



THE EFFECT OF DIFFUSE REFLECTIONS ON NOISE MAPS USING PHONON MAPPING

Bram de Greve*, Tom Willems, Tom De Muer, and Dick Botteldooren

Department of Information Technology, Ghent University
Sint-Pietersnieuwstraat 41, B-9000 Ghent, Belgium
bram.degrev@intec.ugent.be

Abstract

Constructing noise maps for large urban areas is a CPU intensive task. Engineering approximations to outdoor sound propagation are commonly introduced to reduce the computational complexity. In particular, diffuse reflections are often neglected because of their great computational expense. This paper reports on an effort to approximate the diffusely reflected field by using a technique called phonon mapping.

In noise mapping, many techniques can be and are adapted from the field of computer graphics where various techniques to approximate indirect lighting exist. A popular technique in computer graphics is photon mapping. It was introduced in 1995 by Jensen and is known to be an efficient algorithm for simulating indirect lighting. In this paper, it will be shown how this photon mapping algorithm can be adapted to the field of noise mapping to simulate diffuse reflections.

The phonon mapping algorithm consists of two passes. In the first pass, a phonon map is built by bouncing phonons, abstract sound particles, through the environment. On each interaction with a (partially) diffusely reflecting surface, it is stored in the phonon map. The Russian roulette technique is used to determine if the phonon is absorbed, specularly or diffusely reflected. Upon reflection, the phonon is traced through the environment again. This process is continued until the phonon map is sufficiently populated. The second pass is a more traditional backward ray tracer. Rays are emitted from the receivers. Specular reflections are handled as usual but for diffuse reflections, no secondary, reflected rays are cast. Instead, the phonon map is queried for nearby phonons. Their presence is used to estimate the energy that is available from this diffuse reflector.

DEFINITIONS

Directional sound power \mathcal{D} : sound power per unit solid angle $d\omega$. Its units are $\frac{\text{W}}{\text{sr}}$.

Acoustic radiance \mathcal{L} : sound power per unit projected area dA^\perp and per unit solid angle $d\omega$, where $dA^\perp = dA \cos \theta$ is the projected area of dA on the hypothetical surface orthogonal to $d\omega$. Its units are $\frac{W}{m^2 sr}$.

INTRODUCTION

Present methods for constructing noise maps of large areas are complex and computational intensive, even when engineering methods are employed. Classic ray/beam tracers [3, 4, 6] often implement only specular reflections based on image sources. Diffuse reflections are neglected because of their great computational cost, despite the fact that the diffusely reflected field can significantly contribute to urban sound fields. This paper presents an effort to approximate the diffusely reflected field using a technique called *phonon mapping*.

Phonon mapping belongs to the family of geometric acoustics, together with methods like ray tracing and beam tracing, all making the same high frequency approximation. Phonon mapping is a particle tracing technique in which acoustic power quanta, called *phonons*, are bounced through the environment. At diffuse surfaces, incident phonons are stored in a *phonon map*. Later, the phonon map is queried to estimate the exitant *acoustic radiance* \mathcal{L} at the diffuse reflector.

This technique is adapted from the field of computer graphics where it is known as *photon mapping*. It was introduced in 1995 by Jensen [7, 8] as an efficient alternative to *radiosity* [5] and *Monte Carlo ray tracing* [9] for simulating *global illumination*.

The goal of noise mapping is to determine how much acoustic power W reaches a receiver located at \mathbf{r} . It is given by the integral of the incoming directional power \mathcal{D} over all directions:

$$W(\mathbf{r}) = \int_{\Omega} \mathcal{D}(\mathbf{r}, \omega) d\omega \quad (1)$$

In geometric acoustics, this integral is written as the sum over all sound paths that terminate at \mathbf{r} :

$$W(\mathbf{r}) = \sum_k \mathcal{D}_k(\mathbf{r}, \omega_k) \quad (2)$$

This sum is taken over the set of all sound paths from a source to the receiver, via any number of specular or diffuse reflections. It contains an infinite number of sound paths, therefore for practical applications a finite approximation to this sum must be made.

