Geometrical Theory of Diffraction for Modeling Acoustics in Virtual Environments

Category: research

Abstract

Realistic modeling of reverberant sound is crucial in 3D virtual worlds, providing users important cues for localizing sound sources and understanding the geometric and material properties of the environment. Unfortunately, current geometric acoustic modeling systems do not accurately simulate reverberant sound. Instead, they model only direct transmission and specular reflection, while diffraction is either ignored or modeled through statistical approximation. However, diffraction is extremely important for correct interpretation of acoustic environments, especially where the direct path between sound source and receiver is often occluded.

The Geometrical Theory of Diffraction (GTD) extends geometrical optics with diffraction phenomena. When a ray hits a wedge, it produces diffracted rays. Thus, wedges become secondary sources of diffracted rays that in turn may be reflected and diffracted. In this paper, we extend earlier beam-tracing methods to include diffraction, using the GTD for deriving secondary diffracted rays and associated diffraction coefficients. Our main contributions are: 1) a beam tracing algorithm that uses adjacencies stored in a spatial subdivision to find diffracting edges and trace diffraction beams; 2) a method for computing reverberation paths, incorporating specular reflection, transmission, and diffraction; and 3) a real-time auralization system demonstrating that diffraction improves dramatically the audio quality in virtual environments.

Keywords: Spatialized Sound, Sound Visualization, Virtual Environments, Geometrical Theory of Diffraction, Beam Tracing.

1 Introduction

Realistic modeling of the acoustic environment is of crucial importance when a user is immersed in a 3D virtual world. Reverberant sound provides important cues of its dimensions and material properties. In a large cathedral we expect significant reverberation, while in a small room with lush carpeting and heavy drapes we expect a "dry" sound. Furthermore, we depend on accurate spatialization to locate the direction and position of a sound source and to get a sense of immersion [4]. Unfortunately, current geometrical acoustic modeling techniques fail to accurately simulate sound in complex virtual environments because they do not account for diffraction effects except by crude approximations.

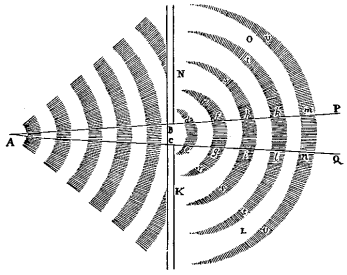

Figure 1: Sketch in Newton's *Principia* (1686) illustrating the concept of diffraction. The wave propagating through the aperture BC is going to "illuminate" regions lying in the geometrical shadow regions from the source A (above line BP and below line CQ).

Diffraction is a form of scattering by objects whose size is of the same order of magnitude as the wavelength of the wave phenomena. Usually neglected for lighting simulation, diffraction is a fundamental aspect of sound propagation, because the wavelength of audible sound is in the range from 0.02 to 17 meters. Without accurate modeling of diffraction around occluding surfaces, we experience abrupt changes in the simulated sound as we turn a corner and the sound source disappears from the line of sight. By adding diffraction, the transition appears smooth. Furthermore, in densely occluded situations, the first wavefront reaching the listener, fundamental for sound localization, is generally a diffracted wave [4]. Thus, accurate modeling of diffraction is extremely important, especially in building interiors and urban environments where the direct path between a sound source and the receiver is often occluded.

Historically, diffraction has been modeled separately from reflection since it was not described by the classical ray theory of wave propagation [5]. It was only in 1962 that Keller proposed the Geometrical Theory of Diffraction (GTD) [20], which extends geometrical optics to include diffraction phenomena. It treats wedges (an edge and its two adjacent surfaces) of the model as secondary sources of diffracted rays, that can propagate into the shadow regions (Figure 1), and that in turn can be reflected and diffracted before reaching the receiver.

We propose an efficient way for generating diffraction paths produced by wedges in complex polyhedral models. There are three major contributions in the paper. The first is an extension to the beam tracing algorithm introduced in [12, 13], using adjacencies stored in a spatial subdivision to quickly find diffracting edges and trace diffraction beams originating from these edges. Second, we derive a method for computing reverberation paths, incorporating specular reflection, transmission, and diffraction according to the GTD. Finally, we describe a real-time auralization system that produces realistic sound during interactive walkthroughs in complex environments.

We demonstrate that 1) beam tracing is an efficient way to find diffraction sequences in densely occluded environments for moderate number of diffraction events, 2) it is possible to construct corresponding propagation paths in real-time, and 3) diffraction improves dramatically the audio quality in virtual environments. To our knowledge, this is the first time the GTD is used for sound rendering and the first time it is used in combination with specular reflections to produce sound at interactive rates in a complex reverberant virtual environment.

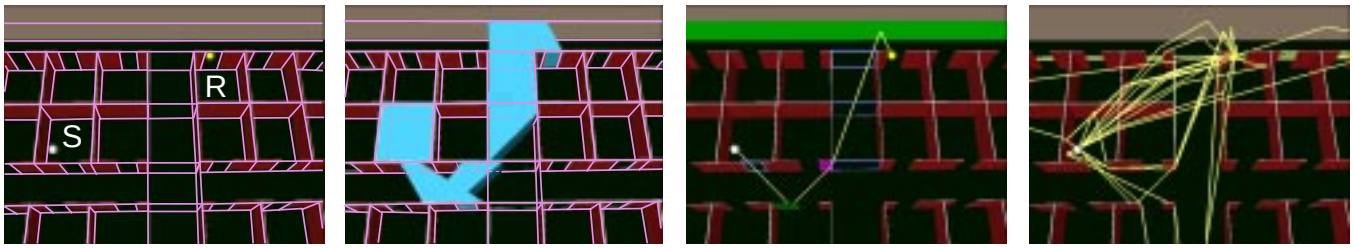

(a) (b) (c) (d)

Figure 2: Overview of our approach: (a) Virtual environment with source S, receiver R, and spatial subdivision marked in pink. (b) Sample reflected and diffracted beam (cyan) containing the receiver. (c) Path generated for the corresponding sequence of faces (green), portals (purple), wedges (magenta). (d) The procedure repeated for all beams containing R.

2 Previous work

There are three major theories for solving diffraction problems: the Huygens-Fresnel diffraction theory, boundary integral representations using the Helmoltz-Kirchoff integral theorem, and the Geometrical Theory of Diffraction.

According to Huygens' construction, every point on a wavefront can be regarded as a source of a secondary disturbance (a spherical wavelet). The wavefront is at any instant defined by the envelope of these wavelets. Fresnel supplemented this principle by adding interference between the wavelets to treat diffraction phenomena [5]. He also defined a subdivision of space between the source and the receiver into concentric ellipsoids with frequency-dependent radii: the Fresnel ellipsoids. By modeling diffraction effects as a loss in signal intensity, Bertoni [3] and Tsingos and Gascuel [38] used Fresnel ellipsoids to determine relevant obstacles at any given frequency. By replacing the binary geometrical visibility by an extended visibility term between 0 and 1, they achieve frequency-dependent sound "muffling". This removes abrupt cuts in simulated sound, producing more pleasing sound. However, such simulations are very crude approximations, limited to direct sound attenuation.

