

Instant Sound Scattering

N. Tsingos¹, C. Dachsbacher¹, S. Lefebvre¹ and M. Dellepiane²

¹REVES-INRIA, Sophia Antipolis, France

²ISTI-CNR, Pisa, Italy

Abstract

Real-time sound rendering engines often render occlusion and early sound reflection effects using geometrical techniques such as ray or beam tracing. They can only achieve interactive rendering for environments of low local complexity resulting in crude effects which can degrade the sense of immersion. However, surface detail or complex dynamic geometry has a strong influence on sound propagation and the resulting auditory perception. This paper focuses on high-quality modeling of first-order sound scattering. Based on a surface-integral formulation and the Kirchhoff approximation, we propose an efficient evaluation of scattering effects, including both diffraction and reflection, that leverages programmable graphics hardware for dense sampling of complex surfaces. We evaluate possible surface simplification techniques and show that combined normal and displacement maps can be successfully used for audio scattering calculations. We present an auralization framework that can render scattering effects interactively thus providing a more compelling experience. We demonstrate that, while only considering first order phenomena, our approach can provide realistic results for a number of practical interactive applications. It can also process highly detailed models containing millions of unorganized triangles in minutes, generating high-quality scattering filters. Resulting simulations compare well with on-site recordings showing that the Kirchhoff approximation can be used for complex scattering problems.

Categories and Subject Descriptors (according to ACM CCS):
I.3.3 [Computer Graphics]: Line and Curve Generation

1. Introduction

Proper modeling of sound propagation is very important for virtual acoustics and virtual reality applications. While virtual acoustics has made enormous progress in the recent years, the environments that can be treated interactively remain quite simple. This is mostly due to the complexity of sound propagation phenomena, in particular the modeling of reflection and diffraction effects. Late reverberation is traditionally rendered by means of statistical techniques or “artificial reverberators” while early scattering is computed using geometrical approaches [FJT02]. Such approaches can model early specular reflection (and possibly edge diffraction) interactively [FCE*98, TFNC01] but are limited to environments of low local complexity. Furthermore, as geometrical acoustics (GA) approaches are valid only for surfaces large compared to the wavelength, it is unclear whether increasing geometrical complexity will provide more accu-

rate results. Hence, surface detail is generally non-existent in GA simulations while it has a strong influence on sound propagation due to the induced scattering. This results in unrealistic effects.

In this paper, we focus on high-quality modeling of first-order sound scattering off complex surfaces using an extended geometrical approach. Our approach handles both diffraction and reflection phenomena in a unified way. Our contributions are:

- The development of an approach to compute the scattering from detailed, dynamic geometry which maps well to programmable graphics hardware (GPUs). This approach is based on a surface integral formulation and uses the Kirchhoff approximation, widely used in acoustics. In particular, we use a *source-view* approach to efficiently find visible surfaces from the source and “mip-mapping” to integrate contributions of all surface samples.
- A level-of-detail approach which reduces the geometry processing for audio rendering while preserving the scat-




Figure 1: *Left: We present an approach to compute sound scattering effects for interactive applications, including diffraction and direct reflections. Center: Our approach uses the GPU to evaluate a scattering integral on all surfaces visible from a sound source. Right: A 13-million triangle model of the scanned façade of a cathedral and close-ups on surface detail. Our approach can compute perceptually accurate scattering effects over the full audible bandwidth in seconds on such complex, unstructured, models that would be intractable otherwise.*

tering behaviour of complex surfaces. This is possible thanks to our use of graphics hardware. We believe that this paradigm shift for geometrical acoustics is of prime importance for interactive applications.

- An interactive rendering framework for realtime simulation of audio scattering effects which greatly enhance the realism of virtual environments.
- An evaluation of the Kirchhoff approximation for auralization of complex off-line simulations.

While only considering first order phenomena, our approach can provide realistic results for a number of practical interactive applications. It can also handle highly detailed models containing millions of triangles. For such models, we demonstrate that our approach can generate full-bandwidth scattering filters usable for high-quality auralization in minutes. To our knowledge, no other existing approach could be used in practice to process such complex environments.

2. Related Work

A wide variety of methods exist to solve acoustic scattering problems including analytic expansions, approximations and numerical methods [FHLB99]. In the following section we review the techniques most directly related to our work.

Geometrical theories:

Geometrical acoustics (GA) is probably the most widespread approach for interactive acoustic modeling. GA is a high-frequency approximation that models sound propagation along ray-paths. The propagation paths can be constructed using techniques such as ray or beam tracing [LSVA07, FCE*98] using spatial data structures to maximize efficiency. For complex dynamic environments, the cost of updating these structures can become in itself a significant bottleneck. In interactive GA applications, the reflection of sound rays is usually modeled purely specular assuming the size of the surfaces is large compared to the wavelength,

which is often not the case in practice. In off-line acoustical simulations, additional lambertian or “glossy” reflections [ISO04,CR05,ZCR06,CDD*06], have been classically used to approximate scattering off complex surfaces, which are then replaced by a flat proxy geometry. However, little work has been devoted to exploring level-of-detail (LOD) approaches in the context of GA [JMT03, WRR04, Sil05], especially how to design general simplification schemes that preserve the correct scattering properties. The framework introduced in this paper brings a possible solution to this issue.

