Precomputed Acoustic Transfer: Output-sensitive, accurate sound generation for geometrically complex vibration sources

Doug L. James   Jernej Barbič
Carnegie Mellon University

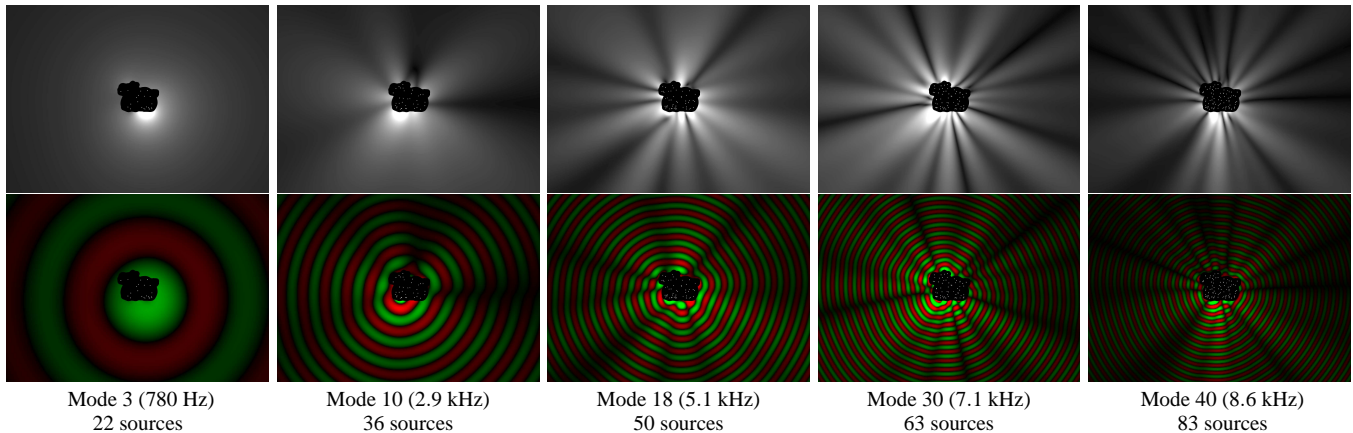Dinesh K. Pai
Rutgers Univ. and Univ. British Columbia

| Mode 3 (780 Hz) | Mode 10 (2.9 kHz) | Mode 18 (5.1 kHz) | Mode 30 (7.1 kHz) | Mode 40 (8.6 kHz) |
| 22 sources | 36 sources | 50 sources | 63 sources | 83 sources |

Figure 1: **Sound field of a vibrating hollow bronze dragon** *(with two holes on bottom) exhibits strong directionality and frequency dependence; (Top) absolute value and (Bottom) real part of complex-valued acoustic pressure field. Equivalent dipole sources (white dots) were placed to achieve* 4% *relative pressure error. Low-frequency radiation fields require few equivalent sources, whereas higher frequency modes with more structured pressure fields typically require more sources. Our Precomputed Acoustic Transfer (PAT) models exploit the fact that equivalent source counts are far smaller than polygon counts for complex geometry, and can therefore accelerate pressure evaluation several thousand times to enable real-time sound rendering.*

Abstract

Simulating sounds produced by realistic vibrating objects is challenging because sound radiation involves complex diffraction and interreflection effects that are very perceptible and important. These wave phenomena are well understood, but have been largely ignored in computer graphics due to the high cost and complexity of computing them at audio rates.

We describe a new algorithm for real-time synthesis of realistic sound radiation from rigid objects. We start by precomputing the linear vibration modes of an object, and then relate each mode to its sound pressure field, or acoustic transfer function, using standard methods from numerical acoustics. Each transfer function is then approximated to a specified accuracy using low-order multipole sources placed near the object. We provide a low-memory, multilevel, randomized algorithm for optimized source placement that is suitable for complex geometries. At runtime, we can simulate new interaction sounds by quickly summing contributions from each mode's equivalent multipole sources. We can efficiently simulate global effects such as interreflection and changes in sound due to listener location. The simulation costs can be dynamically traded-off for sound quality. We present several examples of sound generation from physically based animations.

CR Categories: I.3.5 [COMPUTER GRAPHICS]: Computational Geometry and Object Modeling—Physically based modeling; I.6.8 [SIMULATION AND MODELING]: Types of Simulation—Animation

Keywords: sound synthesis, modal vibration, Helmholtz, acoustic radiation, equivalent sources, source simulation, Trefftz, boundary element method, multipole

1 Introduction

Sounds have always been an important part of computer animation and visual media. There has been significant recent progress in developing algorithms for physically based sound synthesis for computer animation. Using recent algorithms for physically based simulation of sound sources [van den Doel et al. 2001; O'Brien et al. 2001; O'Brien et al. 2002; Dobashi et al. 2003] one can automatically synthesize new sounds synchronized with animations.

However, these did not adequately account for how sound is *radiated* from an object with complicated geometry. The sound's timbre changes in a recognizable and spatially variant way due to interaction with the object's own geometry. This is because audible sounds have wavelengths that are comparable to the size of objects in human environments. Recall that human hearing is very sensitive up to about 10 KHz (3.4cm wavelength), and can reach up to 20 KHz (less than 2cm wavelength). Radiation from a particular vibration mode interacts with the object's geometry in a complicated way to affect its *radiation efficiency* [Cremer et al. 1990], boosting radiation from some frequencies and suppressing others, similar to how an open-ended cylindrical tube enhances resonant frequencies.

These phenomena parallel global illumination effects such as self-shadowing and interreflections. Computing these effects is very expensive but recent work on precomputed radiance transfer [Sloan et al. 2002] has shown how interactive rendering can be achieved by preprocessing the linear response of a single 3D object exposed to (low-frequency) environmental lighting basis vectors.

In this paper, we investigate preprocessing of analogous global *sound* radiation effects for a single vibrating object, showing that these expensive linear acoustic transfer phenomena can be largely precomputed to support efficient real-time sound synthesis. An overview of our approach is illustrated in Figure 2.

Starting from a description of the geometry, material properties, and boundary conditions on a sound producing object, we first precompute its vibration modes, i.e., the mode shapes and frequencies.
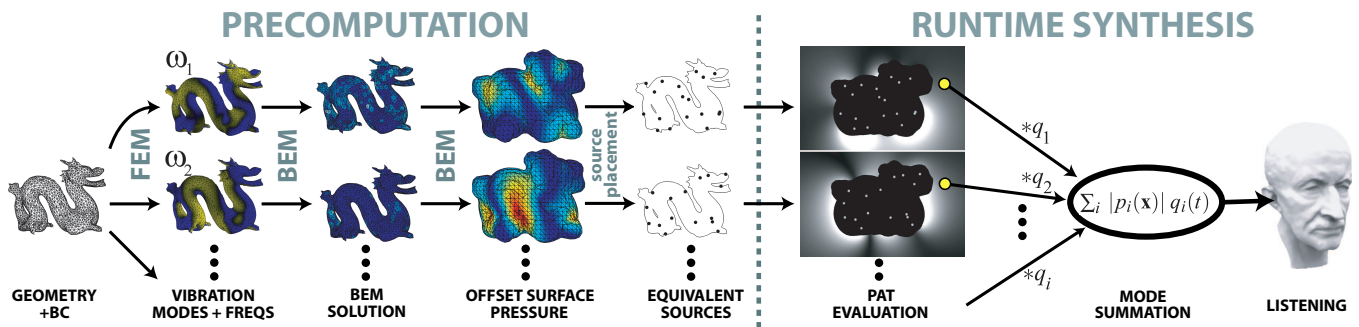
Figure 2: **Overview of Precomputed Acoustic Transfer (PAT)**

Each mode radiates sound over the entire surface of the object; using the boundary element method (BEM) we precompute boundary solution data which can be used to reconstruct the complex, space-variant pressure field (the acoustic transfer function) that accounts for phenomena such as diffraction and interreflection. For runtime performance, we then approximate each transfer function by precomputing a set of equivalent multipole sources that match its pressure values on a fictitious offset surface. Runtime synthesis can then be efficiently performed; when the object is excited by forces from the environment, we only have to compute the excitation of each mode, and compute a weighted combination of contributions from the precomputed sources at the listener's location. Multiple listening locations and real time animation are easily supported, since the equivalent source approximations are cheap to compute, and typically need only be updated at slightly faster than graphics rates.

2 Related Work

Several aspects of sound synthesis have been explored in computer graphics. A significant part of the effort has been in simulating the acoustics of rooms and concert halls, assuming that the sound is provided, for instance, by a recording [Stettner and Greenberg 1989; Funkhouser et al. 1998; Funkhouser et al. 1999; Tsingos et al. 2001; Tsingos et al. 2002; Lokki et al. 2002]. General techniques for specification and synthesis of sounds for animation have been investigated in [Takala and Hahn 1992; Cardle et al. 2003]. Tsingos, et al. [2004] have shown how large numbers of sound sources can be handled using perceptual clustering. A multi-resolution algorithm based on stochastic sampling was proposed in [Wand and Strasser 2004]. Note that in these papers, a source corresponds to an entire waveform, for instance, a recorded voice.