A traditional ray or beam tracer handles only direct paths, or paths with only specular reflections. The sum of their contributions is W_s . Such a ray tracer disregards all sound paths with at least one diffuse reflection. The goal of the phonon mapper presented in this paper is precisely to estimate the sum W_d of the contributions due to these omitted paths. Finally, both results must be summed to get W :

$$W(\mathbf{r}) = W_s(\mathbf{r}) + W_d(\mathbf{r}) \quad (3)$$

FIRST PASS – PHONON TRACING

Emitting Phonons From A Single Source

A number of phonons are emitted from the source with power P until the phonon map is populated with N_{map} phonons. N_{emit} is the total number of phonons emitted and is not known a priori since some phonons do not hit a surface at all and thus are not stored in the phonon map, while others may be stored several times.

All emitted phonons should have the same power. This is important for two reasons: firstly computational resources are not spent on low power phonons which make little contribution to the result, and secondly it reduces the variance of the *radiance estimate* in the gathering pass. Consequently, phonons must be emitted with directional densities proportional to power output. For omnidirectional sources, the phonon emission directions ω must be uniformly distributed over the unit sphere. This is accomplished by transforming a uniform variate (ξ_1, ξ_2) over $[0, 1] \times [0, 1)$ according to:

$$\omega(\xi_1, \xi_2) = \begin{bmatrix} 2\sqrt{\xi_1(1-\xi_1)} \cos 2\pi\xi_2 \\ 2\sqrt{\xi_1(1-\xi_1)} \sin 2\pi\xi_2 \\ 1 - 2\xi_1 \end{bmatrix} \quad (4)$$

The sum of all phonon powers must equal P , so each phonon should have power $P_{ph} = \frac{P}{N_{emit}}$. Since N_{emit} is initially unknown, this can only be computed after the phonon tracing pass, in a normalisation pass over all stored phonons.

Multiple Sources

In case of multiple sources having powers P_i , an additional random variable ξ_3 is used to select an emitting source. Each source is selected with probability $p_i = \frac{P_i}{P_{tot}}$ with $P_{tot} = \sum P_i$, such that a source will emit proportionally many phonons to its (relative) power output. The power of each phonon will be $P_{ph} = \frac{P_{tot}}{N_{emit}}$.

Surface Interaction

Incident sound power at a surface point \mathbf{r} having normal \mathbf{n} is partially reflected back into the scene. The coefficients R_s and R_d correspond to the fraction of specularly and diffusely reflected sound power, respectively¹.

The sound power P carried by phonons hitting a surface must thus be reflected in different directions. One way to accomplish this is *path splitting*: several new phonons are emitted carrying the reflected power: one with the specularly reflected power $R_s P$ and n phonons with diffusely reflected power $\frac{1}{n} R_d P$. Although this technique is straightforward, it has two major drawbacks: it requires an exponentially growing number of paths to be traced, and the phonons in the phonon map will no longer have similar powers, which increases the variance of the radiance estimate in the second pass.

¹In this paper, the spectral and spatial dependency of both R_s and R_d , and the angular dependency of R_s are not explicitly mentioned to simplify the notation. The implementation takes this into account, however.

Instead, another technique called *Russian roulette* [1, 10] is used to determine whether the phonon should be absorbed (discarded) or reflected. A uniform random number over $[0, 1]$ is used to determine which mode of reflection, if any, is chosen:

$$\begin{cases} \xi \in [0, R_s) & \rightarrow \text{specular reflection} \\ \xi \in [R_s, R_s + R_d) & \rightarrow \text{diffuse reflection} \\ \xi \in [R_s + R_d, 1) & \rightarrow \text{absorption} \end{cases} \quad (5)$$

In case of diffuse reflection, the new direction ω is sampled using the diffuse BRDF² component ρ_d . If ρ_d is Lambertian, this function is the constant $\frac{R_d}{\pi}$ and the new direction is uniformly sampled from the entire sphere using (4), hemiflipping it to be in the hemisphere about \mathbf{n} if necessary ($\omega \cdot \mathbf{n} < 0$).

