Acoustic simulation using hierarchical time-varying radiant exchanges

Nicolas Tsingos and Jean-Dominique Gascuel

iMAGIS-GRAVIR/IMAG

1 Introduction

Sound is essential to enhance any visual experience and is becoming an important issue in virtual reality systems, computer animation and interactive computer graphics applications such as video games.

Real-time rendering of sound from a source positioned arbitrarily in space around the listener, often referred to as “3D-sound”, has seen dynamic development in the recent years. Many approaches have been proposed leading to quite impressive results [4, 2, 6, 23, 3]. However, these methods are usually not sufficient to model the complex, environment-dependent phenomena resulting from sound propagation, diffusion or diffraction that are key components of *virtual sound fields*, addressed in this paper.

One of the principal elements of virtual sound rendering is the process of *auralization* or rendering a virtual sound field audible [20, 13]. Auralization involves digital filtering of a “rough”, anechoic sound signal by a filter or *pressure impulse response* proper to a given environment (Figure 1). The pressure impulse response associated with a source/receiver couple may be seen as the sound pressure signal received by the receiver when the source emits a single sound pressure “impulse” (Dirac signal). To compute this digital filter, different phenomena must be taken into account: sound emission, free sound propagation in the medium, occluders that cause reflection, diffusion and diffraction of waves and sound reception. If we assume all these phenomena to be linear, the transformations undergone by the original signal in the environment before it reaches the receiver can be expressed as a convolution product (in temporal domain):

$$y(t) = h(t) * x(t) = \int_0^{+\infty} h(\tau)x(t-\tau)d\tau,$$

where $y(t)$ is the received signal, $x(t)$ is the original signal and $h(t)$ is the pressure impulse response of the environment.

Formerly developed in acoustics to study sound reverberation in concert halls, various computer models are now used in auralization systems. They usually imply significant computational costs which restricts their use in interactive applications to very simple environments and/or ones for which basic sound propagation models are sufficient.

We present a new geometric approach to acoustic simulation. High quality sound fields, including effects of sound global specular and diffuse reflections are using a time-varying and hierarchical radiant exchange technique, similar to hierarchical radiosity used in lighting simulations [29].

2 Previous approaches

In this section we review some previous approaches aimed at generating virtual sound fields, calculating the pressure impulse response and integrating sound and image rendering in interactive graphics applications. We decided to leave it quite detailed to provide enough relevant acoustic background to the reader.

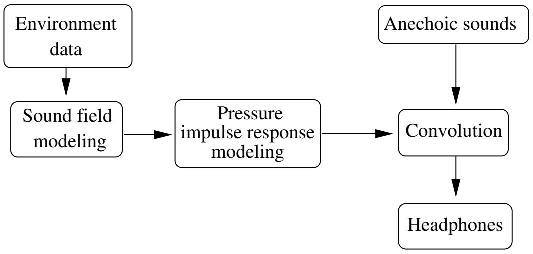

Figure 1: *Overview of the auralization process. The pressure impulse response of the virtual environment is simulated and used to filter a “rough” sound signal before listening.*

2.1 Sound field modeling techniques

Formerly measured on scale models of rooms [14], impulse responses are now calculated using computer models based on geometrical room acoustics [16, 20, 13, 28]. Due to the complexity of directly solving sound propagation equations based on the wave theory [16], several models have been developed using geometric approximation to sound propagation, i.e. representing sound waves by sound rays that propagate in the environment along straight lines. Sound sources and receivers are generally considered as punctual. In this context, two main groups of approaches currently in use are: *image sources* or ray/beam tracing [1, 5, 32, 12, 22, 26, 35, 24] (*wide band methods*) and Monte Carlo particle-tracing [31, 17] (*narrow band methods*).

Wide band methods

Due to the roughness of reflecting surfaces relative to sound wavelengths, in this group of methods sound reflections are assumed to be purely specular. Thus, sound reflected from a planar surface can be considered as sound emitted from a virtual *image source* constructed by mirroring the original source at the surface[2] (Figure 2a). Any order of reflection can be simulated by repeating the mirroring process for the image sources. The energy that arrives at the receiver at a given time is the sum of energies contributed from all sources and valid image sources.

The temporal distribution of reflections gives an “energetic impulse response” or *echogram* of the environment (Figure 3a). The pressure impulse response can be retrieved from the echogram by convolving each spike by a frequency-dependent filter that includes effects of wall attenuation, distance attenuation, source and microphone directivities.

One advantage of this approach is that the obtained energetic impulse response is exact (up to the simulated reflection order and within the reflection model validity). Besides, the pressure impulse response calculation correctly handles any appearing interferences. Moreover, updating the solution remains an easy task for dynamic

[2] Note that a similar method has also been used in computer graphics to add specular reflections to a radiosity solution [30].

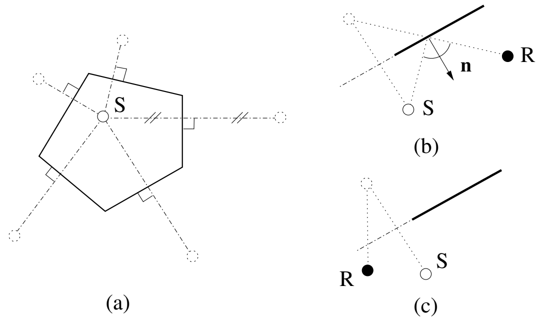

Figure 2: *Image sources method in 2D. Figure (a) shows a sound source (S) and 1^{st} order virtual sources for a pentagon-shaped enclosure. Figure (b) shows a valid virtual source relative to a listening position (M). Figure (c) shows an invalid virtual source since the reflected path between the virtual source and the listening position does not intersect the reflecting patch.*

environments and can be achieved in real-time for the first few orders of reflection [36, 9].

