5th International Meeting
on
Wind Turbine Noise
Denver 28 – 30 August 2013

Physics Based Spatial Acoustics in Virtual Scenes with Application to Wind Farm Noise

Kevin Nelson. GLSV Inc. Houghton MI 49931 kevinn@glsv.com
Steven G. Mattson. GLSV Inc. Houghton MI 49931 stevem@glsv.com

Summary
A technique for characterizing the impact of the environment on wind turbines is presented. The technique is based on Monte Carlo ray tracing which can characterize much larger environments accurately. Historically, the characterization of wind farm noise has been a difficult task involving complex aerodynamic and acoustic interactions between the turbines and the environment. The size of the surrounding environment renders classical numerical techniques such as the finite element method or boundary element analysis intractable. Also, the complex and varying nature of the environment from site to site makes building detailed analytical or empirical models difficult. There is a clear need for a large scale environmental model which can correctly predict the environment's impact on turbine noise sources.

1. Introduction
The goal of this ray tracing technique is to produce physically accurate three dimensional (3D) audio at various receiver points in a given virtual environment. This is accomplished by emitting rays from the receiver into the scene. When rays intersect with a sound source, the time of arrival and total energy across the sound spectrum is accounted for. This information is then used to generate metrics about sound levels and characteristics at the receiver position, as well as to generate three dimensional immersive audio. This audio stream can then be used to objectively and perceptually judge the noise quality of a wind turbine farm.

This technique shows improvement over state of the art for several reasons. First, the quality of data generated by this model can be arbitrarily controlled by the resolution of the representative geometry given to the simulator. Simpler models can be run faster to generate relatively quick estimates on how changing factors in the environment changes the quality of noise produced. More detailed models can be run at a higher cost of setup and computation to accurately predict the final wind farm noise quality. Because all important wave propagation phenomena can be approximated by these techniques, arbitrary environments produce accurate results.

Secondly, this approach allows a user to hear an accurate representation of the wind farm before it is built, providing otherwise difficult to obtain subjective information about the quality of the wind farm sound. Part of the power of this approach is that the software was designed to be synchronized with a visual display, giving a complete representation of the wind farm before it is built.

2. Background

A Monte Carlo based ray tracing algorithm shows promise for modeling acoustic interactions within a large environment. Ray tracing is less computationally intensive than other numerical methods over large domains because of simplifying assumptions made about acoustic wave interactions. Specifically, the ray tracer assumes that waves cannot interact with one another except at the receiver location. Also, an assumption is made that the frequencies in question are relatively high. These assumptions allow much larger environments to be characterized – on the order of several square kilometers. Work is currently being done to adapt ray tracing techniques to urban environments [4], and a ray tracing approach may work well for wind farm sites as well.

2.1 Simulated Physical Phenomena using Linear Rays

Many important wave phenomena can be captured by ray tracing, including spherical losses, atmospheric attenuation, reflections, diffraction, and interference.

2.1.1 Spherical Losses

Spherical losses, or spreading losses, are caused by the wave front expanding which reduces the energy per unit area of the wave. If

$$e(r) = \frac{E_1}{2\,\pi\,r^2}$$

is the intensity at radius r, derived by the total energy emitted from the source E_1 divided over the surface area of a sphere, then

$$e(r) \propto \frac{1}{r^2}$$

which implies that the energy per unit area on the wave front decreases proportionally to the square of the distance traveled.

This effect is captured naturally in the ray tracer. Each ray represents a finite amount of energy, and the density at which rays fall over the target represent the intensity. Because the rays spread outwards as they travel, their density decreases with distance. Therefore, so long as the ray tracer conserves energy, the effect of spherical loss emerges intrinsically from the geometry of the ray propagation.

2.1.2 Atmospheric Attenuation

Atmospheric attenuation is the process by which a mechanical sound wave looses energy because of frictional losses in the medium it is traveling through. Atmospheric attenuation occurs as a function of four variables, which are frequency, distance, humidity, and air pressure.

