

Priority-Driven Acoustic Modeling for Virtual Environments

Patrick Min and Thomas Funkhouser

Department of Computer Science, Princeton University, Princeton NJ, USA

Abstract

Geometric acoustic modeling systems spatialize sounds according to reverberation paths from a sound source to a receiver to give an auditory impression of a virtual 3D environment. These systems are useful for concert hall design, teleconferencing, training and simulation, and interactive virtual environments. In many cases, such as in an interactive walkthrough program, the reverberation paths must be updated within strict timing constraints – e.g., as the sound receiver moves under interactive control by a user. In this paper, we describe a geometric acoustic modeling algorithm that uses a priority queue to trace polyhedral beams representing reverberation paths in best-first order up to some termination criteria (e.g., expired time-slice). The advantage of this algorithm is that it is more likely to find the highest priority reverberation paths within a fixed time-slice, avoiding many geometric computations for lower-priority beams. Yet, there is overhead in computing priorities and managing the priority queue. The focus of this paper is to study the trade-offs of the priority-driven beam tracing algorithm with different priority functions. During experiments computing reverberation paths between a source and a receiver in a 3D building environment, we find that priority functions incorporating more accurate estimates of source-to-receiver path length are more likely to find early reverberation paths useful for spatialization, especially in situations where the source and receiver cannot reach each other through trivial reverberation paths. However, when receivers are added to the environment such that it becomes more densely and evenly populated, this advantage diminishes.

1. Introduction

Spatialized sound greatly enhances the sense of presence in a virtual environment by giving a *spatial impression* of the environment, a “conceptual image of the type, size, and properties of the actual or supposed room”.¹ By simulating real world acoustics more accurately, the “audio realism” of the environment is improved, and as a result the feeling of immersion increases.² Example applications include concert hall design, architectural CAD, teleconferencing, training and simulation, warning and guidance systems, displays for the visually handicapped, film sound editing, and multi-player games.

Geometric acoustic modeling methods spatialize sound by convolution with an impulse response representing the acoustic effect of reverberation paths along which sound waves travel from a source to a receiver, taking into account reflections, transmissions, and diffractions. Of course, there are an infinite number of reverberation paths in an indoor en-

vironment, some of which are psychoacoustically very significant (e.g. direct paths to receivers), while others follow complex sequences of reflections and add little or nothing to a listener’s psychoacoustic comprehension of the environment. The ideal goal of an acoustic modeling system is to compute all the reverberation paths perceptible to a listener. However, in many applications, such as an interactive virtual environment, computational resources are limited, and the computation must be completed in a specified amount of time (e.g., an interactive time-slice). In these cases, the difficult challenge is to compute acoustic models by finding only the most psychoacoustically relevant reverberation paths.

Our approach is to use priority-driven beam tracing. The basic idea is that *beams* representing potential sound paths are traced from each sound source along paths of transmission, diffraction, and specular reflection via a *best-first* traversal of a spatial subdivision of the environment. The primary advantage of this approach is that the most significant beams are considered first, enabling methods based on

progressive refinement, graceful degradation, and dynamic termination criteria. As an example, we use our method in an off-line acoustic simulation program to compute approximate impulse responses containing only the most significant reverberation paths quickly, rather than computing exact solutions, which would take an extremely long time. We also use it in a real-time virtual environment system that updates impulse responses with the highest accuracy possible at interactive rates as a user moves through the environment. The priority-driven approach is critical to the success of this application as it enables good time-critical approximate solutions to be found within the real-time constraints of an interactive system.

Priority-driven beam tracing was first introduced in Funkhouser et al.³ In *this* paper, we study the trade-offs of different beam priority functions and evaluate their use in an acoustic modeling system. The paper is organized as follows. The next section contains a survey of previous work in acoustic modeling. Section 3 provides an overview of our priority-driven beam tracing method, while Sections 4 and 5 describe the motivation and computational methods for our several priority functions. The results of experiments with priority-driven beam tracing in a real-time acoustic modeling system appear in Section 6. Section 7 contains a discussion of the limitations of our method and suggestions for extensions. Finally, Section 8 contains a brief summary and conclusion.

2. Previous work

There has been decades of work in acoustic modeling, especially in the area of concert hall design.⁴ For instance, several commercial CAD systems are available that simulate the reverberations of a 3D architectural model and allow a user to visualize the results interactively.^{5, 6, 7} However, to our knowledge, no prior system traces paths in priority order, nor does any provide time-critical performance guarantees.

To review, prior geometric acoustic modeling methods have been based on the following four approaches: 1) image source methods, 2) ray tracing, 3) radiant exchange methods, and 4) beam tracing.

2.1. Image source methods

Image source methods^{8, 9} compute specular reflection paths by considering *virtual sources* generated by mirroring the location of the audio source over each polygonal surface of the environment. The key idea is that a direct path from each virtual source has the same directionality and length as a specular reflection path. Thus, specular reflection paths can be modeled up to any order by recursive generation of virtual sources. This method is simple for rectangular rooms.⁸ However, in general environments, $O(N^R)$ virtual sources must be generated for R reflections over N polygons. Moreover, for every new receiver location, each of the $O(N^R)$ virtual

sources must be checked to see if it is visible to the receiver, since the specular reflection path might be blocked by another polygon or intersect a mirroring plane outside any polygon.⁹ As a result, this method is practical only for computing very few specular reflections in simple environments.