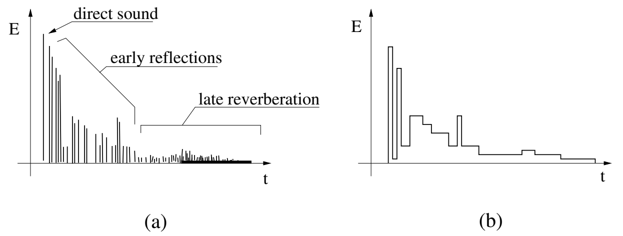

Figure 3: *Temporal distribution of reflections. (a)* Echogram *or "energetic impulse response" representing energy repartition through time as a list of energetic Dirac peaks sent from different sources and reaching the receiving location. Three main parts can be distinguished: direct sound (i.e. sound coming directly from the sound source that reaches receiving location first), early reflections and reverberation or reverberant tail. (b) An echogram integrated over short time-steps.*

However, the number of virtual sources grows exponentially ($N(N-1)^{i-1}$ for i-th order of reflection and N reflecting planes). Besides, not all image sources contribute to the solution (Figures 2b and 2c), calling for validity tests to be performed before summing contributions from virtual sources. The iterative mirroring process can be replaced by a direct construction of a set of valid virtual sources up to a given reflection order using ray/beam-tracing from a source [21, 22, 35] where a fixed number of rays/beams is fired in all direction from the sound source and virtual sources are constructed each time a ray hits a surface.

Another problem arises from the hypothesis on the pure specular sound reflection which rapidly tends to be invalid for complex materials (e.g. audience, often represented by a single surface with a high diffusion coefficient due to multiple reflections and diffraction over seats) or high frequencies [18, 11]. Pure diffuse reflections can be added to the image sources solution by applying statistical approaches as a postprocessing [12, 22, 35, 24]. An original combined ray-tracing/radiant exchanges technique has also been proposed in which a "stochastic" radiant exchange simulation is fed with non specular portion of the energy while a ray-tracing simulation is conducted [19].

Narrow band methods

The second family of methods uses Monte-Carlo particle tracing to determine all sound paths from the sound source to a given listening position [31, 17]. Particles are fired in all directions from the sound source and successive reflected particles (due to both specular and diffuse reflections) are fired until they reach a given "counting volume" (Figure 4). The energy brought to the receiver by a particle is marked in an echogram for the particle's frequency. Simulations are only conducted for a chosen set of frequency bands (usually octave bands) leading to one echogram per frequency band.

These echograms are integrated over short time steps (i.e. energy contribution of the particles are added together) (Figure 3b). For each time step, the spectrum of the pressure signal can be constructed by reading the energy corresponding to the time step in each integrated echogram. The pressure impulse response is obtained by calculating inverse Fourier transforms of these spectra. Since phase information is missing, usually random phase is used for the transform, which is physically disputable, but seems satisfactory from the psychoacoustic point of view [15, 17].

One drawback of this method lies in the need for the time versus quality trade-off in the choice of the sampling resolution. Also, due to the energetic aspect of the simulation, possible interferences are not taken into account during echogram integration leading to over-estimated energy values. Moreover, the method is very compute intensive and the solution is not easily updatable when changes occur in the environment.

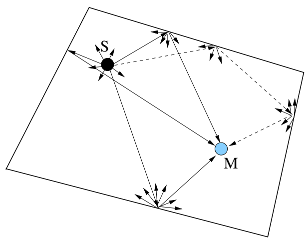

Figure 4: *Particle tracing method in 2D. Sound particles are traced in all directions from sound source position (S) and different propagation paths are collected when reaching a given listening volume (M). The dashed line denotes a path of double reflection.*

Late reverberation

As neither wide nor narrow band methods can simulate high orders of reflection, statistical room acoustics is usually used to simulate the late part of the response or *reverberant tail* (Figure 3a), leading to hybrid geometric/statistical algorithms. This part of the response, although important, does not carry fine psychoacoustic information since it consists of numerous overlapping echoes arriving from all directions. Thus, it is often modeled using colored noise with exponential decay [25].

3 Overview

We present an original sound simulation method which can take into account complex sound propagation phenomena, including global sound reflections through time and interferences. It can treat surfaces with hybrid ideal diffuse and ideal specular reflectance properties. It follows the narrow band approach to sound rendering but offers better control over the solution's quality. Moreover, the algorithm is based on a hierarchical approach. Thus, the resolution of the process can be easily tuned using proposed refinement oracles that direct the computational effort to regions where more detailed calculations are required.

4 Hierarchical sound radiosity

Hierarchical radiosity is widely used in lighting simulation to model global energy exchanges between surfaces [29]. Scene surfaces are subdivided into a hierarchy of patches. Energy transfer between patches is evaluated using some error criteria: if the quality is sufficient patches are linked and exchange energy; otherwise, they are subdivided and quality of the transfer is re-evaluated for the subpatches. Radiosity is stored as a pyramid in a hierarchy, the value on a given patch being the area average of the energy of its subpatches. The patches can thus exchange energy in a coherent way at different levels of the hierarchy.

Hierarchical radiosity presents a number of advantages for sound simulation. In particular:

- It is very well suited for simulating diffuse exchanges. However, although diffusion is an important aspect of sound propagation, sound reflectance functions usually have a strong specular component. Therefore, our approach explicitly takes into account specular reflections. In particular, directivity of first reflections is very important for correct sound imaging (i.e. perception of source positions in space), especially for complex receivers like human ears. We will use an integrated image-sources model to simulate global specular reflections.

- To handle interactive environment modifications, several approaches have been proposed to interactively recompute a radiosity lighting solution when sources, material attributes or geometry are modified [7, 10, 27, 8]. For sound simulation, one important example of scene modification is moving the receivers. Since information is stored on the surfaces, a single gathering operation is sufficient to update the solution. Considering the relatively low resolution of the ear, the costly "final gathering" operation used in radiosity systems can be avoided.