Interactive GA simulations can also be enhanced by introducing diffraction effects from wedges [MPM90] in order to avoid audible discontinuities when the source or a strong specular reflection path becomes occluded from the receiver [TFNC01]. As all GA models, the geometrical theory of diffraction (GTD) assumes edges to be large compared to the wavelength. Increasing geometrical complexity would imply using smaller primitives and eventually would fall outside the validity domain of GA. Experimental studies have shown that low-resolution models might be more appropriate to evaluate acoustical criteria with GA approaches [RSCG99], which is somewhat counter-intuitive. Thus, it is unclear how classical GA+GTD approaches could apply to more realistic scenes, for instance highly detailed models from CAD-CAM or acquired through scanning techniques.

Other approaches, such as the Biot-Tolstoy-Medwin model [SFV99,KN00,CS07], follow a Huygens-Fresnel formalism which states that a wavefront can be seen as a superimposition of many elementary “wavelets”. The impulse response due to scattering from finite-sized wedges can then be accurately computed by integrating small contributions along the edges. A key aspect of this approach is to directly construct the response in time-domain. However, due to its computational complexity it has found limited use in interactive applications.

Finite and boundary element methods:

Finite element methods (FEM) are numerical solutions to the wave (Helmholtz) equation and associated boundary conditions. They are classically solved in frequency domain by subdividing the environment into small elements (voxels) but alternative time-domain formulations can also be used [SRT94]. Of special interest is the Green surface integral formulation. Using this formulation, the pressure solution $P(R)$ to the Helmholtz equation can be obtained using an arbitrary surface Σ surrounding the receiver R [FHLB99]:

$$P(R) = P_0(R) - \int_{\Sigma} (P(U)\nabla G(U,R) - G(U,R)\nabla P(U)) \cdot \mathbf{dS}, \tag{1}$$

where $\mathbf{dS} = \mathbf{n}dS$ (\mathbf{n} unit vector) and $G(U,R) = -e^{ikr}/4\pi r$ is the Green function corresponding to the propagation of a spherical wavefront in free-field (see Figure 2).

$P_0(R)$ is the free-field pressure emitted by the source and the integral term, often called *diffracted field*, is the solution of the homogeneous Helmholtz equation associated with the boundary conditions on the surfaces of the environment. This surface integral is also called the *Helmholtz-Kirchhoff integral theorem*. The Green surface formulation serves as a basis for boundary element methods (BEM) techniques. In this case, Σ corresponds to the surfaces of the environment.

FEM/BEM methods can account for full scattering effects in a unified way and are widely used to compute off-line reference solutions. They are not well suited to interactive applications except for very low frequencies since they require a dense subdivision of space or tessellation of surfaces to properly account for interferences. In the case of BEM, edge-length for surface elements must typically be smaller than $1/4th$ of the wavelength. This makes such approaches very time-consuming at high frequencies or for large-scale problems. They also require carefully-designed meshes with uniform elements to limit errors in the solution. They are thus hard to use with most 3D models, particularly those acquired with scanning or designed by CG artists.




Figure 2: Notations for the Kirchhoff-Helmholtz integral theorem. S and R denote the source resp. receiver.

Kirchhoff approximation:

The Kirchhoff approximation (KA) can be seen as a hybrid strategy between GA and wave acoustics [FHLB99]. It is

based on Eq. 1 but imposes $P(U) = P_o(U)$ and $\nabla P(U) = \nabla P_o(U)$ on the “illuminated” side of the surfaces (visible to the source) and $P(U) = \nabla P(U) = 0$ on the “shadowed” side. As a result, $P(R)$ can be computed by a direct integration but surfaces facing away from the sound source will not contribute to the solution. However, enforcing the value of both P and ∇P on the same surface is not rigorous for the Helmholtz equation. Hence the obtained result is not a true solution to the wave equation and is generally limited to first-order scattering.

Neglecting the contribution from occluded surfaces and higher-order scattering is the major source of error in the KA. As a result, the approximation will degrade at very low frequencies as the scattered component becomes more important in the regions occluded from the source. The KA will also lead to inaccurate results when second-order occlusions/reflections become prominent. Several studies have shown that it can introduce significant errors in the computed scattering from simple flat or randomly-rough surfaces when compared to reference solutions [JM82, Tho87, NNK93, CL93]. In particular, errors were found to be more important at grazing angles and near-field from the surface. However, some of these studies also used additional far field approximations which might also contribute to the observed errors. As stated in [Tho87, EDS01], further work is still required to evaluate the validity of the KA which is still not well established even today.

Despite these limitations, the KA has been widely used to solve off-line scattering/occlusion problems in acoustics [SN81, CI90, CL93, TG98, Emb00, EDS01]. In [SN81] it was shown that the KA can be used to simulate impulse response of first-order reflection/diffraction off rigid panels with good agreement to measured responses. Due to its surface integration and interference treatment, it captures phenomena that cannot generally be modeled using classical GA approaches, providing a continuous first-order sound-field and extended validity range. In this paper, we will also be using the KA and extend the previous work to show that it can be efficiently used to compute very convincing scattering effects off arbitrary complex surfaces.

3. Overview

In Section 4 we show that the Helmholtz-Kirchhoff theorem combined with the Kirchhoff approximation can be used to derive an expression for first-order scattering effects off complex surfaces that compares well with BEM simulations. This formulation is well suited to an implementation using graphics hardware that shares some similarity with the *reflective shadow map* [DS05], introduced to compute interactive first order global illumination effects. Using graphics hardware for audio rendering also opens the opportunity for leveraging classic LOD techniques such as bump or displacement mapping, which we evaluate in Section 5. In Section 6 we present an interactive framework for real-time au-

realization of complex scattering effects. The throughput of our hardware-accelerated approach makes it well suited for calculating impulse responses of complex architectural environments which can be used for off-line auralization. In this context, Section 7 validates the Kirchhoff approximation against real-life recordings. Finally, in Section 8 we discuss the current limitations of our approach before concluding.