The pressure of a ray propagated over a distance can be calculated as a function of its initial pressure as follows:

$$p_t = p_i \, e^{-0.1151 \, \alpha \, s}$$

where s is the distance travelled in meters, and α is attenuation coefficient which is calculated as follows:

$$\alpha = 8686 \, f^2 \left(\left[1.84 \times 10^{-11} \left(\frac{p_a}{p_r}\right)^{-1} \left(\frac{T}{T_0}\right)^{\frac{1}{2}} \right] + \left(\frac{T}{T_0}\right)^{-\frac{5}{2}} \left\{ 0.01275 \left[e^{-\frac{2239.1}{T}} \right] \left[f_{rO} + \left(\frac{f^2}{f_{rO}}\right) \right]^{-1} + 0.1068 \left[e^{-\frac{3352.0}{T}} \right] \left[f_{rN} + \left(\frac{f^2}{f_{rN}}\right) \right]^{-1} \right\} \right)$$

given that

$$f_{rO} = \frac{p_a}{p_r} \left(24 + 4.04 \times 10^4 \, h \, \frac{0.02 + h}{0.391 + h} \right)$$

$$f_{rN} = \frac{p_a}{p_r} \left(\frac{T}{T_0} \right)^{-1/2}$$

where

- f is the frequency of the sound
- h is the molar concentration of water vapor (as a percentage)
- p_r is the reference ambient atmospheric pressure in kilopascals (normally taken to be 101,325 kPa)
- p_i is the initial sound pressure amplitude in Pascals
- p_t is the sound pressure amplitude in Pascals
- p_0 is the reference sound pressure (20 μPa)
- p_a is the ambient atmospheric pressure in kilopascals
- T is the atmospheric temperature in Kelvin
- T_0 is the reference air temperature in Kelvin (normally taken to be 293.15 K)

The attenuation constant formula is derived from how mechanical waves interact with the various gasses in the atmosphere. For a more rigorous derivation, see ISO9613-1

Because atmospheric attenuation is a non-linear property, it must be applied individually across each ray bounce.

2.1.3 Specular and Diffuse Reflections

Reflections occur when the wave interacts with a surface in the environment. They can act differently depending on the porosity of the material and can be classified into two separate groups, which are specular and diffuse reflections.

Specular reflections occur when the wave interacts with a relatively flat surface. In this case energy is propagated according to the law of reflection, which states that *the angle of incidence equals the angle of reflection,* or as follows:

$$\theta_i = \theta_r$$

The reflection occurs in the plane created by the surface normal and the incoming ray.

Specular reflections can be easily implemented in a ray tracing engine. When a ray encounters a surface, the surface's material properties are checked to determine what percentage of the energy it reflects specularly. The outgoing direction is then computed based on the angle from the surface normal and the incoming ray. Finally a new ray is emitted in the correct direction with an energy spectra corrected for the losses in the reflection.



Figure 1: A BRDF determines how an incoming ray is scattered by a material

Unlike specular reflections, diffuse reflections occur when a wave encounters a porous surface (relative to the wavelength). In this case, the ray is scattered across all possible radiating angles weighted by a material property known as the bidirectional reflectance distribution function, or BRDF. The BRDF is a material dependant function which takes the incoming ray, outgoing ray, and surface normal and returns the reflected intensity.

In a sense, specular reflections can be implemented as a subset of diffuse reflections, where the BRDF is defined as a piecewise function as follows:

$$f_r(\overline{d_\iota}, \overline{d_r}, \overline{d_n}) = \begin{cases} s & \overline{d_r} = 2(\overline{d_n} \cdot \overline{d_\iota})\overline{d_n} - \overline{d_\iota} \\ 0 & otherwise \end{cases}$$

where s is a percentage which represents absorptive losses in the material, and $\overline{d_\iota}, \overline{d_r}, \overline{d_n}$ are unit vectors representing the incoming ray, reflected ray, and surface normal respectively. The condition $\overline{d_r} = 2(\overline{d_n} \cdot \overline{d_\iota})\overline{d_n} - \overline{d_\iota}$ is only met if the angle of incidence equals the angle of reflectance within the reflection plane.

