

**Physically Based Sound for Computer Animation and
Virtual Environments**
Acoustic Transfer

Doug L. James
Stanford University

Please note: This document is ©2016 by Stanford University. This chapter may be freely duplicated and distributed so long as no consideration is received in return, and this copyright notice remains intact.

Acoustic Transfer

Doug L. James
Stanford University

TOPICS TO BE COVERED: *Sound pressure; acoustic wave equation; Helmholtz equation; boundary conditions for one-way coupling; transfer functions; solvers and precomputation; multipole expansions; fast evaluation; rendering details.*

1 Introduction: How loud is each mode?

At this point we know how to analyze the small vibrations of solids, and determine their modal oscillations when impacted. While we can use these modal amplitudes to synthesize an approximate sound, we don't really hear the modal amplitudes of the vibrating solids. What we hear are sound pressure waves which are radiated from the vibrating surface into the surrounding air where we are listening. Ignoring this fact leads to the wrong amplitudes for our modes, it just doesn't sound right.

A simple model that captures the different relative importance of each vibration mode is one where the synthesized sound at position \mathbf{x} and time t is given by the following expression

$$\text{sound}(\mathbf{x}, t) = \sum_j a_j(\mathbf{x}) q_j(t) = \mathbf{a}(\mathbf{x})^T \mathbf{q}(t) \quad (1)$$

where the synthesized sound is the weighted linear superposition of the modal oscillations, where each mode j is weighted by an acoustic transfer function $a_j(\mathbf{x})$ which specifies how loud that mode would be at the specified position if it were undergoing a unit amplitude vibration. It turns out that acoustic transfer functions for different modes can vary by hundredfold factors, and lead to very noticeable sound changes. Furthermore they capture complex spatial variations which lead to rich sound fields around the object. These effects are due to the complex interactions of the surface vibrations in the geometric shape with the diffracting pressure waves of a certain wavelength in the surrounding air.

In the remainder of this chapter we will describe the basic mathematical background for sound wave radiation based on the linear wave equation, describe solvers for computing the acoustic transfer functions, as well as discuss representations for efficient evaluation of acoustic transfer functions for sound rendering. Finally, it's worth pointing out that the fundamentals of acoustic transfer will be used when modeling sound from objects that vibrate in more complex ways than linear modal vibrations, such as liquids.

2 Wave Equation Basics

The small pressure fluctuations associated with soundwaves are well approximated by the linear wave equation in many cases of interest to sound rendering. If the acoustic pressure fluctuations are described by $p(\mathbf{x}, t)$ then it must satisfy the following wave equation in the surrounding fluid

$$\nabla^2 p - \frac{1}{c^2} \frac{\partial^2 p}{\partial t^2} = 0, \quad (2)$$

where c is the speed of sound in the fluid (approximately 340 m/s in air), and the Laplacian is given by

$$\nabla^2 p = \frac{\partial^2 p}{\partial x^2} + \frac{\partial^2 p}{\partial y^2} + \frac{\partial^2 p}{\partial z^2},$$

and essentially measures the curvature of the pressure field.

For radiation problems where we were trying to compute the sound given off by a vibrating surface into the surrounding air or fluid, we need to have a boundary condition that relates the surface acceleration to the pressure quantity that we were trying to compute (see Figure 1). It turns out that the boundary condition that we need relates the normal derivative of the pressure (at $\mathbf{x} \in \Gamma$ and t),

$$\partial_n p(\mathbf{x}, t) \equiv \frac{\partial p(\mathbf{x}, t)}{\partial n} \equiv \nabla p(\mathbf{x}, t) \cdot \mathbf{n}(\mathbf{x}), \quad (3)$$

to the surface acceleration via the boundary condition

$$\frac{\partial p}{\partial n} = -\rho a_n(\mathbf{x}, t), \quad \mathbf{x} \in \Gamma, \quad (4)$$

where $a_n(\mathbf{x}, t) = \mathbf{n}(\mathbf{x}) \cdot \mathbf{a}(\mathbf{x}, t)$ denotes the normal acceleration of position \mathbf{x} on the object's surface; and ρ refers to density of the surrounding medium, assumed constant (1.2041 kg/m³ for air at standard temperature and pressure).

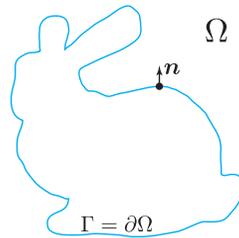


Figure 1: Radiation problem for a vibrating solid with external interface, Γ , in contact with the surrounding acoustic medium, e.g. air, which exists in the unbounded domain, Ω . We seek to compute the acoustic transfer function $p(\mathbf{x})$ in Ω associated with harmonic vibrations of Γ , corresponding to the Neumann radiation boundary condition.

While we could time step the wave equation (to solve an initial value problem) to compute the pressure vibrations produced by a specific sequence of surface vibrations, the computational cost would be prohibitive for sound rendering. That is because time stepping the wave equation on large 3-D volumetric domain at audio rates with small time steps would involve significant memory and computational costs to produce a sound. In practice we can do much better, and avoid the combined space-time complexity by considering the problem in the frequency domain.

3 Frequency-domain Wave Radiation

To approximate the radiated sound waves amplitude due to a modal surface vibration, we can exploit the strong frequency localization of each mode individually. We can also ignore damping effects which tend to very weakly damp modes over a single vibration in most cases, e.g., for a ringing object. Given a harmonic vibration of the surface boundary, the resulting pressure solution to the wave equation will also have a pure harmonic form.

Using complex numbers to represent oscillations: We can represent a unit-amplitude harmonic oscillation at natural frequency ω , using the complex number $e^{+i\omega t} = \cos(\omega t) + i \sin(\omega t)$. By linearity of the wave equation and boundary conditions, we can then use the exponential factor $e^{+i\omega t}$, to represent any time-harmonic structure. Different phases can be introduced using an $e^{i\phi}$ factor, since $e^{+i\omega t} e^{i\phi} = \cos(\omega t + \phi) + i \sin(\omega t + \phi)$.