Physically based synthesis of sound sources has been explored in computer graphics and computer music [Cook 2002]. Most approaches are concerned with sounds emitted by vibrating solids, but see [Dobashi et al. 2003] for aerodynamic sounds. Depending on their need for speed, the approaches can differ significantly.

For interactive simulation, the method of choice has been to directly use the vibration modes (frequencies and corresponding decays) for sound synthesis [van den Doel and Pai 1996; van den Doel et al. 2001; O'Brien et al. 2002; Raghuvanshi and Lin 2006]. Pai et al. [2001] described how these modal frequencies and decays could be measured from real objects. However, these papers treated the sound source as a point, effectively assuming that the listener is far from the object, and did not capture the subtleties of the sound field.

For non-interactive simulation, O'Brien et al. [2001] used FEM to directly simulate vibrations in the time domain. They were the first in computer graphics to simulate the radiated sound field, using the "Rayleigh method." Even though the method is very expensive, they were able to obtain many interesting effects. The "Rayleigh

method" treats each element of the vibrating surface independently, summing each element's contribution together with a visibility test. While faster than boundary element analysis, with speed can come considerable inaccuracy since the method completely ignores reflections and diffraction [von Estorff 2003].

The Helmholtz equation is a standard tool for modeling sound waves radiated from vibrating yet essentially rigid objects, and its numerical solution has been studied extensively. Boundary element methods (BEM) are widely used for acoustic radiation and scattering calculations (see [Ciscowski and Brebbia 1991; von Estorff 2000; Wu 2000]), and numerous production codes exist for engineering analysis. One major drawback of the acoustic BEM computations is that they can have large memory requirements ($O(N^2)$ memory for N boundary elements), and so do not scale gracefully to detailed geometry that is required to resolve higher frequencies [Desmet 2002; von Estorff 2003]. In this respect, Helmholtz variants of the fast multipole method (FMM) have been successful (see [Gumerov and Duraiswami 2005]), providing high-accuracy, memory- and time-scalable solutions for radiation and scattering problems. Unfortunately, such methods are not optimized for pressure evaluation in real-time sound rendering.

In the last decade, a relatively simple approach called the source simulation technique (also called equivalent sources method, method of substitute sources, equivalent-sphere methods, etc.) has been introduced to solve the Helmholtz equation for scattering and radiation problems [Ochmann 1995; Ochmann 1999], and to estimate measured radiation sources [Magalhães and Tenenbaum 2004]. Similar to our PAT-approximation problem, the primary challenge is placement of equivalent sources, which is currently an active topic of research; some recent examples are [Gounot et al. 2005; Pavić 2005; Pavić 2006]. However, these methods solve the (Neumann) radiation problem directly, while we first compute each acoustic transfer function using trusted engineering acoustics methods, e.g., BEM, then use a novel equivalent sources technique to solve an easier pressure (Dirichlet) approximation problem on an offset surface. The latter makes it easier to accommodate complex geometry, thin shells, and polygon soups common in graphics. Our use of an offset pressure surface to support equivalent sources is also closely related to work on faithful rendering of vibration sources by [Johnson et al. 2003]. We focus on radiation produced by a single complex object in isolation, and introduce an efficient, randomized, multilevel source placement algorithm suitable for approximating radiation from geometrically complex models.

In graphics, diffraction effects are usually less important for lighting than for sound since the wavelength of visible light is much shorter than geometric features. Exceptions are reflections from geometric microstructure, such as the bumpy surface of a compact disc, for which Fourier-based precomputation approaches exist [Stam 1999]. Our method is in the same spirit as precomputed

radiance transfer (PRT) [Sloan et al. 2002; Sloan et al. 2003], however, due to mathematical differences in the nature of sound generation, PRT techniques cannot be directly used for acoustics.

3 Preliminary Analysis

Our preprocess uses established techniques of linear modal vibration analysis and acoustic analysis to generate input pressure fields for our PAT approximation algorithm.

3.1 Modal Vibrations

Sound is produced by small vibrations of objects. These vibrations can be approximated as a superposition of a discrete set of *mode shapes* $\hat{u}_k$, each oscillating with angular frequency (2π/period) ω_k and amplitude $q_k(t)$. The vibration's displacement vector is

$$u(t) = Uq(t) = [\hat{u}_1 \cdots \hat{u}_K] q(t) \qquad (1)$$

where $q(t) \in \mathbb{R}^K$ is the vector of modal amplitude coefficients. For acoustics we only need to retain modes with frequencies in the audible range. While modes decay over time due to complicated factors such as mechanical damping and acoustic radiation, we use the simple Rayleigh damping approximation, with constants tuned to the material behavior desired. Modal analysis is standard [Shabana 1990]; we used ABAQUS, a commercial FEM package.

3.2 Acoustic Transfer Functions

Sound radiates from an object's surface S into the surrounding medium, Ω, as pressure waves. Since the modes decay in time slowly relative to frequency, it is more convenient to deal with Fourier transforms of all time dependent quantities. We denote the complex-valued time-harmonic pressure field due to a single vibration mode at a point $\mathbf{x} \in \Omega$ outside the object as $p(\mathbf{x})e^{+i\bar{\omega}t}$, where ω is the frequency of the mode. The wave number, k, for a sound wave of wavelength λ is given by

$$k = \frac{\omega}{c} = \frac{2\pi}{\lambda}, \qquad (2)$$

where c is the speed of sound in the surrounding medium, e.g., air has $c = 343m/s$ at standard temperature and pressure. We refer to the spatial part $p(\mathbf{x}) \in \mathbb{C}$ of each mode's pressure field as the **acoustic transfer function** of that mode. It satisfies the temporal Fourier transform of the wave equation, the Helmholtz equation,

$$\nabla^2 p + k^2 p = 0, \quad \mathbf{x} \in \Omega. \qquad (3)$$

Given $p(\mathbf{x})$, we can approximate the contribution to sound pressure at the ear (up to a phase) by the real-valued quantity, $|p(\mathbf{x})| q_k(t)$, where $q_k(t)$ is this mode's amplitude. Boundary conditions are required to solve (3) for sound radiation from surface vibrations. First, one assumes that pressure decays far away from the boundary (Sommerfeld radiation condition). Second, the normal derivative of pressure on the vibrating object's surface, $S = \partial\Omega$ is given by a Neumann boundary condition

$$\frac{\partial p}{\partial n} = -i\omega\rho v \quad \text{on } S \qquad (4)$$

where ρ is the fluid density, and v is the surface's normal velocity. The latter is specified as $v = i\omega(\mathbf{n} \cdot \hat{u})$, where $\mathbf{n} \cdot \hat{u}$ is the modal displacement in the normal direction.

Solving each mode's exterior radiation problem can be done using any standard numerical acoustics solver. For example, the boundary element method (BEM) is widely used in the engineering community to solve Helmholtz radiation problems for applications such as noise analysis and reduction [Ciscowski and Brebbia 1991]. We used a commercially available acoustic BEM package (Neo Acoustics, Paragon NE, Netherlands). For graphics models, where thin-shell structures are common, e.g., plastic chair or bell,

we recommend using indirect BEM solvers [Ciscowski and Brebbia 1991]. For more complex geometries, very high accuracies, or difficult high-frequency analyses, i.e., where $kR \gg 1$ (R is bounding radius of object), sophisticated solvers are available, such as the fast multipole method (see [Gumerov and Duraiswami 2005]).

Once the BEM problem has been solved on the object boundary, it still remains to evaluate the BEM solution's pressure field $p(\mathbf{x})$ at non-boundary locations $\mathbf{x} \in \Omega$. Unfortunately, this incurs an $O(N)$ cost for an object with N boundary elements, for an $O(NM)$ cost for M modes. Such costs make evaluation noninteractive for all but the simplest models (see Figure 3).