2.1.4 Diffraction

Diffraction is the apparent bending of a wave front around an object. The bending of the wave can be approximated by the Huygens-Fresnel Principal, which states that *each point on a wave front acts as a secondary source of wavelets. The combination of these secondary wavelets produces the new wave front in the direction of propagation.* Fresnel used this principal in combination with a series of zones called 'Fresnel zones' to predict the intensity of the wave at various points.



Figure 2: Fresnel Zone Ellipsoid used to compute a diffraction path

A Fresnel zone is an ellipsoid defined such that the foci of the ellipse are the locations of the source and receiver and the length of the difference between the longest path and the shortest path does not exceed $\lambda/2$, or as follows:

$$\Delta p = p2 - p1 \leq \frac{\lambda}{2}$$

A barrier which does not obstruct the line of sight, but exists inside the Fresnel zone will cause constructive or destructive interference in the amplitude of the signal as the barrier is moved farther or nearer from the line of sight. On the other hand, a barrier which impedes the line of sight will attenuate the signal.

Within a Fresnel zone, the wave path can be interpreted using a secondary source model [2]. Secondary sources are placed at edge points located at the corners of $b1$ and $b2$. When these points are excited by the wave front emitted from A, they emit 'diffraction' waves. The superposition of all paths from A to B represents the total sound transfer and can be used to generate a transfer function from A to B.

2.1.5 Constructive and Destructive Interference

Constructive and destructive interference occur as a result of two different signals reinforcing or diminishing one another. In complex signals with more than one frequency component, only the similar frequency components in each of the two signals interact. This is because separate frequencies are independent of one another. Whether constructive or destructive interference occurs depends on the phase of each of the frequency components.

Interference can be generated by a ray tracing algorithm by the property of superposition. Superposition states that if multiple sound sources are evaluated independently of one another at a specific receiver location, and then their results are summed, the total response will be the same as if both sources were evaluated together. This property allows us to disregard every interaction between the waves generated by different sources within an environment except for the ones that occur at the receiver location.

In order to correctly recreate the interference between the two sources, two pieces of information are needed, which are the travel time from each source to the receiver and the phase of the signals emitted from each source. Once the signals are combined accounting for the phase offset by the travel time, interference occurs automatically.

Because the ray tracer maintains the time of arrival for rays which reach the source, no information necessary for interference is lost in the ray trace of the environment. The actual interference effect is introduced when the source audio is mixed using the information from the ray tracer.

2.2 Simulated Physical Phenomena using Curved Rays

While the acoustic effects above can be modeled using linearly propagating rays, there exist a few acoustic phenomena which change the direction of wave propagation en

route, and therefore cannot be modeled by straight rays. These phenomena include wind and refraction due to temperature gradients.

2.2.1 Wind

Wind conditions affect sound propagation by modifying the local speed of sound which will change the propagation path of the wave front.

Wind can be modeled as a vector field which modifies the local speed of sound. The vector field represents the magnitude and direction of the wind speed at each location in the environment. This field can be added to the global speed of sound to produce a localized value.

2.2.2 Refraction

Refraction can be classified into two distinct effects. When a sharp acoustic impedance change occurs the acoustic wave path becomes discontinuous, bending sharply at the interface between the two impedances. An example of when this occurs is at the interface between air and water. This type of refraction can be modeled using linearly propagating rays, and needs no modification to the ray tracing algorithm.

However, when the impedance change is gradual, the wave front will tend to curve and a different kind of ray is needed. An example of when this occurs is in temperature gradients in the atmosphere. These gradients gradually affect the speed of sound in the atmosphere, causing the propagating sound waves to curve.