Harmonic pressure: Given the harmonic forcing, the acoustic pressure will also have a harmonic form,

$$p(\mathbf{x}, t) = p(\mathbf{x})e^{+i\omega t},$$

where $p(\mathbf{x})$ is a complex-valued pressure field. The amplitude of the pressure wave at \mathbf{x} is given by the modulus, $|p(\mathbf{x})|$. Using Euler's formula we can write $p(\mathbf{x}) = |p|e^{i\phi}$, and see that the other $e^{i\phi}$ factor affects the phase of the wave at \mathbf{x} , since

$$p(\mathbf{x}, t) = p(\mathbf{x}) e^{+i\omega t} = |p(\mathbf{x})| e^{+i(\omega t + \phi(\mathbf{x}))} = |p(\mathbf{x})| \cos(\omega t + \phi) + i |p(\mathbf{x})| \sin(\omega t + \phi).$$

Helmholtz equation: Substituting the harmonic pressure $p(\mathbf{x}, t) = p(\mathbf{x})e^{+i\omega t}$ into the wave equation, and taking derivatives with respect to time, we get

$$0 = \nabla^2 p(\mathbf{x}) e^{+i\omega t} - \frac{1}{c^2} p(\mathbf{x}) \left(\frac{\partial^2}{\partial t^2} e^{+i\omega t} \right) \tag{5}$$

$$= \left(\nabla^2 p(\mathbf{x}) - \frac{\omega^2}{c^2} p(\mathbf{x}) \right) e^{+i\omega t}, \tag{6}$$

and since $e^{+i\omega t}$ is never zero, the term in brackets must be. We thus obtain the frequency-domain wave equation, or *Helmholtz equation*,

$$(\nabla^2 + k^2) p(\mathbf{x}) = 0, \tag{7}$$

where $k = \frac{\omega}{c} = \frac{2\pi}{\lambda}$ is the wave number, which is proportional to the natural frequency, and inversely proportional to the wavelength. The solution to this Helmholtz equation (with suitable boundary conditions) is a complex-valued pressure field defined everywhere outside the radiating object.

Radiation boundary conditions: To solve for the sound wave radiated by the vibrating surface, we need boundary conditions (BCs) on the air domain—both on the object's surface and at infinity (where the waves are going). To convert the time-domain Neumann BC on the object's surface (4) to its harmonic counterpart, we will need an expression for the normal surface acceleration, $a_n(\mathbf{x}, t)$. Consider infinitesimal surface displacements due to a single vibration mode, and its successive derivatives:

$$\mathbf{u}(\mathbf{x}, t) = \bar{\mathbf{u}}(\mathbf{x}) e^{+i\omega t} \implies u_n = \bar{u}_n e^{+i\omega t} \tag{8}$$

$$\mathbf{v}(\mathbf{x}, t) = \frac{d\mathbf{u}}{dt}(\mathbf{x}, t) = i\omega \bar{\mathbf{u}}(\mathbf{x}) e^{+i\omega t} \implies v_n = i\omega \bar{u}_n e^{+i\omega t} \tag{9}$$

$$\mathbf{a}(\mathbf{x}, t) = \frac{d^2\mathbf{u}}{dt^2}(\mathbf{x}, t) = (i\omega)^2 \bar{\mathbf{u}}(\mathbf{x}) e^{+i\omega t} \implies a_n = -\omega^2 \bar{u}_n e^{+i\omega t}, \tag{10}$$

where $\bar{u}_n = \mathbf{n} \cdot \bar{\mathbf{u}}$ is the normal component of the displacement mode. Note that the surface normal $\mathbf{n}(\mathbf{x})$ can be considered as constant for infinitesimal surface vibrations. Therefore a suitable time-independent

radiation boundary condition can be obtained:

$$\frac{\partial p(\mathbf{x}, t)}{\partial n} = -\rho a_n(\mathbf{x}, t) \quad (11)$$

$$\frac{\partial p(\mathbf{x})}{\partial n} e^{i\omega t} = \rho \omega^2 \bar{u}_n e^{i\omega t} \quad (\text{cancelling nonzero } e^{i\omega t} \text{ factor}) \quad (12)$$

$$\frac{\partial p(\mathbf{x})}{\partial n} = \rho \omega^2 \bar{u}_n(\mathbf{x}), \quad (13)$$

and since it specifies $\frac{\partial p}{\partial n}$ it is also referred to as a Neumann boundary condition.

To model physically meaningful wave radiation we also need a boundary condition at infinity that captures the notion that the waves are radiating outwards to infinity from a source at the object, and not the other way around. The *Sommerfeld radiation condition* essentially says that the amplitude of the waves goes to zero as the wave goes to infinity, and also the waves radiate outward from the object. This condition avoids other nonphysical mathematical solutions where the wave approaches the object from a singular solution at infinity. We will skip a precise mathematical statement of this radiation conditions here, and only consider solutions that satisfy these conditions hereafter.

Acoustic Transfer BVP: We refer to the pressure solution $p(\mathbf{x})$ of each mode's radiation problem as the *acoustic transfer function*. For a mode at frequency ω , the transfer function $p(\mathbf{x})$ is defined as the solution to the Helmholtz equation (7), subject to the mode-driven Neumann BC (13), and Sommerfeld radiation condition.

Visualizing solutions: We can visualize the waves produced by a unit amplitude modal vibration by looking at the real or imaginary part of $p(\mathbf{x})e^{i\omega t}$ (see Figure 2).

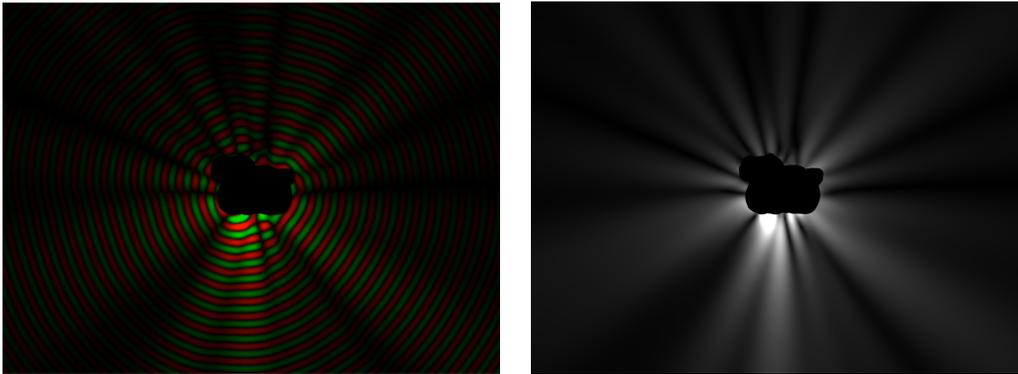


Figure 2: A visualization of the acoustic transfer function representing (left) pressure fluctuations (as $\text{Re}(p(\mathbf{x}))$) due to sound waves radiated from a surface undergoing harmonic vibrations due to modal oscillations. (Right) a visualization of the absolute value of the acoustic transfer function, $|p(\mathbf{x})|$, which represents the sound amplitude due to the unit modal oscillation. These acoustic transfer functions can be precomputed and efficiently represented for rapid evaluation during sound rendering.