4 Approximating Acoustic Transfer

We now describe how we approximate $p(\mathbf{x})$ to support real-time evaluation at costs independent of geometric complexity, i.e., output-sensitive evaluation. We use the boundary element analysis to evaluate the pressure $p(\mathbf{x})$ on an offset surface enclosing the object. To support repeated (real-time) evaluation, we then estimate a set of simple equivalent sources that approximate the offset pressure boundary condition, and therefore the acoustic transfer function in the exterior space. See Figure 3 for a comparison of a resulting PAT pressure approximation to the ground truth BEM solution.
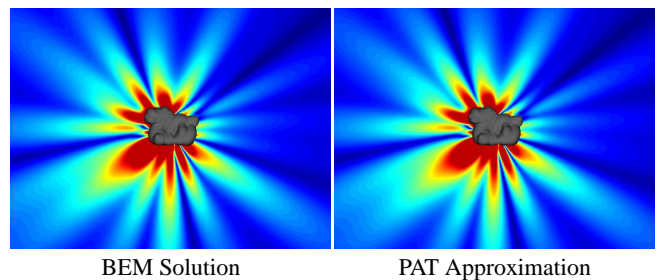

| BEM Solution | PAT Approximation |

Figure 3: **Comparison of BEM to PAT approximation** *(4% error) using exaggerated shading reveals near identical absolute pressure fields, $|p(\mathbf{x})|$ (mode 30 of dragon model shown in Figure 1); each 640x480 raster has 307200 samples. However, evaluating the BEM raster samples took 8.5 hours while PAT only took 4.3 seconds–*a 7000x speedup*. This illustrates the benefit of output-sensitive PAT evaluation: PAT only needs 63 cheap dipole sources to approximate the pressure field to 4% accuracy, whereas BEM must accumulate contributions from all 7689 boundary elements.*

4.1 Background: Spherical Multipole Radiation

Our PAT pressure fields are linear combinations of spherical multipoles $\psi_{lm}(\mathbf{x} - \bar{\mathbf{x}})$, an important class of functions representing radiating waves (of particular wavenumber k). They decay as $\|\mathbf{x} - \bar{\mathbf{x}}\| \to \infty$, and satisfy the free-space Helmholtz equation everywhere except at the source position, $\bar{\mathbf{x}}$, where they are singular. They are given by the product of two complex-valued functions,

$$\psi_{lm}(\mathbf{x} - \bar{\mathbf{x}}) = h_l^{(2)}(kr) Y_{lm}(\theta, \phi), \quad |m| \le l, \; l = 0, 1, \ldots, \qquad (5)$$

where (r, θ, ϕ) are spherical coordinates of the vector, $\mathbf{x} - \bar{\mathbf{x}}$. The spherical harmonics $Y_{lm}(\theta, \phi) \in \mathbb{C}$, widely used in graphics [Sloan et al. 2002], describe the angular variables, and the radial factors are spherical Hankel functions of the 2^{nd} kind,

$$h_l^{(2)}(kr) = j_l(kr) - i y_l(kr) \in \mathbb{C}, \qquad (6)$$

where j_l and y_l are real-valued *spherical* Bessel's functions of the first and second kind [Abramowitz and Stegun 1964]. For example, a spherical radiating wave is simply $\psi_{00} \propto \frac{e^{-ikr}}{r} = \frac{\cos kr}{r} - i\frac{\sin kr}{r}$.

Dipole Sources: Fortunately, one can precompute and evaluate good PAT approximations using only combinations of 4-term dipole sources ($l < n = 2$ or $(l, m) \in \{(0,0), (1,-1), (1,0), (1,1)\}$), simplifying evaluation and implementation. Code for efficient evaluation of dipole sources is provided on the CDROM, and achieves a throughput of more than 3 million dipoles/sec on a Pentium IV 3GHz PC. Combinations of only monopole sources ($l = 0$) tend to be too sensitive to source placement, and are also poorly suited to approximating radiation from thin shells, which is dipole-like locally. Higher-order sources can also be used, but are not analyzed in this paper (although see [Ochmann 1995; Pavić 2006]).

4.2 Equivalent Multipole Sources

Any linear combination of M multipoles of order-n

$$p(\mathbf{x}) = \sum_{q=1}^{M} \sum_{l=0}^{n-1} \sum_{m=-l}^{l} \mathsf{c}_{qlm} \psi_{lm}(\mathbf{x} - \bar{\mathbf{x}}_q) \equiv \sum_{j=1}^{Mn^2} \mathsf{c}_j \psi_j(\mathbf{x}) \qquad (7)$$

also satisfies the Helmholtz equation (and radiation condition), except at source points $\{\bar{\mathbf{x}}_q\}$ where it is singular. Here $\mathsf{c}_{qlm} \in \mathbb{C}$ are coefficients to be determined, as are M and the source points $\{\bar{\mathbf{x}}_q\}_{q=1}^{M}$, and j is a generalized index for (q, l, m). One must also satisfy the boundary condition on the object. Source simulation methods [Ochmann 1995; Pavić 2005] place sources *inside* the object so as to satisfy the (Neumann) boundary condition. Unfortunately, in graphics many objects are essentially thin shells, such as our bell or chair or dragon, and they have no inside to hide singular sources in. We therefore take a different approach.

Defining an offset surface, S_o, that is manifold and closed, and encloses the object geometry provides a fictitious boundary interface that serves two purposes. First, it provides a clear "inside" region within which to place our fictitious sources. Second, it provides a boundary on which to evaluate our precomputed (BEM) solution's pressure values to provide a pressure (Dirichlet) boundary condition for estimation of equivalent multipole sources. We obtain offset surfaces by first computing a distance field to the model geometry, and then using marching cubes to extract an outer iso-surface for a given distance offset [Lewiner et al. 2003]. Resulting surfaces are smooth and of marching cubes resolution; a chair example is shown in Figure 6. The distance field's voxel resolution, h, is chosen small enough to resolve the offset shape, as well as the smallest wavelength λ_{min} analyzed; in our implementation we choose $kh \le 1$, but in practice kh may approach one for high frequencies due to computational limitations. The offset distance δ is chosen large enough to avoid over-fitting during source placement; in practice we set δ to a multiple, e.g., 2, of the largest offset surface mesh's edge length. We use BEM to sample the pressure $\bar{p}(\mathbf{x})$ on the offset surface at N vertex (or centroid) sample positions $(\mathbf{s}_1, \mathbf{s}_2, \ldots, \mathbf{s}_N)$, collectively denoted by $\bar{\mathsf{p}} \in \mathbb{C}^N$. Each vertex (or centroid) sample i also has an effective area, a_i, given by 1/3 of adjacent triangle areas (or triangle area).

Computing Equivalent Source Amplitudes: If one places sources inside the volume enclosed by S_o, such that the Dirichlet BC $p = \bar{p}$ is satisfied on S_o, one satisfies the Helmholtz radiation problem everywhere in the space exterior to S_o; the approximation error is only determined by how well we satisfy the boundary condition, $p(\mathbf{x}) = \bar{p}(\mathbf{x})$, $\mathbf{x} \in S_o$ (see Trefftz methods [Kita and Kamiya 1995; Desmet 2002]). Assuming a set of M unique source positions, $\{\mathbf{x}_1, \ldots, \mathbf{x}_M\}$, we can compute each multipole's expansion coefficients by fitting the multipole expansion's pressure field to the BEM-sampled offset-surface pressure values, $\bar{\mathsf{p}}$, in a least-squares sense. In our implementation, we use weighted least squares to fit offset-surface pressure fields at N locations by solving the over-determined N-by-Mn^2 system of equations:

$$(\mathsf{WV})\mathsf{c} = \mathsf{W}\bar{\mathsf{p}} \quad \Leftrightarrow \quad \mathsf{Ac} = \mathsf{b}, \qquad (8)$$

where $\mathsf{c} \in \mathbb{C}^{Mn^2}$ is the vector of unknown multipole coefficients from (7); rows are scaled by the N-by-N weight matrix $\mathsf{W} = \mathrm{diag}(\sqrt{a_i})$; V is the multipole basis matrix with $\mathsf{V}_{ij} = \psi_j(\mathbf{s}_i)$ being the j^{th} multipole function evaluated at the i^{th} sample position, $\mathbf{s}_i$; and $\bar{\mathsf{p}}_i = \bar{p}(\mathbf{s}_i)$ is the BEM pressure evaluated at $\mathbf{s}_i$. These weighted discrete equations can be seen as arising from the minimization of the squared $\mathcal{L}^2$ pressure error, $\int_{S_o} |p(\mathbf{y}) - \bar{p}(\mathbf{y})|^2 dS_y$.

For numerous or closely placed sources, the A matrix in (8) can be poorly conditioned. We solve the least squares problem in (8) using a truncated singular value decomposition with a small relative singular value threshold of 10^{-6}. Double precision is used to construct and solve the system, and also to evaluate the PAT pressure approximation (7) at runtime.

4.3 Multipole Placement Algorithm

Given multipole positions one can solve (8) to obtain a multi-point multipole expansion that approximates the offset-surface pressure boundary condition. Determining good source positions and how many, however, is nontrivial.

