

Harmonic Shells: A Practical Nonlinear Sound Model for Near-Rigid Thin Shells

Jeffrey N. Chadwick

Steven S. An
Cornell University

Doug L. James




Figure 1: Crash! Our physically based sound renderings of thin shells produce characteristic “crashing” and “rumbling” sounds when animated using rigid body dynamics. We synthesize nonlinear modal vibrations using an efficient reduced-order dynamics model that captures important nonlinear mode coupling. High-resolution sound field approximations are generated using far-field acoustic transfer (FFAT) maps, which are precomputed using efficient fast Helmholtz multipole methods, and provide cheap evaluation of detailed low- to high-frequency acoustic transfer functions for realistic sound rendering.

Abstract

We propose a procedural method for synthesizing realistic sounds due to nonlinear thin-shell vibrations. We use linear modal analysis to generate a small-deformation displacement basis, then couple the modes together using nonlinear thin-shell forces. To enable audiorate time-stepping of mode amplitudes with mesh-independent cost, we propose a reduced-order dynamics model based on a thin-shell cubature scheme. Limitations such as mode locking and pitch glide are addressed. To support fast evaluation of mid-frequency mode-based sound radiation for detailed meshes, we propose *far-field acoustic transfer maps* (FFAT maps) which can be precomputed using state-of-the-art fast Helmholtz multipole methods. Familiar examples are presented including rumbling trash cans and plastic bottles, crashing cymbals, and noisy sheet metal objects, each with increased richness over linear modal sound models.

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; H.5.5 [Information Systems]: Information Interfaces and Presentation—Sound and Music Computing

Keywords: Sound synthesis; thin shells; contact sounds; modal analysis; dimensional model reduction; subspace integration; acoustic transfer; Helmholtz equation

1 Introduction

Linear modal sound models are widely used for rigid bodies in computer animation and virtual environments [van den Doel et al. 2001; O’Brien et al. 2002; Bonneel et al. 2008], and when combined with acoustic transfer models for sound radiation [James et al. 2006] they

can provide convincing physically based sound sources, especially for pure ringing tones such as chimes, bells, or “knocks.” Unfortunately, we lack effective sound models for a broad class of noisy virtual objects: thin shells (objects with thicknesses orders of magnitude smaller than their other dimensions). Thin shells are very common in real and virtual environments, and produce rich and easily recognizable impact sounds: sheet metal objects (trash cans, oil drums, tin roofs, machinery), plastic containers (water bottles), musical instruments (cymbals), etc. Their rich nonlinear vibrations produce proverbial “crashes” and “rumbles” that are poorly approximated by linear modal sound models which lack nonlinear mode coupling. To make matters worse, thin shells are often very loud and important sound sources due to their ability to vibrate and radiate sound so effectively, e.g., consider a metal roof pelted by hail. Alas, their expensive nonlinear dynamics have made thin shells computationally impractical for physically based sound synthesis.

In this paper, we propose an efficient method for synthesizing realistic sounds from thin-shell structures undergoing small but nonlinear vibrations. Given a description of an object’s geometry and material properties, we compute linear vibration modes, then couple these modes together using the nonlinear thin-shell force model. To accelerate nonlinear modal dynamics, we optimize a thin-shell cubature scheme to evaluate reduced-order shell forces at costs independent of the geometric complexity of the model. We show that the complex internal dynamics of thin-shell models can be approximated with sufficient accuracy and efficiency to allow practical synthesis of plausible thin-shell sounds. We also address sound-related locking effects that arise when simulating nonlinear modal dynamics that might produce pitch-glide artifacts in general animations.

Our nonlinear reduced-order dynamics model can synthesize modal vibrations using hundreds of modes, which then drive sound radiation. Unfortunately the estimation of sound wave radiation via prior acoustic transfer models is complicated for two reasons: (1) the nonlinear mode vibrations are no longer linear harmonics, and (2) higher frequency acoustic transfer with high-resolution meshes is expensive to precompute, represent, and evaluate at runtime. First, we observe that nonlinear thin-shell vibrations produced by our animations exhibit frequency-localized modes for which linear frequency-domain radiation models still provide a plausible approximation. Second, we propose **far-field acoustic transfer maps**

(FFAT maps) for fast runtime evaluation of high-frequency acoustic transfer from general modal vibration sources. Our texture-based approach leverages state-of-the-art fast Helmholtz multipole methods to precompute acoustic transfer functions for complex thin-shell (or more general) structures, while delivering the simplicity and speed of texture sampling for runtime transfer evaluation. At runtime, sounds are synthesized by time-stepping the nonlinear reduced-order model to estimate modal amplitudes, which are then multiplied by acoustic transfer values to auralize the thin-shell sound source.

Other Related Work: Thin elastically deformable models have seen widespread use in computer animation, especially for large-deformation simulations of parameterized cloth models [Terzopoulos et al. 1987; Baraff and Witkin 1998]. Plate-like cloth models with flat rest configurations are most common, however recently shell models for objects with intrinsically curved rest configurations have gained attention, e.g., to model clothing with folds and wrinkles [Bridson et al. 2003]. While various thin-shell models exist in the mechanics literature (see [Chapelle and Bathe 2003]), discrete shell models that better meet computer animation needs have appeared recently that use simple bending energy formulations that allow standard cloth solvers to produce convincing large-deformation thin-shell dynamics [Grinspun et al. 2003; Bridson et al. 2003]; we use the elastic thin-shell model summarized in Gingold et al. [2004]. Other works have aimed to make shell simulations faster and better [Cirak et al. 2000; Cirak and Ortiz 2001; Green et al. 2002]. The mathematical structure of bending energy formulations has been exploited and simplified to accelerate near-isometric deformation of thin plates [Bergou et al. 2006] and shells [Garg et al. 2007]. Our work differs in that we do not require large-deformation animations or sophisticated linear system solvers, but rather focus on small mode-related shell vibrations and try to obtain explicit time-stepping costs sublinear in geometric complexity to enable audio-rate sound synthesis. We also prefer mode-based representations since they are preferred for auralization with frequency-domain acoustic transfer models.

Linear modal vibration models are well known in animation [Pentland and Williams 1989; James and Pai 2002]. They have proven effective for procedural sound synthesis, and are increasingly used to simulate rigid-body impact sounds [Adrien 1991; Cook 2002; van den Doel and Pai 1996; van den Doel et al. 2001; O’Brien et al. 2002] in part due to excellent synthesis speed for interactive applications [van den Doel et al. 2001; Raghuvanshi and Lin 2006; Bonneel et al. 2008]. “Modal warping” has been used to approximate large-deformation thin shells for animation [Choi et al. 2007], but such models are of limited use for nonlinear sound synthesis since the underlying linear modal oscillators are uncoupled by design.

