

Propagation of Sound in Two-Dimensional Virtual Acoustic Environments

SÉRGIO ALVARES R. DE S. MAFFRA¹, MARCELO GATTASS¹, LUIZ HENRIQUE DE FIGUEIREDO¹

¹TecGraf/PUC-Rio – Rua Marquês de São Vicente, 225, 22.453-900, Rio de Janeiro, RJ, Brazil
{sam,gattass,lhf}@tecgraf.puc-rio.br

Abstract. This paper describes the implementation of a system that simulates the propagation of sound in two-dimensional virtual environments and is also capable of reproducing audio according to this simulation. The simulation, which is a preprocessing stage, consists in creating direct, specular reflection and diffraction sound beams that are used later for the creation of the actual propagation paths, in real-time. As the sound beams are created in a preprocessing stage, the system treats only sound sources with fixed position and moving receivers.

1 Introduction

For a long time, computational simulation of acoustic phenomena has been used mainly in the design and study of the acoustic properties of concert and lecture halls. Recently, however, there is a growing interest in the use of such simulations in virtual environments in order to enhance users' immersion experience. According to [1], the addition of realistic simulation of acoustic phenomena to a virtual reality system can aid in the localization of objects, separation of simultaneous sound events and in the spatial comprehension of the environment.

Generally, we can say that a virtual acoustic environment must be able to accomplish two tasks: simulating the propagation of sound in an environment and reproducing audio with spatial content, that is, in a way that allows its user to recognize the direction of the incoming sound waves.

To simulate the propagation of sound, one can solve the wave equation [2] using finite and boundary element methods [3]. This approach, however, is not suitable for interactive applications due to its high computational cost. An alternative for these expensive methods is the geometric treatment of the propagation of sound, referred to as *geometrical room acoustics* [4]. As Kuttruff describes it, in geometrical room acoustics the concept of a sound wave is replaced by the concept of a sound ray.

The use of sound rays to simulate the propagation of sound in an environment makes the algorithms created for this purpose very similar to the ones used in the analysis of wireless communication networks [5] and in visualization (hidden surface removal), such as *ray tracing* [6] and *beam tracing* [7]. This means that the same techniques used to speed-up visualization applications can also be used in the simulation of sound propagation, as was shown in [1].

In our system we implemented a beam tracer capable of creating beams of specular reflection and diffraction. Section 2 contains more details on how our system simulates the propagation of sound.




Figure 1: Propagation paths and virtual sound sources

The reproduction of the simulated sound field is made by superposing several virtual sound sources located around the user, as Figure 1 illustrates. On Section 3 we discuss how these virtual sound sources are used to render the simulated sound field.

Figure 1 illustrates how a virtual acoustic environment works. First, propagation paths between the sound source and the receiver are found. These are then used to create the virtual sound sources, one for each propagation path between the source and the receiver.

2 Propagation of sound

There are three basic methods that can be used to enumerate propagation paths comprised of specular reflection and diffraction. Namely, the *virtual source method* [4], ray tracing [5, 6] and beam tracing [1].

The virtual source method is basically an exhaustive enumeration technique. It's main problem is computational effort spent in vain in the generation of a large number of

invalid paths that must be identified and discarded. These invalid paths are created due to the lack of visibility information in the method.

Ray tracing has a well known discretization (*aliasing*) problem: no matter how close rays from the same source are created near their origin, as their distance to the source increases, so does the gap between neighboring rays. The existence of gaps between rays creates a discontinuity in the sound field that can lead to audible artifacts.

Beam tracing fixes the problems of the previous methods by dealing with beams, represented by a region of space, instead of individual rays and by using visibility information in the creation of beams, as we show on the next sections. The disadvantage of the beam tracing technique is the complexity of the geometric primitives and data structures necessary for its implementation. It was this complexity that motivated us in implementing a two-dimensional simulation in order to gain more familiarity with the method before moving to the three-dimensional case.

2.1 Beam representation

We begin our brief explanation of the beam tracing method by describing the representation of beams. As we mentioned before, beams are represented by a region of space. This means that a single beam can represent an infinite number of rays, which eliminates the aliasing that occur in ray tracing.

In our implementation, we have used the same representation for beams used in [7], where beams are represented by a local coordinate system (the *beam coordinate system*) and by a cross-section defined in this coordinate system, as Figure 2 illustrates. The figure shows on the left a beam defined in the global coordinate system (axes x and y) with its local coordinate system (axes x' and y'). On the right it shows the same beam (now in its local coordinate system) and its cross-section (defined by the position of a vertical projection plane (x_p) and an interval located on this projection plane [y_{pi}, y_{pf}]).