- Due to the wide extent of the audible sound frequency spectrum, a hierarchical approach is well suited to tune the computation resolution depending of the frequency. Of course, this implies a computation of a solution for several frequencies. In our method, simulation is simultaneously conducted for different frequencies, usually central frequencies of octave bands (see appendix B), with adaptive patch refinement. We present several oracles that can be used to drive the refinement.

The remainder of this section presents an adaptation of the hierarchical radiosity technique to time-varying radiant exchanges between surfaces for sound simulation. In particular we describe how to transport and store energy as a function of time and how to combine energy exchanged at different hierarchy levels.

4.1 Hierarchical time-varying radiant exchanges

Let us consider time-varying energy transport between diffuse reflecting surfaces. Radiance emitted from a surface at a given time in a given direction consists of the radiance due to self-emissivity of the surface and the sum of incoming radiances emitted from all visible surfaces.

More formally, the time-varying energy transport equation for a given wavelength λ can be written as[3]:

$$L^\lambda(x,t,\theta_0,\phi_0) = L_e^\lambda(x,t,\theta_0,\phi_0) + \frac{\rho_d^\lambda(x)}{\pi}\int_\Omega L_i^\lambda(x,t,\theta,\phi)\cos\theta\, d\omega, \qquad (1)$$

where

(θ_0,ϕ_0)	the direction given in spherical coordinates,
$L^\lambda(x,t,\theta_0,\phi_0)$	the total radiance emitted by a point x on a surface at time t,
$L_e^\lambda(x,t,\theta_0,\phi_0)$	the self-emitted radiance,
$\rho_d^\lambda(x)$	the diffuse reflectance at x, assumed constant in all directions (θ,ϕ),
$L_i^\lambda(x,t,\theta,\phi)$	the incident radiance arriving at x at time t from the direction (θ,ϕ),
Ω	the hemisphere of directions,
$d\omega$	a small solid angle.

This formulation, at first sight identical with the light transport equation used in light radiosity [29], differs from by its temporal aspect. Indeed, the time of energy transport cannot be neglected for sound waves since energy emitted from a point x takes a non-zero time to arrive at point y. We will call this time a *propagation delay*. The following considerations show how the propagation delay influences the energy transport for sound waves.

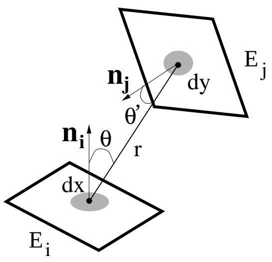

Figure 5: *Notation for energy exchange between two surfaces.*

Consider notation in Figure 5. The propagation delay between x and y can be calculated from the wave speed c, assumed constant over time and space (for sound approximately $334.1 m.s^{-1}$ [16]) and the distance between x and y: $t' = \frac{r}{c}$. Therefore, energy arriving at point x at time t from the point y was emitted from y at time $t - t'$. Equation (1) can be expressed in terms of sound *intensity*, analogical to light radiosity, as follows:

$$I^\lambda(x,t)dt = I_e^\lambda(x,t)dt + \rho_d^\lambda(x)\int_{y\in S} I^\lambda(y,t-t') \frac{\cos\theta\cos\theta'}{\pi r^2} A^\lambda(r) V^\lambda(x,y)\, dy\, dt, \qquad (2)$$

[3] Expressions superscripted with λ denote wavelength-dependent quantities. We shall maintain this notation in the remaining of this document.

where	(see also notations in Figure 5)
$I^\lambda(x,t)dt$	the total energy emitted from the point x at time t during a short emission time dt,
$I_e^\lambda(x,t)dt$	the self-emitted energy during dt,
$I^\lambda(y,t-t')$	the energy emitted from the point y at time $t-t'$,
$A^\lambda(r)$	a medium scattering coefficient,
$V^\lambda(x,y)$	a ¡¡visibility¿¿ term due to diffraction by occluders [33, 34],
S	a set of surfaces in the scene.

The expression under integral in equation (2) describes the intensity contribution from any point y for which the propagation delay to the point x is equal to t'. Taking the integral of this term over all surfaces S in the scene gives the total intensity that arrives at x at time t. This integral can be rewritten by sorting the points y by their propagation delay to x:

$$\int_{\tau=0}^{\infty}\int_{y\in S_\tau} I^\lambda(y,t-\tau)F^\lambda(x,y)dyd\tau, \tag{3}$$

where	
τ	the propagation delay between x and y,
S_τ	the set of points y with the propagation delay to x equal to τ,
$F^\lambda(x,y)$	a point-to-point "form factor" between x and y.

Let us consider now that surfaces are discretized into a set of polygonal patches P_i. We assume that the intensity is spatially constant over a patch and during the exchange:

$$I_i^\lambda(x,t)\equiv I_i^\lambda \qquad {I_e}_i^\lambda(x,t)\equiv {I_e}_i^\lambda.$$

Surfaces are of finite extend, we must take into account a duration of the exchange defined as the difference between the maximum $\tau_{\max}$ and the minimum $\tau_{\min}$ propagation delays between points on a pair of patches.

We can thus rewrite equation 3 as:

$$\begin{aligned} I_i(t) &= {I_e}_i(t) + \\ &= I_j^\lambda(t-T_{ij})\underbrace{\int_{\tau_{\min}}^{\tau_{\max}}\int_{y\in P_j} F_{ij}^\lambda(x,y)\delta(r-c\tau)dyd\tau}_{F_{ij}^\lambda(\tau)}, \end{aligned} \tag{4}$$

where $T_{ij}=\tau_{\min}$ is the minimum propagation delay between the two patches, $\delta(r-c\tau)$ is a Dirac distribution which is non zero only for pairs of points such as the propagation delay is equal to τ and $F_{ij}^\lambda(\tau)$ is the geometric form factor between P_i and P_j.