The Helmoltz-Kirchoff integral theorem is a formalization of the Huygens-Fresnel principle [5, 9]. It expresses the scattered field at any point in space as a function of the field on the surface of the diffracting objects. Mathematically, it can be expressed as a surface integral and solved by numerical methods such as Boundary Element Methods (BEM) [16, 18] that discretize surfaces into patches. BEM allow for very accurate treatment of sound diffraction but are far too compute intensive for interactive sound rendering. Moreover, the extent of the patches must be less than the wavelength to obtain correct interference effects. In some cases the integral can be solved analytically [35], such as for height fields or periodic surfaces, assuming the listener remains far (compared to the wavelength) from diffracting surfaces. However, in the case of virtual environments, neither of these assumptions is usually valid.

The Geometrical Theory of Diffraction (GTD) incorporates diffraction effects into the ray theory of light [20]. The wedges of the model act as secondary sources and generate new diffracted rays. Each diffracted ray is attenuated by a diffraction coefficient in the same way a reflected ray is attenuated by a reflection coefficient. As for reflected rays, diffracted rays follow Fermat's principle: if the propagation medium is homogeneous, the rays follow the shortest path from the source to the receiver, stabbing the diffracting edges. Thus, incident and diffracted rays make equal angle with the edge direction.

GTD has been applied successfully in acoustics for solitary wedges [24, 32, 17]. In particular, Kawai shows that it gives good agreement with measurements for noise barrier assessments [19]. In such applications, a simple geometric model is sufficient and audio rendering is not required. Moreover, most sound barrier simulations are limited to non-reverberant spaces and do not combine effects of reflected and diffracted sound, or do so for only one surface (when ground effects are taken into account). For virtual environments, we are interested in treating complex reverberant spaces for which both reflected and diffracted sound is crucial.

In general polyhedral environments, finding diffracted paths is closely related to finding the shortest 3D Euclidean path avoiding a set of obstacles. Unfortunately, this well-known motion planning problem is NP-hard [6] and has no straightforward solution. For radio wave propagation, Rajkumar *et al.* [34] proposed a ray-tracing technique using the GTD in which diffracted rays are broadcast when a ray "intersects" an edge. This approach is prone to aliasing due to difficulties in detecting intersections between rays and edges, which are both infinitely thin. "Thick" rays must be used to avoid this, leading to other types of sampling errors [27].

Beam tracing is also used for radio wave propagation by [21], who find the ray/edge intersections without sampling errors, but simulate edge diffraction by a set of point sources along the diffracting edge. To reduce spatial aliasing, multiple overlapping beams must be broadcast from the edges. The complexity in managing the large number of diffracted beams precludes this method from real-time applications. However, beam tracing has been shown to be well suited for real-time applications [12, 13], particularly in architectural environments. But, beam tracing needs to be extended to include non-point sources, in particular edge sources, to model diffractions correctly.

3 Overview of our approach

Our approach extends the beam tracing algorithm described in [12, 13] to incorporate diffraction. It is decomposed into four phases (see Figure 2). First, we precompute a spatial subdivision data structure to be used for efficient space traversal. Second, we trace conservative beams inside this structure to find possible propagation sequences from a source to a receiver. Third, from the list of all possible sequences, we compute the valid 3D propagation paths. Finally, we evaluate the contribution of every path to the impulse response and render the audio accordingly.

However, each of the four phases is substantially different from Funkhouser's work: 1) spatial subdivision must be extended to incorporate spatial adjacency information in order to retrieve diffracting wedges efficiently, 2) beam tracing must be extended to handle diffracted beams following the Geometrical Theory of Diffraction (GTD), with an edge span as a single source to avoid broadcasting multiple beams, 3) propagation paths can no longer be constructed with simple methods previously used for specular reflections, but require an optimization process. Finally, 4) for each diffracted path contributing to the impulse response, a diffraction coefficient must be computed each time the path stabs a wedge.

In the following four sections, we further detail each phase.

4 Extended spatial subdivision

The goal of this precomputation phase is to partition space into convex polyhedral cells whose boundaries are aligned with the polygons of the 3D input model and to encode the topological adjacencies of the vertices, edges, faces, and cells in a data structure that enables output-sensitive traversals of 3D space. The key contribution over previous beam tracing methods is that topological information (edge-face adjacencies) is used to construct wedges between adjacent faces. We call the new spatial subdivision a *winged-pair* representation.

The winged-pair data structure stores topological adjacencies in fixed-size records associated with vertices, edges, faces, cells, and face-edge pairs. As such, it is very similar to the well-known winged-edge data structure [2], but it describes topological structures one dimension higher. Our representation is also similar to the facet-edge representation described by Dobkin and Laszlo [8] and to the polyhedral subdivision used by Fortune [10].

Every vertex stores its 3D location and a reference to one attached edge; every edge stores references to its two vertices and one attached face-edge pair; every face stores references to its two cells and one attached face-edge pair; and, every cell stores a reference to one attached face. The face-edge pairs store references to one edge E and to one face F along with adjacency relationships required for topological traversals. Specifically, they store references (*spin*) to the two face-edge pairs reached by spinning F around E clockwise (CW) and counter-clockwise (CCW) (see Figure 3) and to the two face-edge pairs (*clock*) reached by moving around F in clockwise and counter-clockwise directions from E (see Figure 3).

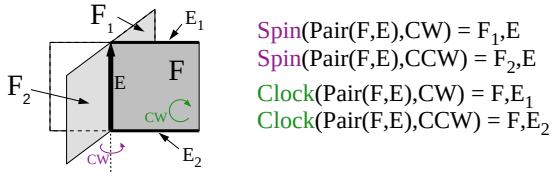

Figure 3: Winged-pair structure.

The winged-pair data structure is crucial for our application. In particular, we find all faces adjacent to a given edge, a critical operation for finding diffracted edges, by following the *spin* references in face-edge pairs.

We build the winged-pair data structure for any 3D model using a Binary Space Partition (BSP) [11] using the method described in [30]. As BSP cells are split by a polygon P, the corresponding winged-pair cells are split along the plane supporting P, and the faces and edges on the boundary of each split cell are updated to maintain a 3-manifold in which every face is convex and entirely inside or outside every input polygon. As faces are created, they are labeled according to whether they are opaque (coincide with an input polygon) or transparent (split free space). The binary splitting process continues until no input polygon intersects the interior of any BSP cell, leading to a set of convex polyhedral cells whose faces are all convex and cumulatively contain all the input polygons.