Rendering a sound: If we are only interested in the amplitude of this sound source at a given position, then we only need to consider the absolute value of the acoustic transfer function $|p(\mathbf{x})|$. Referring to the

earlier equation (1), by identifying the weight $a(\mathbf{x})$ for mode j as its absolute value acoustic transfer function $p_j(\mathbf{x})$, we now have a simple model for sound rendering given by the following expression

$$p(\mathbf{x}, t) = \sum_j |p_j(\mathbf{x})| q_j(t). \quad (14)$$

Given that $\mathbf{q}(t)$ are easy to compute, the remaining challenge remains to efficiently evaluate the acoustic transfer functions.

More sophisticated models of rendering the sound, such as those that include time delays, phase effects, or head related transfer functions (HRTFs), will be considered in a later section. We now consider simple multipole sound sources, then how to solve for acoustic transfer functions, including a practical approach that is easy to implement, and finally techniques for representing the computed transfer functions so that they can be cheaply evaluated for rendering, as well as have low memory requirements if necessary.

Aside (Radiation efficiency): A measure of how well a mode converts its surface vibrational energy into radiated acoustic energy is estimated by the so-called *radiation efficiency* [4]. It can be computed as the ratio of the radiated power Π_{rad} to a power-like measure of the surface vibration given by

$$\sigma = \frac{\Pi_{rad}}{\rho c S \langle v_n^2 \rangle_S} = \frac{\Pi_{rad}}{\rho c S \omega^2 \langle \bar{u}_n^2 \rangle_S} \quad (15)$$

where S is the surface area of the object, and $\langle \bar{u}_n^2 \rangle_S$ is the average of the squared normal displacement over the surface—computing the radiated power Π_{rad} is beyond the scope of our discussion here (see [4] or [9]). An interesting phenomena is that the radiation efficiency of each mode can vary greatly and non-monotonically with increasing mode index (see Figure 3).

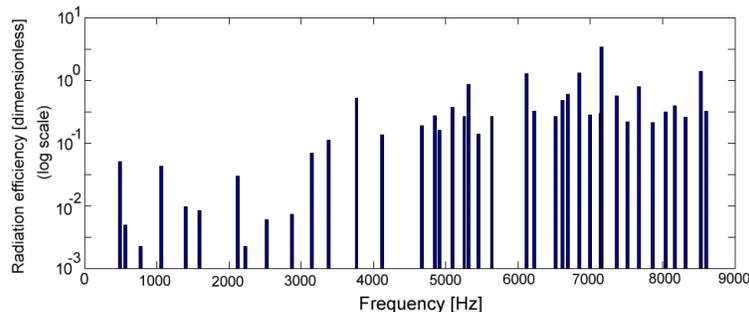


Figure 3: Radiation efficiency for the dragon shell model from [8].

4 Multipole Sound Sources

Multipole pressure fields are an important class of solutions to the wave equation, which are associated with point-like (singular) sound sources. You will find them particularly useful for constructing solvers, and building fast representations of acoustic transfer functions for sound rendering.

Hello Monopole: The simplest multipole field is a *monopole* sound source. It represents expanding spherical waves

$$p(\mathbf{x}) = \frac{e^{-ikr}}{4\pi r} \quad (\text{Monopole source}) \quad (16)$$

due to a point-like pulsation at the source point, \mathbf{x}_0 ; here $r = \|\mathbf{x} - \mathbf{x}_0\|_2$ is the distance from the source location, \mathbf{x}_0 , and $k = \frac{\omega}{c} = \frac{2\pi}{\lambda}$ is the wavenumber. The wave strength falls off as $\frac{1}{r}$, and have wavelength λ . To see this better, multiply by the time-harmonic factor, $e^{+i\omega t}$, to obtain $e^{i(\omega t - kr)}/(4\pi r)$. Observe that for any particular value of the exponential's argument, $\omega t - kr = B \in \mathbb{R}$, we have a spherical wavefront expanding outward at speed c with radius $r(t) = \frac{\omega}{k}t - \frac{B}{k} = ct - \frac{B}{k}$, and falling off in amplitude as $\frac{1}{4\pi r}$.

Like all multipoles, this monopole $p(\mathbf{x})$ satisfies the Helmholtz equation everywhere, except at the source location where a singularity occurs. For the monopole, we have

$$\nabla^2 p(\mathbf{x}) + k^2 p(\mathbf{x}) = \delta(\mathbf{x} - \mathbf{x}_0), \quad (17)$$

where $\delta()$ is the *Dirac delta function*, which is zero everywhere except when its argument is zero (here when $\mathbf{x} = \mathbf{x}_0$), in which case it is infinite; it also has the interesting property that its integral is one for any region that contains the source location,

$$\int_{\Omega} \delta(\mathbf{x} - \mathbf{x}_0) d\Omega_{\mathbf{x}} = \begin{cases} 1, & \mathbf{x}_0 \in \Omega \\ 0, & \mathbf{x}_0 \notin \Omega \end{cases} \quad (18)$$

The Multipole Basis Functions: Higher-order multipole sources have more complex spatial patterns, and are very useful for describing the far-field structure of sound radiation from objects (see Figure 4). In general, there are an infinite number of multipole basis functions to the wave equation, which we can write as

$$S_n^m(\mathbf{x} - \mathbf{x}_0) = h_n^{(2)}(kr) Y_n^m(\theta, \phi), \quad n = 0, 1, 2, \dots, \quad m = -n, \dots, n, \quad (19)$$

where \mathbf{x}_0 is the location of the multipole source, and \mathbf{x} is the evaluation position; the radial vector $\mathbf{r} = \mathbf{x} - \mathbf{x}_0$ is represented in spherical coordinates as (r, θ, ϕ) ; $h_n^{(2)} \in \mathbb{C}$ are spherical Hankel functions of the second kind; and $Y_n^m \in \mathbb{C}$ are spherical harmonics.