2.2. Ray tracing

Ray tracing methods^{10, 11} find reverberation paths between a source and receiver by generating rays emanating from the source position and following them through the environment until an appropriate set of rays has been found that reach a representation of the receiver position. Although these methods are very general and simple to implement, they are subject to aliasing artifacts as the space of rays leaving the source is sampled discretely. Also, ray tracing is very compute intensive, usually taking hours to compute an impulse response in a complex environment. In work very related to our own, Yagel and Meeker¹² have investigated accelerating a ray tracing algorithm for illumination by processing rays in priority order, considering factors such as intensity, object complexity, and surface properties. They conclude that by using a priority-driven method they are able produce an image which is close to a final image much faster than when using traditional ray tracing methods. Because a large number of rays may be in the priority queue at any given time the memory requirements are substantial (they give a maximum of 192,000 rays for a 256 by 256 image). They found that little extra time was required for the priority-driven method (about 5 % for one particular scene).

2.3. Radiant exchange methods

Radiant exchange methods consider every polygon of the environment a potential emitter and reflector of acoustic energy.^{13, 14, 15, 16} Conceptually, for every pair of polygons, A and B , a form factor is computed which measures the fraction of the energy leaving polygon A that arrives at polygon B . This approach yields a set of simultaneous equations which are solved to obtain the radiant energy for each polygon. Priority-driven approaches, such as progressive refinement radiosity¹⁷ and importance-driven radiosity¹⁸ are possible. However, it is very difficult to model acoustic phase accurately with radiant exchange methods without substructuring polygons to a very fine mesh, and thus we choose not to use these methods in our system.

2.4. Beam tracing

Beam tracing methods¹⁹ classify reflection paths from a source by recursively tracing pyramidal beams (i.e., sets of rays) through the environment. A set of pyramidal beams is constructed that completely covers the 2D space of directions from the source. For each beam, polygons are considered for intersection in order from front to back.

As intersecting polygons are detected, the original beam is clipped to remove the shadow region, a transmission beam is constructed matching the shadow region, and a reflection beam is constructed by mirroring the transmission beam over the polygon's plane (see Figure 1). This method has been used in a variety of applications, including acoustic modeling^{20, 21, 22, 23, 24}, illumination^{25, 26, 27, 19, 28}, radio propagation^{29, 30}, and visibility determination.^{31, 32}

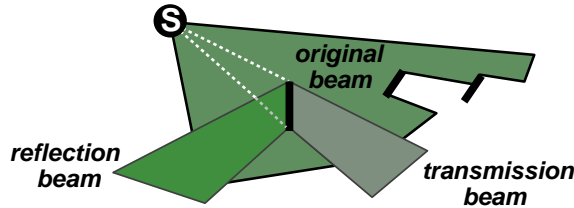


Figure 1: Beam tracing method

The primary advantage of beam tracing is that it takes advantage of spatial coherence, as each beam represents an infinite number of potential ray paths emanating from the source location. It also does not suffer from the sampling artifacts of ray tracing³³ or the overlap problems of cone tracing^{34, 35}, since the entire 2D space of directions leaving the source can be covered by beams exactly. The primary disadvantage is that the geometric operations required to trace a beam through a 3D model (i.e., intersection and clipping) are relatively complex, as each beam may be reflected and/or obstructed by several surfaces. Although several methods have been proposed to make these geometric computations more efficient, including ones based on BSP-trees²⁰, cell adjacency graphs^{31, 30, 21, 3, 32}, layers of 2D triangulations²⁹, and medial axis approximations^{36, 14, 37}, no prior work has evaluated the costs and benefits of re-ordering the beam tracing computation to account for relative priorities.

3. Priority-driven beam tracing method

The priority-driven method presented in this paper is based on the beam tracing method described by Funkhouser et al.^{21, 3} and extended by Tsingos et al.³⁹ Beams are traced from the position of each source along paths of transmission, specular reflection, and diffraction via a best-first traversal of an adjacency graph of polyhedral cells.

Throughout the traversal, a priority queue stores the set of beams to be traced sorted according to a *priority function*. Initially, the priority queue contains only one beam representing the entire space inside the cell containing the source. Then, during each of many iterations, the highest priority beam B is removed from the priority queue and traced into adjacent cells as follows:

- **transmission beams:** For each face F on the boundary of C and intersected by B , a pyramidal *transmission beam*

is formed whose apex is the source of B and whose sides pass through the edges of $F \cap B$. This new beam BT_F represents sound rays traveling along B that transmit through F into the cell which is adjacent to C across F .

- **specular reflection beams:** For each reflective face F intersected by B , a *reflection beam* is created in the same manner, but with the apex mirrored over the plane supporting F . This new beam BR_F represents sound rays traveling along B that reflect specularly off of F and back into cell C .
- **diffraction beams:** For each edge E shared by two non-transparent faces on the boundary of C and intersected by B , a *diffraction beam* is formed whose source is the line segment describing E and whose polyhedral extent contains a cone of potential diffraction paths in the shadow region described by the Geometrical Theory of Diffraction (GTD).^{38, 39} This new beam BD_E represents sound rays traveling along B that diffract around E into a shadow region.

As beams are created, they are assigned priorities and inserted into a priority queue. The process continues until an ending condition is met, such as a limit on the number of beams, or a limit on the number of transmissions, reflections, and diffractions, or the end of a time slice.

An example beam tracing sequence is shown in Figure 2 (note that for clarity beams are shown in two dimensions, whereas they are three-dimensional in our system). The sequence starts in the cell labeled C . Then, the beam is trimmed as it transmits through face u into cell F , mirrored over faces o and p , trimmed at face t as it transmits into cell E , and diffracted into the shadow region of cell D by the lower edge of face s . The resulting beams contain all ray paths traversing this particular sequence of reverberations (one of which is shown as a solid black line).