2.3 Source Modeling of the Turbine

In order to accurately propagate turbine sound through an environment, first an accurate turbine model must be created and integrated into the simulator. Several good turbine source models exist already [5] [7], but a model must be used which operates well within a ray tracing simulator. In order to allow general use, the simulator was designed to accept inputs as moving point sources radiating in a monopole, dipole, or quadripole pattern. Therefore to successfully model turbine noise, the turbine model must be expressed in terms of these sources.

Producing an arbitrary far field pattern with a collection of sources is a difficult problem, but by using simplifying assumptions based on the physical geometry of the turbine, approximations can be made. Lighthill's acoustic analogy is one technique that allows the modeling of sound sources as point sources. More research needs to be done to determine if this framework will be suitable for modeling wind turbine noise sources.

3. Implementation

A Monte Carlo based ray tracing algorithm which can characterize all of the phenomena mentioned above is described in detail.

3.1 Representative Geometry

The first step in characterizing the effects of an environment on noise sources is to build a representative geometry of the environment. While other methods exist, one of the most robust ways to store geometry is in a polygon mesh. The mesh can be significantly reduced in complexity from the actual scene, containing only features that are relevant to the shortest wavelength which is important in the simulation. However, the actual model complexity necessary to gain accurate data from an environment is a matter of engineering judgment.

Many standard commercial tools exist to generate terrain meshes such as TerraSim or Terra Vista. Rather than redesign these tools, this software has been designed with the capability to import meshes from either of these standards.

Once a virtual environment has been created which represents the physical environment in question, it can be loaded into the ray tracer for analysis.

3.2 Ray Tracing

The current ray tracing algorithm works by propagating rays in small incremental time steps and then checking if a ray has intersected with the virtual environment at each step. This interacting of the rays with surfaces captures the effects of wave propagation within the environment. Because the propagation is incremental, rays which curve based off of a set of differential equations can be implemented with relative ease.

Each ray represents an 'energy packet' of sound emanating from a source. Each ray contains spectral information about the relative energies against frequency. The energy based approach gives the ray spectral magnitudes physical meaning which can be measured and understood. The scene then becomes an energy equation which must be balanced, and the ray tracing engine becomes an accountant, keeping track of the energy in the environment. As rays travel, they lose energy based on the damping ratio of the atmosphere. When they interact with a surface, some energy is absorbed and transformed into heat, while the rest is reflected according to the law of reflection.



Figure 3: A ray trace of an environment with one occlusion. The source is marked by the red dots. Non source intersections are marked with a white dot. The receiver is located directly over the bright white patch. A uniform grid of rays was emitted from the receiver, producing the ring pattern.

In order to aid in a physical interpretation, the energy contained in a ray must be spread over a representative area. Then the ray can be interpreted as an intensity, which can be measured in the physical world. This representative area of a ray is inversely proportional to the number of rays within a certain area, or the ray density. An estimate

of the energy in an area can be determined by summing the total energy of all the rays in that area, and then dividing by the ray density in the area. By randomly distributing the rays over some ray distribution, a non-deterministic approximation of the total energy in the area is found. This is a Monte Carlo Integration estimator technique, and can be described as follows:

$$F_N = \frac{1}{N}\sum_{i=1}^{N}\frac{f(X_i)}{p(X_i)}$$

where F_N is the energy over the sampling range, X_i is a set of random variables drawn over an arbitrary PDF $p(x)$ valid over the sampling range, and $f(x)$ is the energy at location x [1]. At this point, the ray tracing method is a valid representation of energy transfer in a wave and may be used to characterize the scene.

In order to save on computation, rays which contain energy below a certain threshold are terminated early. While analytical methods could be used to determine the threshold number, such as the limit of human hearing, in reality it needs to be changed as a trade off based on what works well for the current scene complexity and content.

3.2.1 Adaptive Sampling

Another technique for reducing the computational burden of characterizing large scenes is to use adaptive sampling. By the Monte Carlo estimator above, areas in the scene that vary gradually may be captured with only a few samples. However areas with sharp gradients in energy need denser sampling to be captured correctly. Rather than trace the entire scene at the resolution necessary to capture the sharpest gradient in the system, an adaptive mesh can be utilized.



