SOUND SPATIALIZATION BASED ON FAST BEAM TRACING IN THE DUAL SPACE

Marco Foco, Pietro Polotti, Augusto Sarti, Stefano Tubaro

Dipartimento di Elettronica e Informazione
Politecnico di Milano
Piazza L. Da Vinci 32, 20133 Milano, Italy
marco.foco@fastwebnet.it polotti/sarti/tubaro@elet.polimi.it

ABSTRACT

This paper addresses the problem of geometry-based sound reverberation for applications of virtual acoustics. In particular, we propose a novel method that allows us to significantly speed-up the construction of the beam tree in beam tracing applications, by avoiding space subdivision. This allows us to dynamically recompute the beam tree as the sound source moves. In order to speed-up the construction of the beam tree, we determine what portion of which reflectors the beam "illuminates" by performing visibility checks in the "dual" of the geometric space.

1. INTRODUCTION

Advanced sound rendering techniques are often based on the physical modeling of the acoustic reflections and scattering in the environment. This can be modeled using either numerical methods (finite elements, boundary elements and finite differences) or geometrical methods (image source, path tracing, beam tracing and radiosity). With all such techniques, however, computational complexity may easily become an issue. One approach that enables an efficient auralization of all reflected paths is based on beam tracing [1]. This method is based on a geometric pre-computation of a tree-like topological representation of the reflections in the environment (beam tree) through spatial subdivision techniques. The beam tree is then used for real-time auralization through a simple look-up of which beams pass through the auditory points. It is important to notice, however, that the tree-like structure of the beam reflections makes the approach suitable for early reverberations only, as it prevents us from looping beams and implementing the corresponding IIR structures. In order to overcome this difficulty, we recently proposed an approach [2] that models both early and late reverberation by cascading a tapped delay line (a FIR) with a Waveguide Digital Network (WDN) [3] (an IIR). In this approach, the beam tree is looked-up to generate the coefficients of a tapped delay line for the auralization of early reverberation. In order to account for late reverberation as well, we feed the outputs of the tapped delay lines into WDN, whose parameters are determined through path tracing [2].

One problem of the beam tracing approach is that it assumes that sound sources are fixed and only listeners are allowed to move around. In fact, every time the source moves, the beam tree needs to be re-computed. This may easily become a costly operation, particularly with environments of complex topology, as it is based on spatial subdivision algorithms.

In this paper we propose a novel method that allows us to significantly speed-up the construction of the beam tree, by avoiding space subdivision. This allows us to dynamically recompute the beam tree as the source moves. In order to speed-up the construction of the beam tree, we determine what portion of which reflectors the beam "illuminates" by performing visibility checks in the "dual" of the world space. In order to illustrate the approach, we will assume that the world space be 2-dimensional.

2. TRACING BEAMS IN THE DUAL SPACE

The world space that we consider is made of sources and reflectors. Sources are assumed as point-like, and reflectors are linear segments. We call "active" that region of a reflector that is directly "illuminated" by a source. The active portion of a reflector can be made of one or more active segments (connected sets), due to occlusions. Each one of the active segments generates a beam, which is defined as that bundle of acoustic rays that connect the source with points of the active segment. Each acoustic ray w can be described by an equation of the form $y = ax + b$, where a is the angular coefficient and b is the offset on the y axis, all referred to a world coordinate system (x, y). The dual $\hat{w}$ of that ray is thus a point (a, b) in parameter space. The dual of a point p of coordinates (x, y), on the other hand, is a line $\hat{p}$ in parameter space, given by all pairs (a, b) that correspond to rays passing through p. This means that the dual of a source is a ray.

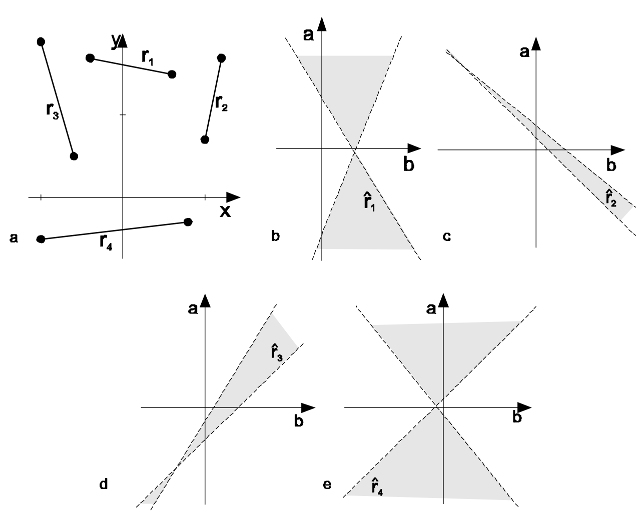

Figure 1: The duals of a set of physical reflector is a collection of strip-like regions in the dual space.