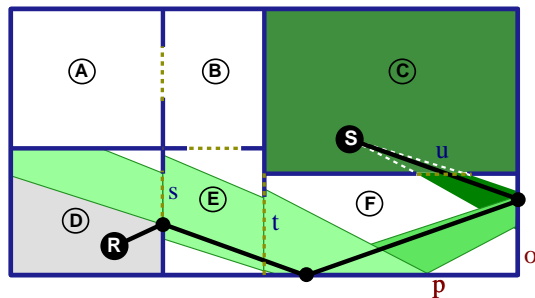


Figure 2: Priority-driven beam tracing

After the beams have been traced, they are used to find reverberation paths from each source to receivers. For every beam containing a receiver, a single reverberation path is constructed by iteratively mirroring the source over reflecting faces in the sequence and by solving for diffraction points according to the Geometrical Theory of

Diffraction.^{38, 39} For each such path, the source-receiver impulse response is updated by adding one pulse, whose delay is determined by l/c (with l the length of the path and c the speed of sound), and whose amplitude is computed by the combined attenuation due to path length, surface absorption, wedge diffractions, source direction, etc. (see Figure 3 for an example impulse response). Finally, the resulting impulse responses are used as convolution filters to spatialize audio signals according to the computed reverberations between sources and receivers in the environment. See ^{21, 3, 39} for more details.

4. Assigning priorities to beams

The primary issue in implementing our priority-driven approach is selecting a “priority function” which orders beams by assigning them a priority value. The function should assign a higher priority to beams which lead to psychoacoustically more significant reverberation paths, e.g., shorter paths with less attenuation. In Funkhouser et al.³ we used a function based on a predicted path length from source to receiver. In this paper, we investigate this priority function and several others, and analyze their performance. Our motivations for choosing priority functions are based on simple psychoacoustic principles, as follows:

- **find many early reverberation paths:**

Following accepted practice of the acoustics literature, we aim to find many early reverberation paths. Early reverberations are the ones that arrive within the first t_{st} milliseconds of the impulse response, when the density of reverberations is low enough that the human ear is able to distinguish individual paths (less than 2,000 reflections per second).⁴⁰ These early reverberations provide a human listener with most of the spatial information about an environment, because of their relatively high strengths, recognizable directionalities, and distinct arrival times^{41, 42, 43, 44} (see Figure 3).

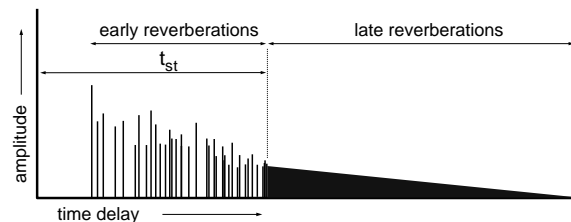


Figure 3: Early and late reverberations.

In the late reverberation phase, when the sound has reflected off many surfaces in the environment, the impulse response resembles an exponentially decaying noise function with overall low power⁴¹ and with such a high density that the ear is no longer able to distinguish them independently.⁴⁰ Consequently, like many other geometric acoustic modeling systems, we focus on computing

only early reverberations, and we model late reverberations with statistical approximations (e.g. ^{45, 46}).

- **find the highest amplitude paths:**

In cases where we are not able to find all early reverberation paths, we aim to find the ones with highest amplitude. By discarding the lowest amplitude pulses, the error in power between the resulting impulse response and the true (real life) impulse response is minimized. Also, higher amplitude reverberations are more likely to mask lower amplitude ones, making the lower amplitude ones inaudible.⁴⁷

5. Priority functions

The goal of our work is to devise a priority function, $f(B)$, that can be used to order the computation of beams efficiently, while capturing the essence of these psychoacoustic principles. In our case, each beam B represents a sequence of transmissions, diffractions, and specular reflections, and thus we can use information related to the predicted length and attenuation of paths represented by the beam to guide our priority-driven algorithm. So far, we have experimented with the following priority functions:

- **traversed path length**

This priority function returns an estimate of the length of the shortest reverberation path traveling along the sequence of transmissions, reflections, and diffractions traversed by the beam. Specifically, for sequences incorporating only transmissions and specular reflections, we can quickly find the length of the shortest possible path by computing the distance from the source of the beam to the closest point on the last cell boundary traversed by the beam. For instance, in Figure 4, the *traversed path length* of the beam b_3 is the length of the line from the source S to the boundary P between cell C and cell F . For sequences also including diffraction, the computation is performed piecewise, starting with the distance from the beam virtual source to the first diffracting edge, adding the distances between every sequential pair of diffracting edges, and ending with the distance from the last diffracting edge to the cell boundary last traversed by the beam. The motivation behind this function is to trace beams in an order expanding approximately spherically from the source, causing earlier reverberation paths to be found before later ones.

- **remaining path length**

This priority function returns an estimate of the straight “as a crow flies” distance between the cell boundary P last traversed by the beam and the receiver closest to P . For an example, consider Figure 5 in which the source is labeled S and the receiver is labeled R , and the last cell boundary traversed by the beam b_3 is labeled P . In this case, the *remaining path length* for beam b_3 is the distance from P to receiver R . The motivation behind this algorithm is to give beams near receivers higher priority. Beams traced




Figure 4: Traversed path length for beam b_3

with this priority function ought to reach regions containing receivers more quickly than the other functions.




Figure 5: Remaining path length for beam b_3

- **potential path length**

This priority function returns a conservative estimate of the *total* length of the shortest reverberation path from the source to a receiver. It is computed as the sum of the *traversed path length* and the *remaining path length* (as described above). The motivation behind this function is to order beams according to how quickly sound traveling along a path inside the beam can reach a receiver. The beams tend to expand in an ellipsoidal pattern from the source, where the foci of the ellipsoid correspond to the locations of the source and the closest receiver.