5 Tracing Diffraction Beams

We trace beams containing potential reverberation paths using an extension to the algorithm by Funkhouser *et al.* [12, 13]. Briefly, we trace beams from each sound source along reverberation paths by traversing the cell-face and face-edge adjacency relationships in the winged-pair structure. The algorithm starts in a cell with an omnidirectional source and a beam representing the entire cell. Then, it visits adjacent cells iteratively, considering different permutations of transmissions, specular reflections and diffractions due to the faces and edges on the boundary of the "current" cell. As each new cell is visited, we update incrementally a polyhedral beam, conservatively including potential reverberation paths along the current traversal sequence.

The novel beam tracing contribution of this paper is a method for tracing beams along paths of diffraction. According to the Geometrical Theory of Diffraction (GTD), an acoustic field, incident upon an edge between two non-coplanar surfaces, forms a diffracted wave that propagates from the intersected part of the edge in a cone-shaped pattern of reverberation paths, as shown in Figure 4.

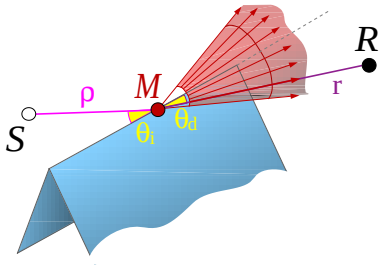

Figure 4: Rays diffracted by a 3D wedge. At oblique incidence, an incoming ray gives rise to a cone of diffracted rays. The angle of the cone θ_d is equal to the angle between the incident ray and the edge θ_i. At normal incidence ($\theta_i = \frac{\pi}{2}$), diffracted rays form a disk of omnidirectional coplanar rays in a plane orthogonal to the edge.

For a given beam B traced through a cell C, we quickly find the diffracting edges as they must be: 1) on the boundary of C, 2) intersected by B, and 3) shared by two faces on the boundary of C that are either non-coplanar or have different acoustic properties (e.g., one is transparent and the other opaque). For each such edge ε, we trace a new diffraction beam B', whose source is $B \cap \varepsilon$ and whose polyhedral extent contains the cone of potential diffraction paths bounded by the solid wedge of opaque surfaces sharing the edge (see Figure 5). This conservatively approximate beam contains all potential paths of sound initially traveling along B and then diffracted by ε. Using the adjacency information in the winged-pair structure, the beams are constructed in constant time.

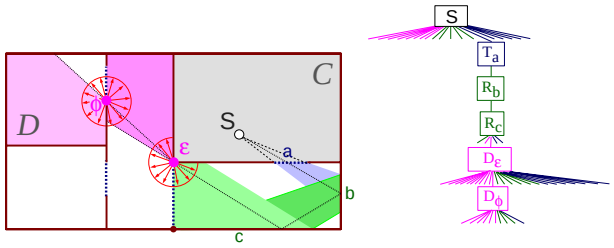

Figure 5: Example beam tree. A beam traced from source S in cell C is transmitted through the transparent portal a, then specularly reflected by surfaces b and c, and finally diffracted twice by edges ε and ϕ reaching cell D. Full cone of diffracted rays is shown for each edge while we trace diffracted beams only in shadow regions.

For efficiency reasons, we trace only a portion of the diffracted beam. Tracing diffracted beams in all directions around the edge would result in a combinatorial explosion of the beam tree. We restrict our diffracted beams to shadow regions (Figure 5) since: 1) we are most interested in providing the acoustic field missing in those regions, and 2) the amplitude of the diffracted sound field is largest in that region and quickly falls off in the illuminated region (see also Figure 7). Even though this approximation results in sound field discontinuities at shadow boundaries, in practice these discontinuities are not perceivable.

As diffraction beams are traced along paths of reflections and transmissions for a sequence of convex polygons, their interior volume can become quite complex, bounded by quadric surfaces due to triple-edge (EEE) events [37]. Rather than modeling these complex volumes exactly, we over-estimate the potential space of paths from each diffracting edge with a bowtie-shaped polytope, as in [36]. We compensate for this approximation later by validating each sequence of reflections and diffractions when constructing explicit reverberation paths.

All beams are traced in priority order, either during an off-line precomputation (as in [12]) or in real-time using multiple asynchronous processes (as in [13]). They are stored in a *beam tree* data structure [15, 12] augmented to contain nodes representing diffraction. An example of a partial beam tree is shown in Figure 5.

6 Construction of diffracted paths

In a third phase, in which the user interactively navigates a simulated observer (receiver) through a virtual environment, *reverberation sequences* to the moving receiver point R are generated via lookup in the beam trees. First, we find the cell containing the receiver point by logarithmic-time search of the BSP spatial subdivision. Then, we check each beam tree node, T, associated with that cell for whether the beam stored with T contains the receiver point. If it does, we have a potential reverberation sequence from the source to the receiver, and the ancestors of T in the beam tree explicitly encode the set of reflections, transmissions, and diffractions through the boundaries of the spatial subdivision. From this sequence of surfaces, wedges and portals, we now construct the corresponding propagation path.

Since we are only considering diffraction and specular reflection and the source and receiver are points, only one propagation path exists for a particular sequence of faces and wedges. This is a result of the angle constraints associated with specular reflection and diffraction. Moreover, according to the GTD, this path is the shortest among all possible paths stabbing the surfaces and edges in the sequence. The difficulty in constructing this path is to find the reflecting points on the surfaces and diffracting points on the edges for the given position of the source and receiver.

For sequences of specular reflection only, the construction of the reflected points is trivial and can be achieved by successively mirroring the source relative to the reflecting surfaces [1].

For sequences of diffraction only, the problem is more difficult and diffraction points on edges cannot be easily determined. One approach may be to successively intersect diffraction cones with the next edge in the sequence. But, assuming we obtain two intersections each time, this could lead to as many as $(2^n - 1)$ tests for a sequence of n wedges. Moreover, deciding if the receiving point lies on a diffraction cone is prone to aliasing. For a hybrid sequence of reflections and diffractions, the problem is even more complicated since intersection of diffraction cones and surfaces results in an infinite number of points.

Hence, it is more efficient to build the path by solving a non-linear system expressing angle constraints between the edges and the incident/diffracted directions [7]. For a single wedge, the diffraction point M on the edge satisfies the equation (see insert in Figure 6):

$$\overrightarrow{MS} \cdot \overrightarrow{E} = \overrightarrow{MR} \cdot (-\overrightarrow{E}),$$

where $\cdot$ denotes the dot product, S is the source point, R the receiving point and $\overrightarrow{E}$ the edge direction vector (all vectors are *unit vectors*). $\overrightarrow{E}$ is oriented such that $\overrightarrow{E} \times \overrightarrow{W} = \overrightarrow{N}$, where $\overrightarrow{W}$ is a vector that lies in the plane of one of the two polygons (arbitrarily chosen) of the wedge, normal to the edge and directed away from the edge. $\overrightarrow{N}$ is the unit normal of the chosen wedge surface (see insert in Figure 8).