4.3.1 Greedy Randomized Source Placement Algorithm

Before presenting our complete multipole placement algorithm optimized for complex geometry and faster convergence (§4.3.2), we first introduce a simple algorithm to illustrate the basic concepts. We note that by using this simple algorithm one can already generate fair quality PAT approximations.
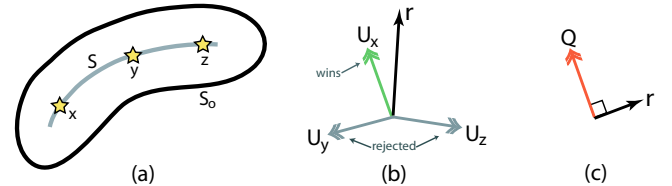

Figure 4: **Greedy selection of multipole positions** *begins with (a) some candidate source positions* $\{\mathsf{x}, \mathsf{y}, \mathsf{z}\}$ *inside or on the object surface S, and a pressure residual* r *on the offset surface* S_o; *(b) Source position* x *is selected since its multipole subspace* $U_\mathbf{x}$ *has the largest residual projection; (c) The residual is made orthogonal to this known subspace* $\mathsf{Q} = U_\mathbf{x}$. *The process is repeated to greedily select additional source positions, updating the subspace* Q *and residual* r *at each iteration.*

Greedy Multipole Placement: We incrementally add sources one by one. Given a set of candidate positions, $\mathcal{X}$, to place the next multipole of order n, we rank each position, $\mathbf{x} \in \mathcal{X}$, based on the ability of its multipole matrix, $\mathsf{V}_\mathbf{x}$, to describe the current offset surface pressure residual, $\mathsf{r} = \mathsf{b} - \mathsf{Ac}$ (see Figure 4). Mathematically, given a unitary basis matrix spanning the multipole's n^2 vectors,

$$\mathsf{U}_\mathbf{x} \equiv \mathrm{basis}(\mathsf{WV}_\mathbf{x}) \in \mathbb{C}^{N \times n^2}, \qquad (9)$$

the *multipole point fitness* is defined as the norm of the residual's subspace projection, $\|(\mathsf{U}_\mathbf{x})^H \mathsf{r}\|_2$, where $()^H$ denotes the matrix Hermitian conjugate. Given a sampling of candidate multipole positions, $\mathcal{X}$, we can select the best position via the largest projection:

$$\mathbf{x}^* = \arg\max_{\mathbf{x} \in \mathcal{X}} \|(\mathsf{U}_\mathbf{x})^H \mathsf{r}\|_2^2, \qquad (10)$$

or, as will be important later for multilevel refinement, select a subset of M best multipole positions. Repeated greedy placement can result in a steady decrease in residual pressure error (see Figure 5).

Candidate multipole positions, $\mathcal{X}$, need only be chosen inside the offset surface, S_o, but to avoid singularities in *any* listening spaces outside the object, we place them inside (or, in the case of shells, on) the original object surface, S. Candidate positions are generated using random sampling; for open shell objects, points
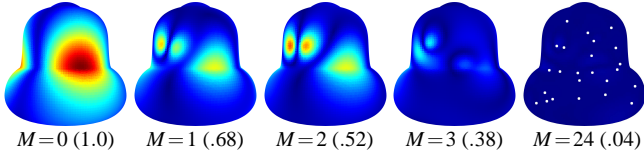
Figure 5: **Greedy multipole placement:** *Plots of the bell's residual offset pressure error are shown at each placement iteration for increasing numbers of sources, M, and decreasing relative residual norms, $\|r\|$ (in brackets). At each iteration, a new multipole position, x^*, is selected from a sampling of candidate positions based on the ability of its multipole basis, U_x, to capture the largest fraction of r. (Far right) Placed source positions.*

are chosen randomly on S; for closed objects rejection sampling (using a sphere-tree accelerated ray-intersection test) is used to find points inside S. Increasing the number of candidate positions, $|\mathcal{X}|$, can improve point selection at the cost of more expensive iterations.

Updating the multipole subspace and residual: After selecting a new multipole position x^*, we need to update the residual r. This is achieved by incrementally updating a unitary basis Q for the space spanned by all multipoles selected, followed by removal of the residual's component in that subspace. The multipole basis matrices can be ill-conditioned, so we use a stable orthogonalization scheme such as modified Gram-Schmidt [Golub and Van Loan 1996]. Given a new multipole position x, we expand the basis Q incrementally and update the residual as follows:

$$\begin{array}{ll} \text{EXPANDSUBSPACEANDUPDATERESIDUAL}(\mathsf{Q}, \mathsf{r}, \mathbf{x}) & \\ 1 \quad \mathsf{Q} \leftarrow [\mathsf{Q}\,|\,\mathsf{Q_x}] \leftarrow \text{MODGRAMSCHMIDT}([\mathsf{Q}\,|\,\mathsf{WV_x}]) & \\ 2 \quad [\mathsf{Q_x}\,|\,\mathsf{r'}] \leftarrow \text{UPDATERESIDUAL}([\mathsf{Q_x}\,|\,\mathsf{r}]) & \\ 3 \quad \mathsf{r} \leftarrow \mathsf{r'} & \end{array} \quad (11)$$

In line 1, modified Gram-Schmidt is used to convert the matrix WV_x into a unitary basis Q_x orthogonal to the previous Q basis. Since the input arguments satisfy $r \perp Q$, it suffices to only subtract r's projection onto Q_x, i.e., $r' \leftarrow r - Q_x Q_x^H r$, which is done in a stable manner using a modified Gram-Schmidt pass (line 2). For a surface with N samples, and a basis with M multipoles of order n, we have $Q \in \mathbb{C}^{N \times Mn^2}$ and $V_x, Q_x \in \mathbb{C}^{N \times n^2}$. Line 1 dominates the cost for large M; each multipole added has cost complexity $O(NMn^4)$. In our implementation we use dipoles ($n=2$), so that the cost is $O(NM)$; the total updating cost of incremental orthogonalization of M dipoles is therefore $O(NM^2)$, as expected. This cost is reasonable for complex models where $N \gg M$, but small M is desirable.

Basic greedy multipole placement algorithm: At each iteration the algorithm chooses greedily from a pool of newly drawn randomized source positions, $\mathcal{X}$, then updates the multipole basis, Q, and the residual. The algorithm terminates when the normalized residual error falls below a tolerance, TOL:

$$\begin{array}{ll} \text{PLACEMULTIPOLES}(TOL, \bar{\mathsf{p}}, S, S_o, ...) & \\ 1 \quad \mathsf{r} \leftarrow \mathsf{W\bar{p}} & \\ 2 \quad \mathsf{r} \leftarrow \mathsf{r}/\|\mathsf{r}\|_2 \quad \text{// init residual} & \\ 3 \quad \mathsf{Q} \leftarrow 0 \quad\quad\quad \text{// init subspace} & \\ 4 \quad \mathcal{Y} \leftarrow \emptyset \quad\quad\quad \text{// init selected points} & \\ 5 \quad \textbf{while } \|\mathsf{r}\|_2 > TOL & \\ 6 \quad\quad \mathbf{x}^* \leftarrow \text{SELECTMULTIPOLEPOSITION}(\mathsf{r}) & \\ 7 \quad\quad \text{EXPANDSUBSPACEANDUPDATERESIDUAL}(\mathsf{Q}, \mathsf{r}, \mathbf{x}^*) & \\ 8 \quad\quad \mathcal{Y} \leftarrow \mathcal{Y} \cup \mathbf{x}^* & \\ 9 \quad \textbf{return } \mathcal{Y} & \end{array} \quad (12)$$

Obviously, the important details are hidden in how SELECTMULTIPOLEPOSITION() picks the new position at each iteration. A simple implementation that captures the essence of our optimized multilevel algorithm is as follows:

$$\begin{array}{ll} \text{SELECTMULTIPOLEPOSITION}(\mathsf{r}) & \\ 1 \quad \mathcal{X} \leftarrow \text{ draw } P \text{ random candidate source positions} & \\ 2 \quad \mathbf{x}^* \leftarrow \arg\max_{\mathbf{x} \in \mathcal{X}} \|(\mathsf{U_x})^H\,\mathsf{r}\|_2 & \\ 3 \quad \textbf{return } \mathbf{x}^* & \end{array} \quad (13)$$