If ρ_d is sufficiently similar to a Lambertian BRDF, uniform sampling can still be used if the power of the outgoing phonon is modified according to the real ρ_d . However, the phonons will no longer have equal powers, increasing the variance of the radiance estimate in the gathering pass.

Multiple Frequencies

In this paper, the spectral dependency of source powers and reflection coefficients are not explicitly taken into account. Its effect is however trivial to include using (1/3) octavebands. Consider a source whose spectral power distribution is given by a number of bands having center frequencies f_k and powers $P_{i,k}$, with $P_i = \sum_k P_{i,k}$. A straightforward approach would be to emit phonons with full spectral information. However, if R_s and R_d are frequency dependent, this complicates surface interaction: Russian roulette requires *scalar* thresholds in (5), and reflection will change the phonon's power spectrum.

A better approach, used in the implementation, is to emit single-frequency phonons. Analogous to selecting one from many sources, f is selected from the set f_k with probabilities $p_k = \frac{P_{i,k}}{P_i}$. This approach leaves the Russian roulette computation unchanged, and the phonon's power is also unchanged upon reflection.

PHONON MAP CONSTRUCTION

In the first pass, each time a phonon hits a surface with $R_d > 0$, it is stored in a list for building the phonon map. It is important to do this independently of the Russian roulette outcome, since the phonon map represents an estimation of the *incident*, not *exitant*, power flux. Each phonon's power P_{ph} , its position \mathbf{r}_{ph} and its incoming direction ω_{ph} are stored.

Once sufficiently many (N_{map}) phonon interactions have been stored, the first pass is terminated and the phonon map must be built. A balanced, adaptive k-d tree [2, 10] is constructed from the stored phonons. This is a compact and efficient data structure that enables locating the M nearest phonons around a point \mathbf{r} in $O(M \log N_{map})$ time on average. The

²BRDF (*Bidirectional Reflectance Distribution Function*) $\rho(\mathbf{r}, \omega_i, \omega_o)$: reflection coefficient as a function of position \mathbf{r} , incident and exitant angles ω_i and ω_o .

cost of storing the tree is $O(N_{map})$, and constructing it (which is performed only once) is $O(N_{map} \log N_{map})$.

SECOND PASS – GATHERING PASS

In the second pass, $W_d(\mathbf{r})$ is evaluated at each receiver. The integral gets approximated as a sum over M random direction ω_k , chosen from a distribution with probability density $p(\omega)$:

$$W_d(\mathbf{r}) = \int_{\Omega} \mathcal{D}_d(\mathbf{r}, \omega) d\omega \approx \frac{1}{M} \sum_{k=1}^M \frac{\mathcal{D}_d(\mathbf{r}, \omega_k)}{p(\omega_k)} \quad (6)$$

If the directions are uniformly distributed over the unit sphere using (4), then $p(\omega) = \frac{1}{4\pi}$, and (6) simplifies to:

$$W_d(\mathbf{r}) \approx \frac{4\pi}{M} \sum_{k=1}^M \mathcal{D}_d(\mathbf{r}, \omega_k) \quad (7)$$

$\mathcal{D}_d(\mathbf{r}, \omega_k)$ is computed by casting a ray (\mathbf{r}, ω_k) into the scene. It intersects a surface at \mathbf{r}' having normal \mathbf{n}' . $\mathcal{D}_d(\mathbf{r}, \omega_k)$ equals the acoustic radiance $\mathcal{L}_o(\mathbf{r}', -\omega_k)$ leaving the surface at \mathbf{r}' towards the receiver. It is evaluated as an integral over the unit hemisphere \mathcal{U} about \mathbf{n}' of the incoming radiance $\mathcal{L}_i(\mathbf{r}', \omega'')$ and the BRDF $\rho(\mathbf{r}', -\omega_k, \omega'')$:

$$\mathcal{D}_d(\mathbf{r}, \omega_k) = \mathcal{L}_o(\mathbf{r}', -\omega_k) = \int_{\mathcal{U}} \rho(\mathbf{r}', -\omega_k, \omega'') \mathcal{L}_i(\mathbf{r}', \omega'') (\omega'' \cdot \mathbf{n}') d\omega'' \quad (8)$$