Figure 3: Surfaces are sampled using hardware rendering from the point of view of the sound source. We evaluate the scattering terms at each pixel before global integration through mip-mapping. In this figure, S and R denote the source resp. receiver.

4. Scattering from Detailed Geometry

We first derive an expression for first-order scattering off a purely reflecting surface using the Kirchhoff approximation. We assume a steady state harmonic point sound source of frequency f . The free-field pressure created by the source at a listening point R at distance d is given by $P_0(R) = P_o/d e^{ikd}$, where P_o is the pressure amplitude of the source, $k = 2\pi f/c$ and c is the speed of sound ($\approx 340 \text{ m}\cdot\text{s}^{-1}$). Using the Kirchhoff approximation and following the conventions of Figure 3, the integrand in Eq. 1 becomes:

$$p(R) = -\frac{P_o e^{ik(\rho+r)}}{4\pi\rho r} \left(\left(ik - \frac{1}{\rho} \right) (\mathbf{n} \cdot \mathbf{v}) + \left(ik - \frac{1}{r} \right) (\mathbf{n} \cdot \mathbf{r}) \right),$$

where \mathbf{r} and \mathbf{v} are unit vectors resp. from the source to the surface-sample and from the surface-sample to the receiver. For additional details, see [Hec98].

Note that this equation does not assume sources or receiver to be in far-field of the scattering surfaces. The diffracted contribution, P_{diff} , received by the listener can thus be interpreted as the difference between the free-field contribution and the portions of the wavefront blocked by the scattering surface Σ :

$$P_{diff}(R) = P_0(R) - \int_{\Sigma} p(R) dS. \quad (2)$$

Following the same interpretation, the purely reflected contribution from Σ can be expressed using an integrand $p'(R)$, similar to the one appearing in Eq. 2. In this case, \mathbf{v} is replaced by its mirror-image at the tangent plane. The total scattered pressure, $P(R)$, at the location of the receiver can

thus be computed as the sum of the three components, free-field direct, diffracted and reflected as (see Fig. 4 for an illustration):

$$\begin{aligned} P(R) &= P_0(R) + \int_{\Sigma} (-p(R) + p'(R)) dS, \\ &= P_0(R) + \int_{\Sigma} \hat{p}(R) dS. \end{aligned} \quad (3)$$

In the more general case of locally-reacting surfaces of finite impedance (i.e., not purely reflecting), the reflection term can be further weighted by the complex-valued reflection coefficient for plane waves [Pie84]. This coefficient will only be valid, however, in far-field from the surface.

Efficient implementation on the GPU

Evaluating Eq. 3 involves two subproblems. First, the integration domain, i.e. all the scattering surfaces visible from the source, must be determined and sampled. In itself, this is a difficult task for complex geometries. Second, the differential contribution of blocked plus reflected wavefronts must be evaluated for all surface samples and summed up. These two tasks perfectly match the operations implemented by the graphics hardware. To maximize efficiency we implemented these two steps using a “source-view” strategy which provides the most natural mapping to the GPU architecture and results in a very straightforward implementation.

By rendering the scene from the location of the sound source, in a way similar to a *shadow mapping* technique in computer graphics, we can sample the set of directly “illuminated” scattering surfaces (see Figure 3). Computing the source view using perspective projections also provides a form of importance sampling strategy by allocating more fragments to surfaces close to the source.

We can then evaluate the integral in Eq. 3 as a sum over all visible fragments i in this view:

$$\hat{P}(R) \approx \sum_i \hat{p}_i(R) dS_i, \quad (4)$$

and $dS_i = (w/rez)^2 (nearDist / (-\mathbf{u} \cdot \mathbf{t}))^2 / (\mathbf{n} \cdot \mathbf{u})$, where \mathbf{t} is the viewing direction, \mathbf{u} the vector from the sample to the viewpoint, \mathbf{n} the normal, $nearDist$ is the view plane distance, rez the rendering resolution and w the width of the view frustum (assuming aspect ratio is 1). All vectors are unit vectors.

Our GPU implementation renders the geometry to a floating-point offscreen render-target. For each rendered pixel, we evaluate the corresponding value of $\hat{p}_i(R) dS_i$. We store the resulting complex-valued number in two of the four color components of each pixel (see Figure 1 center).

The sum over all visible surfaces can then be efficiently computed using hierarchical integration (i.e., “mip-mapping”), classically performed in $\log(rez)/\log(k)$ render passes, where k is the reduction factor. At each pass a $k \times k$ block of values is summed-up to give a single value which will be

recursively integrated in the next pass until the total value of the integral is eventually reached. In our case, we typically used a $4\times$ factor resulting in 5 passes for a 1024×1024 render-target.

To evaluate the integral over multiple frequencies, we use a deferred-shading approach [DS05] and render the necessary geometrical parameters (distances and dot products) only once. However, Eq. 4 must still be fully re-evaluated. For scenarios where the scattering of multiple dynamic sources must be computed at interactive rates, deferred shading can be leveraged to sample the scattering surfaces from a unique location at the expense of some visibility error. Then, the scattering integrals can be computed and accumulated using hardware blending before the final mip-mapping pass. For k sources however, the process requires $k * b$ render passes of the scattering shader (Eq. 4) which currently limits the approach to a small number of sources or frequency components.