Thus, we can derive a discrete version of equation (4):

$$E_i^\lambda[t] = {E_e}_i^\lambda[t] + \rho_i^\lambda \sum_{j=1}^{N} F_{ij}^\lambda I_j^\lambda[t-T_{ij}], \tag{5}$$

where	square brackets indicate that the energy is calculated only for some discrete values of time,
$E_i^\lambda[t]$	the total energy emitted by patch P_i at time t,
${E_e}_i^\lambda[t]$	the energy self-emitted by patch P_i at t,
ρ_i^λ	the diffuse reflectance of patch P_i (constant),
F_{ij}^λ	the form factor between patches P_i and P_j,
T_{ij}	is the minimum propagation delay between between P_i and P_j,
$I_j^\lambda[t-T_{ij}]$	the intensity emitted by P_j at time $t-T_{ij}$.

Equation (5) describes time-varying energy exchanges between two patches P_i and P_j.

The three principal components of a time-varying energy exchange are then:

- A wavelength-dependent form factor F_{ij}^λ defined by:

$$F_{ij}^\lambda = \frac{1}{A_i}\int_{x\in P_i}\int_{y\in P_j}\frac{\cos\theta\cos\theta'}{\pi r^2}A^\lambda(r)V(x,y)dxdy$$

  computed as an average of point-to-polygon form factors for sampled points on Pj. Note that the visibility term is no longer marked as wavelength-dependent. Indeed, during the form factor calculation, additional evaluation of a frequency dependent visibility term would be too expensive. The effects of diffraction are most noticeable for direct sound which has not hit a surface at the time of form factor calculation. Thus, we assume the visibility to be standard light, geometric visibility and compute a diffraction term for direct sound only using an approach similar to [33, 34].

  Point to polygon form factors are also computed by sampling the receiving patch P_i. Taking the following definition for a point-to-point *delta* form factor $\delta F \equiv \frac{\cos\theta\cos\theta'}{\pi r^2}A^\lambda(r)V(x_n,y_m)$ we obtain:

$$F_{ij}^\lambda = \frac{1}{A_i}\sum_{n=1}^{N}\sum_{m=1}^{M}\delta F dx_n dy_m \tag{6}$$

- The minimum propagation delay T_{ij} computed for all couples of *visible* sample points on patch P_i and P_j :

$$T_{ij} = \min_{x\in P_i, y\in P_j}\frac{r}{c}\ , \tag{7}$$

  where r is the distance between x and y, mutually visible points respectively on P_i and P_j.

- The duration of the exchange τ_{ij} defined by:

$$\tau_{ij} = \max_{x\in P_i, y\in P_j}\frac{r}{c} - \min_{x\in P_i, y\in P_j}\frac{r}{c}. \tag{8}$$

Two linked patches exchange energy in an iterative process, each iteration corresponding to an order of reflection. At each iteration oracles are used to evaluate the quality of the energy transfer. If quality constraints are not met, patches are subdivided and links are refined. The process is iterated until it reaches a given order of reflection or until overall energy gathered by all surfaces has decayed by more that 60 dB relative to the energy originally sent by the source (which is the criterion commonly used in room acoustics to measure the duration of the reverberation).

4.2 Echos: representing sound intensity through time

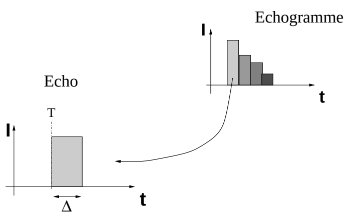

Figure 6: *Representing sound intensity through time.*

Equation 5 provides a discrete space/time formulation for time-varying radiant exchanges. An elementary information that is going to be exchanged between patches is an "echo". It can be seen as an energy distribution during a time interval corresponding to the duration of the exchange and modified by each reflection. An echo arriving on a patch is thus described by three quantities (Figure 6):

- an arrival time, T
- an intensity, I
- a width, Δ

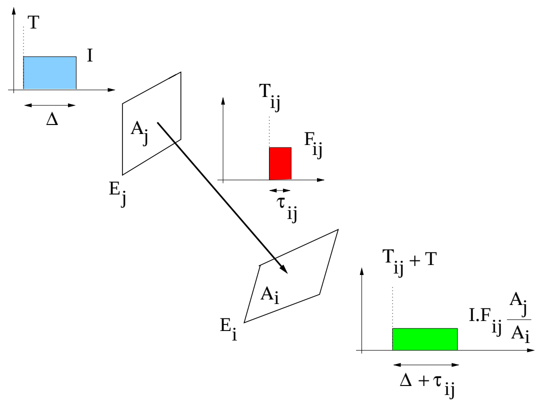

Figure 7: *The temporal energy transport operator. A_i and A_j are the areas of patches P_i and P_j.*

At the beginning of the process punctual sound sources emit a single Dirac echo of unit energy and zero width. These echos are then propagated from patch to patch using the transport operator, defined in equation (5) and graphically illustrated in Figure 7. Each reflection modifies the echo: the energy of the received echo becomes equal to the energy of the shot echo multiplied by the form factor; its arrival time is augmented by a patch-to-patch propagation delay; and its width is extended by the exchange duration.

Equation (5) defines discrete energy exchanges between surfaces. In order to represent the energy received/emitted by a surface through time, we store a list of echos on each patch in echograms as presented in section 2.1. Each patch stores relevant information in three echograms:

- a shooting echogram which stores echos to be shot during the current iteration,
- a gathering echogram in which gathered echos will be stored before the "push-pull", and
- an overall echogram storing all echos gathered through time by the patch *at its level in the hierarchy only*. Note that the energy is not stored in a pyramid to avoid memory overhead.

As in standard hierarchical radiosity algorithms, energy received by a father patch must be "pushed" to its children and then "pulled" to shoot coherent values of energy during the next iteration. This process is similar to the classic "push-pull" process of hierarchical light radiosity. However, due to the temporal nature of the exchange, one extra detail must be taken into account. When an echo has been received by a father patch, its propagation time has to be corrected before it can be reshot by the children to take into account the difference in propagation time between the shooter-to-father path and the shooter-to-child path (T_1 and T_2 in Figure 8). In order to compute this correction, each echo stores a value identifying its shooter. Then, an approximate correction is computed by evaluating the difference between propagation time from the shooter patch to children patches centroids and to the father patch centroid. A similar time correction must be used when refining a shooting patch. The correction can be performed at the moment of the shooting/gathering operation.