Discussion of computational complexity: Given an offset surface with N samples, simply evaluating the $P = |\mathcal{X}|$ multipole basis matrices, $U_x \in \mathbb{C}^{N \times n^2}$, constitutes a large $O(NPn^4)$ cost per iteration. For dipoles ($n=2$), the total cost of selecting a multipole position, expanding the subspace, and updating the residual each iteration is $O(NP + NM)$ flops. Unfortunately, to obtain good multipole positions and a small M, it is desirable to have very large P (ideally $P \gg M$). Caching and reuse of multipole bases is possible, but has undesirable memory requirements.[1] We now provide a simple, low-memory, multilevel approach suitable for large P.
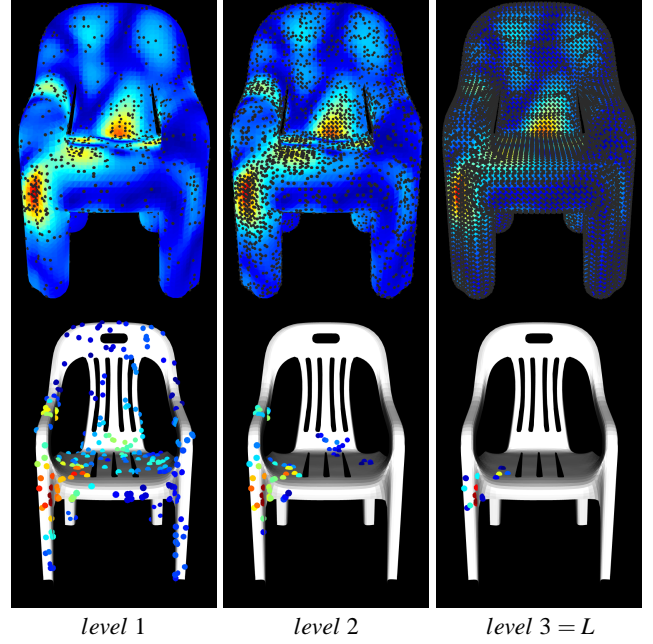

level 1   *level 2*   *level 3 = L*

Figure 6: **Multilevel source placement** ($L=3$ **levels):** *(Top) Offset surface importance-sampled at increasing density based on plotted pressure residual; (Bottom) $P = 256$ candidate source positions are considered on the coarsest level ($\ell = 1$), then progressively thinned by a factor of 4 at each level until only 16 must be evaluated at the finest resolution ($L = 3$).*

4.3.2 Multilevel Acceleration of Source Placement

We can greatly accelerate SELECTMULTIPOLEPOSITION(r) using a simple multilevel source placement strategy. For example, we can rank $4\times$ as many source positions if we estimate their fitness $\|(U_x)^H r\|_2$ using only 1/4 of the offset surface samples. Obviously this is approximate, but even sparse approximations can be expected to cull away the worst candidate positions. In a multilevel setting, given a large set of P candidate positions, each new level culls bad positions while evaluating the fitness of remaining ones with increasingly more samples until the best multipole position is selected from a small handful of promising ones evaluated using all offset surface samples (see Figure 6).

---

[1] This basic algorithm is related to work by Pavić [2005; 2006] on equivalent source simulation for exterior radiation (and scattering) problems with Neumann (pressure derivative) boundary conditions on the object surface: a greedy approach is used to select sources from a dense uniform grid of candidate source positions thereby providing engineering results for 2D radiation problems. To amortize orthogonalization ($V_x \rightarrow U_x$) costs, one could keep a large number of $\{U_x\}_{x \in \mathcal{X}}$ matrices resident in memory, analogous to how Pavić [2005] caches a 2D grid of normalized monopole vectors. However, for complex 3D meshes such an algorithm has prohibitive memory requirements to achieve sufficient accuracy, and restricts the size of the pool, $|\mathcal{X}|$, introducing position sampling bias. In summary, for complex models, storing a large (double precision) Q basis matrix in memory is possible, but it is undesirable to store thousands of candidate $\{U_x\}$ matrices.

Model	Object Surface, S		Offset Surface, S_o				Vibration Modes				BEM Analysis			Wavelength Regime		
	Vtx	Tri	Vtx	Tri	h	δ	Modes	Freq (Hz)	FEM Time	Modes Size	Memory	Solve	Offset	radius R	λ	max kR
Dragon	3861	7689	2617	5230	0.72 cm	1.5 cm	40	490–8600	3m 45s	3.5 MB	171 MB	5h 40m	2h 54m	10 cm	4.0–70 cm	15.8
Rabbit	2562	5120	11100	22196	0.49 cm	5 cm	60	160–999	2m 50s	3.5 MB	101 MB	32m	11h 56m	10 cm	34.3–214 cm	1.8
Bell	3651	7300	14173	28342	1.4 cm	15 cm	85	230–2750	4m 07s	7.1 MB	149 MB	3h 37m	16h 01m	30 cm	12.5–149 cm	15.1
Chair	5246	9581	11060	22124	1.6 cm	5 cm	200	25–2045	11m 04s	25 MB	221 MB	40h	76h	40 cm	16.8–1390 cm	15.0

Table 1: **Model Statistics:** *Triangle meshes of the object surface S are used for both thin-shell FEM analysis, and indirect BEM analysis (rabbit uses direct BEM with CHIEF). Finite element modal analysis is very fast (FEM Time), whereas BEM analysis is substantially slower (but computed in parallel): see BEM analysis memory footprints (Memory), and times to solve all BEM problems (Solve) and evaluate the BEM solutions on the offset surface S_o (Offset). The relationship between the object size, given by its bounding radius (R), and the problem wavelength λ, is summarized by the quantity, $kR = 2\pi R/\lambda$.*

Specifically, consider L levels, with level 1 the coarsest, and level L is the finest level which uses all N offset position samples. Offset surface sample indices on level ℓ are denoted by $\mathcal{S}^\ell$ and are constructed using random sampling such that $|\mathcal{S}^\ell| = \lceil N\alpha^{\ell-L}\rceil$, where α is the ratio between levels–we use $\alpha = 1/4$. Given P candidate multipole positions on level 1, level ℓ has only $\lceil P\alpha^\ell \rceil$ positions. Then given a list of sample indices $\mathcal{S}^\ell$, the level-ℓ sparse multipole position fitness is $\|(U_\mathbf{x}^\ell)^H r^\ell\|_2$, where $r^\ell = (r_i)_{i \in \mathcal{S}^\ell}$ is r restricted to the offset samples, and similarly $U_\mathbf{x}^\ell \in \mathbb{C}^{|\mathcal{S}^\ell| \times n^2}$ is the unitary basis associated with the row restriction of WV to row indices in $\mathcal{S}^\ell$. Our multilevel algorithm for estimating the best multipole position from P candidates is:

SELECTMULTIPOLEPOSITION(r)

$$
\begin{aligned}
&1 \quad \mathcal{X}^1 \leftarrow \text{ draw } P \text{ random candidate source positions} \\
&2 \quad \textbf{for } \text{Level } \ell = 1 \ldots L-1 \\
&3 \quad\quad \mathcal{S}^\ell \leftarrow \text{ select } \lceil N\alpha^{L-\ell}\rceil \text{ offset surface samples} \\
&4 \quad\quad r^\ell \leftarrow \text{ restrict r to indices in } \mathcal{S}^\ell \\
&5 \quad\quad \mathcal{X}^{\ell+1} \leftarrow \lceil P\alpha^\ell \rceil \text{ points in } \mathcal{X}^\ell \text{ with largest } \|(U_\mathbf{x}^\ell)^H r^\ell\|_2 \\
&6 \quad \mathbf{x}^* \leftarrow \arg\max_{\mathbf{x} \in \mathcal{X}^L} \|(U_\mathbf{x})^H r\|_2 \quad // \text{ best on finest level} \\
&7 \quad \textbf{return } \mathbf{x}^*
\end{aligned} \tag{14}
$$

Random sampling of offset surface positions changes at each level (line 3) to avoid persistent bias. To avoid missing structure in the residual r, importance sampling is used to balance samples selected based on area-weighted probability, $prob \propto a_i$, and squared residual error, $prob \propto |r_i|^2 = a_i|\mathbf{p}_i - \bar{\mathbf{p}}_i|^2$. In practice, we observe that sampling only based on area or squared residual distributions results in slightly larger M values than using an equal combination. The memory requirements of the multilevel version are comparable to the single-level case. Each level has the same theoretical cost of α^{L-1} that of a single-level version with the same number of candidate positions. There are L levels to evaluate, so the cost of the multilevel algorithm is $\alpha^{L-1}L$ that of a single-level approach. For example, with $\alpha = 1/4$ an $L = 3$ scheme provides approximately a $5\times$ speedup, whereas $L = 4$ gives a $16\times$ speedup.

5 Results

We now present results for four different objects: a large tin bell; a hollow bronze dragon with holes on the bottom; a plastic chair; and a plastic thin-shell rabbit. Please see accompanying video, and our real-time software demonstration (on CDROM and website). Table 1 gives detailed statistics obtained on a Pentium IV, 3.0 GHz machine using C/C++ code. Our source placement implementation uses Java, and is timed using Sun Java JDK 1.5.

PAT approximations for the dragon model were shown in Figure 1, and exhibit a highly-structured pressure field, except at very low frequencies (more on this later). A comparison of PAT to the ground truth indirect BEM solution was provided in Figure 3 for the dragon model, and is nearly indistinguishable despite being several thousand times faster to evaluate. The radiation efficiency [Cremer et al. 1990] of individual modes exhibits a complicated attenuation structure (see Figure 7) that is perceptually relevant yet missing from traditional modal rendering methods [van den Doel and Pai 1996].