Performances

We implemented Eq. 4 both in software using C++ and using programmable graphics hardware with *OpenGL* and *Cg*. For the interactive demos, we used a *Direct3D* implementation. Example shader code can be found in the supplemental materials. We ran performance evaluation tests of our approach on several CPU and GPU configurations using simple models (sphere and plane), where all surfaces are visible from the source. In the software case, sampling the visibility from the source would decrease the performance of the approach for complex geometries. In contrast, using the GPU gives us almost instant access to the visible surfaces. Table 1 summarizes the timings required to compute the full scattering integral at a single frequency. Our GPU implementation was found to be 20 to 40 times faster than an equivalent C++ implementation with refresh rates in excess of 700Hz for a fully filled 256×256 render target on our most recent hardware.

Validation against BEM simulations

We first conducted comparisons to BEM simulations using a Fast Multipole Method (FMM) [Dar00] to assess the accuracy and limitations of our approach for a given frequency. We used two simple cases: a spherical occluder made of 8192 triangular elements and a square-shaped plate made of 10200 elements. Figure 4 shows the amplitude of the sound pressure when the receiver rotates around the spherical obstacle. Evaluation was done for 360 positions (every degree). In this case, we obtain very good agreement with BEM solution although the accuracy of our approach is slightly reduced at low frequencies. Please, refer to the supplemental material for additional low-frequency comparisons and results for the square-shaped plate.

We also ran comparisons for more complex situations, such as the displacement surface of Figure 5. In this case, second




Figure 4: Polar plot of BEM (magenta) pressure-amplitude solutions compared to our approach (green) as a receiver rotates around a spherical obstacle scattering a 1KHz wave. The radius of the scattering sphere is 1m. Source and receiver are respectively 5 and 10 meters away from the center of the sphere. The plot also shows the unoccluded pressure amplitude (blue) and the amplitude of the scattered component (orange).

order scattering becomes more prominent and our approach introduces more errors at high frequencies. However, it still provides a reasonable estimation.




Figure 5: Scattering patterns for a detailed surface. The figure compares sound pressure levels in a plane medial to the surface obtained by BEM and our approximation. Source is 5m directly above the center of the face and the pressure is plotted at a distance of 10m.

5. Simplified Modeling of Surface Detail

Apart from raw computing power, using the GPU also offers the possibility to leverage level-of-detail schemes originally developed for real-time graphics rendering [COM98]. In this section, we propose a strategy combining normal mapping with displacement correction to model complex surface detail for acoustic scattering calculations.

Displacement surfaces [CCC87] use textures to encode fine-grain surface detail. They can be rasterized in hardware us-

Hardware	256x256			1024x1024			2048x2048		
	kirch(Hz)	mipm(Hz)	total(Hz)	kirch(Hz)	mipm(Hz)	total(Hz)	kirch(Hz)	mipm(Hz)	total(Hz)
Intel Core2 6600 @2.4GHz	-	-	43	-	-	2.7	-	-	0.7
Intel Xeon @3.2GHz	-	-	35	-	-	2.2	-	-	0.55
QuadroFX Go 1400	380	870	270	37	60	23	n/a	n/a	n/a
GeForce 7950GX2	1120	2120	730	141	170	76	39	26	15

Table 1: Performance tests on various hardware platforms. The table shows update rates for the scattering integral computed at single frequency using an increasing number of samples. GPU performance is detailed both for the integral evaluation shader (kirch) and the summation shader (mipm). Values do not include the cost of sampling the visibility from the source.

ing a ray-casting approach [HEGD04, BD06], leading to correct visibility handling at the expense of a higher shader-processing cost. To avoid the cost of ray-casting we propose a simpler approach that accounts for the correct propagation delay but omits accurate visibility calculations. To remove the contribution of backfacing fragments, which should not be contributing to the integral, we use an additional weighting term defined as $-\mathbf{v} \cdot \mathbf{n}$, where \mathbf{v} is the direction from the 3D location to the sound source and \mathbf{n} is the surface normal.

To evaluate our approach, we created several 4x4 meter surface samples from displacement textures. The amplitude of displacement was 0.5 meters. Figure 6 shows example surfaces and scattering impulse responses calculated with true displaced geometry. In this case, source and microphone were directly above the center of the surface respectively 10 and 20 meters away. The impulse responses were computed off-line by evaluating the scattering integral for 8192 linearly spaced frequency components ranging from 0 to 22.05KHz. The result is then obtained by the inverse Fourier transform of the resulting transfer function. A resolution of 1024x1024 was used for the render target. We also created corresponding normal maps from the displaced geometry and performed a scattering calculation using a flat proxy surface with our displacement-corrected normal mapping and standard normal mapping only.

Figure 7 compares the different results. As can be seen, normal maps without displacement correction result in very little difference compared to a flat surface. This demonstrates the importance of the interference phenomena which are paramount in modeling the proper scattering effect. Using our displacement-corrected normal mapping approach results in a much better approximation.

We also evaluated our simplification technique at oblique incidence. Please refer to the supplemental materials for an illustration of the results. At grazing incidence, errors have been found to be more significant due to self occlusions but our approximation remains acceptable.

6. Interactive Rendering Framework

Our approach can be used to render interactive scattering effects using a framework similar to the sound occlusion approach of [TG97b], which we extend to add our more physically-grounded scattering calculation.