The cross-section of a beam is responsible for limiting the area it occupies. Notice, however, that beams are actually structures with infinite area, as the cross-section only limits how open beams are and not how far they can reach. That is, rays defined inside the gray area shown in Figure 2 have infinite length.

An essential part of the beam tracing method, as implemented in our system, is the decomposition of the environment into convex cells. This decomposition permits the efficient traversal of the environment and also allows to limit the range of a beam, as each beam must be associated with a single cell of the environment. This association means that operations realized with a beam are valid only inside the cell it is associated with. The next section defines




Figure 2: Representation of beams

the basic operations that are realized on beams.

Figure 3 illustrates the association between beams and convex cells. Notice that two different beams are created when the original beam (a) strikes the boundary of the first cell. Beam c is a *reflection beam*, created due to the intersection of beam a with an opaque portion of the boundary of the cell. The intersection of the original beam with a transparent portion of the boundary originates a *transmission beam* (beam b), that only differs from the original beam in its cross-section.

2.2 Beam operations

There are two basic operations that are frequently made on beams during a beam tracing algorithm: determining whether a beam contains a point in space and determining the intersection of a beam and a segment of the boundary of a convex cell of the environment. Both operations are based on the projection of a vertex in the cross-section of a beam. This projection is made along the ray defined (in the beam coordinate system) by the origin of the beam and the vertex.

Once the projection is made, it is enough to check if the projected vertex lies inside the interval that limits the cross-section to determine if the vertex being tested is located inside the beam.

The intersection of a beam and a segment of the environment is used in the creation of transmission and reflection beams to determine the cross-section of the new beams. When the newly created beams inherit the position of the projection plane from the beam that originated them, the intersection operation can be greatly simplified. In this case, the only information needed for the creation of the new beams is the interval that results from the intersection of two other intervals: the cross-section of the original beam and the interval defined by the projection of the endpoints of the segment in the projection plane of the original beam.

For more detail on the implementation of the basic operations realized with beams, in two and three dimensions, refer to [8].




Figure 3: Beams and their association to cells




Figure 4: Diffraction beams

2.3 Beam tracing

The beam tracing method has two different stages. The first stage, implemented in our system as preprocessing stage, comprises the construction of the *beam-tree* data structure. The beam-tree is the data structure that links all beams originating from the same source, allowing the construction of the actual propagation paths, which is the second stage of the method. Each node of a beam-tree represents a beam that is linked to its parent beam (the one responsible for its creation). Figure 3 illustrates, along with the association of beams and convex cells, the beam-tree created for the beams illustrated in the figure. The next sections discuss the stages of the beam tracing method in more detail and also how the contribution of each propagation path for the simulated sound field is calculated.

2.3.1 Beam-tree construction

As mentioned in the previous sections, whenever a beam intersects a segment of the environment a new beam is created. This creation involves the computation of the new beam's representation (local coordinate system and cross-section), its insertion in the beam-tree and its association to one of the convex cells of the environment.

The segments intersected by the beams can be either opaque (represented in our figures as continuous line segments) or transparent (represented as dashed line segments). Opaque segments represent the walls of the environment that reflect sound waves, while transparent segments are artificial walls, commonly referred to as *portals* [9], that are inserted in the environment to obtain its convex cell decomposition.

When an opaque segment is intersected, a new reflection beam is created. Its coordinate system can be obtained by reflecting the coordinate system on the line supporting the segment. The cross-section of the new beam can be obtained by performing the intersection of the original beam with the segment (as described in section 2.2). Finally, reflection beams are always associated with the same cell associated with its parent beam. In the case of an intersection with a transparent segment, the new beam inherits the coordinate system of its parent and is associated with

a neighboring cell (the one adjacent through the intersected segment). As with opaque segments, the cross-section of the new beam is obtained by the intersection with the parent beam.

There is also another kind of beam we have neglected to mention until now: *diffraction beams*. Diffraction is the scattering of a wave that happens when it strikes a wedge of the environment. Figure 4 illustrates the diffraction of a wave incident to a wedge. Notice how the scattered wave propagates in all directions around the wedge, forming the figure of a cone. As the scattered wave propagates in all directions around the wedge, the number of beams might explode. To avoid this increase in the number of beams we use the same approximation adopted in [10]: diffraction beams are traced only in the region shadow region of the wedge (the region around the wedge that is not illuminated by the incident beam). This approximation is also shown in Figure 4. The justification for using diffraction beams that only cover the shadow region is the high attenuation of the amplitude of the wave that occurs in the diffraction. In the region around the wedge that is illuminated by the incident wave and, occasionally, also by its reflection, the contribution of the scattered wave can be discarded without great losses to the resulting sound field. Notice that in the shadow region, the only contribution to the sound field is the diffracted wave.