The spherical Hankel functions $h_n^{(2)}(kr)$ represent radial solutions to the wave equation, with the functions of the “second kind” representing out-going waves for us (the “first kind” functions actually represent in-going waves for our $e^{+i\omega t}$ convention). The first few of these functions are as follows:

$$h_0^{(2)}(kr) = \frac{i}{kr} e^{-ikr} \quad (20)$$

$$h_1^{(2)}(kr) = -\frac{kr - i}{(kr)^2} e^{-ikr} \quad (21)$$

$$h_2^{(2)}(kr) = -i \frac{(kr)^2 - 3ikr - 3}{(kr)^3} e^{-ikr} \quad (22)$$

$$h_3^{(2)}(kr) = \frac{(kr)^3 - 6i(kr)^2 - 15kr + 15i}{(kr)^4} e^{-ikr}. \quad (23)$$

Numerical evaluation of $h_n^{(2)}(kr)$ for larger n can be done efficiently using the recurrence relation

$$h_{n+1}^{(2)}(kr) = \frac{2n+1}{kr} h_n^{(2)}(kr) - h_{n-1}^{(2)}(kr). \quad (24)$$

Derivatives required for computing boundary gradients can be obtained with

$$\frac{d}{dz} h_n^{(2)}(z) = \frac{1}{2} \left[h_{n-1}^{(2)}(z) - \frac{h_n^{(2)}(z) + z h_{n+1}^{(2)}(z)}{z} \right]. \quad (25)$$

Spherical harmonics are common in graphics, and we use the complex-valued form given by

$$Y_n^m(\theta, \phi) = \sqrt{\frac{2n+1}{4\pi} \frac{(n-m)!}{(n+m)!}} P_n^m(\cos \theta) e^{im\phi}, \quad m = -n, \dots, n, \quad (26)$$

where P_n^m is the associated Legendre polynomial [1, 14]. These functions can be computed individually, or using recurrence relations. For example, in [15] all $Y_n^m(\theta, \phi)$ are computed recursively in two passes by first computing $Y_m^m(\theta, \phi)$, $m = 0 \dots p_i$ using the recurrence relation between Y_m^m and Y_{m-1}^{m-1} [13], then computing Y_n^m , $n = m+1 \dots \bar{n}_i$ using the recurrence relation between Y_n^m and Y_{n-1}^m .

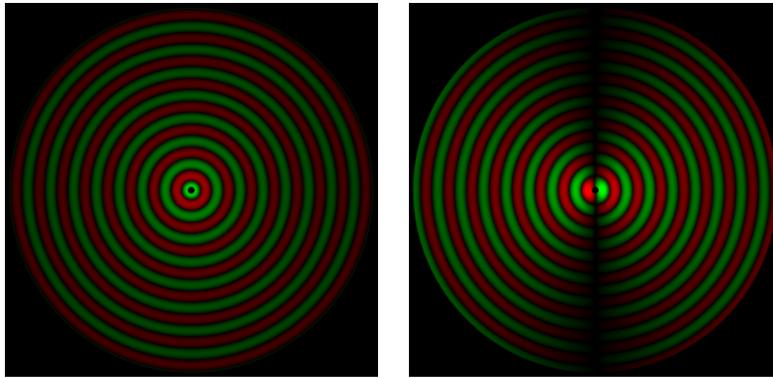


Figure 4: Multipole sound sources: (Left) Monopole spherical source, $S_0^0(\mathbf{x})$, and (Right) one of the dipole basis functions $S_1^m(\mathbf{x})$.

Multipole Expansions: By taking linear combinations of the multipole basis functions, one can construct spherical multipole wave expansions of the form,

$$p(\mathbf{x}; \mathbf{x}_0) = \sum_{n=0}^{\bar{n}} \sum_{m=-n}^n S_n^m(\mathbf{x} - \mathbf{x}_0) c_n^m \quad (27)$$

where $c_n^m \in \mathbb{C}$ are multipole expansion coefficients, and S_n^m are the multipole basis functions (singular, radiating solutions to the Helmholtz equation). We would call this a multipole expansion of order \bar{n} . Given the multipole coefficients, c_n^m , one can efficiently evaluate the approximation at any listening position \mathbf{x} . Recall that these solutions represent outgoing waves which decay as they extend towards infinity, and are precisely the solutions needed for sound rendering problems. Any multipole expansion of the form (27) will satisfy the Helmholtz wave equation everywhere except at the singular source point \mathbf{x}_0 . Later we will find that we can use linear combinations of these multiple sources to describe acoustic transfer functions for efficient rendering, and that we can even use these sources for easy solution of the Helmholtz equation itself. Lastly, we mention that we will sometimes wish to express (27) using a single summation over a generalized summation index j for some basis functions S_j with coefficients c_j as follows,

$$p(\mathbf{x}; \mathbf{x}_0) = \sum_j S_j(\mathbf{x} - \mathbf{x}_0) c_j \quad (28)$$

for some suitable summation limits.

5 Solvers

Now that we know what an acoustic transfer function is and we know the equation that it must satisfy, we can go ahead and compute it. One wrinkle is that the acoustic transfer function is specified on an infinite domain outside the object. We certainly don't want to store an infinite number of function samples, so it's important to represent that solution in a way that is convenient for sound rendering. In this section we will summarize three types of solvers which you might use to precompute acoustic transfer function quantities, then we will revisit the issue of representation in a later section. As you'll see, the representation and solution of these functions is closely related.

5.1 Equivalent Source Method

By far my favorite technique for computing acoustic transfer functions for sound rendering is to use the equivalent source method it's also referred to by different names, such as "source simulation" (by Ochmann [11]), or the method of Trefftz, and is described in various references for the Helmholtz problem. The benefit of this method is that it is easy to implement, can provide sufficient accuracy, and that its solution representations are immediately useful for rapid evaluation and sound rendering. The downsides are that it has trouble achieving reasonable accuracy at higher frequencies, and it only works for volumetric objects—not thin objects, like a cymbal.

The basic idea is that if we place multipole sound sources inside the vibrating object, then they are guaranteed to satisfy the wave equation outside the object as well as the Sommerfeld radiation condition. All of the error in the approximation is then in the boundary condition on the vibrating surface of the object. By estimating multipole coefficients and possibly specific locations to place the multipoles, we can minimize the boundary condition error and therefore obtain an approximation to the acoustic transfer function. Since these approximations typically involve modest numbers of multipoles, they can also be efficiently evaluated during sound rendering, and they do not have large storage requirements.