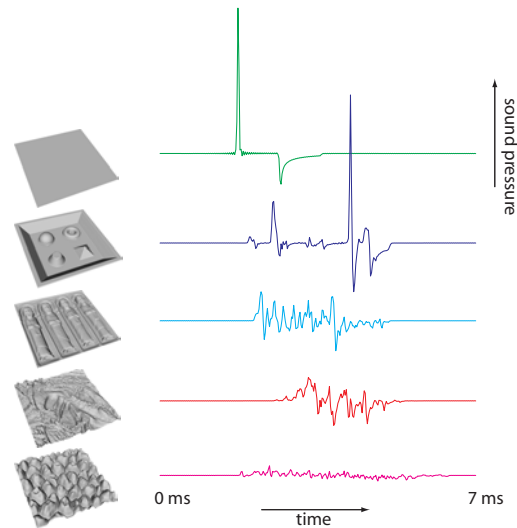


Figure 6: Responses from different 4x4m surface samples. Each surface is composed of 131072 triangles and generated from displacement maps. Note the secondary scattering component due to the finite extent of the flat surface on the top row (green curve) and the increasingly “diffusing” nature of the surfaces from top to bottom.

Our current audio-visual rendering pipeline performs both visual rendering and calculation of the audio scattering coefficients on the GPU. In order to support interactive audio rendering in fully dynamic environments, we compute the scattering integral over a small number of frequencies, typically 10 to 20, using uniform spacing on a Bark scale to provide a better match to auditory perception [Moo97]. This set of frequency subbands b will be selectively modulated by a coefficient defined by:

$$\alpha_b(P) = |P_b(R)|, \quad (5)$$

where $P_b(R)$ is the sum of the direct and scattered components (see Eq. 3) calculated at the center frequency of band b .

All audio processing is performed asynchronously on the CPU. To auralize the effect of scattering we re-equalize the signals based on the coefficients α_b computed on the GPU. We process the input audio signal using short time-frames of 1024 samples with 50% overlap at CD quality. We trans-

form the signal into the Fourier domain and directly multiply each complex coefficient by a scalar factor, which we interpolate from the small set of subband values α_b . This simplified approach does not explicitly model the propagation delay between the direct and indirect components but will still capture the coloration due to scattering [Hal01]. While appropriate for occlusions, it might result in erroneous rendering of reflections off distant surfaces for which a distinct echo would be perceived. However, psychoacoustical evidence suggests that this approximation will hold for delays up to 50 ms (≈ 17 m difference in path-length) but this threshold generally depends on the level of the reflected sound and the nature of the signal (see [Ber96], Ch. 14). As illustrated in the accompanying video, this approach still results in convincing renderings for a number of practical situations. In particular, we show our technique applied to a prototype driving simulation (see Figure 1) which is a challenging application due to fast moving sources and scatterers of different sizes as well as complex road-side geometry. To demonstrate the flexibility of our approach we compute the scattering of the car's engine from the roadside, tunnels, and other cars, as would be heard through the open windows. The environment in this case comprises 95000 triangles. We used a rendering of each side of the car to sample all possible scattering surfaces and provide a stereophonic rendering effect. We also provide examples of interactive auralization of occlusion by a dynamic obstacle and reflections off a complex surface deformed in real-time. The effect of the scattering can be fully appreciated in these examples.

All GPU operations were implemented in *Direct X* and run with frame rates in excess of 100Hz on a *GeForce 8800 GTX*. The GPU could easily handle all the graphics rendering, calculation of the deferred shading buffers and subsequent integration passes for each of the 10 frequencies we used. Note that the graphics rendering also includes per-pixel bump mapping on the road surface and rendering of an equalizer display for illustration purposes. To better balance GPU load between graphics and audio related render passes, we perform calculations for the left and right side of the road every other frame. View-frustum culling was used to optimize rendering but any other acceleration technique can be applied. In our current implementation, 512×512 render targets were used for audio scattering calculations and seemed appropriate to avoid aliasing problems. We believe that the scattered audio component enhances the sense of presence in the environment compared to the direct sound alone since it is tightly coupled to the surrounding geometry. This effect can be appreciated in the accompanying video.

7. Extension to Off-line Simulations

Our approach could also be used to compute high-quality impulse responses suitable for off-line auralization of outdoor acoustic problems. As a first step towards this goal we computed scattering filters for large-scale real-world situa-

tions and compared them directly to corresponding recordings in the field. Due to the scale/complexity of the models, we believe that no other approach could be used at this time to obtain such results. Please, refer to the video or the supplemental audio files provided with the paper to evaluate the quality of our simulations.

An interesting example is the Kukulcan temple, a Maya staircase-pyramid located in Chichén Itzá, Mexico (see video and supplemental materials for an illustration). The stairs of this pyramid act as a sound diffraction grating. They reflect a particular chirped echo which has been the object of a number of studies [DDBL04, Bil06]. For additional information, we refer the reader to <http://www.ocasa.org/MayanPyramid.htm>. We modeled a 856 polygons virtual replica of the pyramid, on which we applied our scattering algorithm (see supplemental material). The same model was used for all frequencies and did not require frequency-dependent adaptation of the tessellation.

A full-bandwidth (0-22KHz) transfer function was computed in 2.45 Hz increments (8192 frequencies) in 92 sec. on a *Pentium 4 3.4GHz* workstation using a *Cg/openGL* implementation running on a *GeForce 8800 GTX* GPU. We used a render target of size 1024×1024 . In this case, we also included modeling of atmospheric scattering to better account for high-frequency attenuation over long distances. This was simply achieved by weighting the integrand in Eq. 3 by an additional scalar factor computed according to [ISO93]. The obtained filter was then applied to the handclap used in the on-site recording available at the above URL. Figure 8 (left) compares the result of our simulation to the recording. Although the recording contains significant environmental noise, it can be seen that our algorithm convincingly captures the chirped echo from the stairs.