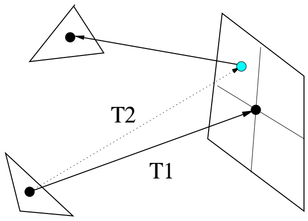

Figure 8: *Handling hierarchical propagation delays. Energy received at time $T1$ by a patch should not be reshot at time $T1$ from its children. A time correction $T2 - T1$ is added before reshooting the echo from the children.*

In order to accelerate all inserting/deleting echos operations, echograms are structured as binary trees of echoes, sorted by their arrival time on the patch.

4.3 Space-frequency refinement

We propose to use two *oracles* to control the spatial subdivision of surfaces in the scene. Since they depend on data that is itself wavelength-dependent (e.g. form factors), they lead to different subdivisions at different frequencies.

- A standard BF_ϵ oracle which refines patches that receive the highest energy.
- A consistency oracle T_ϵ that guarantees that the hypothesis on the surfaces shooting energy in phase is valid (and so is the energetic formulation). It chooses for refinement the patches for which the exchange duration is high. Figure 9 shows an example of such a subdivision for different frequency bands. The subdividing criterion is:

$$\tau_{ij} > T_\epsilon \frac{c}{\lambda}$$

Note that at high frequencies, this oracle tends to be ineffective, calling for the use of time corrections as presented in the previous section.

4.4 Echo Fusions

In order to handle the overwhelming number of echoes and be able to simulate quickly higher orders of reflection we designed an echo merging mechanism such that several echoes can be combined in a new one. The main problem in merging echos together is to guarantee that possible interferences are taken into account in order not to over estimate the energy. Two cases of echo merging can occur: coherent echos and incoherent echos merging.

Coherent echo merging occurs when two (or more) echos reach a surface coming from the same shooter with an arrival time difference *a priori* inferior from the period $T = c/\lambda$ of the signal. In this case we can reconstruct the two complex pressure distributions associated with each echo and sum them up to get a new echo.

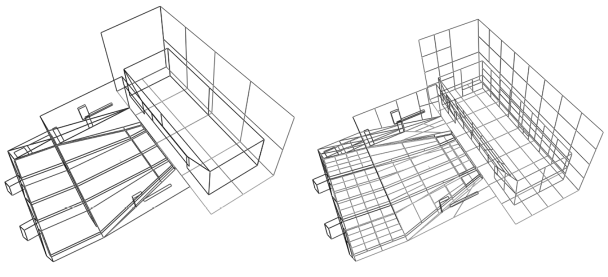

(a) (b)

Figure 9: *Example of frequency refinement on T_ϵ oracle at (a) 50 Hz, (b) 100 Hz. Value of the oracle is 1.0*

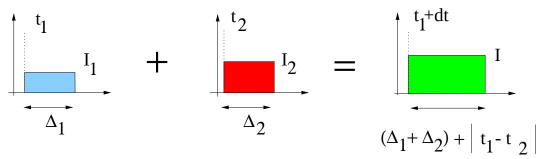

Figure 10: *Summing echos together*

Each echo represents a complex pressure distribution through time $P = Ee^{i\omega t}$, where E is the square root of the intensity of the echo [4]. We can thus write:

$$P_1 = E_1 e^{i\omega t_1}, P_2 = E_2 e^{i\omega t_2},$$

where $\omega = \frac{2\pi c}{\lambda}$ and

$$P_1 + P_2 = P_1(1 + \frac{P_2}{P_1}) = Ee^{(\omega(t_1+dt))},$$

where

$$E = \sqrt{E_1{}^2 + E_2{}^2 + E_1 E_2 \cos(\omega(t_2 - t_1))}$$

and

$$dt = \frac{1}{\omega} atan(\frac{E_2 \sin(\omega(t_2 - t_1))}{E_1 + E_2 \cos(\omega(t_2 - t_1))}).$$

We use a parameter m to merge coherent echoes that reach a surface within a time-step defined by $\delta = mT$. Effects of coherent echo merging efficiency and energy conservation are illustrated on figures 11 and 12.

4.5 Handling specular reflections

Surface reflectances for sound usually have a strong specular component. Thus, in order to be able to simulate this essential specular term we propose to extend our method to take into account combined ideal specular and ideal diffuse reflectors. To achieve this goal, we use an image-source like approach. We extend the data structure used to represent the echos in order to store a point of origin. When echos are gathered from a specular surface, we use this origin point to construct a mirrored image according to the plane of the shooting patch (see Figure 13). A point-to-element extended

[4] intensity is proportional to the mean square pressure over time: $I = \overline{p^2}/\rho_0 c$ where $\rho_0 c$ is the characteristic impedance of the medium ($414 kg.m^{-2}.s^{-1}$ for air)

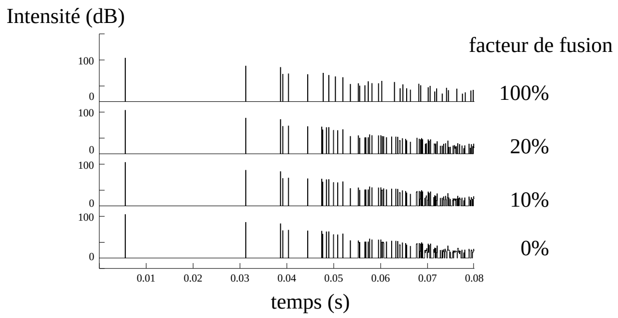

Figure 11: *Comparison of four echograms resulting from simulations with different coherent echo merging factors. Note that the energy does not tend to be overestimated by the merging process.*