Consider the case where we have several multipole sources located inside the object, possibly at different positions, and of various expansion orders, but with as yet undetermined expansion coefficients. We can write this pressure solution model as a sum over all multipole basis functions used,

$$p(\mathbf{x}) = \sum_{j=1}^N \psi_j(\mathbf{x}) c_j, \quad (29)$$

assuming there are N basis functions in total. Here the $\psi_j(\mathbf{x})$ basis functions are none other than $S_n^m(\mathbf{x} - \mathbf{x}_0)$ for some particular values of n , m , and source location \mathbf{x}_0 . Rewriting this in a vector form, we can say that

$$p(\mathbf{x}) = \boldsymbol{\psi}(\mathbf{x}) \cdot \mathbf{c}, \quad (30)$$

where $\mathbf{c}^T = [c_1, c_2, \dots, c_N]$ and $\boldsymbol{\psi}(\mathbf{x})^T = [\psi_1(\mathbf{x}), \psi_2(\mathbf{x}), \dots, \psi_N(\mathbf{x})]$. Here the undetermined coefficients are represented by the complex valued vector \mathbf{c} . We can determine \mathbf{c} by minimizing the error in the surface boundary condition. There are many different ways to formulate this optimization problem which have different properties (for example, see work by Ochmann [11, 12]). A simple and effective way is to consider the boundary condition at numerous uniformly sampled locations \mathbf{x}_i on the boundary Γ . By minimizing

the error in the Neumann boundary condition (13) at these points, we obtain many linear constraints on the coefficients \mathbf{c} . Specifically we equate the normal derivative of our pressure expansion on the boundary,

$$\frac{\partial p(\mathbf{x})}{\partial n} = \sum_{j=1}^N \frac{\partial \psi_j(\mathbf{x})}{\partial n} c_j = \frac{\partial \psi}{\partial n}(\mathbf{x}) \cdot \mathbf{c} \quad (31)$$

to the BC RHS at each of M sample points to obtain the complex-valued system of equations,

$$\frac{\partial \psi}{\partial n}(\mathbf{x}_1) \cdot \mathbf{c} = \rho \omega^2 \bar{u}_n(\mathbf{x}_1), \quad (32)$$

$$\frac{\partial \psi}{\partial n}(\mathbf{x}_2) \cdot \mathbf{c} = \rho \omega^2 \bar{u}_n(\mathbf{x}_2), \quad (33)$$

$$\vdots \quad (34)$$

$$\frac{\partial \psi}{\partial n}(\mathbf{x}_M) \cdot \mathbf{c} = \rho \omega^2 \bar{u}_n(\mathbf{x}_M). \quad (35)$$

In matrix form we have a rectangular system of equations

$$\mathbf{A} \mathbf{c} = \mathbf{b} \quad (36)$$

where \mathbf{A} is an M -by- N matrix of multipole basis function normal derivatives, and \mathbf{b} is a complex-valued M -vector of the sampled $\rho \omega^2 \bar{u}_n(\mathbf{x})$ boundary values. By choosing sufficiently many boundary samples (large enough M) we obtain an overdetermined linear system for the multipole coefficients which we can solve using a least squares method. Since the system matrix \mathbf{A} can be highly ill-conditioned, appropriate numerical methods must be used. One effective technique is to use a truncated singular value decomposition (TSVD) to compute an approximate \mathbf{C} vector while discarding small singular values [5, 8].

For higher frequency vibration modes, one typically finds that increasing numbers of multipole coefficients are required to obtain the same relative error [8]. This is due to the increased difficulty and higher wavelength structure of these higher frequency acoustic transfer functions. Fortunately in practice one does not require high accuracy solutions to these wave radiation problems to provide convincing results for sound rendering. There are several reasons for this. First, a seemingly high numerical error such as 10% relative error may only result in a small sound pressure air when considered on the logarithmic decibel scale. Furthermore many of these errors can be difficult to discern because it is not just a single mode that is often listen to but the superposition of many modes, and therefore errors in an individual mode can be masked to some extent. The reason for allowing larger errors in the boundary conditions, is that more detailed structure in the boundary conditions can be associated with near field detail in the pressure field which does not radiate effectively into the far field. Since most of the far field structure is dominated by the lowest order multipole moments, this can be acceptable in practice.

5.2 Boundary Integral Solvers

Perhaps the most widely used technique for solving the linear wave equation for exterior radiation problems is to use a boundary integral formulation where singular sources are placed on the boundary of the object itself. Such boundary integral formulations are often solved using discretizations such as the boundary element method (BEM) [3, 10], which is a workhorse method of computational acoustics. Several open-source BEM implementations are available, e.g., BEM++ (<http://www.bempp.org>).

Special care is required to deal with thin objects, such as shells which are common in sound rendering since they often are very efficient radiators. Alternate boundary integral formulations are used for these cases, and can be mathematically complicated to implement. Furthermore, since discretizations of the boundary integral equations such as using the boundary element method, lead to dense coefficient matrices with $O(N^3)$ solution costs for N boundary degrees of freedom, more advanced solution techniques are required to process detailed geometry and/or models with many radiating degrees of freedom. The most popular method for solving high degree of freedom radiation problems are those based on the fast multipole method for the Helmholtz equation [6]. Several commercial implementations are widely available, e.g., the FastBEM implementation is used in [2, 15]. In all of these boundary integral formulations, one inputs the vibration boundary data for the Neumann boundary condition, and obtains the resulting pressure solution (or related data) on the boundary. One can then use this boundary data to evaluate the solution anywhere, the cost of the evaluation may still be large without an efficient representation, as discussed later.

5.3 Spatial Discretizations

A more direct way to approximate the wave equation solutions around the object is to actually mesh the volumetric domain itself. Solvers based on the finite difference method, or finite volume method, or finite element method, are common, and can be used to solve the Helmholtz equation. We refer the interested reader to the following excellent books [7, 10]. These techniques provide sufficient accuracy, but can be challenging to implement correctly due to several factors. Meshing of the domain outside the vibrating solid is an additional complication, not present for the previous methods. In practice only a small region around the object can be meshed due to the cubic growth in the number of elements required with increasing domain volume. Furthermore, finer meshes are required to resolve finer wavelengths at higher frequencies. To avoid reflections from the boundary where the mesh ends, techniques based on perfectly match layers, infinite elements, or other non-reflecting boundaries must be used (see [7]). Essentially these techniques cause the radiated wave to dissipate due to strong damping in this outer region. To evaluate the solution at an arbitrary location in the exterior listening region, some other representation is required since it is usually undesirable to store the solution at all volumetric grid points for all modes, and the limited meshed region is also a problem.