Let us now consider a reflector, which is a segment in world space. In order to determine its dual, we can consider the duals $\hat{p}$ and $\hat{q}$ of its two extremes p and q, respectively. Such points will correspond to two lines $\hat{p}$ and $\hat{q}$ in the dual space. As the reflector r is made of all points in between p and q, its dual $\hat{r}$ will be made of all the lines in between $\hat{p}$ and $\hat{q}$. In other words, the dual of a reflector is a beam-like region. We recall, however, that an active segment of a reflector, together with a virtual source s, completely specifies a beam. The dual of this beam can be determined as the intersection between the dual $\hat{r}$ of the active segment (a beam-like region) and the dual $\hat{s}$ of the virtual source (a line). In conclusion, the dual of a beam is a segment, just like the dual of a segment (reflector) is a beam.

Pushing this formalism even further, we can define active reflecting segments at infinity (just like in projective geometry) and beams at infinity, which are those beams that are not reflected by any surface. If we have a collection of reflectors in world space, their duals will form a collection of beam-like regions (see Fig. 1).

We would now be tempted to say that the intersection between the dual of a source and the dual of the reflectors is the dual of the branching beams, therefore it can be used to determine the beam tree. This would be true if there were no occlusions between reflectors (i.e. if all points of the reflectors were active). In fact, mutual occlusions cause the duals of reflectors to overlap each other, therefore in order to determine the branching of the beams we need first to figure out the layering order of such regions, according to mutual occlusions (i.e. "which region eats which").

The layering order is not the only problem, as we need to construct the tree of all possible reflections in the environment while keeping track of the occlusions. In order to do so, we propose an iterative approach that starts from the source (root) and tracks the branching of all beams in dual space. At each iteration we consider a beam at a time and we determine how this beam is split into various sub-beams as it encounters reflectors on its way. This splitting characterizes the branching of the beam tree.

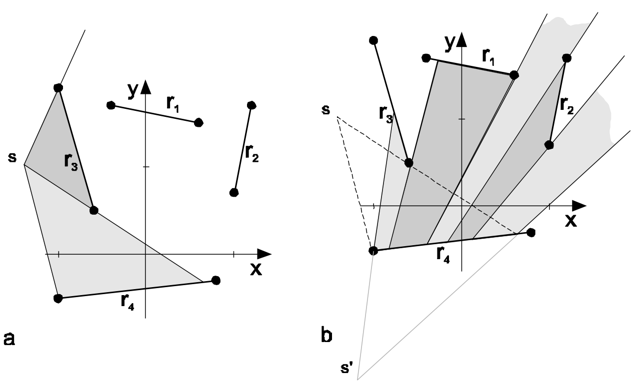

Figure 2: Beam tracing in world-space. a. The source s illuminates the reflector r_3 and partially the reflector r_4. b. The beam that describes the reflection from r_4 (s' is the virtual source for this reflection). In this figure sub-beams illuminating r_3, r_1 and r_2 are shown.

At the first iteration we consider the source and the physical reflectors (Fig. 2a). From this source a number of beams will depart, each corresponding to an active segment of the various reflectors. Such active segments can be determined quite straightforwardly by tracing all rays that depart from the source. At this point we consider all the beams one by one and figure out how each of them branches out as it encounters the reflectors (Fig. 2b). Each beam will keep branching out until its cumulative attenuation or its aperture angle falls beneath some pre-assigned threshold. Let us consider a beam originating from a virtual source s. If we want to characterize how this beam branches out, start from the active reflector segment r_0 (Fig. 3a) that defines that beam (aperture) and we perform a change of reference frame in order for its extremes p_0 and q_0 to fall in the coordinates $(0, 1)$ and $(0, -1)$, respectively (Fig. 3b).