The BRDF is split into specular and diffuse components ρ_s and ρ_d . The specular part is non-zero only for the perfectly reflected direction, and is evaluated by tracing the ray $(\mathbf{r}', \omega_k - 2\omega_k \cdot \mathbf{n}')$ and evaluating (9) again at the new intersection.

$$\mathcal{L}_o(\mathbf{r}', -\omega_k) = R_s \mathcal{L}_i(\mathbf{r}', \omega_k - 2\omega_k \cdot \mathbf{n}') + \int_{\mathcal{U}} \rho_d(\mathbf{r}', -\omega_k, \omega'') \mathcal{L}_i(\mathbf{r}', \omega'') (\omega'' \cdot \mathbf{n}') d\omega'' \quad (9)$$

The diffuse term is estimated by querying the phonon map. In the neighbourhood of \mathbf{r}' , the n nearest phonons $(P_{ph}, \mathbf{r}_{ph}, \omega_{ph})_j$ are found with $\|\mathbf{r}_{ph} - \mathbf{r}'\| \leq r_{max}$ and $n \leq n_{max}$. The phonons provide information about the incident power flux P_i , which must be converted to incident radiance \mathcal{L}_i :

$$\mathcal{L}_i(\mathbf{r}', \omega'') = \frac{d^2 P_i(\mathbf{r}', \omega'')}{(\omega'' \cdot \mathbf{n}') dA d\omega''} \quad (10)$$

$$\int_{\mathcal{U}} \rho_d(\mathbf{r}', -\omega_k, \omega'') \mathcal{L}_i(\mathbf{r}', \omega'') (\omega'' \cdot \mathbf{n}') d\omega'' \approx \frac{1}{A} \sum_{j=1}^M \rho_d(\mathbf{r}', -\omega_k, -\omega_{ph,j}) P_{ph,j} \quad (11)$$

where A is the area of the circular neighbourhood in which the n phonons are found. If r is the largest distance between \mathbf{r}' and the n phonons, then $A = \pi r^2$. In case of a Lambertian surface, this simplifies to:

$$\int_{\mathcal{U}} \rho_d(\mathbf{r}', -\omega_k, \omega'') \mathcal{L}_i(\mathbf{r}', \omega'') (\omega'' \cdot \mathbf{n}') d\omega'' \approx \frac{R_d}{A\pi} \sum_{j=1}^M P_{ph,j} \quad (12)$$




Figure 1: SPL and error in receivers $\mathbf{r}(x, 0, z_0)$ with $z_0 = 10$ m; plane $z = 0$ with $R_d = 0.95$, $R_s = 0$; source $\mathbf{s}(0, 0, z_0)$ with $P = 0$ dB; $M = 225$, $n_{max} = 20$, $r_{max} = 10$ m.

Stratified Sampling

When evaluating (6) with directions ω_k generated using (4) with (ξ_1, ξ_2) directly sampled from $\Lambda = [0, 1] \times [0, 1]$, the distribution of directions over the full sphere isn't as good as one may think. The samples will tend to clump, causing some areas to have too many samples, while others are undersampled.