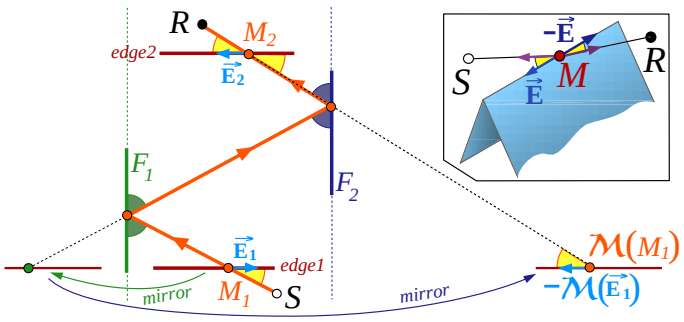

Figure 6: Construction of hybrid diffracted and specular reflected paths (top-view). We express the angle constraints (marked in yellow) using mirror-images of the edges. The insert shows the case of a single diffraction.

For a sequence of n wedges, a system of n equations is obtained:

$$\begin{cases} \overrightarrow{M_1S} \cdot \overrightarrow{E_1} & = & \overrightarrow{M_1M_2} \cdot (-\overrightarrow{E_1}) \\ \overrightarrow{M_2M_1} \cdot \overrightarrow{E_2} & = & \overrightarrow{M_2M_3} \cdot (-\overrightarrow{E_2}) \\ & \vdots & \\ \overrightarrow{M_nM_{n-1}} \cdot \overrightarrow{E_n} & = & \overrightarrow{M_nR} \cdot (-\overrightarrow{E_n}) \end{cases} \quad (1)$$

Parameterizing the edges, $M_i = P_i + t_i\overrightarrow{E_i}$ (where P_i is a reference point on edge i, for example, the midpoint), system (1) can be rewritten in terms of n unknowns t_i and solved within a specified tolerance using a non-linear system solving scheme. In our implementation we used Broyden's scheme, which has a complexity $O(n^2)$ in the number of unknowns [33].

In the case of a hybrid specularly reflected and diffracted path, we have to include specular reflection constraints. If we add a sequence of k mirror surfaces F_i between two consecutive edges, we can use successive mirror images of the first edge to rewrite the previous system as:

$$\begin{cases} \overrightarrow{M_1S} \cdot \overrightarrow{E_1} & = & \overrightarrow{\mathcal{M}(M_1)M_2} \cdot (-\overrightarrow{\mathcal{M}(E_1)}) \\ \overrightarrow{M_2\mathcal{M}(M_1)} \cdot \overrightarrow{E_2} & = & \overrightarrow{M_2M_3} \cdot (-\overrightarrow{E_2}) \\ & \vdots & \end{cases} \quad (2)$$

where the operator $\mathcal{M}$ expresses mirroring according to the k successive mirrors F_i in front order from F_1 to F_k (see Figure 6).

We solve the system assuming all edges to be infinite. Then we perform two validity checks. First we ensure that all diffraction points lie inside the finite edge spans, just as it is usually done for specular reflection. If a diffraction point lies outside an edge span, we declare the path invalid. Next we check for visibility, in particular that paths go through the transparent portals (see Figure 2(c)), since the beams are overestimating.

7 Auralization of diffracted paths

In order to render the virtual sound field, we compute a digital filter, the *impulse response* of the environment [22, 28]. The impulse response reproduces correct reverberation effects when convolved with "dry" (anechoic) input signals. Whenever we find a valid path between the source and the receiver, we compute its contribution to the impulse response. Each path makes a contribution to the final

signal, with a delay proportional to the path length and an amplitude proportional to the combined attenuation of the diffraction and reflection coefficients of the encountered wedges and surfaces.

In this section we focus on the evaluation of diffraction coefficients and the diffracted acoustic field. For acoustic reflection coefficients and the reflected field, we refer the reader to classic texts such as [25, 31] and previous work on real-time auralization [12].

GTD provides an analytic expression of the diffraction coefficient only for infinite wedges. Keller's original expression suffers from severe singularities making it useless for simulation purposes (see Figure 7). Koujoumjan and Pathak solved this problem by introducing the Uniform GTD [23], and a new expression for the diffraction coefficient. In our implementation, we use Uniform GTD expression described in Appendix A.

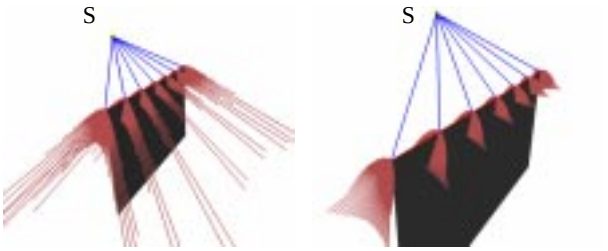

Figure 7: Visualization of diffraction coefficients, calculated by GTD and Uniform GTD, for a "flat-wedge" diffraction from a point source. The cones of diffracted rays (red) are represented for several incidence directions (blue), with segments of length equal to the diffraction coefficient modulus. Note the singularities in Keller's expression (left) which approaches infinity near shadow/reflection boundaries. The Uniform GTD expression (right) is well-defined everywhere.

For a monochromatic harmonic wave in complex notation, the acoustic pressure field diffracted by a wedge can be expressed in terms of the incident field on the edge, $\mathcal{E}_{\text{incident}}(M)$, as:

$$\mathcal{E}_{\text{diffracted}}(R) \quad = \mathcal{E}_{\text{incident}}(M)\ D\ A(r,\rho)\ e^{-ikr}, \qquad (3)$$

where R is the receiver location, M is the diffraction point on the edge, $k = 2\pi/\lambda$ is the wave number, λ is the wavelength, $A = \sqrt{\rho r/(\rho + r)}$ is a scalar distance attenuation term along the diffracted path, and D is the diffraction coefficient, calculated according to the Uniform GTD. The complex exponential represents phase variation along the propagation path (see Figure 8 for other notation). The diffraction coefficient, D, is a complex number accounting for amplitude and phase changes due to diffraction and depends on both the incident and diffracted direction on the edge and the corresponding angles α_i and α_d (Figure 8).

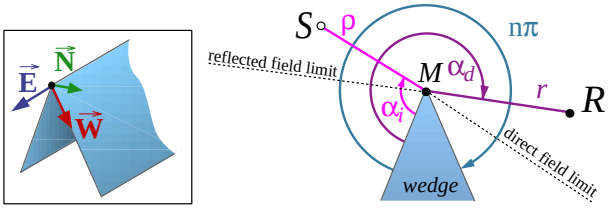

Figure 8: Notations and edge-fixed coordinate system definition for the Uniform GTD wedge diffraction coefficient. The choice of the wedge surface used as reference is arbitrary and will not change the value of the diffraction coefficient.