- **attenuation**

This priority function returns the minimum attenuation of sound traveling along the beam. We model the minimum attenuation by dividing the *traversed path length* by the product of the attenuation coefficients for all reflecting and transmitting surfaces along the beam (attenuation due to diffraction is ignored for this priority function since it depends strongly on a specific receiver location). The motivation behind this priority function is to find the reverberation paths with highest amplitude.

- **breadth first**

Our final method is to create new beams in breadth first order. Specifically, we use a priority function that returns the number of cell boundaries traversed by a beam. Note

that for this function a FIFO queue can be used to implement the beam ordering, rather than a priority queue.

6. Results

We incorporated all these priority functions in the acoustic modeling system of Funkhouser et al.³, extended to trace beams of diffraction.³⁹

To evaluate the functions' performance in a time-critical environment, we ran a series of real-time tests with each priority function, in which a receiver travels along a predefined trajectory in a 3D virtual environment, and a source is stationary in a central location (see Figure 6). During each test, a new set of beams was computed emanating from the source position two times per second. Each test took 45 iterations, or 22.5 seconds, and was run on an SGI Onyx using two 195 MHz R10000 processors, with one processor dedicated to computing the beams, and the other constructing reverberation paths and computing impulse responses. For comparison we included a priority function that returns random values.

The 3D virtual environment used is shown in Figure 6. It consists of about 50 rooms and 1800 polygons, and the BSP spatial subdivision contains about 800 cells. We set the reflectivity value of each surface in the test environment to the same value (80 %).



Figure 6: test environment

The sound source was positioned in a large central room (at the location labeled S in Figures 6 and 7), and a receiver traveled along the trajectory shown as a thick line in Figure 7. This path was chosen because it includes receiver locations both nearby (e.g., in the same room) and far from the source (e.g. down the hall). Different *zones* of the trajectory can be characterized by the relative position of source and receiver: the boundaries between these zones are shown as short black bars on the trajectory in Figure 7, and as vertical lines in the graphs of Figures 8 and 10. Zone 3 represents situations where the source and receiver are in the same room, while zones 1 and 5 represent situations where they are far apart. Zones 2 and 4 fall in between, representing situations where the receiver is in an adjacent room or hallway.

We compare priority functions based on (1) the number of early reverberations, and (2) their power in the computed

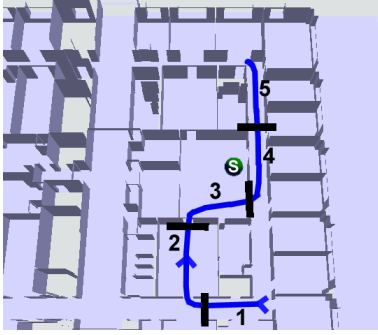


Figure 7: source position and receiver trajectory, divided into 5 zones

impulse responses. Of course, to be able to classify reverberations as early or late, we must have a reasonable estimate of t_{st} , i.e. the time when the pulse density in an impulse response reaches 2,000 per second. To compute such an estimate for every step along the test trajectory, we first generated a “reference” impulse response for each receiver position along the trajectory, by computing up to 2 million beams at each step, using the *potential path length* priority function. We then counted the number of pulses in each 1 ms interval after t_d , in steps of 0.5 ms. The start of the first 20 consecutive intervals with 2 pulses was used as the early reverberation cutoff time (t_{st}). Figure 8 shows for each step the delay time of the first arriving sound (t_d), the 1 ms intervals with 2 pulses (as dots), and the resulting values for t_{st} . The delay time corresponding to the potential path length of the highest priority unprocessed beam (i.e. the beam at the head of the priority queue after 2 million beams had been processed) was larger than t_{st} at every timestep, which guarantees that we found all early reverberations.

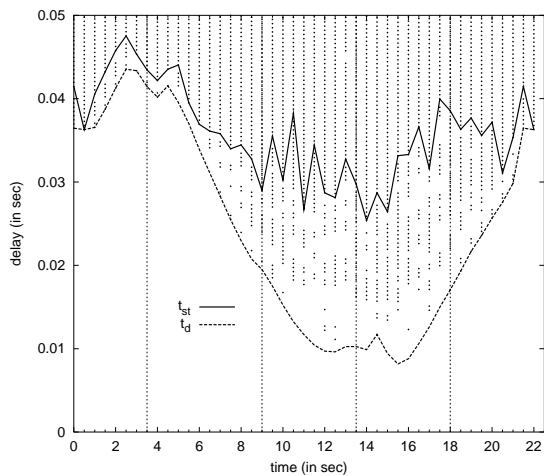


Figure 8: delay time of the first arriving sound (t_d) and cutoff time for early reverberations (t_{st})

We compute the power of the early reverberations of an impulse response (in dB) by:

$$10 \cdot \log \sum_{i=1}^n a_i^2$$

where n is the number of pulses and a_i the amplitude of the i th pulse. For example, an impulse response with one pulse of amplitude 1.0 has a power of 0 dB. An impulse response with no pulses was set to have a power of -60 dB.

6.1. Effect of priority function

Table 1 shows for each function in each zone the average number of early reverberations found, and the average over the whole trajectory in the rightmost column. Table 2 shows the average power of those early reverberations for each function and each zone, and the average over the whole trajectory. The top line in both tables shows these values for the reference impulse response, which contains all early reverberations.