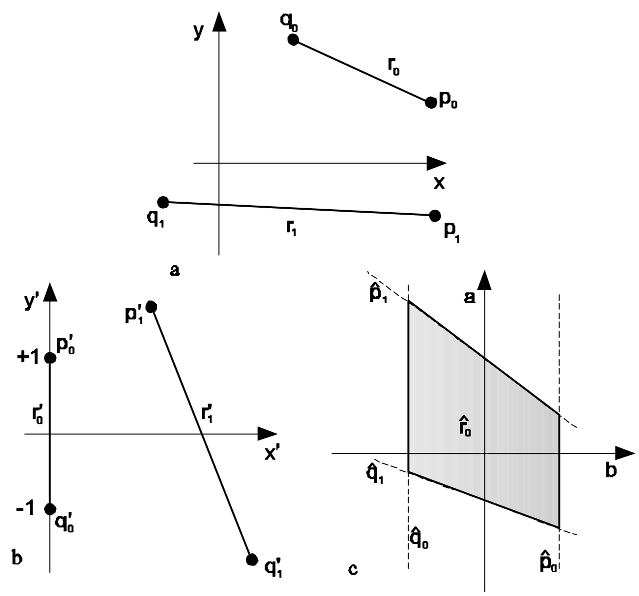

Figure 3: A reflective segment: a. in the world space, b. after normalization relative to r_i, c. in the dual space

This way, the dual of the considered aperture will be the reference strip $-1 \leq b \leq 1$ in parameter space (see Fig. 3c). The intersection in the dual space between the refence strip and the representation oof the othe considered reflection is the area named $\hat{r}_1$ that represents the visibility of r'_1 from r'_0.

In Fig. 4 and 5 we can see some other examples of a reflector visibility from a reference segment.

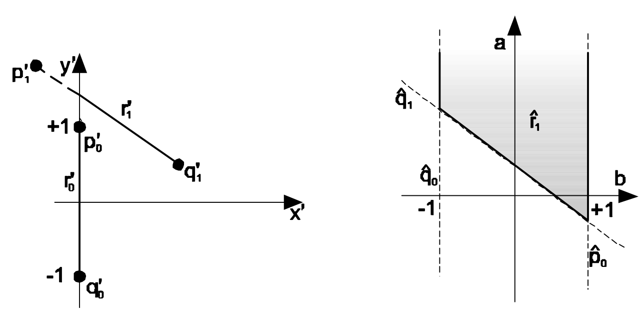

Figure 4: The visibility of r'_1 with respect of r'_0. The visibility is limited to the right half-plane of the reference reflector.

Let us now consider again Fig. 1, which has four physical reflectors. Let us assume that, as the iterations go alone, r_4 is found to be an active reflector, and we want to show how to determine the

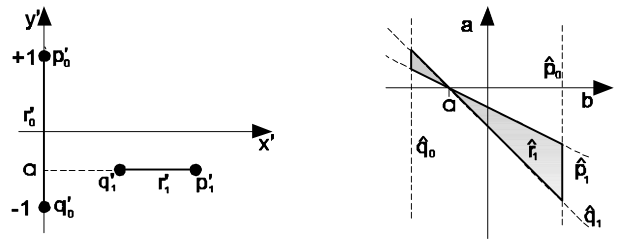

Figure 5: Another particular case of visibility. Varying the observation point on the reference segment, the two different sides of the reflector r'_1 are seen. This is clear from the observation of the dual space.

branching of the beams. Therefore we will consider r_4 as an active aperture, and a coordinate change is used to normalize its representation in the space domain. We will name the new coordinate system (x^4, y^4) (superscript denotes which reflector is currently the reference for the normalization), and the new representation of the other reflectors will become r_1^4, r_2^4 and r_3^4. In the dual space we obtain the set of regions of Fig. 6. Notice that r_3^4 partially occludes r_1^4. In the dual space, when the areas that represents two reflectors overlap, ther is an occlusion. By analyzing the slopes of the dual representation of the reflector endpoints, it is possible to define in the correct way the relative visibility.