The deficiency of monopole sources for thin-shell radiators is illustrated in Figure 8 for the chair model. The number of dipole sources required to approximate various models are shown in Figure 9, and support our output-sensitive approximation claim. Larger approximation errors lead to fewer dipole sources, and improved real-time rendering rates; see Table 3 for PAT precomputation and real-time rendering performance results. The multilevel source placement algorithm is analyzed in Table 2, and illustrates that multilevel placement is more efficient at scanning large numbers of candidate sources to generate compact (low M) dipole approximations than single-level scanning.
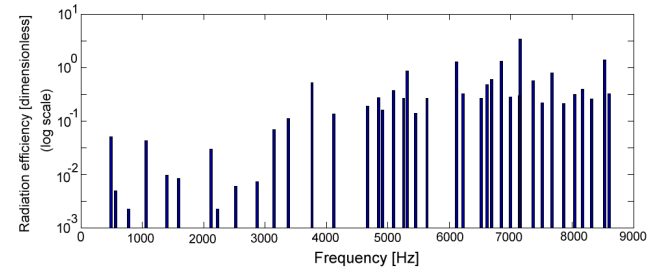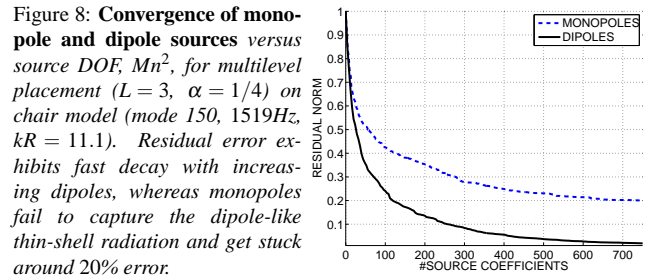

Figure 7: **Radiation efficiency of dragon model** *illustrates that some modes radiate thousands of times more effectively than others. In particular, several low frequencies are suppressed, thereby illustrating that mechanically "dominant" base vibration modes can be less important for sound generation. Even "low accuracy" PAT approximations, e.g., 20% error, can inherit the dramatic radiation efficiency effects.*

Figure 8: **Convergence of monopole and dipole sources** *versus source DOF, Mn^2, for multilevel placement ($L = 3$, $\alpha = 1/4$) on chair model (mode 150, 1519Hz, $kR = 11.1$). Residual error exhibits fast decay with increasing dipoles, whereas monopoles fail to capture the dipole-like thin-shell radiation and get stuck around 20% error.*

Real-time sound synthesis is achieved by computing each mode's PAT absolute pressure amplitudes, $|p(\text{ear})|$. If the object undergoes rigid body motion, listening positions are transformed into the object frame, and $|p|$ interpolated along the listening trajectories; we use a fixed rate of 250 Hz. Offline animations linearly interpolate absolute PAT values at intermediate times, whereas our real-time interactive demonstrations (see CDROM) linearly interpolate $|p|$ with a $1/250$ second delay. Runtime evaluation of modal vibration amplitudes, q, are done using an IIR digital filter [James and Pai 2002]. The resulting sound is computed by summing each mode's absolute PAT pressure, scaled by the modal amplitude, $|p|q$, for each vibration mode. Animations involving rigid body dynamics were time-stepped at audio rates or higher (see Figure 10). Planar ground
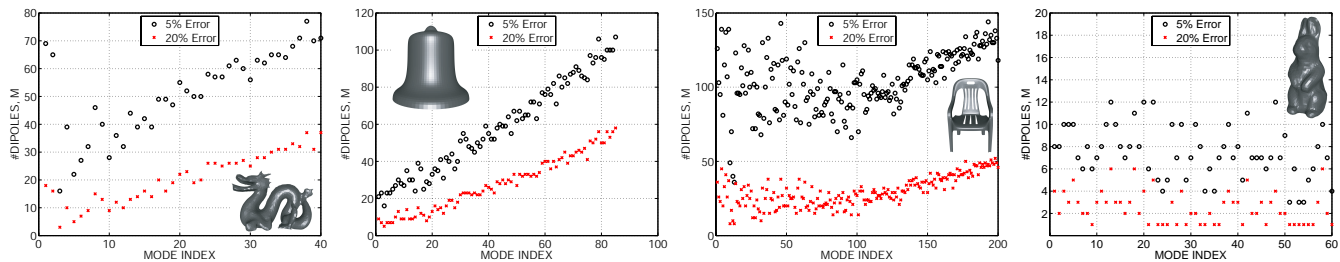
Figure 9: **Number of multipole sources per mode** *computed using the 3-level placement algorithm (L = 3, P = 256) at two approximation tolerances (TOL = 0.05 and 0.20). Output-sensitivity is illustrated by the fact that the number of dipoles needed to approximate each mode are far fewer than the thousands of triangles on the radiating surface. All examples exhibit erratic variations with mode index, e.g., due to greedy placement and varied mode structures, but have a gradual increase in dipole counts at increasing frequency (dragon, bell, chair). The rabbit's simple dipole approximations are an exception that illustrates that it is in the low-frequency regime (small kR) and nearly a dipole source itself.*

| | $L=1$ | | | $L=2$ | | | $L=3$ | | | $L=4$ | | |
	$<M>$	PLACEMULTIPOLES	SVD	$<M>$	PLACEMULTIPOLES	SVD	$<M>$	PLACEMULTIPOLES	SVD	$<M>$	PLACEMULTIPOLES	SVD
$P=16$	78	0.55m	0.42m	80	0.49m	0.91m	79	0.47m	1.2m	–	–	–
$P=64$	73	0.92m	0.32m	71	0.63m	0.29m	74	0.49m	0.41m	75	0.49m	0.53m
$P=256$	71	2.5m	0.27m	68	1.4m	0.27m	69	0.84m	0.31m	70	0.47m	0.35m
$P=1024$	67	7.9m	0.24m	66	4.6m	0.26m	69	2.0m	0.31m	66	0.89m	0.30m
$P=2048$	67	17m	0.29m	66	8.7m	0.24m	67	3.5m	0.31m	66	1.4m	0.32m

Table 2: **Multilevel source placement** *compared for a range of levels L, and candidate source positions, P, for the plastic chair at modest approximation accuracy (10% error). Values for number of dipoles per mode M, and timings of PLACEMULTIPOLES() and computation of equivalent source coefficients (SVD) are averaged over 3 modes (1, 100, 200). The SVD-based solve for source coefficients is based on Intel's MKL library implementation of LAPACK double-precision complex SVD driver. Note that the subspace generated during placement could be reused to reduce SVD solve costs in an optimized implementation. All timings are on a single Opteron 280 core, with PLACEMULTIPOLE() implemented in Java (dipole evaluation cost ≈ 0.39μsec/dipole).*

| Model | 5% Error | | | 20% Error | | |
	Dipoles	Precomp	Eval Rate	Dipoles	Precomp	Eval Rate
Dragon	2056	0.23 hr	1818 Hz	822	0.14 hr	4053 Hz
Rabbit	450	0.27 hr	5899 Hz	149	0.23 hr	20413 Hz
Bell	4979	3.8 hr	799 Hz	2380	1.6 hr	1632 Hz
Chair	20864	19 hr	190 Hz	5958	1.8 hr	574 Hz

Table 3: **PAT precomputation and real-time evaluation rates** *for high (5%) and low (20%) accuracy equivalent source approximations. Total dipole counts (Dipoles), precomputation times for all modes (Precomp), and real-time evaluation rates (Eval Rate) are given (using Pentium IV 3.0GHz). Note that in practice, PAT need only be evaluated at a few hundred Hz, e.g., 250 Hz, and not at the audio sample rate (44100 Hz). All models were constructed with the same multilevel source placement settings (L=3, P=256).*

contacts were resolved using a simple damped linear spring penalty model, with contact forces driving the vibration model. Real-time evaluation performance is easily achieved for our examples. Table 4 gives animation statistics. Comparisons to other sound renderers were also made via careful implementations of the Rayleigh renderer of [O'Brien et al. 2001], ground truth absolute values of acoustic transfer pressure using BEM, and the unscaled sum $\sum_i q_i$ of [van den Doel and Pai 1996]. We also provide a comparison to the traditional far-field ($\|\mathbf{x}\| \gg R$), low-frequency ($kR \ll 1$) monopole approximation [Cremer et al. 1990]

$$|p| = \rho\omega|Q|/(4\pi r), \quad r \gg R, \quad kR \ll 1, \quad (15)$$

where $Q = \int_S v\, dS$ is the so-called volume velocity (compare to similar model in [O'Brien et al. 2002]), and this model sounds nearly identical to the rabbit PAT approximation. Note that (15) yields *zero values* for open double-sided models due to the definition of Q; we provide a single-sided monopole approximation for the thin-shell dragon model using only the outer surface. Sound interactions with the ground were ignored in all renderers. Please see accompanying video. Comparisons show that the modal renderer clearly suffers from a lack of directionality phenomena, especially for highly direction examples such as the swinging bell. One exception is the rabbit model, which is in the low-frequency regime ($kR \approx 1$) thereby making it nearly a monopole source.
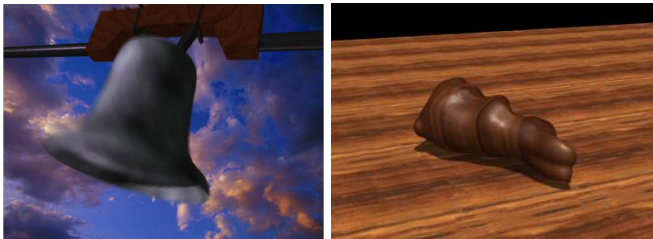
Swaying tin bell