5.4 Testing It Out: Analytical Test Cases

When implementing your own acoustic transfer solver, such as the equivalent source method described earlier, or using a third-party code, it's important to debug your implementation on some known test cases. Fortunately it's easy to cook up endless test cases. The easiest thing to do is to simply use a multipole source with known coefficients placed at some location *inside* a closed mesh, such as a sphere (or a bunny, or whatever you have lying around). This multipole source gives you a pressure field everywhere $p(\mathbf{x})$. You can then evaluate the input Neumann boundary condition data by numerically evaluating $\partial_n p(\mathbf{x})$ at suitable boundary locations. Since you know the pressure values of your chosen multipole source on the boundary and anywhere in the exterior domain, you will be able to evaluate the accuracy of values computed by your solver, and hopefully detect any implementation bugs. Note that you should avoid placing your test sources too close to the boundary, since they can cause numerical singularities which can be difficult to resolve and might cause other computational problems.

6 Representations for Sound Rendering

The solvers described so far compute answers which may only give the pressure indirectly, e.g., in terms of an integral on the object’s boundary, or only in a small domain around the object, or via a multipole expansion which can be evaluated outside the object. Clearly given the infinite size of the domain where sound is being radiated, one cannot sample it everywhere. We now describe techniques for representing the acoustic transfer functions that are suitable for rapid evaluation during sound rendering. Keep in mind that acoustic transfer functions typically need to be evaluated hundreds or thousands of times, per object, per mode, per listening position. Keeping these evaluation costs down is important for practical implementations, especially when many objects are being evaluated.

6.1 Equivalent-Source Representations: A free lunch

The equivalent source method provides an all-in-one technique for solving for acoustic transfer functions, and generating a practical representation. The downside of this representation is that it can be difficult to obtain high-accuracy approximations, especially at higher frequencies, without resorting to numerous multipole coefficients in the solver. Increased coefficient numbers are also prohibitive due to the expensive dense least-squares solution methods employed.

A second deficiency of multi-point multipole expansions for the equivalent source method, is that that technique does not apply to objects which are not volumetric in nature, e.g., thin shells which do not have interior volumes to place internal multiple sources. This limitation was overcome in “Precomputed Acoustic Transfer” [8] by first solving the Helmholtz problem using BEM, then evaluating the pressure $p(\mathbf{x})$ on an enclosing offset surface, and then fitting a many-point dipole approximation to the boundary data (see Figure 5). The method is complicated, but achieves very low dipole approximations to the radiation fields, and supported real-time rendering in 2006.

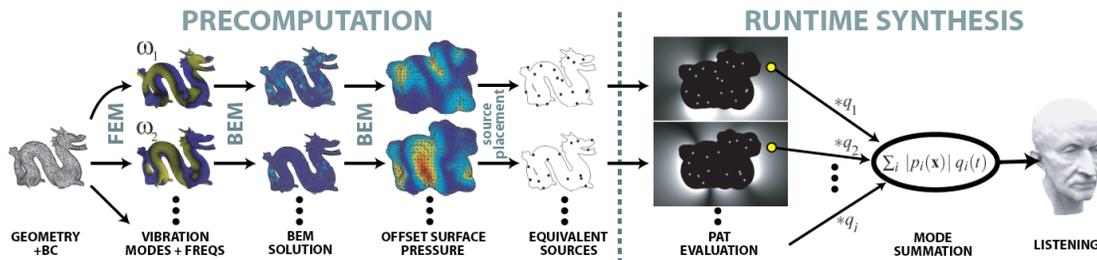


Figure 2: Overview of Precomputed Acoustic Transfer (PAT)

Figure 5: Overview of “Precomputed Acoustic Transfer” [James et al. 2006] which approximated radiation using a multipoint multipole approximation to BEM pressure data on an offset surface.

6.2 Boundary Integral Representations

Solutions computed using boundary integral methods can be evaluated to compute the pressure solution at any point outside the object, e.g., at the listening position. The basic idea of evaluating the acoustic transfer function in the exterior domain using boundary data can be demonstrated using the so-called Kirchhoff integral on a closed volumetric surface, Γ . The acoustic transfer pressure at any point \mathbf{x} in the exterior

domain outside the vibrating object is given by the *Kirchhoff integral formula*

$$p(\mathbf{x}) = \int_{\Gamma} \left[G(\mathbf{x}; \mathbf{y}) \frac{\partial p}{\partial \mathbf{n}}(\mathbf{y}) - \frac{\partial G}{\partial \mathbf{n}}(\mathbf{x}; \mathbf{y}) p(\mathbf{y}) \right] d\Gamma_{\mathbf{y}} \quad (37)$$

where $G(\mathbf{x}; \mathbf{y})$ is the free-space Helmholtz Green's function

$$G(\mathbf{x}; \mathbf{y}) = \frac{e^{-ik\|\mathbf{x}-\mathbf{y}\|}}{4\pi\|\mathbf{x}-\mathbf{y}\|}, \quad (38)$$

and $\frac{\partial G}{\partial \mathbf{n}}$ is its normal derivative. Given the outward-pointing normal derivative of the pressure field $\partial_n p$ on Γ (e.g., given by the Neumann boundary condition), and the pressure value p on Γ , one can evaluate this integral to obtain $p(\mathbf{x})$ in the exterior region. Different integral formulations can use different boundary data types, and more sophisticated integral formulations are required to handle geometrically thin objects, such as thin shells.

Such an approach can provide accurate evaluation of the solution, but it can be expensive for boundary integrals over geometrically complex surfaces. For example, if these boundary integral evaluations sum over N element contributions, the evaluation of *each mode* has an inherent $O(N)$ cost, and will require $O(N)$ storage.

For improved evaluation performance, more sophisticated implementations based on the fast multipole method for the Helmholtz equation can be used to evaluate pressure in the exterior domain using the method's tree data structure. Unfortunately, such implementations can be very complicated, and worse they involve additional memory overhead which can be problematic when considering storage of all mode transfer functions for many simulated objects.