Equation (3) is applied successively for every diffracting wedge along the path, the diffracted field of one wedge becoming the incident field on the next wedge or on a reflecting surface.

To achieve interactive rates, we implemented a prototype frequency independent system using scalar sound amplitude rather than complex sound pressure. We use a scalar version of Equation (3):

$$\mathcal{A}_{\text{diffracted}} = \mathcal{A}_{\text{incident}}\ \hat{D}\ A(\rho, r),$$

where $\mathcal{A}_{\text{diffracted}}$ and $\mathcal{A}_{\text{incident}}$ are scalar amplitudes, and $\hat{D}$ is a "diffractivity" coefficient, computed as the modulus of the diffraction coefficient at a representative frequency (e.g., 1 kHz). We could also use an average of the coefficient over the audible frequency spectrum.

8 Experimental results

The 3D data structures and algorithms described in the preceding sections have been implemented in C++ and run both on Silicon Graphics and PC/Windows computers. We have integrated them into a prototype system that allows a user to move through a virtual environment interactively, while images and spatialized audio are rendered in real-time according to the user's simulated viewpoint.

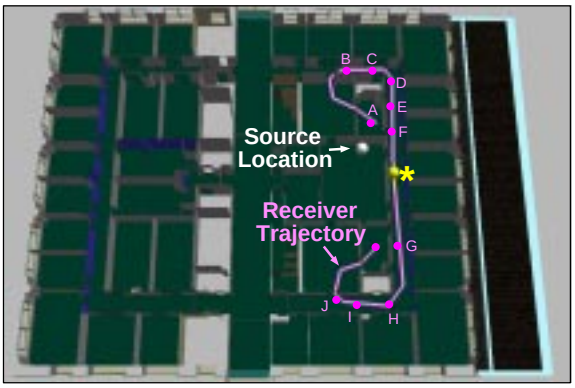

Figure 9: Experimental setup.

To test if our beam tracing approach is practical for modeling diffraction in typical virtual environments, and to evaluate the benefits of incorporating diffraction into real-time auralization, we ran a series of tests computing reverberation paths both with and without diffraction. During each test, we used the experimental setup in Figure 9. The virtual environment is a 3D model with 1,762 polygons representing one floor of a building (Soda Hall). For simplicity, we assume that every polygon in the 3D model is 80% reflective and acoustically opaque (no transmission). Before each test, we traced 100,000 beams in breadth-first order from a stationary sound source (located at the white dot in Figure 9) and stored them in a beam tree data structure (see Section 5). Then, as a receiver moves at three inch increments along the trajectory (the pink line in Figure 9), we compute reverberation paths from source to receiver and update an impulse response (see Sections 6 and 7). All the tests were run on a Silicon Graphics Onyx2 workstation using one 195MHz R10000 processor.

This test sequence was executed three times, once for each of the following beam tracing constraints:

1. **Specular reflection only:** We traced 100,000 beams along paths of specular reflection, with no diffraction. The results represent the state-of-the-art prior to this paper [12, 13].
2. **Diffraction only:** We traced 100,000 beams along paths of diffraction (around silhouette edges into shadow regions), with no specular reflections.
3. **Both specular reflection and diffraction:** We traced 100,000 beams along paths representing arbitrary permutations of specular reflection and diffraction.

Figure 10 shows plots with the number of reverberation paths found between source and receiver (the top plot) and the log power[1] of resulting impulse responses (the bottom plot), for each receiver location along the test trajectory. Note that since traversal is breadth-first with a fixed number of beams, the paths obtained with both diffraction and reflection are not the sum of the paths obtained by reflection and diffraction alone (Figure 10(a)).

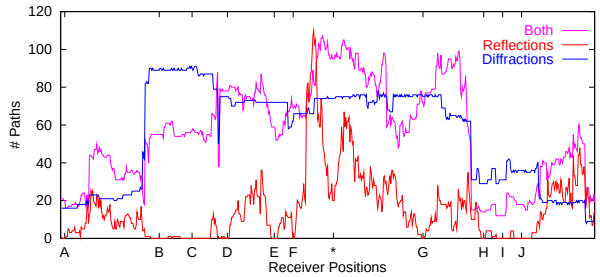

(a) Number of reverberation paths.

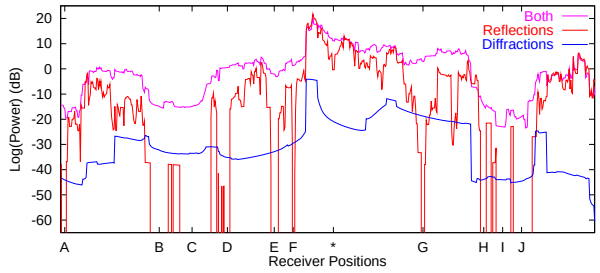

(b) Log power of computed impulse response.

Figure 10: Plots showing the number of reverberation paths found (top) and the log power of the corresponding impulse responses (bottom) for every receiver position along the test trajectory for specular reflection only (red), diffraction only (blue), and both specular reflection and diffraction (purple). Note how adding diffraction gives a smoothly varying sound level as the receiver moves.

From these plots, we confirm that specular reflection alone is not adequate to produce realistic spatialized sound in typical virtual environments. The red curves in Figure 10 show that the number of reverberation paths and the power in the corresponding impulse responses varied dramatically with small changes in receiver location. The cause of this effect is easily understood by examining the top-left image of Figure 11(a), which shows the set of beams traced from the source location during the test with only specular reflection. Note the pattern of thin beams crossing the hallways. As the receiver walks down the hallway along the test trajectory, s/he moves in and out of distinct reflection beams, leading to sharp discontinuities in the computed reverberations. Even worse, there are several locations along the receiver trajectory with no specular reflection paths, and thus the power of auralized sound drops to zero for short periods. These locations are marked with capital letters on the horizontal axis of the plots in Figure 10 and along the pink trajectory line in Figure 9, and they correspond to the visible "holes" in the beam coverage in Figure 11(a). These "dead-zones" are particularly troublesome to users during a walkthrough, as they clearly fail to match real-world acoustical experiences.

In contrast, we note that tracing beams only along paths of diffraction leads to smoothly varying reverberations (the blue

[1] log power is computed as $10 \cdot \log \sum_{i=1}^{n} a_i^2$ where n is the number of reverberation paths and a_i the amplitude along the ith path.