Our second, more challenging example models the scattering of the façade of the *Duomo* on the *Piazza dei Miracoli* in Pisa, Italy, also famous for its leaning bell-tower (see Figure 1). We used a detailed model of the cathedral obtained from time-of-flight laser scanning and containing 13 million triangles (a resolution of about 2 cm). Figure 8 (right) compares a simulation with an on-site recording of a handclap. As can be seen, the approach gives satisfying results although some components, probably due to higher-order scattering or reflections from the ground (which was not acquired) are missing. The computing time for the solution was similar to the one of the pyramid since we use a deferred shading approach. Creating the necessary renderings required only 2 additional seconds using the same hardware.

8. Discussion and Limitations

We demonstrated several successful applications of our approach. We now discuss some remaining concerns.

High-frequency and surface aliasing




Figure 7: Comparison of true displaced geometry with a proxy flat quadrilateral enhanced with normal-map only or combined normal/displacement maps. Source and receiver are respectively 10 and 20 m directly above the center of the face. Note how the normal-map alone has little effect on the obtained response.

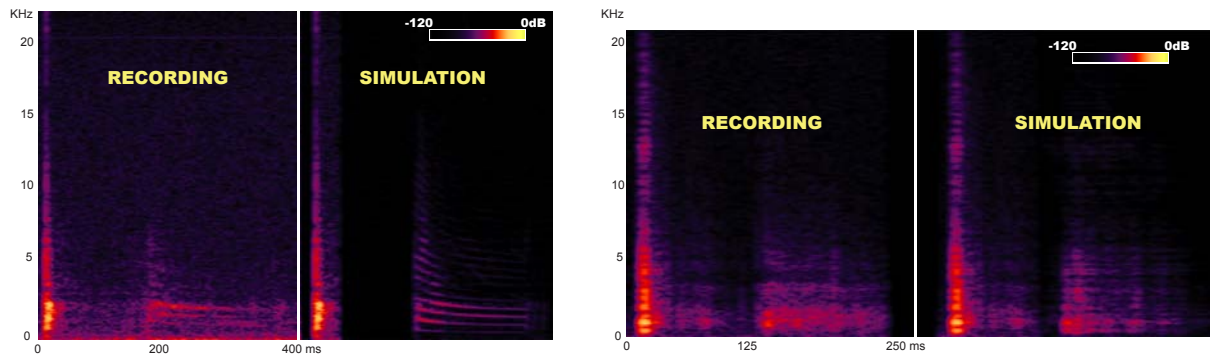


Figure 8: Comparison between spectrograms of a simulation and an on-site recording for the Kukulkan temple (left) the façade of the Duomo (right). The simulated responses are convolved by the handclap of the original recordings.

Our *source-view* approach can be prone to aliasing due to insufficient sampling at high frequencies or to a bad estimation of the dS_i in Eq. 4. A possible solution to this problem could be to implement the algorithm in surface space by rendering each polygon of the scene in a 2D surface “atlas”. Fragments visible from the source have to be determined through the use of a shadow-map. Such an approach would offer several advantages. First, it would allow for precise control of the surface sampling which could be adapted to each primitive. Second, it could be used to treat omnidirectional sources without requiring six renderings to construct an environment map as it is currently the case with our approach. Although resulting in reduced performance in our preliminary implementation, this alternative approach could be useful for off-line high-quality simulations.

Higher-order scattering

Our approach remains limited to first order scattering. As a result it cannot be used to compute reverberation. It also does not account for occlusions of first-order reflected waves.

This is likely to have an influence at high frequencies. A straightforward solution would be to use an additional shadow map from the receiver to account for self occlusions. However, as can be seen in Figure 5, this would not hold at low frequencies for which it seems better to ignore occlusions. One possible approximation would be to determine a suitable blend between the unoccluded and the occluded solution depending on the frequency. For higher-order propagation between objects, which would be especially relevant for indoor simulations, we believe our approach can be used to enhance GA simulations with realistic pre-computed surface scattering functions or filters similar to Figures 5 and 6. This could be achieved in combination with an image-source/beam-tracing technique but would probably better suit a radiosity-like framework to account for non-specular transfer between surfaces. In this context, our approach could also be used to compute the impulse response of the *form-factors* [TG97a].

Simplified frequency-domain processing

Reconstructing an accurate impulse response from sparse data is a very challenging problem which is equally difficult in time or frequency domain. Using a frequency-domain approach is key to our efficient implementation on the GPU since the contribution of all samples can be integrated regardless of their propagation time to the receiver. However, using only a small number of frequency bands introduces an approximation to the true filtering since it makes it difficult to account for phase effects, e.g. due to the propagation delays from the scattering surfaces. For distant surfaces, the delay of the scattered component relative to the direct sound has to be accounted for separately. We believe that an hybrid time-frequency approach could be derived by partitioning the integral into different sub-regions which would contribute separately to the global solution. This could be achieved in world-space by pre-segmenting the scene into different object layers or directly in source-view by stopping our hierarchical integration at a given depth. This latter solution would preserve the efficiency of our current approach. Average or minimum delays could then be computed for each region and used to explicitly account for the propagation time prior to filtering. Better spatial audio rendering could also be achieved by computing an average direction of incidence to the listener for each region.