There is a tremendous amount of work on the nonlinear vibrations of plates and shells in engineering [Nayfeh and Mook 1979; Mousaoui and Benamar 2002] and related structure borne sound [Cremmer et al. 1990]. Important examples are nonlinear vibrations in musical instruments [Fletcher 1999], such as (the notoriously difficult to simulate) gongs and cymbals [Chaigne et al. 2005]. Simple all-pass nonlinear passive filters have been used to mimic nonlinear mode-coupling effects [Pierce and Van Duyne 1997]. For nonlinear modal analysis, linear eigenmodes are often used for nonlinear subspace integration of dynamics to resolve weak material nonlinearities in small-strain configurations and resolve mode-mode coupling [Bathe 1996] (for a good discussion on nonlinear mode coupling in beams and plates see [Malatkar 2003]). Nonlinear normal modes [Nayfeh and Nayfeh 1995; Touzé et al. 2004] have been used to describe a nonlinear vibration mode’s shape as a linear superposition of other linear modes, i.e., to resolve nonlinear mode coupling. However, likely due to the high computational complexity of simulating nonlinearly coupled modal models (often $O(r^4)$ or

worse), we are unaware of sound synthesis results in the literature that demonstrate results comparable to ours, i.e., with several hundred fully coupled modes. We achieve this by extending cubature optimization techniques of An et al. [2008]; their method estimates volumetric cubature schemes, and was even used to evaluate nonlinear modal shell vibrations for sound synthesis, but the thin shell had to be modeled using several hundred thousand tetrahedral elements. We extend cubature schemes to handle thin shells more efficiently, and obtain cubature-based reduced-order modal models with $O(r^2)$ time-step complexity for r nonlinearly coupled modes.

The only physically based sound rendering work in graphics that addresses nonlinear object vibrations is O’Brien et al. [2001]. They use an explicitly integrated large-deformation finite element model to simulate nonlinear vibrations of objects using small time-step sizes, and a time-domain ray-based “Rayleigh method” (a.k.a. direct propagation) to approximate sound radiation. Interesting results were obtained for short animations with large deformations and buckling. Unfortunately simulation times were on the order of a day (circa 2001) for models of rather modest geometric complexity (<2000 tetrahedra), and the expensive radiation model is known to have limited accuracy [Desmet 2002]. Bilbao has considered energy conserving finite difference discretizations and time-stepping schemes for nonlinear plates to generate plausible sounds [Bilbao 2008]. In contrast to these works, we focus on efficient sound models for nonlinearly forced thin-shell vibrations: we develop efficient subspace integration techniques for geometrically complex nonlinear modal models which drive frequency-domain acoustic transfer models that are consistent with frequency-domain wave radiation.

The most closely related work on sound radiation to ours is “Precomputed Acoustic Transfer” [James et al. 2006], which augmented a linear modal sound model with a multipole-based Helmholtz approximation of each mode’s acoustic transfer function, $p(\mathbf{x})$. The approach enables real-time computation of vibration and sound radiation for geometrically complex objects. Unfortunately due to the increasing complexity of mid- to high-frequency radiation (higher kL values) [Desmet 2002], increasingly complex multipole approximations are required at higher frequencies; such models are costly to precompute, costly to evaluate at runtime, and become increasingly inaccurate at difficult higher frequencies. There is significant work in the acoustics community on reconstructing acoustic quantities (such as sound source multipole expansion coefficients) from often sparse acoustic pressure measurements, e.g., for near-field acoustic holography, but most methods are limited to low- to mid-frequency problems [Vorlander 2007; Wu 2008]. Our far-field acoustic transfer (FFAT) maps provide a simple method for constant-time transfer evaluation which handles low- to high-frequency far-field radiation complexity by exploiting (1) fast multipole boundary element methods from acoustics [Liu 2009; Shen and Liu 2007], and (2) texture-based far-field expansions that are well suited to capturing rapid angular variations with simple radial structure.

Finally, it is possible in principle to simulate fully nonlinear vibrations coupled to 3D (nonlinear) sound fields. For instance, two-dimensional simulations of a nonlinear plate coupled to acoustic fluid have been modeled using the nonlinear Euler equations and used to produce nonlinear far-field radiation [Frendi et al. 1994]. Unfortunately such approaches are currently computationally unappealing for 3D audio-rate simulations of nonlinear sound.

2 Nonlinear Modal Vibrations for Thin Shells

2.1 Background: Thin-Shell Dynamics

We consider discrete thin-shell models on manifold with boundary triangle meshes with N_Δ triangles, and N_v vertices. Following a

suitable discretization via the finite element (or other) method we obtain an N -dimensional system of ordinary differential equations,

$$\mathbf{M}\ddot{\mathbf{u}} + \mathbf{D}\dot{\mathbf{u}} + \mathbf{f}_{int}(\mathbf{u}) = \mathbf{f}_{ext} \quad (1)$$

where $\mathbf{u} = \mathbf{u}(t) \in \mathbb{R}^N$ are mesh vertex displacements, $\dot{\mathbf{u}}$ are velocities, $\ddot{\mathbf{u}}$ are accelerations, $\mathbf{M} \in \mathbb{R}^{N \times N}$ is the mass matrix; $\mathbf{f}_{int}(\mathbf{u})$ describes nonlinear internal thin-shell forces; and \mathbf{f}_{ext} are time-dependent external forces such as gravity or contact forces. For lightly damped small vibrations, we use linear Rayleigh damping [Bathe 1996] with $\mathbf{D} = \alpha\mathbf{M} + \beta\mathbf{K}$, with \mathbf{K} the stiffness matrix (Jacobian of \mathbf{f}_{int}) evaluated at $\mathbf{u} = 0$.

Without loss of generality, we use the elastic shell model of [Gingold et al. 2004], in part because it is based on physical material parameters that simplify parameter tuning for sound synthesis. The deformation strain energy, $E(\mathbf{u})$, is an integral over the surface of the strain energy density, $W(\mathbf{u}; X)$, where X is a material position on the undeformed surface. The strain energy density is decomposed in two parts

$$W = W_m + W_b \quad (2)$$

where W_m is the membrane strain energy density which penalizes tangential stretching or compression; and W_b is the bending energy density which resists bending away from the rest configuration. The membrane and bending strain energy densities are defined in terms of per-triangle strain tensors:

$$W_m = \frac{Yh}{2(1-\nu^2)} (1-\nu) \text{tr}(\epsilon_m^2) + \nu \text{tr}(\epsilon_m)^2 \quad (3)$$

$$W_b = \frac{Yh^3}{24(1-\nu^2)} (1-\nu) \text{tr}(\epsilon_b^2) + \nu \text{tr}(\epsilon_b)^2 \quad (4)$$

where h is the shell thickness, Y is Young’s modulus, and ν is Poisson’s ratio (see Gingold et al. [2004] equations for $W_{membrane} = W_m$ and $W_{bending}^{Koiter} = W_b$). The membrane and bending strains are ϵ_m and ϵ_b , respectively; ϵ_m is a 3×3 tensor which varies as triangle edges deviate from their rest lengths; and ϵ_b is a 3×3 tensor which changes as the dihedral angles between a triangle and its three neighbors vary from the corresponding dihedral angles in the shell’s rest pose (see [Gingold et al. 2004] for definitions). Consequently, the strains and strain energy density, W , can be defined as piecewise constant over shell triangles, with each triangle’s W_m value a function of the triangle’s 3 vertex positions, whereas its bending strain energy, W_b , is a function of 6 vertices (see Figure 2).