Figure 4: An adaptive ray sampling technique. The red circle represents the object being sampled.

The adaptive sampling technique used assumes that sound energy varies slowly over the surface of an object, but that the energy can vary rapidly at the edges of objects. Because of this, it is designed to search out object edges. An initial low resolution grid of rays is shot out from the listener. For adaptive sampling to work properly, it is important that each object in the scene is hit by at least one ray. The algorithm then looks for neighbor rays that did not intersect the same object. If a ray is found which meets this criterion, the area around the ray is subdivided and sampled. This algorithm is then recursively executed until an arbitrary resolution is met.

3.3 Echogram and Impulse Response Generation

When the rays intersect a sound source their energy and time of arrival are recorded in a data structure called an echogram. An echogram is a two dimensional data structure in time and frequency that records the 'echoes' as they are perceived by the source. The rays are grouped into bins in frequency and time to reduce the memory footprint of the recorded data. In the implemented version of the software, the time axis of an echogram is divided into 10ms increments and the frequency axis is divided into third octave bands. Research shows that human perception between distinct noises is greater than 10ms and that human hearing is divided in greater increments than third octave bands, so this is a reasonable way to divide the data [12] [13].



Figure 5: A representation of an echogram. The front to back axis represents frequency. The right to left axis represents time.

A sphere is defined around the receiver and is binned in elevation and azimuth into discrete 'windows'. This allows directional audio to be generated. One echogram is generated for each bin on the receiver which characterizes the wave propagation path to one source. Therefore, to completely characterize a scene, the number of echograms needed is equal to the number of bins around the receiver times the number of sources in the scene, times the number of receivers in the scene.



Figure 6: The bins (or windows) around the receiver (in this case a human head)

One set of echograms is recorded for each bin on this sphere. When the sound is characterized in a later step, it is the combined sound from each of these windows being transformed against the data in the corresponding echogram which produces the physically accurate 3D effect. One set of echograms are generated for each sound source in the scene.

Once a series of echograms has been generated, impulse responses need to be created from those echograms. Impulse responses contain information about how the environment modifies a source sound. They are applied to the audio data by the process of convolution.

Impulse responses differ from echograms in that an echogram is a 'coarse' two dimensional structure containing a sample only every 10ms where as an impulse response in a 'fine' single dimensional structure sampled at the system audio rate (usually 44100 Hz).

To generate an impulse response, the echogram is processed in frequency slices. For each frequency in the echogram, a time varying waveform is extracted. This waveform

represents the magnitude of the envelope of energy in the impulse response for that frequency. This envelope is used to modulate noise with unity spectral power in the given frequency band, and no power elsewhere. One band of noise is modulated for each frequency bin in the echogram.

Once frequency bands have been modulated, they each represent a narrowband impulse response at a given frequency. To generate a broadband impulse response, the narrowband responses are summed together. This broad band impulse response, or scene impulse response, represents the transfer function between the source and the receiver, and it can be used to transform the (anechoic) source sound to appear as though it emanated from the virtual environment.

3.4 Data Visualization

After the model has been run, data about the environment may be collected in several forms. By placing multiple receivers in the environment in a grid based pattern, information about the relative loudness at various points may be computed from the scene impulse response functions.

To actually auralize the scene for visualization, a representative sample of wind turbine noise is needed. This can either be recorded from the real world or synthesized using the wind turbine model. By calibrating the system, an accurate 3D representation of the wind farm noise can be played over a standard PC audio card. This is accomplished using a separate piece of software, called the 'audio streamer.'

4. Conclusion

The goal of this paper was to outline a Monte Carlo based ray tracing technique which has the ability to characterize the noise quality generated from wind farms. It was shown that a ray tracer can produce accurate results significantly faster than other numerical techniques, and can be adapted to various environmental conditions easier than analytical or empirical models. All major phenomena associated with mechanical wave propagation can be accounted for using the ray tracing technique.