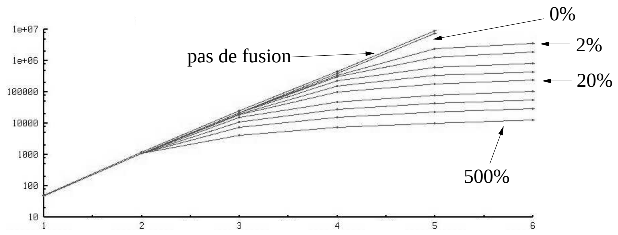

Figure 12: *Comparison of the number of echoes function of the order of reflection for different values of the coherent echo merging threshold m. From 0% (top curve) to 500% (bottom curve).*

form factor is computed and a new echo is transfered to the gathering patch. To simplify the process, when echos come from surfacic elements we use the centroid of the element as the point of origin. This operation is similar to a “three-point transport”. Of course if the specular reflector also has a diffuse component, a “standard” echo is also exchange using the extended form factor between the source and gathering patch. This approach allows to maintain special echos which have zero duration and represent pure specular paths of reflection, which are very important since they will be responsible for the strongest reflections. We do not store in that case the intensity of the echo but the product of all reflection coefficients along the path. Since we also store the propagation time we can thus re-compute the intensity emitted by the equivalent point sound source.

This source-image approach is conducted with each element. If elements are sub-patches of the same surface, several copies of the same source-image will be generated in the case of pure specular paths. We can however eliminate those duplicates when performing the push-pull operation by specifying that echos of zero duration and same origin point can be merged together.

5 Implementating the approach

As in common hierarchical radiosity approaches, there are basically two ways of implementating the process. The approach we chose is to refine using the energetic oracle. Since the energy is decreasing with time, resolution of the process gets more and more coarse. This allows us to get fine results for beginning times of the simulation and more coarse, but fast ones for the late reverberation. Implementation using this approach is also very straightforward.

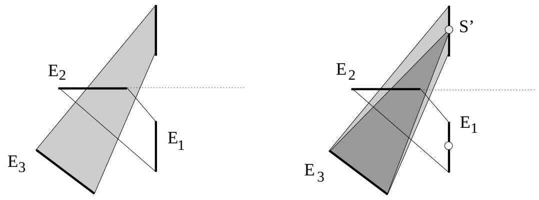

(a) (b)

Figure 13: *Using the center of the element to build an image-source for specular exchanges. The center of the origin element E_1 of the echo is mirrored according to the source element E_2. An extended form factor is then computed between the obtained point (S') and the receiving element E_3.*

Moreover, since it is very time consuming to perform the push-pull operation on every element, we propose to maintain an integral value of all the intensity that has been exchanges during the time of one iteration. This value is computed as the sum of the energies of the echos for each echogram. We perform a standard push-pull operation on those values and use them to evaluate the refinement oracles. When the appropriate level of exchange has been selected, the push-pull operation is performed on the shooting echograms of the shooter's hierarchy.

However, the major drawback is that the refinement operator does not use a global estimate of the solution. Another possibility is then to perform a complete simulation until convergence using links of highest level in the hierarchy. The obtained solution is then coarse but covers the complete duration of the simulation. If a refinement operation is performed at this time, the oracle can better estimate the links through which energy has most transited over time. But the drawback then is that all the echoes that were transmitted through the old links must be canceled before being reshot through the new ones. Once again, the cancellation operation can be quite costly since it involves to reshoot each echo with a negative intensity value.

6 Application to acoustic simulations and data visualization

In this section we present examples of acoustic simulations and acoustic data visualization aimed at validating our simulation technique in the case of concert hall design applications. A few approaches have already been presented to visualize acoustic data [31, 24]. They are based on ray-tracing techniques which seems less suitable for visualization since no direct information is available on surfaces. Thus sampling of the surfaces must be done as a postprocessing [24]. In our approach, since surfaces are discretized for the radiosity process, visualization of relevant acoustic data can be easily integrated in the simulation. By displaying energy that transited on the links designer can easily check for problems in the hall design (lack of lateral energy for example) (see figure 14). Once the time-varying radiosity solution is computed, an animated rendering of energy reaching surfaces through time can also be displayed in real-time to analyze the wave front propagation (see figure 16). Figures 14 and 15 present other examples of visualization in an acoustic simulation context.

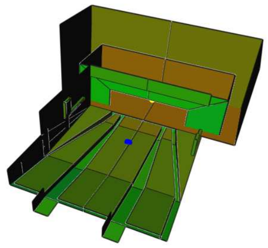

(a)

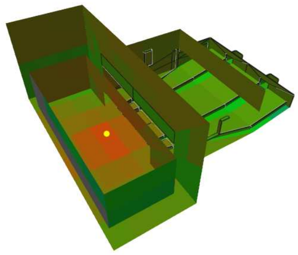

(b)

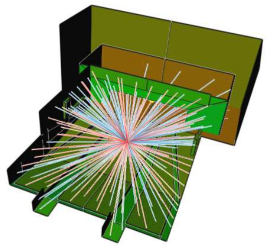

(c)

Figure 14: *Visualization of acoustic simulation data. Sound source is displayed as yellow spheres, microphones as blue ones. Figure (a) and (b) show overall energy received by the surfaces through time. (c) is a display of the overall energy that transited through links reaching the microphone. This kind of visualization may allow the designer to detect problems related with spatial distribution of the energy.*

7 Conclusion

We presented a new solution to acoustic simulations based on a radiant exchange approach adapted to take into account time-varying phenomena. This approach is well suited to simulate global diffuse reflections and has been modified to handle specular reflections which are of primary importance for sound. Moreover, we presented an extension that allows for tuning the temporal complexity of the process while taking into account interferences. The resulting solution is stored on the surfaces of the environment. Thus, it is *quasi* listener-independent. Only a final gathering step is needed when the receiving point is moving, which offers the possibility to use this method for interactive walkthroughs. The solution also provide interesting ways of visualizing spatial and time-varying phenomena.