9. Conclusion

Our approach casts the problem of sound scattering in a framework based on the GPU for all geometric computations. There are three main advantages to this usage of the GPU. First, the raw power of rasterization, thanks to the massively parallel pipeline which allows the determination of visible surfaces at speeds orders of magnitude faster than previous methods. Any type of primitives that can be rasterized can be used with our approach, for instance direct point-based representations from scanned data. Second, the ability to use a mip-mapping strategy to compute the scattering integral extremely efficiently. Third, the ability to use effective level of detail mechanisms such as displacement mapping to treat extremely complex geometry as textured billboards. We believe that this paradigm shift for geometrical acoustics is of prime importance for interactive applications.

Our approach simulates a continuous first-order sound field for complex dynamic environments which would be very difficult, if not impossible, to obtain with concurrent GA techniques. As a result, our approximate interactive rendering allow the addition of convincing and artefact-free sound scattering effects to interactive virtual environments, adding a sense of realism which was previously impossible.

The advantages cited above also allow efficient computation of high-quality sound scattering effects at a scale which was previously impossible. Our results show for the first time that the Kirchhoff approximation can be successfully used for off-line auralization in very complex environments. On our simpler validation examples, we could typically compute

32000 integrals during the time required to obtain a BEM solution for the same scene. Hence, improving further on the accuracy of our approach to bring it closer to BEM could be of tremendous interest for off-line acoustics simulations.

Finally, with a suitable extension to vibrating surfaces, integration with techniques such as the precomputed acoustic transfer (PAT) [JBPO6] would allow much faster solutions, and would render the PAT approach more attractive.

References

- [BD06] BABOUD L., DÉCORET X.: Rendering geometry with relief textures. In *Graphics Interface '06* (2006).
- [Ber96] BERANEK L. L.: *Concert and Opera Halls: How They Sound*. Published for the Acoustical Society of America through the American Institute of Physics, 1996.
- [Bil06] BILSEN F. A.: Repetition pitch glide from the step pyramid at Chichen Itza. *J. of the Acoustical Society of America*, 120 (2006), 594.
- [CCC87] COOK R. L., CARPENTER L., CATMULL E.: The reyes image rendering architecture. *SIGGRAPH Comput. Graph.* 21, 4 (1987), 95–102.
- [CDD*06] COX T., DALENBACK B., D'ANTONIO P., EMBRECHTS J., JEON J., MOMMERTZ E., VÖRLANDER M.: A tutorial on scattering and diffusion coefficients for room acoustic surfaces. *Acta Acustica united with Acustica* (Jul 2006), 1–15.
- [CI90] CHEN J., ISHIMARU A.: Numerical simulation of the second-order Kirchhoff approximation from very rough surfaces and a study of backscattering enhancement. *J. Acous. Soc. of America*, 4 (Oct 1990), 1846–1850.
- [CL93] COX T., LAM Y.: Evaluations of methods for predicting the scattering from simple rigid panels. *Applied Acoustics* 40 (1993), 123–140.
- [COM98] COHEN J., OLANO M., MANOCHA D.: Appearance-preserving simplification. In *SIGGRAPH '98: Proceedings of the 25th annual conference on Computer graphics and interactive techniques* (New York, NY, USA, 1998), ACM Press, pp. 115–122.
- [CR05] CHRISTENSEN C., RINDEL J.: A new scattering method that combines roughness and diffraction effects. *Forum Acousticum, Budapest, Hungary* (2005).
- [CS07] CALAMIA P., SVENSSON U.: Fast time-domain edge-diffraction for interactive acoustic simulations. *EURASIP Journal on Applied Signal Processing, special issue on Spatial Sound and Virtual Acoustics* (2007).
- [Dar00] DARVE E.: The fast multipole method: numerical implementation. *J. Comp. Physics* 160 (2000), 195–240.
- [DDBL04] DECLERCQ N. F., DEGRIECK J., BRIERS R., LEROY O.: A theoretical study of special acoustic effects caused by the staircase of the El Castillo pyramid at the maya ruins of Chichen-Itza in Mexico. *J. of the Acoustical Society of America*, 116 (2004), 3328.
- [DS05] DACHSBACHER C., STAMMINGER M.: Reflective shadow map. *Proceedings of I3D'05* (2005).
- [EDS01] EMBRECHTS J., D-ARCHAMBEAU, STAN G.: Determination of the scattering coefficient of random rough diffusing surfaces for room acoustics applications.