It follows that the deformation energy is given by the integral over the undeformed surface, S :

$$E(\mathbf{u}) = \int_S W(\mathbf{u}; X) dS_X = \sum_{i=1}^{N_\Delta} A_i W_i(\mathbf{u}) \quad (5)$$

where W_i is the piecewise constant strain energy density for triangle i , and A_i is its area. The desired internal thin-shell force is

$$\mathbf{f}_{int} = \nabla_{\mathbf{u}} E(\mathbf{u}) = \sum_{i=1}^{N_\Delta} A_i \nabla_{\mathbf{u}} W_i(\mathbf{u}). \quad (6)$$

which gathers both membrane and bending contributions.

2.2 Reduced-order Thin-Shell Dynamics

We now describe the nonlinear mode-coupled dynamics model used for sound synthesis.

Linear Modal Analysis Basics: As a first step, we compute r linear eigenmodes of the thin-shell system (1) using standard methods [Shabana 1990; Bathe 1996]. We compute undamped vibration modes by solving the generalized eigenvalue problem,

$$\mathbf{K}\mathbf{u}_j = \omega_j^2 \mathbf{M}\mathbf{u}_j, \quad j = 1 \dots r, \quad (7)$$

where the j^{th} displacement eigenmode, \mathbf{u}_j , corresponds to a vibration at the j^{th} (smallest but nonzero) natural frequency, ω_j . We compute the first r modes corresponding to nonzero eigenvalues (“rigid body modes” with zero eigenvalues are discarded). The eigenvectors are assembled in the mode matrix, $\mathbf{U} = [\mathbf{u}_1 \mathbf{u}_2 \dots \mathbf{u}_r] \in \mathbb{R}^{N \times r}$ where N is the number of degrees of freedom in the system ($r \ll N$). Without loss of generality we “mass normalize” the eigenvectors so that

$$\widetilde{\mathbf{M}} = \mathbf{U}^T \mathbf{M} \mathbf{U} = \mathbf{I} \quad \text{and} \quad \widetilde{\mathbf{K}} = \mathbf{U}^T \mathbf{K} \mathbf{U} = \mathbf{\Lambda} = \text{diag}(\omega^2). \quad (8)$$

The linear modal model assumes vertex displacements are given by a linear superposition of mode shapes, $\mathbf{u}(t) = \mathbf{U}\mathbf{q}(t)$, where $\mathbf{q}(t) \in \mathbb{R}^r$ are the generalized modal coordinates.

Nonlinear Subspace Integration: To obtain nonlinear coupling between linear modes we employ dimensional model reduction [Bathe 1996; Krysl et al. 2001] by substituting $\mathbf{u} = \mathbf{U}\mathbf{q}$ into the full-dimensional equations of motion (1), and premultiplying by \mathbf{U}^T to project into the r -dimensional modal subspace:

$$\ddot{\mathbf{q}} + \widetilde{\mathbf{D}}\dot{\mathbf{q}} + \widetilde{\mathbf{f}}_{int}(\mathbf{q}) = \widetilde{\mathbf{f}}_{ext} \quad (9)$$

where

$$\widetilde{\mathbf{D}} = \alpha\mathbf{I} + \beta\mathbf{\Lambda} \quad (10)$$

$$\widetilde{\mathbf{f}}_{int}(\mathbf{q}) = \mathbf{U}^T \mathbf{f}_{int}(\mathbf{U}\mathbf{q}) \quad (11)$$

$$\widetilde{\mathbf{f}}_{ext} = \mathbf{U}^T \mathbf{f}_{ext}. \quad (12)$$

This equation provides nonlinear mode coupling when integrated (see Figure 3). Unfortunately explicit time-stepping of these nonlinear equations at audio rates (44.1 kHz; $\Delta t \approx 2.3 \times 10^{-5}$) is quite expensive due to $\widetilde{\mathbf{f}}_{int}$ evaluation, which involves subspace-projection of \mathbf{f}_{int} in (6) via a gather over N_Δ triangles—an expensive $O(r N_\Delta)$ cost per audio timestep.




Figure 2: Stencil of triangle element




Figure 3: **Nonlinear mode coupling:** The nonlinear and linear modal dynamics of $q_{200}(t)$ are compared for the ride cymbal’s response to a single metal ball impact (the first video example). The nonlinear mode exhibits rich dynamics, and strong coupling to lower frequency modes.

In the linearized small-deformation case this bottleneck disappears since the modal equations decouple: $\mathbf{f}_{int} = \mathbf{K}\mathbf{u}$ so that $\widetilde{\mathbf{f}}_{int} = \mathbf{\Lambda}\mathbf{q}$. Each mode can be integrated efficiently using an IIR filter [Hamming 1983] (e.g., using filter coefficients in [James and Pai 2002]) so that each explicit time-step has only $O(r)$ cost, but nonlinear mode coupling is lost.

2.3 Thin-Shell Cubature Scheme

We accelerate $\widetilde{\mathbf{f}}_{int}$ evaluation by extending the volumetric approach of An et al. [2008] to optimize cubature schemes for thin shells. We use a compound cubature scheme that evaluates both bending and stretching force integrals simultaneously thereby sharing vertex-deformation computations. We make the approximation

$$\widetilde{\mathbf{f}}_{int}(\mathbf{q}) = \mathbf{U}^T \mathbf{f}_{int}(\mathbf{U}\mathbf{q}) = \sum_{i=1}^{N_\Delta} A_i \mathbf{g}_i(\mathbf{q}) \quad (13)$$

$$\approx \sum_{i \in \mathcal{e}} w_i \mathbf{g}_i(\mathbf{q}), \quad (14)$$

where $\mathbf{g}_i(\mathbf{q}) \equiv \mathbf{U}^T \nabla_{\mathbf{u}} W_i(\mathbf{U}\mathbf{q})$, and w_i are cubature weights, and \mathcal{C} are a set of integers defining cubature triangles (see Figure 4).




Figure 4: Illustration of cubature scheme: (Left) Trash can (200 modes) with 800-feature cubature scheme; (Right) close-up of triangle-flap features.

Cubature Training: Since the strain energy density (and thus force density) defined in the shell model is piecewise-constant over a triangle element, it suffices to only consider one cubature point per triangle during the optimization pre-process. However, at run-time the triangle and its three edge flaps must be reconstructed for each cubature point, as the bending force is dependent on the state of neighboring triangles. We also considered alternate cubature schemes, including optimizing separate weights for the two force integrals, but this did not significantly affect convergence behavior. To generate training samples for cubature optimization, we randomly sample a Gaussian for each mode to get physically plausible training poses [An et al. 2008]. Convergence plots are shown in Figure 5.