In the algorithm, a new diffraction beam must be created whenever a beam intersects a wedge of the environment, which happens when it intersects two consecutive segments, one opaque and the other transparent.

Regarding the construction of beam-trees, there is only one more consideration: the termination criteria for the construction of the tree. The most natural criterium for terminating the expansion of a branch of the beam-tree is an auditive criterium, that is, beams shouldn't be created when the sound becomes inaudible. Limiting the maximum number of beams created and the maximum number of reflections and diffractions in each branch are also commonly used.

2.3.2 Propagation path construction

The second stage of the beam tracing method is the one executed in real-time and is responsible for the creation of

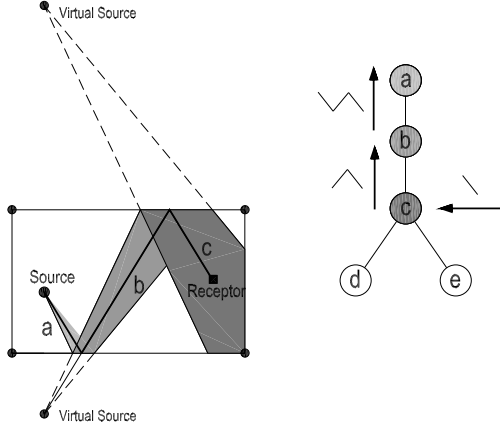


Figure 5: Constructing propagation paths on a beam-tree

the actual propagation paths between the sound source and the receiver. As we mentioned before, each beam stored in the beam-tree contains a reference to its parent beam. Therefore, given any beam b , it is possible to traverse the beam-tree, passing through all ancestor beams of b until the sound source is reached. It is by making this traversal that one can build an actual propagation path between a source and a receiver.

In order to create the propagation paths, the position occupied by the receiver must be determined (which is undetermined during the construction of the beam-tree). Once its position is known, to create the propagation paths the beams that contain the receiver must be identified. This identification can be performed quite efficiently by examining all the beams associated with the convex cell that contains the receiver. Notice that for each beam that contains the receiver, a different propagation path can be built, as for each beam there is a different path on the beam-tree that leads to the source.

The construction of a propagation path is illustrated in Figure 5. The figure shows a rectangular environment with three different beams and the resulting beam-tree. Since the beam c contains the receiver, it is the starting point of the path towards the sound source. Also, note that each node along this path contributes with a segment to the propagation path (the intermediary propagation paths are shown next to the arrows that indicate the path along the beam-tree).

2.3.3 Attenuation and delay

Once the propagation paths have been found, the contribution of each path to the resulting sound field must be calculated. As our main interest is not the rigorous analysis of acoustic phenomena, we have adopted several simplifications in the computation of the contribution of each propaga-

tion path.

As in [1], we disregard phase information when computing the amplitude of the wave reaching the receiver and the phase change due to reflections, that is modelled as a frequency independent constant factor (α in the expression below). Phase changes due to diffraction are also ignored. Given the complexity of the evaluation of more rigorous formulations for diffraction, like the *Uniform Geometrical Theory of Diffraction* [11] or the *Directive Line Source Method* [12], we have adopted an approximation ($\delta(\theta)$, defined in the formulas below) that we believe captures the essence of the effects caused by diffraction, that is the growing attenuation suffered by the amplitude of the scattered wave as it goes deeper into the shadow region around a diffracting wedge. This approximation was obtained through a curve-fitting approach, using diffraction charts present in [10].

The formula below contains the expression used to compute the amplitude of the wave at the receiver. The propagation path modelled in the formula has suffered n_{ref} reflections, n_{diff} diffractions and has length r . P_0 is the initial amplitude of the wave. The diffraction attenuation term of the formula receives as parameter an angle θ that measures how deep into the shadow region the propagation path is.

$$P = P_0 \frac{\alpha^{n_{ref}} \prod_{i=1}^{n_{diff}} \delta(\theta_i)}{r}$$

$$\delta(\theta) = \frac{1}{1 + K\theta^n}$$

$$K = 130$$

$$n = 1.66$$

The time delay associated with a propagation path is given by r/c , where c is the velocity of propagation of sound.