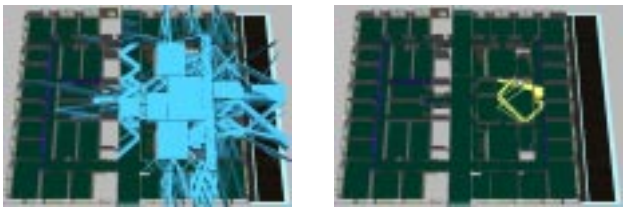

(a) Specular reflection only.

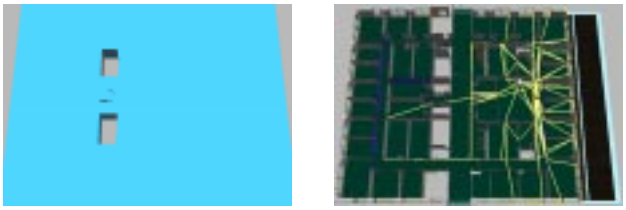

(b) Diffraction only.

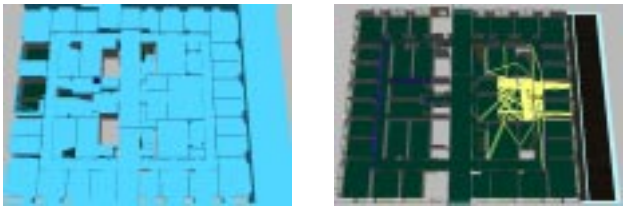

(c) Both specular reflection and diffraction.

Figure 11: Images of 100,000 beams traced in breadth-first order (cyan) and reverberation paths (light yellow) from source (white sphere) to receiver location `*' (dark yellow sphere). Note how diffracted beams fill the entire space uniformly in contrast with specular beams.

curves in Figure 10). The reason can be seen in Figure 11(b). Diffraction beams tend to cover larger volumes of space than specular beams; and, in our test, they cumulatively cover all reachable parts of the 3D environment (the light gray regions in the middle of the model correspond to elevator and wiring shafts unreachable by sound). Accordingly, the computed diffraction paths are very "stable", continuously tracking the receiver as s/he moves down the hallway. In addition, diffraction paths often contain the shortest reverberation path, critical for source localization. Moreover, diffraction coefficients computed with the GTD vary smoothly as the receiver moves, and thus we encounter no "drop-outs" or sudden discontinuities as the receiver moves during the second test.

Interestingly, although many diffraction paths reach the receiver during the second test, the most direct path tends to contribute almost all the reverberant power. For instance, consider the receiver location `*' highlighted in yellow in Figures 9 and 10. The receiver has just passed by an open doorway near the source and has moved behind an opaque wall. At this point, there is no direct sound path from the source. But, there is a very short diffraction path, which bends around an edge of a nearby door frame. That path appears as a solitary blue spike in Figure 12. The other 73 diffraction paths, which are shown in Figure 11(b), arrive much later and their contribution is negligible since they are very long paths.

Looking at the results of the third test (the purple curves in Figure 10), we see that the power varied quite smoothly between successive receiver locations, as one would expect. The nice features of combining specular reflections and diffractions can be seen in Figure 12(c). In that case, the impulse response contains not only the shortest (diffracted) path from the source (the left-most spike), but it also has many high-power early reverberations not found in the other tests because they are reflections of previously diffracted

waves, or diffractions of previously reflected waves. Those paths are going to provide a quasi-direct sound in regions where the actual direct path is occluded.

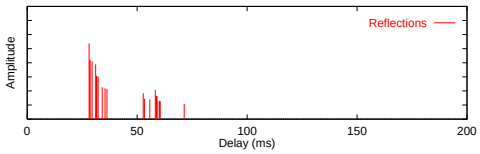

(a) Specular reflections only.

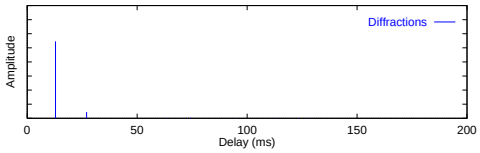

(b) Diffractions only.

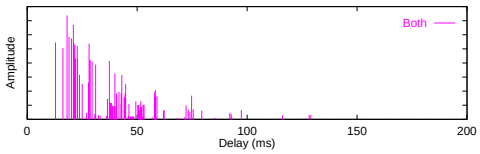

(c) Both specular reflections and diffractions.

Figure 12: Impulse responses corresponding to the reverberation paths computed in each test for the receiver location `*'. Note how diffraction introduces new contributions in the beginning of the response (c), that are going to provide a quasi-direct contribution while the actual direct path is occluded.

We also gathered computational statistics during the three tests (Table 1). The first two columns contain compute rates. Specifically, Column 2 shows the rate (in beams/second) at which beams are traced from the stationary source location in each test. Column 3 shows the rate (in paths/second) at which reverberation sequences are computed and processed to form reverberation paths to the moving receiver location, including calculation of reflection and diffraction coefficients. The next three columns show the average number of transmissions through transparent cell boundaries (Trans), specular reflections (Refl), and diffractions (Diff) along the computed paths in Columns 4-6, respectively. Finally, the right-most column (Update Time(s)) shows the time (in seconds) required to update the impulse response for each new receiver location in each test.

Test Name	Compute Rates		Path Statistics			Update Time (s)
	Beams/s	Paths/s	Trans	Refl	Diff	
Specular	5,028	12,020	5.2	4.6	0.0	0.012
Diffract	3,882	908	10.5	0.0	5.6	0.081
Both	4,500	3,311	5.6	2.0	1.6	0.059

Table 1: Beam tracing and path generation statistics.

We conclude that tracing paths of diffraction is quite practical for interactive virtual environment applications. Although diffraction increases the time required to trace beams and construct reverberation paths, the system can still update impulse responses at interactive time steps (e.g., every 59ms) with our method. More to the point, listening to sound auralized with impulse responses incorporating diffraction, we find that our simple frequency independent prototype renderer provides a surprisingly good impression of the environment. Please review the examples on the enclosed video.

Although space limitations preclude presentation of more detailed quantitative results in this paper, we note that the combinatorial growth of beam trees with both specular reflection and diffraction is very large. In these experiments, tracing 100,000 beams in breadth-first order we are able to model all permutations containing up to 2 specular reflections and 2 diffractions while using around 100MB of memory. However, for instance, tracing all permutations with up to 4 specular reflections and 4 diffractions exceeds the physical memory of our machine (2GB). For this reason, we usually rely upon priority-driven beam tracing methods to prune constructed beam trees to contain only the most significant reverberation paths [13].

9 Discussion

Our approach is limited to polyhedral environments and polygonal wedges. GTD can be extended to treat curved wedges, but finding diffracted paths becomes more difficult since they can contain pieces of geodesic curves on the surface of the curved wedge.