↓func zone→	1	2	3	4	5	avg
reference	9.7	8.4	34.1	36.2	22.5	22.2
potential pl	7.0	4.7	22.1	24.0	12.2	13.9
traversed pl	1.6	0.0	11.9	22.4	10.2	9.2
breadth first	2.9	3.0	17.4	17.2	5.5	9.1
attenuation	3.0	0.7	12.2	19.0	8.6	8.6
remaining pl	0.0	0.5	9.7	6.4	0.0	3.3
random	0.0	0.0	3.0	2.6	0.2	1.1

Table 1: Number of early reverberations found per zone, and overall average

↓func zone→	1	2	3	4	5	avg
reference	-21	-20	-9	-7	-14	-14
potential pl	-35	-28	-8	-9	-17	-19
traversed pl	-38	-59	-14	-9	-17	-28
breadth first	-33	-27	-9	-9	-18	-19
attenuation	-38	-39	-10	-9	-17	-22
remaining pl	-59	-52	-10	-16	-53	-38
random	-59	-59	-27	-20	-49	-43

Table 2: Power of early reverberations found per zone in dB, and overall average

Overall, examining the right-most columns of Table 1 and Table 2, we find that the *potential path length* function finds the most early reverberations with the highest power on average. *Remaining path length* and *random* perform worst of all. *Remaining path length* only considers estimated distance to the receiver, and as a result finds many late reverberations and few early ones, because it tends to cause the beams to

travel to the receiver quickly, which then spawn new beams only in the room the receiver is in. The other priority functions have performance in between these two extremes.

Even though the difference in power between *potential path length* and *breadth first* is small, the impulse responses are substantially different. For example, Figure 9 shows two impulse responses from when the receiver is in zone 5 (at $t = 20$), one generated using *potential path length*, and one using *breadth first*. Their respective powers are -18.5 and -20.8 dB, a difference of only 2.3 dB. However, the impulse response created using *potential path length* has 16 more early reverberations, adding significantly to the acoustic impression.

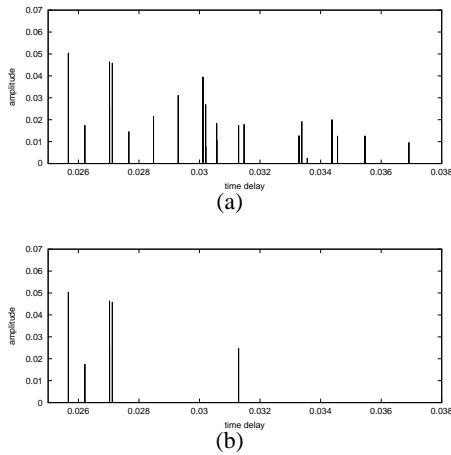


Figure 9: example impulse responses in zone 5 (at $t = 20$), created using (a) potential path length, and (b) breadth first

6.2. Effect of receiver location

We find that the relative performance of priority functions varies depending on receiver location. If we look at the results for each zone individually, specifically in the far away zones 1 and 5, we find that the benefits of incorporating an estimated distance to both source and receiver in the priority function (as in *potential path length*) increase with greater distances between them.

The graph in Figure 10 shows for each priority function (except *remaining path length* and *random*) the number of early reverberations found at every step along the receiver trajectory. Note that *potential path length* finds the most early reverberations at every step, with the largest relative difference when the receiver is far away (in zone 1).

Priority functions that do not use the receiver position cause beams to be created equally in all directions, and since the number of beams grows exponentially with increasing reverberations, they need many more beams to reach the receiver as it moves further from the source. For example, Figure 11 shows 4,000 beams created using the *traversed path*

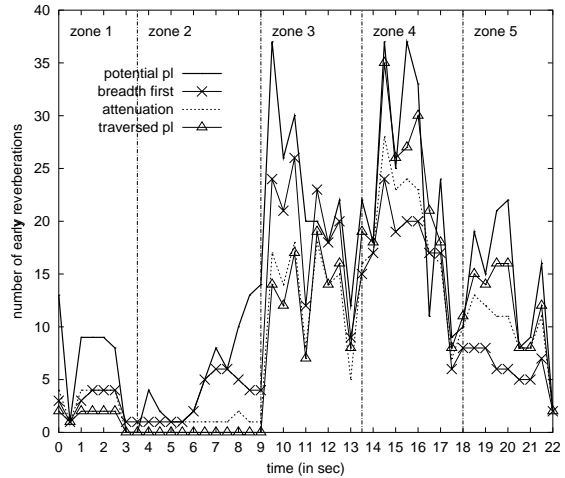


Figure 10: Number of early reverberations for each step along receiver trajectory

length and the *potential path length* function when the receiver is at the end of the trajectory. Notice that the *potential path length* function mostly examines beams in the direction of the receiver, thus making more effective use of available computational resources.

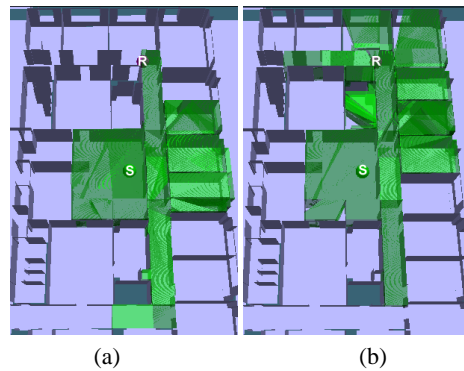


Figure 11: 4,000 beams created using (a) traversed path length, and (b) potential path length, with the receiver at the end of the trajectory

6.3. Effect of more beams

We expect that given enough beams, every function will eventually find all early reverberations. To determine the dependency of the functions' performance on the number of beams computed, we tested several priority functions with two different receiver locations, one near the source and one far away, and limited the number of beams to 250, 500, 1,000, 2,000, 4,000, 8,000, 16,000, 32,000, 64,000, and 128,000. Figures 12 (a) and (b) show for *potential path*

length, breadth first, attenuation, and traversed path length how many early reverberations were found, with in (a) the receiver in the same room as the source, and in (b) the receiver at the start of the trajectory. The number of early reverberations found in the reference impulse responses was 19 and 17 respectively. With the source and receiver in the same room, each function quickly finds all early reverberations (Figure 12 (a)). But when the receiver is further away, only *potential path length* finds all relatively quickly, while the other functions spend most of their time computing use-less beams (Figure 12 (b)).