Plastic rabbit

Hollow bronze dragon

Figure 10: **Rigid body animations** *were generated of the dragon and rabbit models falling on the ground, and the bell swaying back and forth. Dynamics and penalty-based contact forces were integrated at audio rates (44100 Hz). See the video for comparisons to other rendering techniques.*

6 Conclusions and Discussion

We have described a fast method for synthesizing sound radiation from geometrically complex vibrating objects. Our Precomputed Acoustic Transfer (PAT) functions are based on accurate approximations to Helmholtz equation solutions generated by standard numerical methods. We introduced an algorithm for constructing equivalent source approximations that enable real-time sound synthesis in physically based animation. Since the number of low-order multipoles required to approximate each vibration mode's acoustic transfer function is independent of the model's geometric complexity, our method exhibits output-sensitive evaluation costs, and is suitable for interactive applications.

Model	Duration	# of |p| evals	BEM Render	PAT Rendering					Rayleigh Rendering	
				PAT Render	Speedup	Dipoles	Throughput	Cost/Dipole	with ray-casting	w/o ray-casting
Dragon (4%)	2.50 s	625	46 min	0.38 s	7263×	2350	1480 Hz	0.29 μsec	8h 42m	7m 40s
Bell (5%)	5.00 s	1250	97 min	1.53 s	3803×	5051	793 Hz	0.24 μsec	15h 20m	51m 58s
Rabbit (5%)	4.00 s	1000	65 min	0.24 s	16250×	686	4166 Hz	0.35 μsec	–	–

Table 4: **Animation Statistics:** *Sound is synthesized at 44100 Hz, with PAT and BEM samples generated (all modes) at 250 Hz and linearly interpolated. The duration and number of PAT/BEM |p| evaluations is provided, along with BEM and PAT evaluation times (BEM Render; PAT Render) and the speedup resulting from PAT (speedup="BEM Render"/"PAT Render"). The total number of dipoles per object (Dipoles), and effective PAT throughput, and cost per dipole evaluation are given. Rayleigh rendering was performed for the dragon model (see video), and timings are given for versions with and without octree-accelerated ray-casting visibility tests on each triangle radiator.*

Boundary element analysis and offset surface evaluation are currently the most expensive part of our preprocess. Fine mesh resolutions are required to resolve small wavelengths (e.g., $6e_{max} < \lambda_{min}$ for maximum edge length e_{max}) and vibration modes, which limits the range of frequencies that can be analyzed [Desmet 2002]. For a fixed maximum frequency, larger objects are more difficult because the mesh needs to be more detailed. For example, the frequency range of the large chair is particularly restricted. Evaluating the BEM solution on the offset surface is also quite expensive because of the detailed meshes required. Faster analyses are needed to support the complex geometry and large audible frequency ranges needed in graphics. Comparing PAT evaluation to fast Helmholtz multipole evaluation costs would be interesting.

Our multipole placement algorithm has general application in engineering acoustics in solving the Neumann radiation problem in the special case of *closed volumetric objects*, e.g., rabbit. Applying our algorithm could remove expensive BEM analyses from the PAT pipeline, proceeding directly to equivalent source representations, possibly reaching higher frequency problems ($kR \gg 1$). Variants for nonclosed shell geometries, such as extensions of hybrid FEM-Trefftz methods [Desmet 2002] are interesting. Higher order sources and higher frequencies should be investigated for PAT.

Our greedy source placement algorithm does not yield optimal source positions, and its ability to obtain minimal source configurations is unknown. The significant variations in M for the chair example (see Figure 9) suggest room for improvement. Optimization can be used to further refine multipole placements [Ochmann 1995], however our preliminary experiments did not yield significantly better results. Improved multipole source and offset surface sampling strategies may help.

PAT can provide a cheap and accurate approximation to Helmholtz radiation from a single object, but at the cost of neglecting the environment. Scattering interactions between multiple PAT objects, analogous to light interactions between PRT models [Sloan et al. 2002], might be approximated. While the method of images can be used to approximate very simple interactions, e.g., with a floor, the simple source-based nature of PAT models is ideal for incorporation into existing general-purpose sound propagation techniques [Funkhouser et al. 1998; Funkhouser et al. 1999; Tsingos et al. 2001; Tsingos et al. 2002]. The PAT approximation involves steady-state frequency analyses, and transient effects can be important, e.g., for very large objects. Exterior radiation problems were considered here, but interior problems may present special challenges.

Level of detail rendering of PAT models has not been addressed, although simply blending between different error PAT approximations is possible. Perceptually based progressive techniques such as [Tsingos et al. 2004] are ideal candidates for simplifying the equivalent source models, however care must be taken since the linear combination of sources can be ill-conditioned. Compression and hardware rendering are also important areas to study.

Ignoring solid-fluid coupling can be a good approximation for many objects in air, however it may be poor for underwater applications where fluid density is higher, or very thin shells. Modal and boundary element analysis can be coupled and solved simultaneously, significantly increasing complexity.

Doppler effects can be important for high speed motion, and have not been addressed here. However, we note that the phase part $\theta(\mathbf{x})$ of the PAT pressure field $p(\mathbf{x}) = A(\mathbf{x})e^{i\theta(\mathbf{x})}$ could be used to approximate Doppler frequency shifts at low velocities ($|\dot{\mathbf{x}}| \ll c$). For example, a local linear analysis implies a frequency shift of $\Delta\omega = (\nabla_{\mathbf{x}}\theta) \cdot \dot{\mathbf{x}}$, where $\dot{\mathbf{x}}$ is the velocity of the listening position.

Techniques exist for using psychoacoustic models to select a subset of modes based on masking thresholds (see [van den Doel et al. 2002; Raghuvanshi and Lin 2006]) and could be applied to PAT to reduce runtime rendering costs. However, this would result in laboriously precomputed PAT modes being discarded. Ideally such approaches could avoid expensive acoustic analysis for culled modes in the first place. Tools for a priori estimation of a mode's radiation efficiency, such as so-called "radiation modes" could be useful [Cunefare and Currey 1994]. Interestingly, reality-based modal sound models [Pai et al. 2001; van den Doel et al. 2001], although lacking nontrivial spatial variations, obtain these expensive radiation efficiency effects "for free." User studies could help understand the speed-accuracy trade-off for PAT approximations; comparisons in the video suggest that very fast PAT approximations of modest accuracy, e.g., 20%, may be sufficient.

Our method assumes both small deformations and small fluid pressure fluctuations, and these are violated by objects undergoing large deformations. Large deformations can also produce complex aeroacoustic effects, e.g., a cracking whip. It would be desirable to extend PAT to large-deformation modal vibration models, such as [Barbič and James 2005]. For general deformable simulations, simplified sound approximations such as the Rayleigh-based renderer [O'Brien et al. 2001] may be more practical. For some deformations it might be possible to interpolate PAT functions in reduced dimensions similar to deformable PRT [James and Fatahalian 2003], or parameterize equivalent sources like [Sloan et al. 2005].

References

ABRAMOWITZ, M., AND STEGUN, I. A. 1964. *Handbook of Mathematical Functions with Formulas, Graphs, and Mathematical Tables*. Dover, New York.

BARBIČ, J., AND JAMES, D. 2005. Real-Time Subspace Integration for St. Venant-Kirchhoff Deformable Models. *ACM Transactions on Graphics 24*, 3 (Aug.), 982–990.

CARDLE, M., BROOKS, S., BAR-JOSEPH, Z., AND ROBINSON, P. 2003. Sound-by-Numbers: Motion-Driven Sound Synthesis. In *2003 ACM SIGGRAPH / Eurographics Symp. on Computer Animation*, 349–356.

CISCOWSKI, R. D., AND BREBBIA, C. A. 1991. *Boundary Element methods in acoustics*. Computational Mechanics Publications and Elsevier Applied Science, Southampton.