Our current implementation is limited to diffraction by exterior corners (silhouette edges). A more accurate simulation should also treat interior corners but doing so would result in a combinatorial explosion of the number of beams. Besides, diffraction by interior wedges will produce strong exaggerated echoes only for high frequencies and a listener close to the wedge (see [31] p. 500 for a quantitative discussion on this subject).

One could argue that diffuse reflection could produce a smoothing of the sound field by scattering rays in all directions. Although our beam tracing technique can be easily extended to trace diffuse reflection beams (i.e., beams reflected into an entire half-space), it remains disputable if diffuse reflection is a proper model for acoustics. Even if waves are scattered in several directions from a surface, it is very unlikely for the scattering to be isotropic. That explains the difficulty to find diffuse acoustic reflection coefficients for surfaces [26]. Moreover, it is unclear how to correctly auralize diffuse reflection.

Finally, our real-time auralization system is still frequency independent and limited in audio resolution (11 kHz). However, hardware that is now becoming widely available on PCs, makes it possible to perform real-time frequency dependent convolution using different coefficients for several frequency bands at compact disc quality.

Apart from the virtual acoustics applications, we believe that GTD can be successfully applied to computer graphics problems. Modeling of reflectance functions (BRDFs) [39, 14, 35] from surfaces acting as diffraction gratings could be treated using this theory. Its major advantage is that, because it is based on the ray-theory of light, it can be integrated in a ray-tracing framework, removing limitations of previous related work [35].

10 Conclusion

In this paper we introduce a novel, efficient technique for incorporating diffraction effects in interactive audio simulations for virtual environments. Relying upon the Geometrical Theory of Diffraction (GTD), we extend a beam-tracing approach to construct combined diffracted and reflected paths. We demonstrate that it is possible to construct these paths in real-time and that diffraction dramatically improves the realism and quality of the audio experience.

This is the first instance where: 1) GTD is used to produce sound at interactive rates in a complex environment, 2) GTD is used for auralization, and 3) realistic diffraction is introduced in virtual environments. By simulating diffraction, we have removed the disturbing ”cuts” in the audio that occur when a sound source is occluded by an acoustically opaque surface, and made it possible to localize occluded sound sources.

Based on our initial experiences with this system, diffraction is essential for realistic acoustical modeling, and we feel this work

makes a significant step towards realistic acoustics in interactive virtual environments. We are continuing this research, aiming towards increased sense of presence in virtual environments by investigating the perceptual trade-offs between visual and acoustical effects. We are also investigating the GTD for computing BRDFs for surfaces acting as diffraction gratings.

A Computing diffraction coefficients

The Uniform GTD [23, 29] expresses the acoustic diffraction coefficient for an infinite wedge as:

$$\begin{aligned} D(n, k, \rho, r, \theta_i, \alpha_i, \alpha_d) = & -\frac{e^{-i\frac{\pi}{4}}}{2n\sqrt{2k\pi}\sin\theta_i} \\ & \Bigg[\tan^{-1}\left(\frac{\pi+(\alpha_d-\alpha_i)}{2n}\right) F\left(kLa^+(\alpha_d-\alpha_i)\right) \\ & + \tan^{-1}\left(\frac{\pi-(\alpha_d-\alpha_i)}{2n}\right) F\left(kLa^-(\alpha_d-\alpha_i)\right) \\ & + \Bigg\{ \tan^{-1}\left(\frac{\pi+(\alpha_d+\alpha_i)}{2n}\right) F\left(kLa^+(\alpha_d+\alpha_i)\right) \\ & + \tan^{-1}\left(\frac{\pi-(\alpha_d+\alpha_i)}{2n}\right) F\left(kLa^-(\alpha_d+\alpha_i)\right) \Bigg\} \Bigg], \end{aligned} \tag{4}$$

where (see also Figures 4 and 8)

- k is the wave number. $k = 2\pi/\lambda$,
- n is such that the exterior wedge angle is $n\pi$,
- ρ is the source to diffraction point distance,
- r is the receiver to diffraction point distance,
- θ_i is the azimuthal angle (in the plane of the edge) between the edge vector $\overrightarrow{E}$ and the incident direction. $\theta_i \in [0, \pi]$,
- α_i is the elevation angle between vector $\overrightarrow{W}$ and the incident direction. $\alpha_i \in [0, n\pi]$,
- α_d is the elevation angle between vector $\overrightarrow{W}$ and the diffracted direction. $\alpha_d \in [0, n\pi]$,

and where:

$$F(X) = 2i\sqrt{X}e^{iX}\int_{\sqrt{X}}^{+\infty} e^{-i\tau^2}\,d\tau, \tag{5}$$

$$L = \frac{\rho r}{\rho + r}\sin^2\theta_i, \tag{6}$$

$$a^{\pm}(\beta) = 2\cos^2\left(\frac{2\pi nN^{\pm}-\beta}{2}\right), \tag{7}$$

$N^{\pm}$ is the integer that satisfies more closely the relations:

$$2\pi nN^+ - \beta = \pi \quad \text{and} \quad 2\pi nN^- - \beta = -\pi \tag{8}$$

Several approximations exist in the related literature, useful for implementation of Eq. 4. In particular, relations (8) reduce to:

$$N^+ = \begin{cases} 0 & \text{for } \beta \le \pi(n-1) \\ 1 & \text{for } \beta > \pi(n-1) \end{cases}, \qquad N^- = \begin{cases} -1 & \text{for } \beta < \pi(1-n) \\ 0 & \text{for } \pi(1-n) \le \beta \le \pi(1+n) \\ 1 & \text{for } \beta > \pi(1+n) \end{cases}, \tag{9}$$

and Kawai [19] gives an approximate rational expression for the integral in Eq. (5):

$$\begin{aligned} &\text{for } X < 0.8: \; F(X) = \sqrt{\pi X}\left(1 - \frac{\sqrt{X}}{0.7\sqrt{X}+1.2}\right)e^{i\frac{\pi}{4}}\sqrt{\frac{X}{X+1.4}} \\ &\text{for } X \ge 0.8: \; F(X) = \left(1 - \frac{0.8}{(X+1.25)^2}\right)e^{i\frac{\pi}{4}}\sqrt{\frac{X}{X+1.4}} \end{aligned} \tag{10}$$

Cotangent terms in Equation (4) are still singular at a reflection or shadow boundary and cannot be evaluated numerically at these boundaries. However, in the vicinity of such a boundary we can express the terms $\alpha_i \pm \alpha_d$ as $\beta = 2\pi nN^{\pm} \mp (\pi - \varepsilon)$. The coefficient is continuous and its value can be computed using [23]:

$$\tan^{-1}\left(\frac{\pi\pm\beta}{2n}\right)F\left(kLa^{\pm}(\beta)\right) \simeq ne^{-i\pi/4}\left(\sqrt{2\pi kL}\,\mathrm{sgn}(\varepsilon) - 2kL\varepsilon e^{-i\pi/4}\right),$$

where $\mathrm{sgn}(\varepsilon) = 1$ if $\varepsilon > 0$ and -1 otherwise.