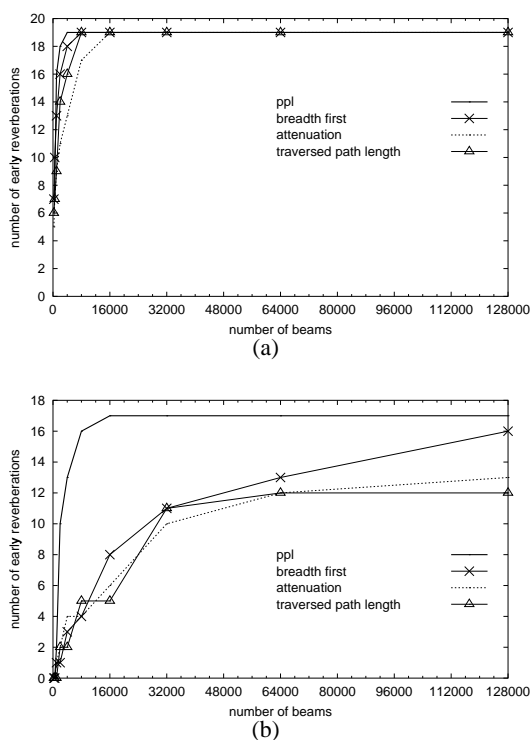


Figure 12: Number of early reverberations found as a function of the number of beams, (a) receiver in same room as source, (b) receiver at start of trajectory

From this we conclude that when the source and receiver are relatively far apart, that the convergence rate of our beam tracing method for finding all early reverberations varies dramatically with different priority functions. Examining the curve in Figure 12 (b), we see that the *potential path length* function outperforms the others significantly, and we expect it to be the only one to find all early reverberations in a practical amount of processing time.

6.4. Effect of priority function overhead

Here we examine the relative overhead of maintaining a priority queue, and the impact of computing priority values on

the rate at which beams can be created. Table 3 shows for the first test what percentage of the total time was spent on queue management, what percentage on beam creation (which includes the priority value computation), and how many beams per second were computed. Note that the *breadth first* priority function computes no priority value, so it does not incur priority queue overhead.

function	queue time %	beam time %	beams/sec
potential pl	2.9	97.1	7,481
breadth first	1.6	98.4	9,143
traversed pl	3.1	96.9	7,861
attenuation	3.0	97.0	8,264
remaining pl	3.1	96.9	6,955
random	2.8	97.2	7,093

Table 3: Percentage of total time spent on priority queue and on creating beams, and beams computed per second

As can be seen from the *queue time %* (column 2 in Table 3), the relative cost of maintaining a priority queue is small. The overhead for computing priority values is fairly large: the *breadth first* function creates over 20 % more beams per second than *potential path length*. Nevertheless it yields fewer early reverberations. The low beam rates for *remaining path length* and *random* are probably caused by the fact that on average each beam spawns fewer new beams than the beams created by the other functions. We conclude that most of the cost of using priority functions is in the computation of priority values, but that this is time well spent, as it results in more early reverberations.

6.5. Effect of multiple receivers

All tests discussed so far included one source and one receiver, enabling *potential path length* to direct every beam towards that one receiver. However, if more receivers are added, they all “compete” for the beams emanating from the source, and thus we expect that the benefits of *potential path length* will diminish. To test this hypothesis, we ran a test four times with one source and 2, 4, 6 and 8 stationary receivers respectively (their locations are marked in Figure 13(b)). In each test, beams were computed from the source for 0.5 seconds using *potential path length*.

We find that the relative advantages of *potential path length* decrease as the environment becomes more densely populated with receivers. In our tests, the total number of early reverberations found to all of the 2, 4, 6, and 8 receivers were 26, 53, 64, and 62 respectively. The growth is less than linear, and even decreases from 6 to 8 receivers. The reason is a combination of two factors. First, *potential path length* has to compare the distances between a beam and every receiver for every beam, and thus the number of beams which

can be computed in 0.5 seconds decreases with an increasing number of receivers (3,741, 3,654, 3,476, and 3,327 beams for 2, 4, 6, and 8 receivers respectively). Second, the beams traced with *potential path length* become less directed towards specific receivers. For example, Figure 13 shows the beams created in 0.5 seconds by *potential path length* for 2 and 8 receivers. Note that in Figure 13(b) the beams are spread more evenly around the source. In the limit, when there are many receivers in the environment, *potential path length* resembles *traversed path length*, and the beams tend to extend equally in all directions. We conclude that the potential path length function is most useful if the environment is sparsely populated.




Figure 13: Beams created in 0.5 seconds by potential path length for (a) 2 and (b) 8 receivers

6.6. Summary of results

Overall, we conclude that priority functions which incorporate estimates of the distance traveled so far from a source *and* the remaining distance to a receiver find more early reverberations, especially when the source and receiver cannot reach each other through trivial reverberation paths. This advantage diminishes as more potential receivers are added to the system and the environment becomes more densely and evenly populated.