$O(r^2)$ **Force Evaluation:** In practice, we can construct cubature schemes to modest accuracies (e.g., 10% relative error) which have the number of cubature samples, $n = |\mathcal{C}| = O(r) \ll N_\Delta$, so that the thin-shell cubature scheme can approximate $\mathbf{f}_{int}(\mathbf{q})$ at $O(r^2)$ cost. More important than complexity is that the scheme is fast in practice. Achieving full nonlinear coupling of r modes at $O(r^2)$ cost makes offline sound synthesis practical for several hundred modes (see Table 2). To understand this cost, note that each $\mathbf{g}_i(\mathbf{q})$ r -vector in (14) can be evaluated in $O(r)$ flops using three steps: given the triangle element’s 6 stencil vertices, V (see Figure 2), we (1) evaluate the 6 vertex displacements, $\mathbf{u}_V = \mathbf{U}_V \mathbf{q}$ in $O(r)$ flops; (2) evaluate the 6 vertex forces $\mathbf{f}_V = \nabla_{\mathbf{u}_V} W_i(\mathbf{u}_V)$ in $O(1)$ flops, then (3) project the vertex forces into the r -dimensional subspace, $\mathbf{g}_i = \mathbf{U}_V^T \mathbf{f}_V$, in $O(r)$ flops.

3 Limiting Artificial “Pitch Glide”

Rigid body animations can generate a wide range of contact impulses, including violent impacts that would dent or damage the shell in reality. Unlike the linear modal model which simply gets proportionally louder with increased impulse strength, we need to take precautions to ensure that the nonlinear modal model operates in a valid energetic range to avoid “mode locking” [Bathe 1996]. Intuitively speaking, hard forcing can produce mode locking when the model has insufficient degrees of freedom to deform due the limitations of the linear modal basis. Without energetic limits, listeners may perceive a nonphysical “pitch glide” artifact during contact, during which vibration frequencies start out very high, but then glide down to their natural frequencies (see Figure 6). While “pitch glide” occurs naturally for some objects, e.g., Chinese opera gongs [Fletcher 1999], locking in low-frequency modes can also artificially increase the effective numerical stiffness of coupled high-frequency modes, e.g., by stretching the shell, which introduces pitch-glide artifacts (also see Figure 8). Numerically we would also like to avoid locking since that nonlinearity can introduce severe time-step restrictions; in our implementation, we use a fixed-rate timestep to avoid temporal artifacts, so smaller timesteps are particularly undesirable. In a full degree-of-freedom model, the

elements would be free to bend yet still avoid excess stretching, so the frequencies remain roughly constant. Unfortunately, full simulations are often impractically expensive and numerically unstable compared to reduced simulation, so we propose a simple technique for limiting pitch-glide artifacts in practice.

Impulse Limiter: To avoid forcing the nonlinear oscillator to high energies, we use a simple per-timestep impulse filter that bounds the system’s vibrational energy. Given the desired velocity impulse, $\Delta \dot{\mathbf{q}} = \Delta t \tilde{\mathbf{f}}_{ext}$, we introduce an (as yet unknown) scale factor, $0 \leq \alpha \leq 1$, and only apply the scaled impulse, $\alpha \Delta \dot{\mathbf{q}}$. We estimate the total energy (kinetic + potential) before the impulse as,

$$\mathcal{E} = \dot{\mathbf{q}}^T \dot{\mathbf{q}} + \mathbf{q}^T \Lambda \mathbf{q}, \quad (15)$$

and the α -parameterized post-impulse energy as,

$$\mathcal{E}'(\alpha) = (\dot{\mathbf{q}} + \alpha \Delta \dot{\mathbf{q}})^T (\dot{\mathbf{q}} + \alpha \Delta \dot{\mathbf{q}}) + \mathbf{q}^T \Lambda \mathbf{q}. \quad (16)$$

We set a maximum allowed energy, $\mathcal{E}_{max} = \sigma M$ by specifying σ , the maximum energy per unit mass, where M is the object’s total mass; in our implementation, we use the same σ value for all scene objects (typically $\sigma \approx 1 - 4$). We determine the impulse limiter’s α value using Algorithm 1.

Algorithm 1: Compute α for impulse limiter

```

1 begin
2   if  $\mathcal{E}'(1) \leq \mathcal{E}_{max}$  or  $\mathcal{E}'(1) < \mathcal{E}$  then
3     return  $\alpha = 1$ ;
4   else if  $\mathcal{E} > \mathcal{E}_{max}$  then
5     return  $\alpha = 0$ ;
6   else
7     solve  $\mathcal{E}'(\alpha^*) = \mathcal{E}_{max}$  for  $\alpha^* \in [0, 1]$ ;
8     return  $\alpha^*$ ;
9 end
```

4 Mapping Far-Field Acoustic Transfer

We now describe a general method for approximating acoustic transfer using far-field acoustic transfer (FFAT) maps. Please see Appendix A (and [James et al. 2006]) for background on acoustic-transfer-based sound rendering of modal models. Our approach involves three steps: (1) for each mode we recompute detailed pressure samples on concentric exterior spherical surfaces using commodity Helmholtz boundary integral solvers, then (2) we precompute a low-order Laurent expansion for each outgoing ray direction; then (3) at runtime we can evaluate a low-order expansion of any mode’s transfer value at a far-field listening position using $O(1)$ operations. This data-driven approach avoids the use of multi-point multipole expansions (as in James et al. [2006]) which can be complex to fit and evaluate for higher frequency radiation [Wu 2008].

Far-Field Acoustic Transfer (FFAT) Maps: To accelerate render-time acoustic transfer evaluation, we propose a simple approximation to the pressure field, $p(\mathbf{x})$, motivated by the far-field, asymptotic M -term series expansion [Gumerov and Duraiswami 2005]

$$p(\mathbf{x}) \sim h_0(kR) \sum_{j=1}^M \frac{\Psi_j(\theta, \phi)}{(kR)^{j-1}} \quad (17)$$

for $\mathbf{x} = \text{spherical}(\theta, \phi, R)$, where h_0 is the monopole-like 0th-order spherical Hankel function of the first kind, $h_0(kR) = -\frac{ie^{-ikR}}{kR}$, and the Ψ_j functions describe the angular dependence of the pressure field. In practice we seek an M -term expansion where M is very small, e.g., $M = 1 \dots 4$.




Figure 5: Cubature training convergence plots reveal that $\sim 10\%$ error is often obtained after $n = 4r$ cubature features, which is higher than the volumetric modal models in [An et al. 2008].




Figure 6: Illustration of locking-related “pitch glide:” Spectrograms are shown for virtual cymbal sounds resulting from different impulse magnitudes. As the impulse magnitude grows, one can see an increasingly noticeable frequency drop at the beginning of the spectrograms.