It is possible to counteract this, thereby also reducing the variance of the result, by using *stratified sampling* [10]. The sampling space Λ is split into $M_1 \times M_2$ equal *strata* Λ_{jk} (13), and each sample $(\xi_1, \xi_2)_{jk}$ is taken from a uniform distribution over Λ_{jk} . This way, samples are much more evenly spread over the sampling space and no directions are severely undersampled.

$$\Lambda_{jk} = \left[\frac{j}{M_1}, \frac{j+1}{M_1} \right) \times \left[\frac{k}{M_2}, \frac{k+1}{M_2} \right) \quad (13)$$

APPLICATIONS AND RESULTS

Validation

A source with power $P = 0$ dB is positioned at $\mathbf{s}(0, 0, z_0)$ with $z_0 = 10$ m, above a very large Lambertian reflector with $R_d = 0.95$ and $R_s = 0$. Using the phonon mapper, (6) is evaluated for receivers at positions $\mathbf{r}(x, 0, z_0)$. As a reference solution, (7) is evaluated using a direct solution for $\mathcal{D}_d(\mathbf{r}, \omega_k)$ with $M = 80000$:

$$\mathcal{D}_d(\mathbf{r}, \omega_k) = \mathcal{L}_o(\mathbf{r}', -\omega_k) = \frac{R_s}{\pi} \frac{P}{\|\mathbf{s} - \mathbf{r}'\|^2} \frac{(\mathbf{s} - \mathbf{r}') \cdot \mathbf{n}'}{\|\mathbf{s} - \mathbf{r}'\|} \quad (14)$$

Figure 1 shows the sound level and error for two phonon mapping simulations against the reference. It can be seen that the photon mapper slightly overestimates the levels. This is due to an underestimation of A in (11). A is a too aggressive a bound for the phonons.

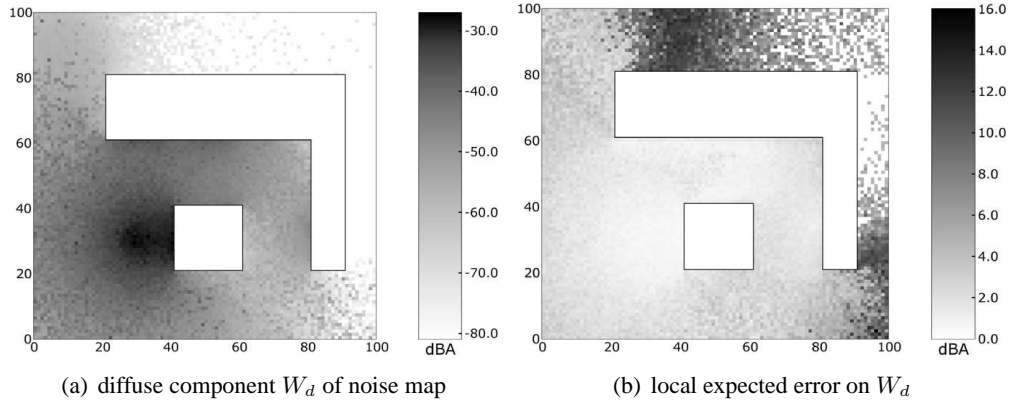


Figure 2: Diffuse component W_d of noise map and local expected error for one source $\mathbf{s}(30, 30, 4)$ with $P = 0$ dB; receivers $\mathbf{r}(x, y, 4)$ with $\Delta x = \Delta y = 1$ m; ground & buildings $R_d = 0.05$, $R_s = 0$; $N_{map} = 40000$, $M = 100$, $n_{max} = 20$, $r_{max} = 10$ m.

Variance Statistics

Monte Carlo simulations often suffer from noise or variance in the estimate. For the phonon mapper as described in this paper, this is mainly influenced by the number M of rays cast in the gathering pass. The local expected error at each point of the map is estimated using the standard deviation of the results of 20 different simulation runs with different random seeds, since the average will converge to the exact solution. The simulation consists of one source, 4 m above a ground surface, and in the vicinity of some buildings. Figure 2 shows the result of one run and the local expected error for $M = 100$. It can clearly be seen that the largest errors are found where the source is occluded by an obstacle, since fewer phonons can be sampled. Fortunately, the sound level significantly drops in such regions, reducing the importance of the error.