The benefit of priority driven beam tracing is strongest when the receiver is not too close to the source. When the receiver moves out of the source room, the *potential path length* function finds substantially more reverberation paths. As the receiver moves further away, it eventually enters a region where only *potential path length* finds reverberation paths. Figure 14* provides a visualization of this effect: part of the test environment was divided into a grid of 1 by 1 meter squares, and the number of reverberation paths from the source to a receiver in the center of each square was counted, using 3,741 beams for *potential path length* and 4,572 for *breadth first* (corresponding to 0.5 seconds compute time). Note that because *potential path length* is dependent on receiver position, the beams were recomputed

for each receiver position. The lighter a square, the more reverberation paths were found to the receiver position in that square. It can be seen that the patterns in the source room are similar. However, in adjacent rooms and the hallway they start to differ, and further away only *potential path length* finds early reverberations.

7. Limitations

In this section, we discuss limitations and possible future extensions of our system.

Our acoustic modeling system is a research prototype, and it has several limitations. First, the implementation does not include sophisticated models for acoustical reflectance distribution functions or directionality of audio sources and receivers. For instance, we currently represent the reflectance of each surface with an angle-independent and frequency-independent absorption coefficient. But, since we compute reverberation paths explicitly, adding these auralization features is relatively straight-forward.

Second, our beam tracing methods are only practical for coarse 3D models without highly faceted surfaces, such as the ones often found in acoustic modeling simulations. The difficulty is that beams get fragmented by cell boundaries as they are traced through a cell adjacency graph. For this reason, we are *not* optimistic that our implementation can be easily adapted to model light transport for photorealistic image synthesis. However, the algorithms probably can be applied effectively in other application areas concerned with simulation of wave phenomena with long wavelengths (e.g., radio propagation).

8. Conclusion

In this paper we have evaluated several priority functions as used in a priority-driven beam tracing method for spatializing sound in virtual environments. The advantage of using a priority-driven method is that computational resources are spent on the psychoacoustically most significant beams, while the overheads of priority management seem relatively insignificant.

We found that functions which incorporate an estimate of the distance traveled so far from a source, and of the remaining distance to a receiver find more early reverberations, especially when the source and receiver cannot reach each other through trivial reverberation paths (e.g. a direct path). When receivers are added to the environment such that it becomes more densely and evenly populated, this advantage diminishes.

In future work, we plan to experiment with other psychoacoustically motivated priority functions, perhaps related to source volume, receiver direction, or source sound onset. We also hope to perform psychoacoustic experiments to help determine the impact of different priority functions on sound localization and sense of presence in a virtual environment.

Acknowledgements

The authors would like to thank Perry Cook and Georg Essl for their valuable comments. We also appreciate the support and motivation provided by Ingrid Carlbom, Gary Elko, Mohan Sondhi, Jim West, and Gopal Pingali of Bell Laboratories. This work was supported by a fellowship from the Alfred P. Sloan Foundation.