COOK, P. 2002. *Real Sound Synthesis for Interactive Applications*. A.K. Peters.

CREMER, L., HECKL, M., AND UNGAR, E. 1990. *Structure Borne Sound : Structural Vibrations and Sound Radiation at Audio Frequencies*, 2nd ed. Springer, January.

CUNEFARE, K. A., AND CURREY, M. N. 1994. On the exterior acoustic radiation modes of structures. *The Journal of the Acoustical Society of America 96*, 4 (October), 2302–2312.

DESMET, W. 2002. Mid-frequency vibro-acoustic modelling: challenges and potential solutions. In *Proceedings of ISMA 2002*, vol. II.

DOBASHI, Y., YAMAMOTO, T., AND NISHITA, T. 2003. Real-Time Rendering of Aerodynamic Sound Using Sound Textures Based on Computational Fluid Dynamics. *ACM Trans. on Graphics 22*, 3 (July), 732–740.

FUNKHOUSER, T., CARLBOM, I., ELKO, G., PINGALI, G., SONDHI, M., AND WEST, J. 1998. A Beam Tracing Approach to Acoustic Modeling for Interactive Virtual Environments. In *Proceedings of SIGGRAPH 98*, Computer Graphics Proceedings, Annual Conference Series, 21–32.

FUNKHOUSER, T. A., MIN, P., AND CARLBOM, I. 1999. Real-Time Acoustic Modeling for Distributed Virtual Environments. In *Proceedings of SIGGRAPH 99*, 365–374.

GOLUB, G., AND VAN LOAN, C. 1996. *Matrix Computations*, third ed. The Johns Hopkins University Press, Baltimore.

GOUNOT, Y., MUSAFIR, R. E., AND SLAMA, J. G. 2005. A Comparative Study of Two Variants of the Equivalent Sources Method in Scattering Problems. *Acta Acustica united with Acustica 91*, 5, 860–872.

GUMEROV, N., AND DURAISWAMI, R. 2005. *Fast Multipole Methods for the Helmholtz Equation in Three Dimensions*. Elsevier Series in Electromagnetism. Elsevier Science, March.

JAMES, D. L., AND FATAHALIAN, K. 2003. Precomputing Interactive Dynamic Deformable Scenes. *ACM Trans. on Graphics 22*, 3, 879–887.

JAMES, D. L., AND PAI, D. K. 2002. DyRT: Dynamic Response Textures for Real Time Deformation Simulation With Graphics Hardware. *ACM Transactions on Graphics 21*, 3 (July), 582–585.

JOHNSON, M. E., LALIME, A. L., GROSVELD, F. W., RIZZI, S. A., AND SULLIVAN, B. M. 2003. Development of an efficient binaural simulation for the analysis of structural acoustic data. In *8th Intl. Conf. on Recent Advances in Structural Dynamics*.

KITA, E., AND KAMIYA, N. 1995. Trefftz method: An overview. *Advances in Engineering Software 24*, 89–96.

LEWINER, T., LOPES, H., VIEIRA, A. W., AND TAVARES, G. 2003. Efficient implementation of Marching Cubes' cases with topological guarantees. *Journal of Graphics Tools 8*, 2, 1–15.

LOKKI, T., SAVIOJA, L., VÄÄNÄNEN, R., HUOPANIEMI, J., AND TAKALA, T. 2002. Creating Interactive Virtual Auditory Environments. *IEEE Computer Graphics & Applications 22*, 4 (July-August), 49–57.

MAGALHÃES, M. B. S., AND TENENBAUM, R. A. 2004. Sound Sources Reconstruction Techniques: A Review of Their Evolution and New Trends. *Acta Acustica united with Acustica 90*, 199–220.

O'BRIEN, J. F., COOK, P. R., AND ESSL, G. 2001. Synthesizing Sounds From Physically Based Motion. In *Proceedings of ACM SIGGRAPH 2001*, 529–536.

O'BRIEN, J. F., SHEN, C., AND GATCHALIAN, C. M. 2002. Synthesizing Sounds from Rigid-Body Simulations. In *ACM SIGGRAPH Symposium on Computer Animation*, 175–181.

OCHMANN, M. 1995. The Source Simulation Technique for Acoustic Radiation Problems. *Acustica 81*.

OCHMANN, M. 1999. The full-field equations for acoustic radiation and scattering. *Journal of the Acoustical Society of America 105*, 5 (May).

PAI, D. K., VAN DEN DOEL, K., JAMES, D. L., LANG, J., LLOYD, J. E., RICHMOND, J. L., AND YAU, S. H. 2001. Scanning Physical Interaction Behavior of 3D Objects. In *Proc. of ACM SIGGRAPH 2001*, 87–96.

PAVIĆ, G. 2005. An Engineering Technique for the Computation of Sound Radiation by Vibrating Bodies Using Substitute Sources. *Acta Acustica united with Acustica 91*, 1, 1–16.

PAVIĆ, G. 2006. A Technique for the Computation of Sound Radiation by Vibrating Bodies Using Multipole Substitute Sources. *Acta Acustica united with Acustica 92*, 1, 112–126.

RAGHUVANSHI, N., AND LIN, M. C. 2006. Interactive Sound Synthesis for Large Scale Environments. In *SI3D '06: Proceedings of the 2006 symposium on Interactive 3D graphics and games*, ACM Press, New York, NY, USA, 101–108.

SHABANA, A. A. 1990. *Theory of Vibration, Volume II: Discrete and Continuous Systems*, first ed. Springer-Verlag, New York, NY.

SLOAN, P.-P., KAUTZ, J., AND SNYDER, J. 2002. Precomputed Radiance Transfer for Real-Time Rendering in Dynamic, Low-Frequency Lighting Environments. *ACM Transactions on Graphics 21*, 3 (July), 527–536.

SLOAN, P.-P., HALL, J., HART, J., AND SNYDER, J. 2003. Clustered Principal Components for Precomputed Radiance Transfer. *ACM Transactions on Graphics 22*, 3 (July), 382–391.

SLOAN, P.-P., LUNA, B., AND SNYDER, J. 2005. Local, Deformable Precomputed Radiance Transfer. *ACM Transactions on Graphics 24*, 3 (Aug.), 1216–1224.

STAM, J. 1999. Diffraction Shaders. In *Proceedings of SIGGRAPH 99*, Computer Graphics Proceedings, Annual Conference Series, 101–110.

STETTNER, A., AND GREENBERG, D. P. 1989. Computer Graphics Visualization For Acoustic Simulation. In *Computer Graphics (Proceedings of SIGGRAPH 89)*, vol. 23, 195–206.

TAKALA, T., AND HAHN, J. 1992. Sound rendering. In *Computer Graphics (Proceedings of SIGGRAPH 92)*, vol. 26, 211–220.

TSINGOS, N., FUNKHOUSER, T., NGAN, A., AND CARLBOM, I. 2001. Modeling Acoustics in Virtual Environments Using the Uniform Theory of Diffraction. In *Proceedings of ACM SIGGRAPH 2001*, Computer Graphics Proceedings, Annual Conference Series, 545–552.

TSINGOS, N., CARLBOM, I., ELBO, G., KUBLI, R., AND FUNKHOUSER, T. 2002. Validating Acoustical Simulations in Bell Labs Box. *IEEE Computer Graphics & Applications 22*, 4 (July-August), 28–37.

TSINGOS, N., GALLO, E., AND DRETTAKIS, G. 2004. Perceptual Audio Rendering of Complex Virtual Environments. *ACM Transactions on Graphics 23*, 3 (Aug.), 249–258.

VAN DEN DOEL, K., AND PAI, D. K. 1996. Synthesis of shape dependent sounds with physical modeling. In *Intl Conf. on Auditory Display*.

VAN DEN DOEL, K., KRY, P. G., AND PAI, D. K. 2001. FoleyAutomatic: Physically-Based Sound Effects for Interactive Simulation and Animation. In *Proceedings of ACM SIGGRAPH 2001*, 537–544.

VAN DEN DOEL, K., PAI, D. K., ADAM, T., KORTCHMAR, L., AND PICHORA-FULLER, K. 2002. Measurements of Perceptual Quality of Contact Sound Models. In *Intl. Conf. on Auditory Display*, 345–349.

VON ESTORFF, O., Ed. 2000. *Boundary Elements in Acoustics: Advances and Applications*. WIT Press, Southhampton, UK.

VON ESTORFF, O. 2003. Efforts to Reduce Computation Time in Numerical Acoustics – An Overview. *Acta Acustica united with Acustica 89*, 1–13.

WAND, M., AND STRASSER, W. 2004. Multi-resolution sound rendering. In *2004 Eurographics Symp. on Point-Based Graphics*, Eurographics.

WU, T. W., Ed. 2000. *Boundary Element Acoustics: Fundamentals and Computer Codes*. WIT Press, Southhampton, UK.