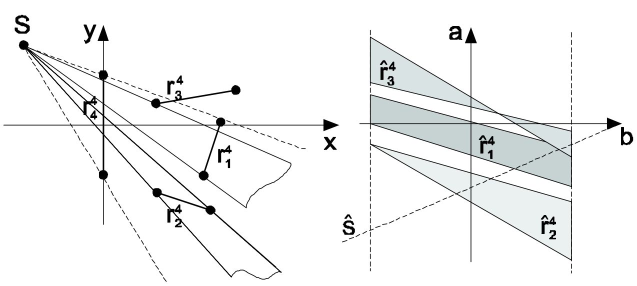

Figure 6: Beam tracing in the dual space. The intersection between $\hat{s}$ and the colored regions generates a collection of segments, which are duals of the branching beams.

A particular ordering situation between reflectors occurs when the sorting between two (or more) reflectors is ambiguous with respect to a third reflector (see Fig. **??**). This happens when the reflector's order changes with respect to a moving point on the reference reflector. This could be a problem when working in metric space, but it becomes easy to solve in the dual space. In fact, in the dual space the visibility information is obtained by partitioning the reference strip ($-1 \leq b \leq 1$), by looking at the intersections of the segments that define the reflector's endpoints.

As a last foremark it is important to notice that, without loss of generality, the source can always be placed on the left of the normalized aperture, as we can always replace a (virtual) source with its specular with respect to the aperture. As we can see in Fig. 6, the scan of the rays that define a source beam corresponds to the scan of the points of its dual along its extension.

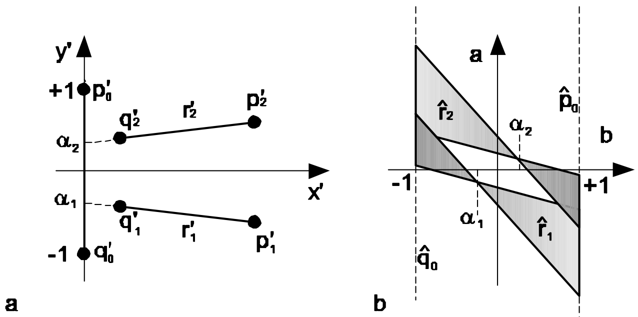

Figure 7: The sorting problem: a.world-space sorting is ambiguous; b. solution in dual-space.

3. IMPLEMENTATION

In order to test the method that we propose, we implemented a software application that includes a 2D CAD for creating test environments, and an auralization system based on both a traditional beam tracer and a dual-space beam tracer. Auralization is based on tapped delay lines, and includes HRTF and cross-fading between beam configurations.

The geometries considered by this application are only two-dimensional, but the method can be extended to the 3D space. The basic structure of our auralization system is shown in Fig. 8.

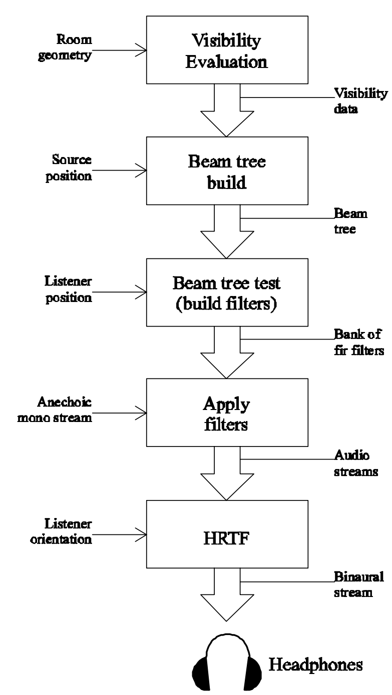

Figure 8: A schematic view of the auralization algorithm embedded in the beam tracing application.

In a first step, the algorithm analyses the geometric description of the reflectors of the environment, and it extracts all the necessary information on visibility by working in the dual of the geometric space. As this step depends only on the geometry of the reflectors, as long as the environment is fixed, this operation can be performed beforehand, as a preprocessing step.

The location of the sound source and the visibility information can now be used to construct the beam tree, or to update it every time the source moves.