- Acta Acustica united with Acustica* 87 (June 2001), 482–494.
- [Emb00] EMBRECHTS J.: Simulation of first and second-order scattering by rough surfaces with a sound-ray formalism. *J. of Sound and Vibration* 229, 1 (June 2000), 65–87.
- [FCE*98] FUNKHOUSER T., CARLBOM I., ELKO G., PINGALI G., SONDEHI M., WEST J.: A beam tracing approach to acoustic modeling for interactive virtual environments. *ACM Computer Graphics, SIGGRAPH'98 Proceedings* (July 1998), 21–32.
- [FHLB99] FILIPPI P., HABAUT D., LEFEVRE J., BERGASSOLI A.: *Acoustics, basic physics, theory and methods*. Academic Press, 1999.
- [FJT02] FUNKHOUSER T., JOT J., TSINGOS N.: Sounds good to me ! Computational sound for Graphics, VR, and Interactive systems. *Siggraph 2002 course #45* (2002).
- [Hal01] HALMRAST T.: Sound coloration from (very) early reflections. *ASA, Acoustical Society of America Meeting, Chicago* (June 2001).
- [Hec98] HECHT E.: *Optics, Chapter 10, pp. 501-507*. 3rd edition, Addison Wesley, 1998.
- [HEGD04] HIRCHE J., EHLERT A., GUTHE S., DOGGETT M.: Hardware accelerated per-pixel displacement mapping. *Proc. of Graphics Interface'04, Canadian Human-Computer Communications Society* (2004), 153–158.
- [ISO93] ISO: Acoustics - Attenuation of sound during propagation outdoors - Part 1: Calculation of the absorption of sound by the atmosphere. *International Organization for Standardization, ISO 9613-1* (1993).
- [ISO04] ISO: Acoustics - Sound-scattering properties of surfaces - Part 1: Measurement of the random-incidence scattering coefficient in a reverberation room. *International Organization for Standardization, ISO 17497-1* (2004).
- [JBP06] JAMES D. L., BARBIČ J., PAI D. K.: Precomputed acoustic transfer: Output-sensitive, accurate sound generation for geometrically complex vibration sources. *ACM Transactions on Graphics (SIGGRAPH 2006)* 25, 3 (Aug. 2006).
- [JM82] JEBSEN G., MEDWIN H.: On the failure of the Kirchhoff assumption in backscatter. *J. Acous. Soc. of America*, 5 (Nov 1982), 1607–1611.
- [JMT03] JOSLIN C., MAGNENAT-THALMANN N.: Significant fact retrieval for real-time 3D sound rendering in complex virtual environments. *Proc. of VRTST 2003* (October 2003).
- [KN00] KEIFFER R., NOVARINI J.: A time-domain rough surface scattering model based on wedge diffraction: Application to low-frequency backscattering from two-dimensional sea surfaces. *J. Acous. Soc. of America*, 1 (Jan 2000), 27–39.
- [LSVA07] LENTZ T., SCHRÖDER D., VORLÄNDER M., ASSENMACHER I.: Virtual reality system with integrated sound field simulation and reproduction. *EURASIP Journal on Advances in Signal Processing* 2007 (2007), Article ID 70540, 19 pages. doi:10.1155/2007/70540.
- [Moo97] MOORE B. C.: *An introduction to the psychology of hearing*. Academic Press, 4th edition, 1997.
- [MPM90] MCNAMARA D., PISTORIUS C., MALHERBE J.: *Introduction to the Uniform Geometrical Theory of Diffraction*. Artech House, 1990.
- [NNK93] NORTON G., NOVARINI J., KEIFFER R.: An evaluation of the Kirchhoff approximation in predicting the axial impulse response of hard and soft disks. *J. Acous. Soc. of America*, 6 (June 1993), 3094–3056.
- [Pie84] PIERCE A.: *Acoustics. An introduction to its physical principles and applications*. 3rd edition, pp. 107-111, American Institute of Physics, 1984.
- [RSCG99] RINDEL J., SHIOKAWA H., CHRISTENSEN C., GADE A. C.: Comparisons between computer simulations of room acoustical parameters and those measured in concert halls. *Joint meeting of the Acoustical Society of America and the European Acoustics Association, Berlin, 14-19 March* (1999).
- [SFV99] SVENSSON U. P., FRED R. I., VANDERKOOY J.: Analytic secondary source model of edge diffraction impulse responses. *J. Acoust. Soc. Am.* 106 (1999), 2331–2344.
- [Sil05] SILTANEN S.: Geometry reduction in room acoustics modeling. *Master Thesis, Helsinki University Of Technology, Department of Computer Science Telecommunications Software and Multimedia Laboratory* (September 2005).
- [SN81] SAKURAI Y., NAGATA K.: Sound reflections of a rigid plane panel and of the "live-end" composed by those panels. *J. Acous. Soc. of Japan*, 1 (Jan 1981), 5–14.
- [SRT94] SAVIOJA L., RINNE T., TAKALA T.: Simulation of room acoustics with a 3D finite difference mesh. *Proceedings of Intl. Computer Music Conf. (ICMC94)* (Sept. 1994), 463–466.
- [TFNC01] TSINGOS N., FUNKHOUSER T., NGAN A., CARLBOM I.: Modeling acoustics in virtual environments using the uniform theory of diffraction. *ACM Computer Graphics, SIGGRAPH'01 Proceedings* (Aug. 2001), 545–552.
- [TG97a] TSINGOS N., GASCUEL J.: A general model for the simulation of room acoustics based on hierarchical radiosity. *technical sketch, in visual proceedings of SIGGRAPH'97, Los Angeles, USA* (Aug. 1997).
- [TG97b] TSINGOS N., GASCUEL J.-D.: Soundtracks for computer animation: sound rendering in dynamic environments with occlusions. *Proceedings of Graphics Interface'97* (May 1997), 9–16.
- [TG98] TSINGOS N., GASCUEL J.-D.: Fast rendering of sound occlusion and diffraction effects for virtual acoustic environments. *Proc. 104th Audio Engineering Society Convention, preprint 4699* (May 1998).
- [Tho87] THORSOS E.: The validity of the Kirchhoff approximation for rough surface scattering using a gaussian roughness spectrum. *J. Acous. Soc. of America*, 1 (Jan 1987), 78–92.
- [WRR04] WANG L., RATHSAM J., RYHERD S.: Interactions of model detail level and scattering coefficients in room acoustic computer simulation. *Intl. Symp. on Room Acoustics, a satellite symposium of ICA, Kyoto, Japan* (2004).
- [ZCR06] ZENG X., CHRISTENSEN C., RINDEL J.: Practical methods to define scattering coefficients in a room acoustics computer model. *Applied Acoustics* (2006).