We produce least-squares estimates of Ψ_j for a given M as follows. Given a coordinate system defined at the center of the object (we use the center of mass), we define a fixed set of angular directions, $\Theta_l = (\theta_l, \phi_l)$, and a set of radii R_1, \dots, R_K , $K > M$ (see Figure 7). Using a Helmholtz boundary integral solver, we rasterize a reference $p(x)$ solution at all (R_i, Θ_l) locations; we use the *FastBEM Acoustics* implementation (www.fastbem.com) of the fast multipole boundary element method [Liu 2009; Shen and Liu 2007]. In our examples, we use $K = 6$ shells to estimate between 1 to 4 Ψ_j maps, rasterizing $\theta \in [0, \pi]$ into T values, and $\phi \in [0, 2\pi]$ into $2T$ values where T is a function of wave number. Given an object radius of \bar{R} , we select geometrically larger shell radii, $R_i = 2.5\bar{R}\gamma^{i-1}$ (we used $\gamma = 1.54$). Using this pressure data, each angular direction, Θ_l , has the following complex-valued least-squares problem,




Figure 7: Geometry of FFAT Map Estimation

$$\sum_{j=1}^M \frac{h_0(kR_i)}{(kR_i)^{j-1}} \Psi_j(\Theta_l) = p(R_i, \Theta_l) \quad \Leftrightarrow \quad \sum_{j=1}^M A_{ij} \Psi_j = p_i.$$

or $\mathbf{A}\Psi = \mathbf{P}$. Since each angular direction has the same \mathbf{A} matrix, we solve all directions simultaneously. Different rows of \mathbf{P} have different radii, and can therefore differ greatly in magnitude. Therefore we use a weighted least squares approach that normalizes by the RMS magnitude of each \mathbf{P} row. Specifically, let \mathbf{W} be a diagonal weighting matrix with $\mathbf{W}_{ii} = 1/\|\mathbf{P}_i\|$, then we solve the weighted least-squares problem, $\mathbf{WA}\Psi = \mathbf{WP}$ using TSVD to obtain $\Psi = (\mathbf{WA})^\dagger (\mathbf{WP})$. Each row in Ψ can be extracted and stored as a floating-point texture for subsequent evaluation.

Adaptive FFAT map resolution: In our implementation, all FFAT maps are computed via a uniform sampling of angle $(\theta-\phi)$ space. That is, given some number of θ divisions T , we store FFAT map terms and precompute samples on each spherical shell at $2T^2 + 2$ angular positions. In general, the acoustic transfer function’s angular complexity increases with modal frequency (see figures 12 and 13) suggesting that some modes require greater angular FFAT map resolution than others. We propose to exploit simple linear dependence between angular resolution T and wave number k ; $T = \lceil ck + d \rceil$. In our experiments, we found that a base resolution of $d = 15$ and a slope of $c = 2.25m$ was sufficient to adequately capture the angular complexity of acoustic transfer functions.

Discussion of frequency localization: We have assumed that the nonlinear modal vibrations are approximately time-harmonic with frequencies similar to the linear modal vibrations. While this is true for weak forcing, it is far less true for hard forcing (see Figure 8). Nevertheless, we believe that linear frequency-domain acoustic transfer provides a cheap yet plausible sound model. The alternative evaluation of 3D time-domain wave radiation is significantly more expensive (and less appealing) than the $O(r)$ /object runtime evaluation of FFAT map transfer.

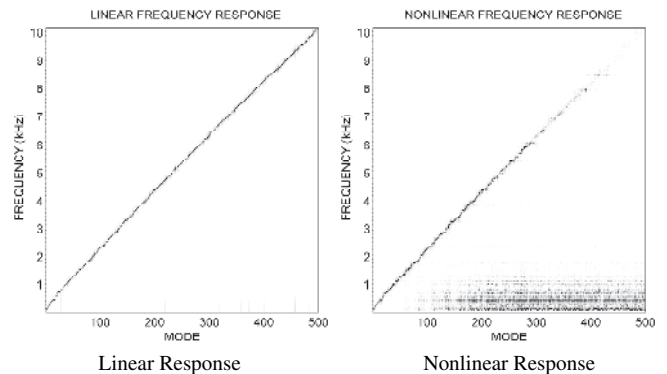


Figure 8: Frequency spectrum of $q(t)$ for a hard cymbal “crash” (first example in video): The linear model (Left) illustrates that each modal coordinate $q_i(t)$ is strongly localized in the frequency domain about its modal frequency, ω_i , whereas the nonlinear modal model (Right) exhibits a more complex response that is frequency-localized for lower-frequency modes, but higher modes become increasingly coupled to low-frequency modes—a possible sign of modal locking after this strong forcing.

5 Results

We now describe numerical and sound experiments for several models and multibody collision scenarios. Please see our accompanying video for all animation and sound rendering results. Model statistics are provided in Table 1. Representative timings are given in Table 2.

Model	L (m)	tri	vtx	N	modes	freq (kHz)	material	ν	Y (GPa)	h (mm)	α	β (10^{-9})	n_{cuba}	Error _{cuba}	kL	Δt (s)
Trash Can	0.75	77536	38833	116499	200	0.071 – 4.43	Steel	0.30	190	2	0.5	75	800	10.3%	0.98 – 61	1/44100
Trash Lid	0.55	34312	17286	51858	200	0.112 – 6.79	Steel	0.30	190	2	0.5	75	800	11.5%	1.1 – 68	1/44100
Water Bottle	0.46	28658	14418	43254	300	0.116 – 3.59	Polycarb.	0.37	2.4	2.25	0.5	400	900	10.7%	0.98 – 48	1/44100
Recycling Bin	0.61	109568	54945	164835	300	0.062 – 2.21	Polycarb.	0.37	2.4	5	4.0	300	1200	15.7%	0.70 – 30	1/44100
Cymbal	0.50	61952	31104	93312	500	0.061 – 9.94	Bronze	0.33	124	0.7	1.0	6.25	1500	10.7%	0.57 – 92	1/88200

Table 1: Model Statistics

We provide sound comparisons between four cases:

1. **Nonlinear/Transfer (“Harmonic Shells”):** Nonlinear modal vibrations with acoustic transfer (FFAT maps, or fast multipole method evaluation).
2. **Linear/Transfer:** Linear modal vibrations with acoustic transfer. This case typically sounds plausible, but misses characteristic nonlinear “crash” and “rumble” effects, and amplitude-based timbre variations.
3. **Linear/Monopole:** Linear modal vibrations with the low-frequency, far-field monopole radiation model (equation (15) in [James et al. 2006]; also used in [Bonneel et al. 2008]). Lacking both nonlinear vibrations and acoustic transfer, this case usually sounds quite unrealistic.
4. **Nonlinear/Monopole:** For comparison, we also render nonlinear modal vibrations with the far-field monopole model. While the vibrations are nonlinear, without acoustic transfer the sound quality is poor.