Of course, many points could be improved in our method. Extensions include further acoustic validation of the method. Taking into account a "phase" factor in the form factors could significantly im-

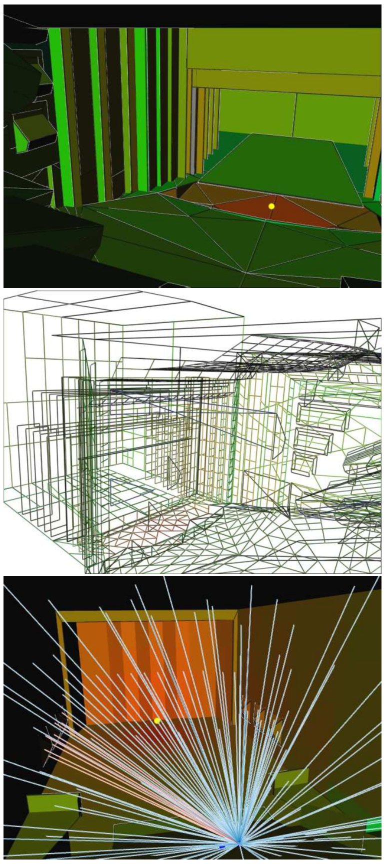

Figure 15: *Visualization of acoustic simulation data in the Opera de la Bastille in Paris and in another smaller hall.*

prove the solution if patches large compared to the wavelength are used. We are also thinking of extending the hierarchical approach to take into account clusters of surfaces and perform incremental update of the solution. We could also think of separating the calculation of the images-sources of the energy exchange process. Since we recompute the image-sources each time we exchange the echos, we should improve a lot the efficiency of our approach.

A Surface reflectance for sound

Usually, wide band data of complex material absorption filters are not available, due to the difficulty of the measurements. In practice the most common coefficients are Sabine's diffuse field absorption coefficients α^λ given in octave bands. Thus, reflected energy coefficient is $(1-\alpha^\lambda)$. Another coefficient β^λ defines the proportion of energy which is diffusely reflected. Thus diffuse reflection coefficient is $(1-\alpha^\lambda)\beta^\lambda$ and the specular reflection coefficient is $(1-\alpha^\lambda)(1-\beta^\lambda)$. Below, we present a few examples of reflectance coefficients measured in a real concert hall. Note the high diffusion aspect of some materials:

Material	125 Hz	250 Hz	500 Hz	1 kHz	2 kHz	4 kHz
Audience (seats)	0.30	0.45	0.50	0.60	0.65	0.70
Back wall (carpet)	0.10	0.10	0.10	0.09	0.09	0.07
Ceiling (stucco)	0.25	0.30	0.30	0.35	0.35	0.40
Stage (parquet floor)	0.05	0.05	0.04	0.03	0.01	0.01

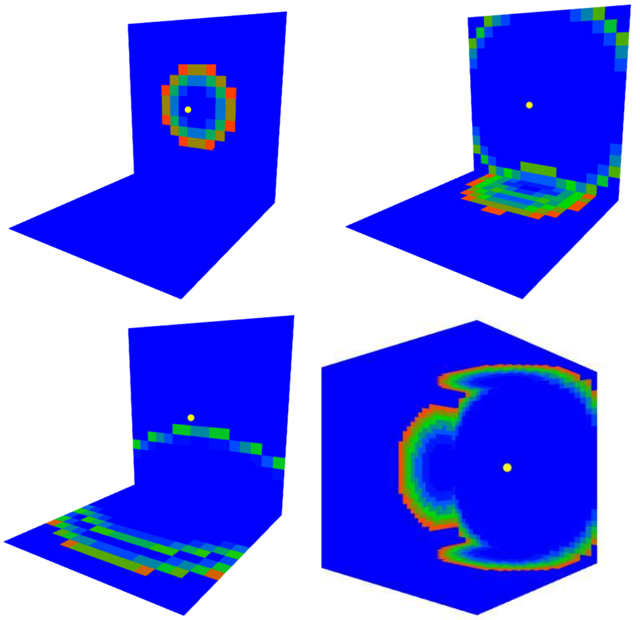

Figure 16: *Visualization of a wavefront propagation with one order of reflection by displaying our radiosity solution through time. Note the two secondary wavefronts due to pure specular reflections on walls on the second and third picture. The radiosity solutions were obtained with 685 elements and 132610 links for the two faces and 8193 elements/6188 links for the cube.*

B Digital signal processing issues

Reconstructing the pressure impulse response

Remember that our simulation process is a narrow band approach. At the end of the simulation our solution consists of several echograms for different octave bands. Reconstructing the pressure impulse response involves three steps:

- For each echogram, reconstruct the pressure data stored in each echo to obtain a one wide band pressure impulse response. Retrieving the wide band pressure information stored in an echo can be achieved simply by replacing the echo by a given time "window function" defined on the width δ of the echo and whose integral on this time-step is the square root of the energy of the echo (remember that the energy is proportional to the square root of the pressure). In our current implementation, we use a window function f_δ defined by:

  $$f_\delta(x) = 1 - \sin\left(\frac{2\pi x}{\delta}\right),$$

  for $x \in [-\pi, \pi]$.

  The pressure impulse response is then given by:

  $$P(t) = \sum_i \frac{\sqrt(E_i)}{\delta_i} f_{\delta_i}(t - \bar{t}_i)$$

- Filter the obtained wide band pressure impulse responses by the corresponding band pass filter to keep only the valid part of the response (remember that all values in echos were computed for a specific frequency band).

- Add the obtained “narrow band” pressure impulse responses together to get the final wide band pressure impulse response of the environment.