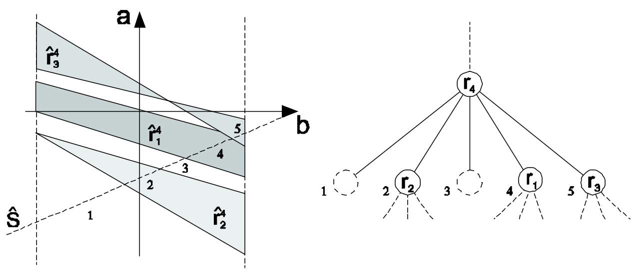

Figure 9: Building a tree level

Fig. 9 refers to the construction of one level of the beam tree (a reflection on the r_4 reflector, as already shown in Fig. 6). In this portion of the beam tree, we describe the splitting of the generating beam (the one that is incident to reflector r_4) into five sub-beams (numbered from 1 to 5). Sub-beams 1 and 3 are "unbounded" beams (they keep propagating undisturbed), while sub-beams 2, 3 and 5 are "bounded", as they end up onto reflectors r_2, r_1 and r_3, respectively. The process is repeated for each one of the "bounded" sub-beams, until the number of reflections reaches a pre-fixed maximum or until the energy associated to the beam becomes negligible.

The beam-tree lookup block of Fig. 8 generates the filter taps according to what beams reach the listener's location. This update operation is fast and can be performed frequently [1]. Through this lookup process we obtain the direction of arrival (DOA), the path length and the list of the reflectors encountered so far. This phase requires all of the tree nodes and leaves to be visited, and the relative beam to be tested for listener's inclusion. If the listener's location falls within the beam area, the DOA and the path length are computed. From there we can go down the beam tree towards the root, in order to obtain the list of reflectors encountered along the propagation of the wavefront.

Since it is impossible to produce one stream for every DOA, we grouped together angular intervals in order to generate a limited number of audio streams, one per interval. Each one of these angular intervals will be attributed a separate delay line (a FIR filter). The taps are computed considering all path lengths and relative attenuations.

The filter bank whose parameters are generated at the previous step, constitutes the auralization algorithm, which generates the directional streams from an anechoic (dry) source. These streams are then mixed for stereo headphone auralization using a HRTF. In our implementation we used OpenAL [7] to mix 16 virtual "equivalent sources" placed in circle around the listener (to simulate 16 discretized DOAs). The listener's head orientation is accounted for only in this last step.

In the following table we summarise what kind of computation is required for source-listener geometry changes.

Type of motion	Recomputation required
Source	beam tree re-computation
Listener's location	beam tree lookup, filter bank update
Listener's orientation	no re-computation required

Every time a re-computation is required, the taps in the delay lines are changed, and this can cause some clicking sounds to be produced in the output. This can be avoided by smoothly interpolating parameters in place of just switching them. A simple method to avoid clicks is a linear interpolation in the audio streams generated by using the "old" and "new" filters (see Fig. 10).

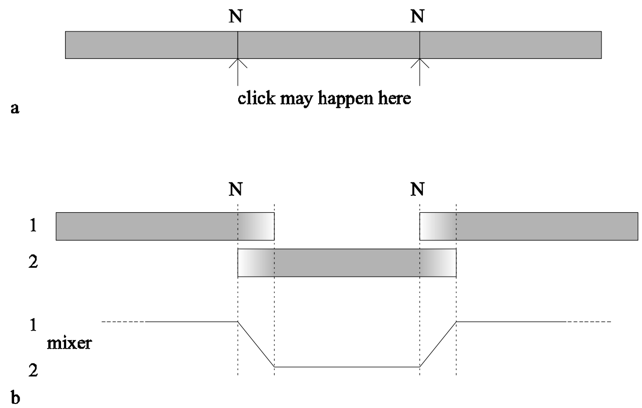

Figure 10: Mixing audio streams to avoid clicks.

4. PERFORMANCE EVALUATION

In order to test the effectiveness of the proposed solution, we compared our implementation (dual space approach) with an optimized one that we developed, which works in geometry space but does not require space subdivision. In spite of the computational efficiency of the approach based working in metric space, the performance improvement of our approach in the dual space is quite apparent.

In the geometry space, the complexity of tracing a beam tree level without space subdivisions depends on n, the number of reflectors. In the dual space, the complexity depends on m ($< n$), the number of "visible" reflectors from the reference. In most cases, such as a typical office configuration, m depends only on the local geometry, and it is independent from n. For example, in a system of square rooms the average visible faces is 6-7, but in a structure characterized by long corridors the number of visible faces can be higher.