Implementation Details: We precompute dominant linear vibration modes using Matlab’s generalized eigenvalue solver (using ARPACK’s shift-and-invert spectral transformation). To improve our graphics model’s mesh quality for vibration and radiation analysis, we remesh the shells (using GNU GTS). We exploit mode-level parallelism to precompute acoustic transfer models (fast multipole solves, and FFAT map estimation) on a 16-node cluster (8-core, 2.66GHz, 8GB, Xeon X5355 processor nodes). All animations are performed using rigid body dynamics with vibration models defined in the appropriate rigid body frame [Shabana 2005]. Collisions are detected using a rigid sphere-tree bounding volume hierarchy, and resolved using a linear Kelvin-Voigt penalty contact model. Rigid body dynamics are time-stepped using symplectic Euler at rates sufficient to resolve penalty contact forces; modal vibrations are time-stepped (explicit subspace Newmark) at audio rates (e.g., 44100Hz); and acoustic transfer is evaluated at 1000Hz along the two-ear listening trajectory. A simple contact damping model is used to damp vibrations of objects in ground contact. In our simulation pipeline, we first simulate rigid-body motion and dump subspace force impulses to disk, then in a second pass we compute modal vibrations and synthesize sound at the listening position. Each object’s final sound is computed as a linear superposition of modal contributions with a simple HRTF model, $H(\omega, \mathbf{x})$ [Brown and Duda 1998]; $sound(\mathbf{x}, t) = \sum_{k=1}^r |H(\omega_k, \mathbf{x})| |p_k(\mathbf{x})| q_k(t)$. Graphics frames were rendered using Pixar’s RenderMan. All floating point computations were performed using double precision.

EXAMPLE (Cymbal): We modeled a large ride cymbal (50cm diameter, bronze), which is known to be a challenging example for modal vibrations [Chaigne et al. 2005]. The linear modal model of a ride cymbal produces a very clean tone that sounds more like a smaller crash cymbal, and it is unable to produce the proverbial “crash” sound as well as the nonlinear model. Both models sound very plausible with acoustic transfer.

EXAMPLE (Trash Can with Lid): The nonlinear modal model produces a dramatic improvement in the sound of the trash can (and lid) relative to the linear modal model; the nonlinear model produces a characteristic “crashing” sound, whereas the linear model makes a

“ding” sound. The trash can also has very interesting acoustic transfer functions; its FFAT maps reveal intricate structure, partly due to the trash can’s side-reinforcing ribs, and very loud values near its opening (see Figure 13). The spilling (spinning & rolling) of the trash can lid has a more distinctive sound than the linear model.

EXAMPLE (Water Bottle): We modeled a round 5-gallon water bottle out of polycarbonate plastic, and tuned damping parameters by comparing to informal experiments. The nonlinear sound model captures a characteristic drum-like fluttering after impact better than the linear sound model. The complex structure of the FFAT maps are shown in Figure 12. A comparison to a real water bottle impact experiment is provided in the accompanying video, and produces a qualitatively similar sound.

EXAMPLE (Plastic Recycling Bin): While less dramatic than other examples, the nonlinear model captures a familiar “wobbling” sound which is missing from the linear model.

MULTIBODY EXAMPLES: We simulated several multibody collision scenarios to demonstrate the feasibility of “Harmonic Shells” for computer animation (see Figure 9, and video results).

Stability: Unlike linear modal models which can be stably integrated with IIR filters, our nonlinear subspace vibration model can suffer time-stepping instabilities. Fortunately, subspace integrators are typically more stable than their unreduced counterparts [Krystl et al. 2001]. We observed that our explicit subspace Newmark integrator was stable at audio rates (44.1 kHz) for all examples, except the cymbal which we integrated at 88.2 kHz. In contrast, traditional explicit Newmark integration required an exceedingly small timestep to be stable, e.g., the water bottle required 11.025 MHz rates (or $250 \times$ the 44.1 kHz rate).

COMPARISON (reduced vs. unreduced simulation): For validation, we compared water bottle impact sounds from our reduced-order model ($\Delta t = 1/44100s$) to those of nonlinear vibrations simulated in a full, unreduced setting via an explicit Newmark integrator ($\Delta t = 1/11025000s$ for stability)—implicit Newmark (with full Newton solves) was less competitive in our experiments. Although the unreduced model produced richer tones at higher amplitude impacts, both sounds were comparable and more interesting than pure linear vibrations. Unfortunately, while the reduced-order model took roughly 17.1 minutes to compute 1 second of sound ($1026 \times$ slower than real time), the unreduced approach took 89.8 hours per second of sound ($323,000 \times$ slower than real time).

Details of unreduced computation: Given the unreduced displacement \mathbf{u} of the object, modal amplitudes are obtained via projection with the basis $\mathbf{q} = \mathbf{U}^T \mathbf{u}$ so that any of the previously discussed radiation methods may be applied. Given that the simulated model is unconstrained, we take steps to avoid rigid body motions in the unreduced simulation as these will result in errors in the modal projection. A $N \times 6$ “rigid basis” matrix \mathbf{U}_R is constructed out of the rigid modes (those corresponding to eigenvalue 0) computed in equation 7. This is used to produce a 6×6 “rigid mass” matrix $\mathbf{M}_R = \mathbf{U}_R^T \mathbf{M} \mathbf{U}_R$. Given the external forces acting at each time step, the component of acceleration resulting in rigid motion is identified as $\mathbf{U}_R \mathbf{M}_R^{-1} \mathbf{U}_R^T \mathbf{f}_{ext}$ and subtracted from the total acceleration vector. This allows the mesh to vibrate freely in place

Model	Modes r	Modal Analysis	Cubature Precomp.	Timestep Cost	Simulation Cost (per second of audio)	FFAT Precomp. (average time/mode)	FFAT Eval (all modes, $M = 4$)	FFAT Storage (floats, $M = 1$)
Trash Can	200	569 s	2.49 hr	16.1 ms	714 s	109.2 min	0.151 ms	56 MB
Trash Lid	200	170 s	1.87 hr	14.6 ms	642 s	85.5 min	0.151 ms	113 MB
Water Bottle	300	314 s	4.31 hr	23.6 ms	1026 s	25.6 min	0.227 ms	54 MB
Recycling Bin	300	2332 s	9.65 hr	27.8 ms	1224 s	48.0 min	0.227 ms	25 MB
Cymbal	500	1155 s	3.88 hr	44.3 ms	3900 s	318 min	0.378 ms	512 MB

Table 2: Representative Timings: All timings are for a single 2.66GHz Xeon X5355 processor core, except “Cubature Precomp” which used 8 cores.



Figure 9: Multibody collision scenarios were simulated for (from left to right) a cymbal with metal balls, multiple cymbals, two trash cans, a trash can and lid, and polycarbonate water bottles—as well as the teaser image (Figure 1).

without undergoing rigid translation and rotation.

COMPARISON (different cubature errors): We simulated nonlinear modal models with different cubature errors to informally demonstrate their respective sound behavior. See the video for a comparison of the trash can simulated with cubature errors (and timestep costs) of 15.3% (10.5ms), 10.3% (16.1ms), 6.1% (27ms), and using brute-force subspace integration [Krysl et al. 2001] we simulated 0% (166.7ms). We find that even cubature schemes with large relative error provide a significant qualitative improvement in sound quality over the linear model. Our cubature schemes are chosen to provide a tradeoff between sound quality and evaluation speed.