6.3 Spatial Discretizations and the Kirchhoff Integral Representation

Volumetric discretizations on restricted domains, such as with the finite difference method with a perfectly matched layer, are not immediately suitable representations for rendering since one can not evaluate the solution everywhere in the exterior domain. Also, explicit storage of volumetric sound fields does not scale to support hundreds of modes for multiple objects. In practice these solutions can be used to compute a different representation for rendering. The most common technique is based on the Kirchhoff integral formula (37) mentioned earlier. For example, given the volumetric solution in some region, you can extract the pressure and its normal derivative on a closed N-triangle iso-surface surrounding the object. The acoustic transfer function pressure $p(\mathbf{x})$ can then be evaluated for rendering using the Kirchhoff integral method mentioned earlier. Therefore the volumetric field values can be discarded once the values on the Kirchhoff iso-surface have been extracted. Again, the cost of evaluating the Kirchhoff integral can be high for detailed boundary meshes, and therefore a faster approach may be desired. Yet another approach is to convert the volumetric solution to a multipole expansion, as follows.

6.4 Multipole Source Representation

Perhaps the simplest and most practical representation for vibrating sound sources is simply a multipole expansion, with as many coefficients are required to obtain sufficient accuracy [15]. The method also has the added benefit that level-of-detail techniques can be used to only render as many multipole orders as needed to obtain sufficient accuracy, which is useful for very distant sources. These expansions can be estimated in different ways, and we will describe an approach based on the Kirchhoff integral here. Given

the pressure and pressure derivative on an iso-surface surrounding the object, the values of the multipole coefficients for the expansion can be computed as described in [15], as follows. One can use a free-space Green's function identity to expand it in a series of singular and regular basis functions [6],

$$G(\mathbf{x}; \mathbf{y}) = ik \sum_{n=0}^{\infty} \sum_{m=-n}^n S_n^m(\mathbf{x} - \mathbf{x}_0) R_n^{-m}(\mathbf{y} - \mathbf{x}_0), \quad (39)$$

where S_n^m are the singular spherical Helmholtz basis function described previously; R_n^m is the non-singular basis function,

$$R_n^m(\mathbf{r}) = j_n(kr) Y_n^m(\theta, \phi) \quad (40)$$

where $j_n \in \mathbb{R}$ are the *spherical Bessel functions*. The expansion (39) is done about a fixed point \mathbf{x}_0 , which must satisfy $\|\mathbf{x} - \mathbf{x}_0\| > \|\mathbf{y} - \mathbf{x}_0\|$ to ensure that (39) converges absolutely and uniformly, e.g., one can place \mathbf{x}_0 at the object's center of mass. Substituting (39) into the Kirchhoff integral formula (37), we can rearrange terms to obtain the multipole expansion of $p(\mathbf{x})$:

$$\begin{aligned} p(\mathbf{x}) &= \int_{\Gamma} ik \left[\sum_{n=0}^{\infty} \sum_{m=-n}^n S_n^m(\mathbf{x} - \mathbf{x}_0) R_n^{-m}(\mathbf{y} - \mathbf{x}_0) \frac{\partial p}{\partial \mathbf{n}}(\mathbf{y}) \right. \\ &\quad \left. - p(\mathbf{y}) \sum_{n=0}^{\infty} \sum_{m=-n}^n S_n^m(\mathbf{x} - \mathbf{x}_0) \frac{\partial R_n^{-m}}{\partial \mathbf{n}}(\mathbf{y} - \mathbf{x}_0) \right] d\Gamma_{\mathbf{y}} \\ &= \sum_{n=0}^{\infty} \sum_{m=-n}^n S_n^m(\mathbf{x} - \mathbf{x}_0) c_n^m \end{aligned}$$

where the multipole coefficients, c_n^m , are given by the formula (which can be evaluated numerically)

$$c_n^m = ik \int_{\Gamma} \left[R_n^{-m}(\mathbf{y} - \mathbf{x}_0) \frac{\partial p}{\partial \mathbf{n}}(\mathbf{y}) - p(\mathbf{y}) \frac{\partial R_n^{-m}}{\partial \mathbf{n}}(\mathbf{y} - \mathbf{x}_0) \right] d\Gamma_{\mathbf{y}}. \quad (41)$$

Here $\partial_{\mathbf{n}} p$ and p are specified as input, e.g., from the boundary integral solver output, or from a volumetric discretization, etc.

The major benefit of using the multipole expansion representation is that it is simple to evaluate the solution using the expression, it provides a compact storage representation which often requires far fewer coefficients than there exists in the boundary geometry, and therefore it is more compact than storing data on the boundary or a Kirchhoff surface. Downsides of this representation are that it can require many coefficients to approximate more complicated fields, especially for objects which are highly nonspherical in nature. The approximation of the sound field can also be poor close to the object, i.e., in the near field. Multipoint multipole expansions can be used to deal with these deficiencies, and these representations can be estimated using essentially the equivalent source method described earlier. For example in the precomputed acoustic transfer paper, a multipoint multipole expansion was estimated for objects which had thin structures by 1st computing the pressure solution on a surrounding isosurface, then using an adaptive multipoint multipole expansion method to generate an approximate representation. While these techniques do not provide high engineering accuracy for higher frequencies which necessitate more sophisticated methods such as the fast multiple method, they can provide sufficient accuracy for sound rendering. Another benefit of these approaches is that they can be hardware accelerated, using multicore or GPU architectures, as described in [15].

6.5 Far-Field Acoustic Transfer (FFAT) Maps

The multipole expansion is particularly effective far from the object, especially with low-frequency modes. However as we get closer to the object, and consider higher frequency modes, the detailed structure of the sound field can be difficult to represent using only multipolar expansions of limited complexity. By observing that most of the complexity occurred in the angular variations, we proposed a texture mapping like approach to encode the structure and support fast evaluation in [2]. The mathematical basis for the approach is as follows. The AtkinsonWilcox theorem [10](p.209) states that the sound field at any point \mathbf{x} lying entirely outside of a circumscribing sphere, which itself encloses all radiating scattering sources, can be written as an expansion of the following form

$$p(\mathbf{x}, k) = e^{-ikr} \sum_{n=1}^{\infty} \frac{f_n(\theta, \phi)}{r^n}. \quad (42)$$

Here all angular structure is contained in the 2D texture-like functions, $f_n(\theta, \phi)$. In contrast, the usual multipole expansion essentially expands the angular functions using spherical harmonics, however by estimating the angular functions using a raster approximation, a cheap method for evaluation can be obtained. This approach forms the basis of the “FFAT Map” technique described in [2]. Using only a couple terms, or even a single term, provides good approximations for sound rendering. The major shortcoming of this approach is that because of the detailed angular structure of these acoustic transfer functions, especially in the near field, high-resolution angular texture maps, can be required to resolve details, and can result in very large memory requirements for objects with many modes. Despite this shortcoming, the method can provide real-time transfer evaluation, and is probably the fastest method for transfer evaluation currently.