References

- J. Blauert, *Spatial hearing: the psychophysics of human sound localization*. MIT Press, revised ed., (1997).
- N. Durlach and A. Mavor, "Virtual reality scientific and technological challenges", tech. rep., National Research Council, Washington DC, (1995). National Academy Press.
- T. A. Funkhouser, P. Min, and I. Carlbom, "Real-time acoustic modeling for distributed virtual environments", in *Proc. SIGGRAPH 1999*, (Los Angeles), pp. 365–374, (August 1999).
- H. Kuttruff, *Room Acoustics*. London, England: Elsevier Science, 3rd ed., (1991).
- Bose Corporation, "Bose modeler", Framingham, MA, <http://www.bose.com>.
- CATT, "Catt-acoustic." <http://www.netg.se/~catt>. Gothenborg, Sweden.
- G. M. Naylor, "Odeon - another hybrid room acoustical model", *J. Applied Acoustics*, **38**(2-4), pp. 131–144 (1993).
- J. Allen and D. Berkley, "Image method for efficiently simulating small-room acoustics", *J. Acoust. Soc. Am.*, **65**(4), pp. 943–950 (1979).
- J. Borish, "Extension of the image model to arbitrary polyhedra", *J. Acoust. Soc. Am.*, **75**(6), pp. 1827–1836 (1984).
- U. Krockstadt, "Calculating the acoustical room response by the use of a ray tracing technique", *J. Sound and Vibrations*, **8**(18), (1968).
- T. Whitted, "An improved illumination model for shaded display", *Communications of the ACM*, **23**(6), pp. 343–349 (1980).
- R. Yagel and J. Meeker, "Priority-driven ray tracing", *Journal of Visualization and Computer Animation*, **8**(1), pp. 17–32 (1997).
- C. M. Goral, K. E. Torrance, D. P. Greenberg, and B. Battaile, "Modeling the interaction of light between diffuse surfaces", in *Proc. SIGGRAPH 1984*, pp. 213–222, (July 1984).
- T. Lewers, "A combined beam tracing and radiant exchange computer model of room acoustics", *J. Appl. Acoustics*, **38**(2-4), pp. 161–178 (1993).
- G. Moore, *An Approach to the Analysis of Sound in Auditoria*. PhD thesis, Cambridge University, Cambridge, UK, (1984).
- N. Tsingos and J.-D. Gascuel, "A general model for simulation of room acoustics based on hierarchical radiosity", in *SIGGRAPH 1997 Visual Proc.*, (August 1997). Technical Sketches.
- M. Cohen, S. Chen, and J. W. D. Greenberg, "A progressive refinement approach to fast radiosity image generation", in *Proc. SIGGRAPH 1988*, pp. 75–84, (1988).
- B. E. Smits, J. R. Arvo, and D. H. Salesin, "An importance-driven radiosity algorithm", in *Proc. SIGGRAPH 1992*, (Chicago), pp. 273–282, (July 1992).
- P. Heckbert and P. Hanrahan, "Beam tracing polygonal objects", in *Proc. SIGGRAPH 1984*, pp. 119–127, (July 1984).
- N. Dadoun, D. Kirkpatrick, and J. Walsh, "The geometry of beam tracing", in *Proc. of the Symposium on Computational Geometry*, (Baltimore), pp. 55–61, (June 1985).
- T. A. Funkhouser, I. Carlbom, G. Elko, G. Pingali, M. Sondhi, and J. West, "A beam tracing approach to acoustic modeling for interactive virtual environments", in *Proc. SIGGRAPH 1998*, (Orlando), pp. 21–32, (July 1998).
- M. Monks, B. M. Oh, and J. Dorsey, "Acoustic simulation and visualization using a new unified beam tracing and image source approach." Meeting of the Audio Engineering Society, (November 1996).
- U. Stephenson and U. Kristiansen, "Pyramidal beam tracing and time dependent radiosity", in *Proc. 15th Intern. Congress on Acoustics*, pp. 657–660, (June 1995).
- J. P. Walsh and N. Dadoun, "What are we waiting for? the development of godot, ii." presented at the 103rd Meeting of the Acoustical Society of America, Chicago, (April 1982).
- J. Chuang and S. Cheng, "Computing caustic effects by backward beam tracing.", *The Visual Computer*, **11**(3), pp. 156–166 (1995).
- G. Ghazanfarpour and J. M. Hasenfratz, "A beam tracing with precise antialiasing for polyhedral scenes", *Computer & Graphics*, **22**(1), (1998).
- E. A. Haines, "Beams o' light: confessions of a hacker." Frontiers in Rendering Course Notes, SIGGRAPH 1991, (1991).
- M. Watt, "Light-water interaction using backward beam tracing", in *Proc. SIGGRAPH 1990*, vol. 24, pp. 377–385, (1990).
- S. Fortune, "Algorithms for prediction of indoor radio propagation", tech. rep., Bell Laboratories, (1996). A partial version of this paper appeared in *Applied Computational Geometry, Towards Geometric Engineering*, proc. of the FCRC 1996 Workshop in conjunction with WACG 1996, Philadelphia, PA, May 1996, pp. 157–166.
- S. Fortune, "Topological beam tracing", in *Proc. 15th Annual Symposium on Computational Geometry*, pp. 59–68, (1999).
- C. Jones, "A new approach to the 'hidden line' problem", *The Computer Journal*, **14**(3), pp. 232–237 (1971).
- S. Teller, *Visibility Computations in Densely Occluded Polyhedral Environments*. PhD thesis, Computer Science Division (EECS), UC Berkeley, (1992). Also available as UC Berkeley technical report UCB/CSD-92-708.
- H. Lehnert, "Systematic errors of the ray-tracing algorithm", *J. Appl. Acoustics*, **38**(2-4), pp. 207–221 (1993).
- J. Amanatides, "Ray tracing with cones", in *Proc. SIGGRAPH 1984*, pp. 129–135, (1984).
- J. Vian and D. van Maercke, "Calculation of the room response using a ray tracing method", in *Proc. of ICA Symposium on Acoustics and Theater Planning for the Performing Arts*, (Vancouver, CA), pp. 74–78, (1986).
- P. Kreuzgruber, P. Unterberger, and R. Gahleitner, "A ray splitting model for indoor radio propagation associated with complex geometries", in *Proc. of 43rd IEEE Vehicular Technology Conf.*, pp. 227–230, (1993).
- A. Rajkumar, B. Naylor, and L. Rogers, "Predicting rf coverage in large environments using ray-beam tracing and partitioning tree represented geometry", *Wireless Networks*, (1995).
- J. B. Keller, "Geometrical theory of diffraction", *Journal of the Optical Society of America*, **52**(2), pp. 116–130 (1962).
- N. Tsingos, T. Funkhouser, A. Ngan, and I. Carlbom, "Geometrical theory of diffraction for modeling acoustics in virtual environments", tech. rep., Bell Laboratories, (March 2000).
- L. Cremer and H. A. Müller, *Principles and applications of room acoustics*, vol. 1. Applied Science, (1978).
- D. R. Begault, *3-D sound for virtual reality and multimedia*. Academic Press, (1994).
- W. Hartmann, "Localization of sound in rooms", *J. Acoust. Soc. Am.*, **74**(5), pp. 1380–1391 (1983).
- S. H. Nielsen, "Auditory distance perception in different rooms", *J. Audio Eng. Soc.*, **41**(10), pp. 755–770 (1993).
- W. Wagenaars, "Localization of sound in a room with reflecting walls", *J. Audio Eng. Soc.*, **38**(3), (1990).
- W. Ahnert, "Ears auralization software", *J. Audio Eng. Soc.*, **41**(11), pp. 894–904 (1993).
- L. Savioja, J. Huopaniemi, T. Huotilainen, and T. Takala, "Real-time virtual audio reality", in *Proc. ICMC 1996*, (Hong Kong), pp. 107–110, (August 1996).
- B. C. Moore, *An introduction to the psychology of hearing*. Academic Press, 4th ed., (1997).
- W. Furrer, *Room and building acoustics and noise abatement*. London: Butterworth & Co., (1964).