COMPARISON (with/without energy limitation): To demonstrate pitch glide, we artificially increase the magnitude of forces acting on the water bottle by a factor of 5. See the video for a comparison of this scenario with and without the impulse limiter (§3).

COMPARISON (FFAT vs Fast-Multipole-Method Error): Please see Table 3 and Figure 10 for FFAT map accuracy demonstrations. Our video provides animated comparisons: FastBEM required $\sim 17h15m$ to synthesize all-mode transfer for the trash-can animation, and $\sim 13h13m$ for the water bottle animation; for both animations, FFAT Map evaluation required under a second.

COMPARISON: Different FFAT map expansion orders, M , are shown in Figure 11 for the highest frequency mode of the trash can. We use at most 4-maps/mode ($M = 4$) in all of our rendered examples. Convergence is obtained for increasing M values; an error analysis is provided in Table 3. Please see the video for comparisons; in practice, similar sounds are obtained for all M values, suggesting that even one texture map ($M = 1$) is sufficient.

6 Conclusion

We have presented a practical method for generating plausible impact sounds for thin shells. By leveraging reduced-order modeling, we can produce nonlinear modal models that enable simulation of hundreds of vibration modes with fully coupled nonlinear modal dynamics. We proposed a method to limit impact magnitudes and overcome pitch-glide artifacts. Compared to linear modal sound models, our objects produce more characteristic “crashing” and “rumbling” sounds.

We also proposed *FFAT maps*, a fast texture-based approximation of each mode’s far-field acoustic transfer function that captures complex spatial structure, and are more generally applicable than to just thin shells. They exploit the observation that transfer functions exhibit complex angular structure, but possess strong radial

coherence. By using an accurate transfer solution, e.g., from a fast multipole solver, we produced far-field acoustic transfer map approximations that enable fast run-time evaluation of the transfer function to very low tolerances (e.g., 1%). In practice, we found that even a single FFAT map texture per mode ($M = 1$) produced excellent results.

Limitations and Future Work: There are numerous ways to improve this result in future work. Simulating the full range of audible “all-frequency” sound poses numerous challenges, not only for pre-computing thousands of vibration modes but also for coupling them together. The $O(r^2)$ complexity of the nonlinear modal model reflects the intrinsic complexity of simulating r coupled modes, but a near-linear-time subspace force algorithm would be a big breakthrough for all-frequency nonlinear sound synthesis, especially since the modal amplitudes are needed for transfer-based rendering. In all cases, we would have liked to have used more vibration modes to produce “all-frequency” sound renderings; large models can also require many modes. The shell-based cubature schemes exhibit slower convergence rates than volumetric models [An et al. 2008], and higher accuracy and more scalable methods are required for faster and/or more complex models. We use a simplified rigid-body contact model, but collision processing should account for object vibrations (ideally in a reduced-order manner [James and Pai 2004]) to properly capture chattering.

Beyond thin shells, how to devise efficient methods for evaluating all-frequency nonlinear vibration-based sound is an open problem. Modal locking reflects the limitations of the linear modal shape model, and we need more expressive shape models to handle difficult nonlinear and noise-like phenomena; a fully developed (chaotic) cymbal crash is currently beyond the capability of our reduced-order vibration model, although the cymbal’s acoustic transfer model appears plausible. Preliminary experiments with thin-shell models using nonlinear shape functions based on modal warping did not provide a significant improvement [Choi et al. 2007]. More generally, we need methods to evaluate accurate sound for large-deformation animations. Buckling is also a challenging nonlinear phenomenon which can produce significant sound radiation, e.g., crumpling paper.

Our FFAT maps can provide high-accuracy acoustic transfer, but would benefit from precomputation techniques for improved spatial sampling, estimation, and texture compression—FFAT maps can be “big.” By construction, our current FFAT models are less accurate for near-field listening positions, which may be a problem for some applications. Far-field listening positions should also include time delay effects, which can be complicated for dynamic objects.

Model	Mode	Freq (Hz)	kL	θ Resolution (T)	Average Relative Error (%)				Median Relative Error (%)			
					M=1	M=2	M=3	M=4	M=1	M=2	M=3	M=4
Trash Can	0	71	0.90	18	1.47	6.73	0.95	1.04	0.54	7.2	0.91	1.03
	50	1880	25.82	93	19.51	7.11	1.28	0.69	9.14	3.40	0.67	0.44
	100	2823	38.79	132	44.44	28.36	20.00	5.62	31.26	12.88	9.27	2.66
	150	3698	50.81	168	67.21	35.20	8.27	1.76	29.18	12.22	4.37	0.63
	199	4433	60.80	198	53.37	42.64	17.05	3.14	28.64	16.94	9.56	1.27
Trash Lid	0	112	1.13	20	3.12	8.47	1.36	0.78	2.52	9.04	.14	0.78
	50	2296	23.13	110	36.12	21.91	6.41	1.74	34.89	21.18	5.64	1.64
	100	3997	40.27	180	51.54	24.36	18.07	2.36	51.76	21.22	16.10	1.62
	150	5736	57.79	252	14.09	3.96	1.32	0.47	8.53	2.96	0.75	0.32
	199	6791	68.22	295	60.40	23.39	31.04	2.94	62.28	12.82	27.17	1.39
Water Bottle	0	116	0.98	20	251.35	16.00	1.90	0.032	264.86	11.37	1.67	0.018
	75	1321	11.13	70	26.86	11.80	2.98	0.54	13.07	3.51	1.49	0.35
	150	2197	18.51	106	11.14	6.69	0.86	0.16	8.30	5.77	0.30	0.10
	225	2906	24.49	135	14.03	8.41	1.33	0.31	8.73	4.99	0.65	0.25
	299	3593	30.27	164	32.26	15.84	3.33	0.45	15.60	8.15	1.34	0.26
Recycling Bin	0	62	0.70	18	10.85	3.53	0.40	0.43	6.47	3.46	0.32	0.35
	75	820	9.17	49	13.66	3.79	1.05	0.65	7.45	1.73	0.64	0.34
	150	1329	14.85	70	8.54	3.72	0.87	0.65	6.52	2.30	0.67	0.49
	225	1791	20.01	89	7.36	2.89	0.70	0.35	5.03	1.37	0.46	0.24
	299	2209	24.68	107	12.67	5.03	2.07	0.54	9.13	3.05	12.10	0.36

Table 3: Comparison of FFAT map to fast multipole solver $|p(\mathbf{x})|$ pressure values illustrate that very low relative errors ($\approx 1\%$) can be achieved using a 4-term FFAT map expansion. Errors were computed for representative raster images (see Figure 11 for a specific example).