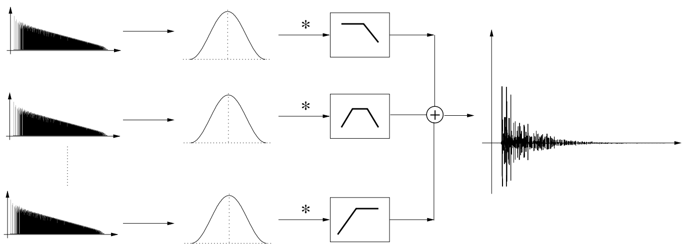

Figure 17: *Converting echograms in octave bands to a pressure impulse response.*

References

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

[2] D.R. Begault. Challenges to the successful implementation of 3D sound. *Journal of the Audio Engineering Society*, 39(11):864–870, November 1991.

[3] Durand R. Begault. *3D Sound for Virtual Reality and Multimedia*. Academic Press Professional, 1994.

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

[5] J. Borish. Extension of the image model to arbitrary polyhedra. *Journal of the Acoustical Society of America*, 75(6), 1984.

[6] D.A. Burgess. Real-time audio spatialization with inexpensive hardware. Technical Report GIT-GVU-92-20, Georgia Institute of Technology, 1992.

[7] Shenchang E. Chen. Incremental radiosity: an extension of progressive radiosity to an interactive image synthesis system. *ACM Computer Graphics, SIGGRAPH'90 Proceedings*, 24(4):135–144, August 1990.

[8] G. Drettakis and F.X. Sillion. Interactive update of global illumination using a line-space hierarchy. *ACM Computer Graphics, SIGGRAPH'97 Proceedings*, pages 57–64, August 1997.

[9] S.H. Foster, E.M. Wenzel, and R.M. Taylor. Real-time synthesis of complex environments. *Proc. of the ASSP (IEEE) Workshop on Application of Signal Processing to Audio and Acoustics*, 1991.

[10] D.W. George, F.X. Sillion, and D.P. Greenberg. Radiosity redistribution for dynamic environments. *IEEE Computer Graphics and Applications*, pages 26–34, July 1990.

[11] C.H. Haan and F.R. Fricke. An evaluation of the importance of surface diffusivity in concert halls. *Applied Acoustics*, 51(1):53–69, 1997.

[12] R. Heinz. Binaural room simulation based on an image source model with addition of statistical methods to include the diffuse sound scattering of walls and to predict the reverberant tail. *Applied Acoustics*, 38:145–159, 1993.

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

[14] M. Kleiner, R. Orlowski, and J. Kirszenstein. A comparison between results from a physical scale model and a computer image source model for architectural acoustics. *Applied Acoustics*, 38:245–265, 1993.

[15] H. Kuttruff. On the audibility of phase distorsions in rooms and its significance for sound reproduciton and digital simulation in room acoustics. *Acustica*, 74(1), 1991.

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

[17] K.H. Kuttruff. Auralization of impulse responses modeled on the basis of ray-tracing results. *Journal of the Audio Engineering Society*, 41(11):876–880, November 1993.

[18] Y.W. Lam. The dependance of diffusion parameters in a room acoustics prediction model on auditorium sizes and shapes. *Journal of the Acoustical Society of America*, 100(4):2193–2203, October 1996.

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

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

[21] T. Lewers. A combined beam tracing and radiant exchange computer model of room acoustics. *Applied Acoustics*, 38, 1993.

[22] J. Martin, D. van Maercke, and J.P. Vian. Binaural simulation of concert halls : A new approach for the binaural reverberation process. *Journal of the Acoustical Society of America*, 94:3255–3263, December 1993.

[23] Henrik Møller. Fundamentals of binaural technology. *Applied Acoustics*, 36:171–218, 1992.

[24] M. Monks, B.M. Oh, and J. Dorsey. Acoustic simulation and visualisation using a new unified beam tracing and image source approach. *Proc. Audio Engineering Society Convention*, 1996.

[25] J.A. Moorer. About this reverberation business. *Computer Music Journal*, 23(2), 1979.

[26] J.M. Naylor. Odeon - another hybrid room acoustical model. *Applied Acoustics*, 38(1):131–143, 1993.

[27] J. Nimeroff, J. Dorsey, and H. Rushmeier. Implementation and analysis of an image-based global illumination framework for animated environments. *IEEE Transactions on Visualization and Computer Graphics*, 2(4), December 1996.

[28] J.D. Polack. Playing billiards in the concert hall: The mathematical foundations of geometrical room acoustics. *Applied Acoustics*, 38:235–244, 1993.

[29] François X. Sillion and C. Puech. *Radiosity and Global Illumination*. Morgan Kaufmann Publishers inc., 1994.

[30] F.X. Sillion and C. Puech. A general two pass method integrating specular and diffuse reflection. *ACM Computer Graphics, SIGGRAPH'89 Proceedings*, 23(3), July 1989.

[31] Adam Stettner and Donald P. Greenberg. Computer graphics visualisation for acoustic simulation. *ACM Computer Graphics, SIGGRAPH'89 Proceedings*, 23(3):195–206, July 1989.

[32] Tapio Takala and James Hahn. Sound rendering. *ACM Computer Graphics, SIGGRAPH'92 Proceedings*, 28(2), July 1992.

[33] Nicolas Tsingos and Jean-Dominique Gascuel. Soundtracks for computer animation: sound rendering in dynamic environments with occlusions. *Proc. of Graphics Interface'97*, 1997.

[34] Nicolas Tsingos and Jean-Dominique Gascuel. Fast rendering of sound occlusion and diffraction effects for virtual acoustic environments. *Proc. 104th Audio Engineering Society Convention*, May 1998.

[35] D. van Maercke and J. Martin. The prediction of echograms and impulse responses within the Epidaure software. *Applied Acoustics*, 38(1):93–114, 1993.

[36] E.M. Wenzel and S.H. Foster. Realtime digital synthesis of virtual acoustic environments. *Computer Graphics (Proc. ACM Symposium on Interactive 3D Computer Graphics)*, 24(2):139–140, March 1990.