To get an idea of the overall expected error, the distribution of the local error over the map is examined. The 90th and 95th percentiles measure how well the largest differences behave, while the 50th percentile is a measure of the average spatial error. Figure 3 shows these percentiles. It can clearly be seen that the expected error decreases with M , however at a linearly increasing cost. It can also be seen that stratified sampling decreases the expected

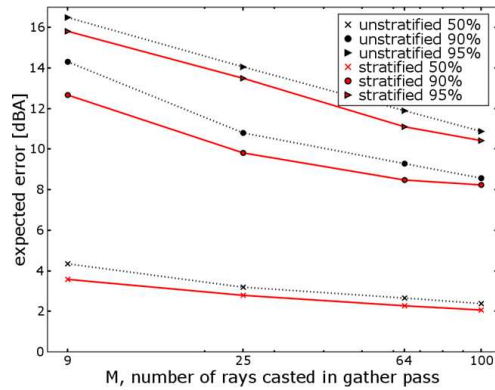


Figure 3: percentiles of expected error distribution for various M (number of rays cast in gather pass), stratified vs. unstratified sampling of rays. Other parameters as in figure 2.

error for the same M , without incurring significant additional computational cost.

POSSIBLE IMPROVEMENTS

In this paper, the diffraction of sound paths is not taken into account for simplicity of presentation. A polygonal beam tracer [3, 4] can easily accommodate this effect in W_s for paths without diffuse reflection. However, the phonon mapper must be extended to handle paths with both diffraction and diffuse reflection. Based on Fresnel zones, Russian roulette may be applied to subject phonons to diffraction.

CONCLUSIONS

In this paper, it is shown how the *photon mapping* technique is adapted to the field of noise mapping for efficiently simulating the effect of diffuse reflections. The new *phonon mapping* technique is orthogonal to traditional ray/beam tracing, which handles the strictly specular paths.

The phonon mapper consists of two passes: a phonon tracing pass, where phonons are propagated in an environment to form the phonon map, and a gathering pass where, at each receiver, the incident diffusive sound field is reconstructed from the phonon map.

The variance of the result is mainly influenced by the number and distribution of gathering rays, the latter of which can be significantly improved using stratified sampling.

REFERENCES

- [1] J. Arvo, D. Kirk, "Particle Transport and Image Synthesis", SIGGRAPH '90, 63–66 (1990)
- [2] J. L. Bentley, "Multidimensional Binary Search Trees Used for Associative Searching". Comm. of the ACM **18** (9), 509–517 (1975)
- [3] B. de Greve, "3D polygonal beam tracer for acoustic simulations" (in Dutch), Master thesis, Ghent University, Belgium (2003)
- [4] T. De Muer, "Policy Supporting Tools for Urban Noise Assessment", PhD thesis, Ghent University, Belgium (2005)
- [5] C. M. Goral, K. E. Torrance, D. P. Greenberg, B. Battaile, "Modeling the interaction of light between diffuse surfaces", SIGGRAPH '84, 213–232 (1984)
- [6] P. S. Heckbert, P. Hanrahan. "Beam tracing polygonal objects". SIGGRAPH '84, 119–127 (1984)
- [7] H. W. Jensen, N. J. Christensen, "Photon Maps in Bidirectional Monte Carlo Ray Tracing of Complex Objects", Computers & Graphics, **19** (2), 215–224 (1995)
- [8] H. W. Jensen, "Global Illumination using Photon Maps", Rendering Techniques '96, Eds. X. Pueyo and P. Schröder, Springer-Verlag, 21–30 (1996)
- [9] J. T. Kajiya, "The Rendering Equation", SIGGRAPH '86, 143–150 (1986)
- [10] M. Pharr, G. Humphreys, *Physically Based Rendering: From Theory to Implementation*. (Morgan Kaufmann, San Francisco, ISBN 0-1255-3180-X, 2004)