In the test we performed with the geometry space approach, the time used to rebuild the beam tree depends quadratically on the number of reflector in the world, while our method (dual space analysis) grows almost in a linear way (Fig. 11).

These tests have been executed on a modular set of rooms connected by not aligned doorways: an example of this model (5 rooms, 20 reflectors) is shown in the program's screenshot in Fig. 12.

5. CONCLUSIONS AND FUTURE WORK

In this paper we showed a novel, very effective algorithm to trace beams in the 2-dimensional space to be used for virtual acoustics. The computational efficiency has been achieved by exploiting particular properties of the representation in the dual space.

We also showed that the above procedure, implemented on a standard PC platform, turns out to be very effective, and the com-

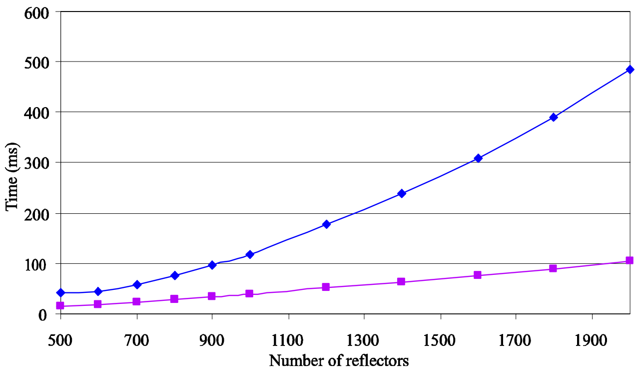

Figure 11: Execution time vs. number of reflectors in the scene for the standard beam tracing (diamonds), proposed method (squares).

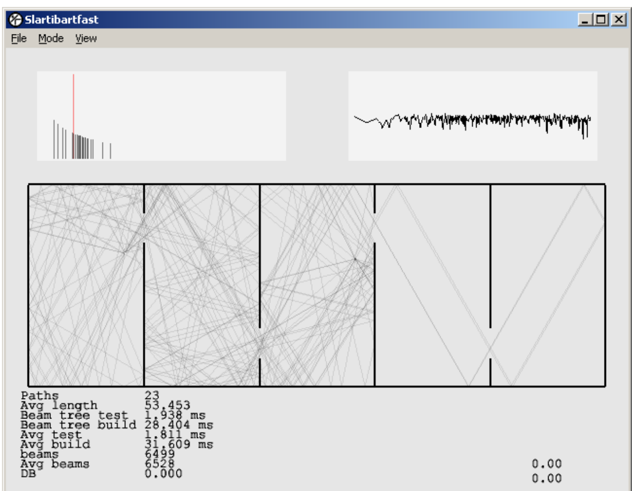

Figure 12: Test application screenshot. The room's model is the 5-rooms test model (20 reflectors).

putational saving compared with standard beam tree methods increases as the complexity of the environment grows.

We are currently working on a 3-dimensional version of this algorithm, to be used with more complex environments.

6. REFERENCES

[1] T. Funkhouser, I. Carlbom, G. Elko, G. Pingali, M. Sondhi, J. West, "Beam Tracing Approach to Acoustic Modeling for Interactive Virtual Environments". *Computer Graphics (SIGGRAPH '98)*, Orlando, FL, July 1998, pp. 21–32.

[2] Augusto Sarti and Stefano Tubaro: "Efficient geometry-based sound reverberation", *X EUSIPCO*, Toulouse, France, 2002.

[3] J.O. Smith, "Principles of digital waveguide models of musical instruments", in *Applications of digital signal processing to audio and acoustics*, edited by M. Kahrs and K. Brandenburg, Kluwer, 1998, pp. 417–466.

[4] S. Teller and P. Hanrahan, "Global visibility algorithms for illumination computations", *in Computer Graphics (SIGGRAPH '93 Proceedings)*, pp. 239–246.

[5] P.S. Heckbert and P. Hanrahan, "Beam Tracing Polygonal Objects", in *Computer Graphics (SIGGRAPH '84 Proceedings)*, pp. 119–127.

[6] L.Savioja, "Modeling Techniques for Virtual Acoustics", PhD dissertation, .Espoo 1999.

[7] OpenAL Documentation, available via *http://www.openal.org/*.