For a more detailed explanation on the calculation of the attenuation and delay suffered by a sound wave, refer to [8].

2.4 Cell partitioning of the environment

As stated previously, the decomposition of the environment into convex cells is an essential part of the beam tracing algorithm. It simplifies the representation of beams, which can be modelled as infinite areas. It also defines an order among the occluders of the environment allowing the implementation of efficient visibility queries and efficient traversal of the environment, which are essential for an efficient creation of beams [8].

The decomposition of an environment into convex cells is usually made using binary space partitions (BSP) [1, 13, 14, 10, 9]. The disadvantage of this technique is the occasional generation of decompositions with a large number of cells. When a large number of cells is created unnecessarily, the large number of portals (or transparent segments) in the decomposition can cause a large increase in the number of beams traced, since whenever a beam crosses a portal, a new transmission beam is created [8].

To avoid this increase in the number of beams traced, we have used a different method, based on one of the techniques in [9] to create cellular decompositions for bidimensional environments. This technique consists in using a Constrained Delaunay Triangulation [15] algorithm to obtain the decomposition. The triangulation, however, also has a large number of cells, not solving the problem of the unnecessary increase in the number of beams. We have avoided this problem by removing edges of the triangulation. Once the triangulation is built, its transparent edges (portals) are sorted in decreasing length order and then removed from the triangulation once it is determined that its removal will not create a concave cell in the decomposition [8].

Figure 6 illustrates the results obtained by the techniques described in the section when applied to the model of a real residential building. As the figure illustrates, the result obtained in this example by the simplification of a triangulation is much superior to the one obtained by the BSP technique.

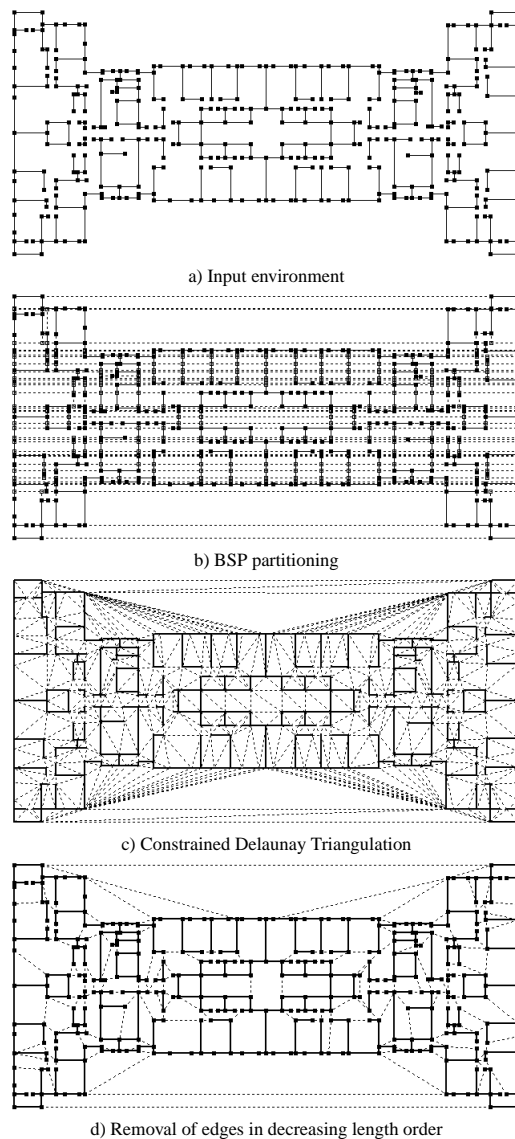
3 Auralization

Auralization is a term created to describe the rendering of sound fields, in analogy to visualization. Many systems for the auralization of sound fields have been developed along the years. A good overview of such systems and of how the localization of sound sources by human beings occur can be found in [16] and [8].

The rendering of the simulated sound field is accomplished by using several virtual sound sources. Being a somewhat lengthy subject, we leave the description on how such virtual sources are implemented to the references above.

In our system we delegate the creation of these virtual sound sources to the DirectX library [17], which offers several algorithms that implement virtual sound sources and also supports different reproduction systems, like headphones and several arrangements of loudspeakers.