Once a characterization is completed, it can be used to produce immersive 3D audio which can be used by non-technical personnel to judge the noise quality emanating from the wind farm. This software was designed to be synchronized with a visual display providing a fully immersive and comprehensive representation of the wind farm before it is built.

To date, a ray tracing engine which correctly recreates the direct sound path and specular reflections has been built and tested. Other phenomena as described in this paper are being implemented in an ongoing effort.

The audio streamer is implemented and tested, which allows for connection to a head tracking system. The system adjusts for head movement in real time, providing a fully immersive experience.

References

1. M. Pharr and G. Humphreys, *Physically Based Rendering: From Theory to Implementation, 2nd Edition* (Morgan Kaufman, Burlington, Massachusetts, 2010)
2. Okada, Masashi, Takao Onoye, and Wataru Kobayashi. "A Ray Tracing Simulation of Sound Diffraction Based on the Analytic Secondary Source Model." *Audio, Speech, and Language Processing, IEEE Transactions on* 20.9 (2012): 2448-2460.
3. Tsingos, N.; Carlbom, I.; Elko, Gary; Kubli, R.; Funkhouser, T., "Validating acoustical simulations in the Bell Labs Box," *Computer Graphics and Applications, IEEE* , vol.22, no.4, pp.28,37, Jul/Aug 2002
4. Steven G. Mattson, et al. "Physics Based Spatial Acoustics in Virtual Scenes" Proceedings of Nose-Con 2013. In Press
5. Wei Jun Zhu. "Modelling Of Noise From Wind Turbines." Diss. Mechanical Department, DTU
6. Lauterbach, Christian, Anish Chandak, and Dinesh Manocha. "Interactive sound rendering in complex and dynamic scenes using frustum tracing." *Visualization and Computer Graphics, IEEE Transactions on* 13.6 (2007): 1672-1679.
7. Bass, Jeremy, Penny Dunbabin, and Mike Anderson. "A CRITICAL APPRAISAL OF WIND FARM NOISE PROPAGATION FINAL REPORT." (2000).
8. Funkhouser, Thomas, et al. "A beam tracing method for interactive architectural acoustics." *The Journal of the Acoustical Society of America* 115 (2004): 739.
9. Kapralos, Bill, Michael Jenkin, and Evangelos Milios. "Sonel mapping: a probabilistic acoustical modeling method." *Building Acoustics* 15.4 (2008): 289-313.
10. Bill Gardner and Keith Martin, MIT Head Related Transfer Function Database, http://sound.media.mit.edu/resources/KEMAR.html
11. Wenzel, Elizabeth M. "Analysis of the role of update rate and system latency in interactive virtual acoustic environments." *Audio Engineering Society Convention* 103. 1997.
12. David R. Perrott, "Studies in the perception of auditory motion", in *Localization of Sound: Theory and Applications,* Amphora Press, CND-Groton, 1982, pp. 169—193, isbn 0-940728-03-6.
13. David R. Perrott and Juliana Tucker, "Minimum audible movement angle as a function of signal frequency and the velocity of the source", *J. Acoust. Soc. Am.* **83**(1988), no. 4, 1522—1527, issn 0001-4966.
14. Jens Blauert, Hilmar Lehnert, Jörg Sahrhage, and Holger Strauss, An interactive virtual-environment generator for psychoacoustic research. I: Architecture and implementation, *Acustica* 86(2000), 94—102, issn 0001-7884.
15. Christoph Pörschmann, 3-D Audio in Mobile Communication Devices: Methods for Mobile Head-Tracking. *JVRB - Journal of Virtual Reality and Broadcasting*, 4(2007), no. 13.
16. J. Arvo, P. Dutre, A. Keller, H. W. Jensen, A. Owen, M. Pharr, and P. Shirley, "Monte Carlo Ray Tracing: Siggraph 2003 Course 44," Course Notes from SIGGRAPH 2003, San Diego, California, US (2003).