7 Rendering Details

Now that we have discussed acoustic transfer functions how to compute them and how to represent them, we will describe how they are used in sound rendering a little more.

Rigid body motion and moving listeners: The acoustic transfer function is evaluated at location \mathbf{x} in the object’s frame of reference. To evaluate the function at runtime, one must first transform the listening position into the object frame of reference to evaluate the virtual listening position. In practice this must be done rapidly enough (e.g., hundreds of times per second) to resolve the variations in the acoustic transfer function along the trajectory. Aliasing can result if the function is not sampled fast enough. In many implementations, the acoustic transfer function is sampled at a fixed rate, such as 600 Hz, to obtain sufficient resolution and avoid artifacts. Multipole representations provide cheap approximations which support fast evaluation [8, 15].

Time delays: So far we have assumed that the acoustic transfer function is evaluated at a simultaneously specified position \mathbf{x} and time t . However for listening positions sufficiently far from the object, time delays can be important to include. A simple model which is often used is to assume that the listener there at position \mathbf{x} hears the sound from an earlier time, $t' = t - c\|\mathbf{x} - \mathbf{x}_0\|$, where $\|\mathbf{x} - \mathbf{x}_0\|$ is the distance from the listener to the object’s center of mass (or suitable representative location).

Auralization and head related transfer functions (HRTF): Simply evaluating the acoustic transfer function at the position of the listener's two ears, in order to evaluate stereo sound, does not usually provide good-quality results. In practice using a head related transfer function (HRTF) can improve the quality of rendered stereo sounds, especially for headphone listening. A simple model of HRTF assumes that the sound is attenuated in a way dependant on the direction of approach of the sound to the ears, and also the frequency. For modal sounds, one can attenuate each mode contribution separately by multiplying each modal component by both the acoustic transfer function, and the HRTF factor.

8 Discussion

Limitations: The frequency-domain acoustic transfer functions provide a simple and practical way to approximate wave radiation attenuation effects for modal vibration and other frequency-localized vibrations. Unfortunately these frequency-dependent transfer functions fail to capture some effects. Objects undergoing rapid motion can result in more complex fluid flow effects which distort the wave radiation. Perhaps the most simple illustration is of a rapidly spinning object, e.g., an object impacts the ground and spins rapidly while ringing. In such cases the orientation of the object can change significantly during several periods of oscillation, and therefore simply looking up the acoustic transfer function in a body-attached spatial field is a poor approximation of the true radiation pattern. Furthermore evaluating the acoustic transfer function along such a rapidly spinning trajectory can result in significant aliasing artifacts which are nonphysical, and can lead to undesirable sound artifacts. A further limitation of the acoustic transfer model is that it is only correct for pure-harmonic vibrations, and lacks proper handling of transients, and attenuation effects due to modal damping. This latter deficiency is particularly noticeable if you have high-frequency modes with high damping. Time-domain methods may be better suited to approximate these sound radiation patterns, especially for objects with concavity's and cavities which can cause sound waves to echo inside.

References

- [1] M. Abramowitz and I. A. Stegun. *Handbook of Mathematical Functions with Formulas, Graphs, and Mathematical Tables*. Dover, New York, 1964.
- [2] Jeffrey N. Chadwick, Steven S. An, and Doug L. James. Harmonic shells: A practical nonlinear sound model for near-rigid thin shells. *ACM Transactions on Graphics*, 28(5):119:1–119:10, December 2009.
- [3] R.D. Ciskowski and C.A. Brebbia. *Boundary element methods in acoustics*. Springer, 1991.
- [4] L. Cremer, M. Heckl, and E.E. Ungar. *Structure-Borne Sound: Structural Vibrations and Sound Radiation at Audio Frequencies*. Springer, 2nd edition, January 1990.
- [5] G.H. Golub and C.F. Van Loan. *Matrix computations*, volume 3. Johns Hopkins Univ Pr, 1996.
- [6] N.A. Gumerov and R. Duraiswami. *Fast multipole methods for the Helmholtz equation in three dimensions*. Elsevier Science, 2004.
- [7] F. Ihlenburg. *Finite element analysis of acoustic scattering*, volume 132. Springer Verlag, 1998.
- [8] Doug L. James, Jernej Barbic, and Dinesh K. Pai. Precomputed acoustic transfer: Output-sensitive, accurate sound generation for geometrically complex vibration sources. *ACM Transactions on Graphics*, 25(3):987–995, July 2006.

- [9] Timothy R. Langlois, Steven S. An, Kelvin K. Jin, and Doug L. James. Eigenmode compression for modal sound models. *ACM Transactions on Graphics (Proceedings of SIGGRAPH 2014)*, 33(4), August 2014.
- [10] S. Marburg and B. Nolte. *Computational acoustics of noise propagation in fluids: finite and boundary element methods*. Springer Verlag, 2008.
- [11] M. Ochmann. The source simulation technique for acoustic radiation problems. *Acta Acustica united with Acustica*, 81(6):512–527, 1995.
- [12] M. Ochmann. The full-field equations for acoustic radiation and scattering. *The Journal of the Acoustical Society of America*, 105:2574, 1999.
- [13] W.H. Press, S.A. Teukolsky, W.T. Vetterling, and B.P. Flannery. *Numerical Recipes: The art of scientific computing*. Cambridge University Press, 2007.
- [14] Peter-Pike Sloan, Jan Kautz, and John Snyder. Precomputed radiance transfer for real-time rendering in dynamic, low-frequency lighting environments. *ACM Transactions on Graphics*, 21(3):527–536, July 2002.
- [15] Changxi Zheng and Doug L. James. Rigid-body fracture sound with precomputed soundbanks. *ACM Transactions on Graphics*, 29(4):69:1–69:13, July 2010.