To auralize the simulated sound field, we create a different sound source for each propagation path found between the source and the receiver. The position of these sound sources is determined, as in a polar coordinate system, by the length of the propagation path and the incidence angle of the wave at the receiver (which is determined by the last



Partitioning	Cells	Vertices	Occluders	Portals
a	—	380	378	0
b	584	883	851	615
c	724	380	378	725
d	206	380	378	207

Figure 6: Comparison of cell partitioning methods

segment of the propagation path). The attenuations due to the reflections and diffractions suffered by the wave along a propagation path were simulated by adjusting the volume of the virtual sound source.

We used a 5.1 surround sound system [18] and headphones as our test reproduction systems.

4 Results

In our tests we obtained the same qualitative results obtained in the literature, such as an exponential growth in the number of traced beams with the increase of the number specular reflections in each propagation path [1] and the acceleration of this growth [10] with the addition of diffractions.

We also noticed that the addition of diffraction beams to the propagation paths results in smoother sound fields [10], what we believe validates the approximation used to evaluate the attenuation of the sound wave due to diffractions. We also noticed that the addition of diffraction can dramatically improve the coverage of the environment.

Regarding the performance of the algorithm, our initial tests indicate that it is suitable for the construction of a large number of propagation paths in real-time.

More detail on the tests performed and on the results obtained in this work can be found in [8]. As examples of the results obtained, we have included a few images generated by our program. Figure 7 illustrates a few propagation paths computed for the residential building example. Figure 9 illustrates the sound fields obtained with and without diffraction, where the smoother sound field (with diffraction) can be viewed quite clearly. And finally, Figure 8 illustrates the sound intensity levels computed for the residential building example.

5 Conclusions and future work

The main objective when we started this work was to obtain more familiarity with a subject previously unknown to us. Given the accordance of our results to the ones existing in the literature and the excellent performance of the algorithm implemented we believe to have successfully implemented a simple virtual acoustic environment.

Our system can still be extended in several ways. The transmission of sound through walls and the use of different materials for the occluders of the environment can be easily implemented. We can also use more physically correct models to evaluate the attenuation due to diffractions and reflections. Another possible extension is the modification of the beam tracing algorithm to treat moving sound sources, as was made in [13] with the use of parallel processing.

Currently, we are working on extending our system to handle 2.5D environments, environments defined by the vertical sweep of environments defined in the plane. We

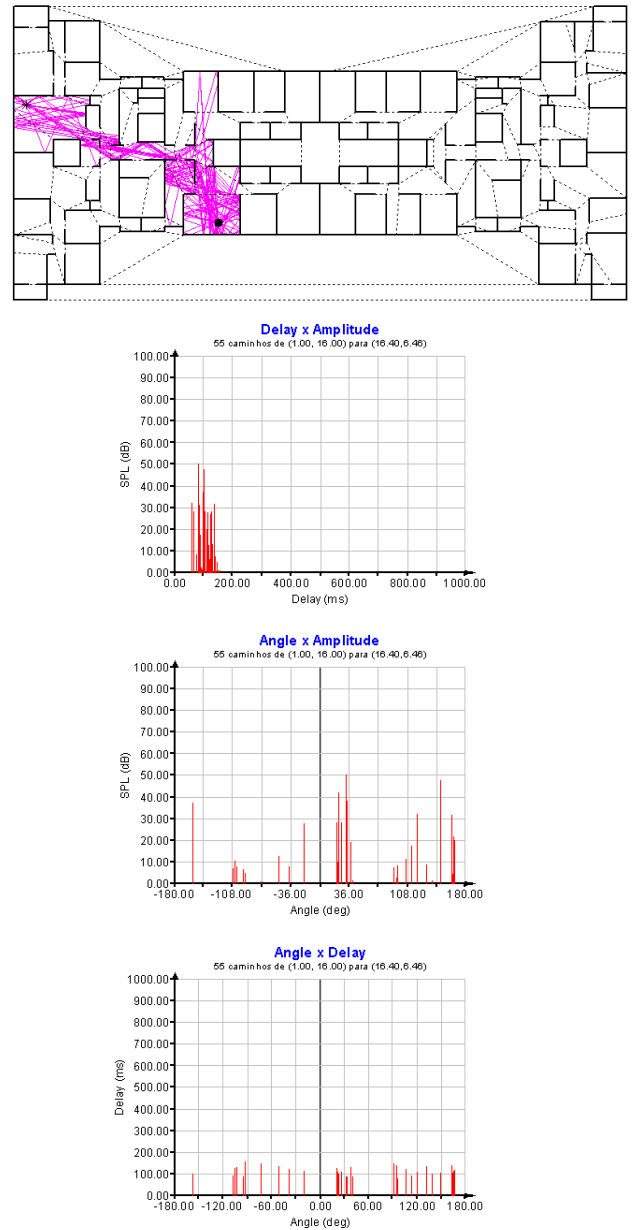


Figure 7: Propagation paths computed for the residential building example

are also studying the possibility of using our algorithm in a computer game that is being currently developed at PUC-Rio and applications of beam tracing for the propagation of radio signals. This application can probably help determining the coverage of a wireless network, helping on the design of such networks and in determining how secure a wireless network is.