References

[1] J. Allen and D. Berkley. Image method for efficiently simulating small room acoustics. *J. of the Acoustical Society of America*, 65(4), 1979.

[2] B. Baumgart. Winged edge polyhedron representation. Technical Report AIM-179 (CS-TR-74-320), Computer Science Department, Stanford University, Palo Alto, CA, 1972.

[3] H. Bertoni, Ed. Coverage prediction for mobile radio systems operating in the 800/900 MHz frequency range. *IEEE Trans. on Vehicular Technology (Special Issue on Mobile Radio Propagation)*, 37(1), 1988.

[4] J. Blauert. *Spatial Hearing : The Psychophysics of Human Sound Localization*. M.I.T. Press, Cambridge, MA, 1983.

[5] M. Born and E. Wolf. *Principles of Optics*. 7th edition, Pergamon Press, 1999.

[6] J. Canny and J. Reif. New lower bound techniques for robot motion planning problems. *Proc. 28th IEEE Symposium on Foundations of Computer Science*, p. 49-60, 1987.

[7] J. Choi, J. Sellen, and C.-K. Yap. Precision-sensitive shortest path in 3-space. *Proc. 11th ACM Symposium on Computational Geometry*, p. 350-359, 1995.

[8] D. Dobkin and M. Laszlo. Primitives for the manipulation of three-dimensional subdivisions. *Algorithmica*, 4(1):3-32, 1989.

[9] P. Filippi, D. Habault, J. Lefevre, and A. Bergassoli. *Acoustics, basic physics, theory and methods*. Academic Press, 1999.

[10] S. Fortune. Topological beam tracing. In *Proc. 15th Symposium on Computational Geometry*, p. 59-68, 1999.

[11] H. Fuchs, Z. Kedem, and B. Naylor. On visible surface generation by a priori tree structures. *Computer Graphics (SIGGRAPH 80)*, 14(3):124-133.

[12] T. Funkhouser, I. Carlbom, G. Elko, G. Pingali, M. Sondhi, and J. West. A beam tracing approach to acoustic modeling for interactive virtual environments. *Computer Graphics, (SIGGRAPH 98)*, p. 21-32.

[13] T. Funkhouser, P. Min, and I. Carlbom. Real-time acoustic modeling for distributed virtual environments. *Computer Graphics, (SIGGRAPH 99)*, p. 365-374.

[14] J. Gondek, G. Meyer, and J. Newman. Wavelength dependent reflectance functions. *Computer Graphics (SIGGRAPH 94)*, p. 213-220.

[15] P. Heckbert and P. Hanrahan. Beam tracing polygonal objects. *Computer Graphics (SIGGRAPH 84)*, 18(3):119-127.

[16] D. Hothersall, S. Chandler-Wilde, and M. Hajmirzae. Efficiency of single noise barriers. *J. of Sound and Vibration*, 146(2):303-322, 1991.

[17] T. Isei, T. Embleton, and J. Piercy. Noise reduction by barriers on finite impedance ground. *J. of the Acoustical Society of America*, 67(1):46-58, 1980.

[18] P. Jean. A variational approach for the study of outdoor sound propagation and application to railway noise. *J. of Sound and Vibration*, 212(2):275-294, 1998.

[19] T. Kawai. Sound diffraction by a many sided barrier or pillar. *J. of Sound and Vibration*, 79(2):229-242, 1981.

[20] J. Keller. Geometrical theory of diffraction. *J. of the Optical Society of America*, 52(2):116-130, 1962.

[21] S. Kim, B. Guarino, T. Willis, V. Erceg, S. Fortune, R. Valenzuela, L. Thomas, J. Ling, and J. Moore. Radio propagation measurements and prediction using thee-dimensional ray tracing in urban environments at 908 MHz and 1.9 GHz. *IEEE Trans. on Vehicular Technology*, 48:931-946, 1999.

[22] M. Kleiner, B. Dalenbäk, and P. Svensson. Auralization - an overview. *J. of the Audio Engineering Society*, 41(11):861-875, 1993.

[23] R. Kouyoumjian and P. Pathak. A uniform geometrical theory of diffraction for an edge in a perfectly conducting surface. *Proc. of IEEE*, 62:1448-1461, 1974.

[24] U. Kurze. Noise reduction by barriers. *J. of the Acoustical Society of America*, 55(3):504-518, 1974.

[25] H. Kuttruff. *Room Acoustics (3rd edition)*. Elseiver Applied Science, 1991.

[26] Y. Lam. A comparison of three diffuse reflection modeling methods used in room acoustics computer models. *J. of the Acoustical Society of America*, 100(4):2181-2192, 1996.

[27] H. Lehnert. Systematic errors of the ray-tracing algorithm. *Applied Acoustics*, 38, 1993.

[28] H. Lehnert and J. Blauert. Principles of binaural room simulation. *Applied Acoustics*, 36:259-291, 1992.

[29] D. McNamara, C. Pistorius, and J. Malherbe. *Introduction to the Uniform Geometrical Theory of Diffraction*. Artech House, 1990.

[30] B. Naylor. Constructing good partitioning trees. *Graphics Interface'93*, p. 181-191.

[31] A. Pierce. *Acoustics. An introduction to its physical principles and applications*. 3rd edition, American Institute of Physics, 1984.

[32] A. Pierce. Diffraction of sound around corners and over wide barriers. *J. of the Acoustical Society of America*, 55(5):941-955, 1974.

[33] W. Press, S. Teukolsky, W. Vetterling, and B. Flannery. *Numerical Recipes in C, 2nd edition*. Cambridge University Press, New York, USA, 1992.

[34] A. Rajkumar, B. Naylor, F. Feisullin, and L. Rogers. Predicting RF coverage in large environments using ray-beam tracing and partitioning tree represented geometry. *Wireless Networks*, 2(2):143-154, 1996.

[35] J. Stam. Diffraction shaders. *Computer Graphics, (SIGGRAPH 99)*, p. 101-110.

[36] S. Teller. Computing the antiumbra cast by an area light source. *Computer Graphics (SIGGRAPH 92)*, 26(2):139-148.

[37] S. Teller. *Visibility Computations in Densly Occuded Polyhedral Environments*. PhD thesis, Computer Science Div., University of California, Berkeley, 1992.

[38] N. Tsingos and J.-D. Gascuel. Soundtracks for computer animation: sound rendering in dynamic environments with occlusions. *Graphics Interface'97*, p. 9-16.

[39] S. Westin, J. Arvo, and K. Torrance. Predicting reflectance functions from complex surfaces. *Computer Graphics (SIGGRAPH 92)*, 26(2):255-264.