Acknowledgements

We would like to PUC-Rio and CAPES for the support that made this work possible. We would also like to thank TecGraf for the resources allocated during this work.

References

- [1] Thomas Funkhouser, Ingrid Carlbom, Gary Elko, Gopal Pingali, Mohan Sondhi, and Jim West. A Beam Tracing Approach to Acoustic Modeling for Interactive Virtual Environments. *Computer Graphics*, 32 (Annual Conference Series):21–32, 1998.
- [2] Philip M. Morse and K. Uno Ingard. *Theoretical Acoustics*. Princeton University Press, 1986.
- [3] Mendel Kleiner, Bengt-Inde Dalenbäck, and Peter Svensson. Auralization – An Overview. In *Journal of the Audio Engineering Society*, pages 861–875, 1993.
- [4] Heinrich Kuttruff. *Room Acoustics*. Spon Press, third edition, 2000.
- [5] Manuel F. Catedra and Jesus Perez. *Cell Planning for Wireless Communications*. Artech House, 1999.
- [6] James D. Foley, Andries van Dam, Steven K. Feiner, and John F. Hughes. *Computer Graphics: Principles and Practice*. Addison-Wesley Publishing Company, Inc., second edition, 1996.
- [7] Paul S. Heckbert and Pat Hanrahan. Beam Tracing Polygonal Objects. In Hank Christiansen, editor, *Computer Graphics (SIGGRAPH '84 Proceedings)*, volume 18, pages 119–127, 1984.
- [8] Sérgio Alvares Rodrigues de Souza Maffra. Propagação de som em ambientes acústicos virtuais bidimensionais. Master's thesis, Departamento de Informática - Pontifícia Universidade Católica do Rio de Janeiro, 2003.
- [9] Seth J. Teller. *Visibility Computations in Densely Occluded Polyhedral Environments*. PhD thesis, Dept. of Computer Science, University of California at Berkeley, 1992.
- [10] Nicolas Tsingos, Thomas Funkhouser, Addy Ngan, and Ingrid Carlbom. Modeling Acoustics in Virtual Environments Using the Uniform Theory of Diffraction. In Eugene Fiume, editor, *SIGGRAPH 2001, Computer Graphics Proceedings*, pages 545–552, 2001.
- [11] D. A. McNamara, C.W.I. Pistorius, and J. A. G. Matherbe. *Introduction to The Uniform Geometrical Theory of Diffraction*. Artech House, 1990.
- [12] Penelope Menounou, Ilene J. Busch-Vishniac, and David T. Blackstock. Directive Line Source Model: a new model for sound diffraction by half planes and wedges. *Journal of the Acoustical Society of America*, 107(6):2973–2986, June 2000.
- [13] Thomas A. Funkhouser, Patrick Min, and Ingrid Carlbom. Real-Time Acoustic Modeling for Distributed Virtual Environments. In Alyn Rockwood, editor, *Siggraph 1999, Computer Graphics Proceedings*, pages 365–374, Los Angeles, 1999. Addison Wesley Longma.
- [14] Patrick Min and Thomas Funkhouser. Priority-Driven Acoustic Modeling for Virtual Environments. *Computer Graphics Forum*, 19(3).
- [15] M. de Berg, M. van Kreveld, M. Overmars, and O. Schwarzkopf. *Computational Geometry: Algorithms and Applications*. Springer Verlag, second edition, 2000.
- [16] Thomas Funkhouser, Jean-Marc Jot, and Nicolas Tsingos. Sounds Good To Me. In *SIGGRAPH 2002 Course Notes*, 2002.
- [17] DirectX. <http://www.microsoft.com/directx>.
- [18] Francis Rumsey. *Spatial Audio*. Music Technology Series. Focal Press, 2001.




Figure 8: Sound intensity levels for the residential building example




Figure 9: Sound intensity level fields without (left) and with (right) diffraction