Figure 10: Comparison of sound pressures, $|p(\mathbf{x})|$, between BEM (left) and FFAT map approximations (right) for various “trash can” modes. In all cases, the 4-term FFAT map ($< 6\%$ average error) results both look and sound essentially same. Error values are given in Table 3.




Figure 11: Comparison of FFAT maps of different order for the trash can (mode 199). (Far Left) Exact transfer field evaluated using the fast multipole method. The remaining figures ($M = 1 \dots 4$) show the result of optimizing the FFAT models with different M values (# maps = M), and their average pointwise relative errors for the $|p(\mathbf{x})|$ rasters. The single-term approximation would provide enough accuracy for real-time applications using linear modal models.

Nonlinear vibrations can introduce other significant frequency contributions into each mode’s vibration especially for high-frequency modes (see Figure 8), so more complex nonlinear radiation models are needed at high frequencies—multi-frequency mode radiation models may provide more realistic sound. We have only considered single-object sound transport to leverage precomputation, and it still remains to include multi-object and environment scattering for correct auralization.

Perceptually based rendering methods are desperately needed to strike a favorable balance between model accuracy and simplicity. For example, our thin-shell transfer maps exhibit increasingly complex structure at high frequencies and are extremely difficult to compute without sophisticated fast multipole solvers. Arguably it is hard to hear all of this structure in sound renderings, and we therefore desire perceptually based precomputation and rendering

techniques which display only what is necessary, and avoid computing what is not.

Acknowledgements: We would like to thank the anonymous reviewers for helpful feedback. This work was supported in part by the National Science Foundation (CAREER-0430528, HCC-0905506), the Alfred P. Sloan Foundation, and generous donations by Pixar, Intel, Autodesk, and also Advanced CAE Research, LLC for the *FastBEM Acoustics* solver. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation or others.

A Background on Acoustic Transfer

Given the dynamics of a single mode, $q(t)$, we can estimate its sound pressure contribution at the listener’s location, \mathbf{x} , by approximating the modal vibrations as time-harmonic with fixed angular frequency, ω , so that $q(t) \propto e^{+i\omega t}$. In that case, the complex-valued pressure field due to this single vibration mode is given by $p(\mathbf{x})e^{+i\omega t}$. The spatial part of the modal pressure field, $p(\mathbf{x})$, is referred to as the *acoustic transfer function* for that mode [James et al. 2006] and it satisfies the frequency-domain Helmholtz wave equation, $(\nabla^2 + k^2)p(\mathbf{x}) = 0$, $\mathbf{x} \in \Omega$, (18)

in the object’s exterior domain, Ω ; here k is the *wave number*, $k = \frac{\omega}{c} = \frac{2\pi}{\lambda}$, where c is the speed of sound in air ($c = 343\text{m/s}$ at STP), and λ is the sound wavelength. To obtain a solution to (18), Neumann boundary conditions are imposed on the vibrating object’s surface, $\frac{\partial p}{\partial n}(\mathbf{x}) = -i\omega\rho v_n(\mathbf{x})$ on $S = \partial\Omega$, where the normal surface velocity is $v_n(\mathbf{x}) = i\omega(\mathbf{n} \cdot \mathbf{u}(\mathbf{x}))$ and $\mathbf{u}(\mathbf{x})$ is the modal displacement at $\mathbf{x} \in S$; also for radiation problems, $p(\mathbf{x})$, must satisfy a Sommerfeld radiation condition at infinity.




Figure 12: FFAT Maps (Water Bottle)




Figure 13: FFAT Maps (Trash Can)

References

- ADRIEN, J.-M. 1991. The missing link: Modal synthesis. In *Representations of musical signals*. MIT Press, Cambridge, MA, USA, 269–298.
- AN, S. S., KIM, T., AND JAMES, D. L. 2008. Optimizing cubature for efficient integration of subspace deformations. *ACM Trans. Graph.* 27, 5, 1–10.
- BARAFF, D., AND WITKIN, A. P. 1998. Large steps in cloth simulation. In *Proceedings of SIGGRAPH 98*, Computer Graphics Proceedings, Annual Conference Series, 43–54.
- BATHE, K.-J. 1996. *Finite Element Procedures*, second ed. Prentice Hall.
- BERGOU, M., WARDETZKY, M., HARMON, D., ZORIN, D., AND GRINSPUN, E. 2006. A quadratic bending model for inextensible surfaces. In *Eurographics Symposium on Geometry Processing*, 227–230.
- BILBAO, S. 2008. A family of conservative finite difference schemes for the dynamical von Karman plate equations. *Numerical Methods for Partial Differential Equations* 24, 1.
- BONNEEL, N., DRETTAKIS, G., TSINGOS, N., VIAUD-DELMON, I., AND JAMES, D. 2008. Fast Modal Sounds with Scalable Frequency-Domain Synthesis. *ACM Transactions on Graphics* 27, 3 (Aug.), 24:1–24:9.
- BRIDSON, R., MARINO, S., AND FEDKIW, R. 2003. Simulation of clothing with folds and wrinkles. In *2003 ACM SIGGRAPH / Eurographics Symposium on Computer Animation*, 28–36.
- BROWN, C. P., AND DUDA, R. O. 1998. A Structural Model for Binaural Sound Synthesis. *IEEE Trans. on Speech and Audio Processing* 6, 5.
- CHAIGNE, A., TOUZÉ, C., AND THOMAS, O. 2005. Nonlinear vibrations and chaos in gongs and cymbals. *Acoustical Science and Technology* 26, 5, 403–409.
- CHAPELLE, D., AND BATHE, K. 2003. *The finite element analysis of shells*. Springer.
- CHOI, M., YONG WOO, S., AND KO, H. 2007. Real-Time Simulation of Thin Shells. In *Computer Graphics Forum*, vol. 26, Blackwell Publishing Ltd, 349–354.
- CIRAK, F., AND ORTIZ, M. 2001. Fully C^1 -conforming subdivision elements for finite deformation thin-shell analysis. *Internat. J. Numer. Methods Engrg.* 51, 813–833.
- CIRAK, F., ORTIZ, M., AND SCHRODER, P. 2000. Subdivision surfaces: A new paradigm for thin-shell finite-element analysis. *Internat. J. Numer. Methods Engrg.* 47, 2039–2072.

- 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.
- DESMET, W. 2002. Mid-frequency vibro-acoustic modelling: challenges and potential solutions. In *Proceedings of ISMA 2002*, vol. II.
- FLETCHER, N. 1999. The nonlinear physics of musical instruments. *Reports on Progress in Physics* 62, 5, 723–764.
- FRENDI, A., MAESTRELLO, L., AND BAYLISS, A. 1994. Coupling between plate vibration and acoustic radiation. *Journal of Sound and Vibration* 177, 2, 207–226.
- GARG, A., GRINSPUN, E., WARDETZKY, M., AND ZORIN, D. 2007. Cubic shells. In *ACM SIGGRAPH Symposium on Computer Animation*, 91–98.
- GINGOLD, Y., SECORD, A., HAN, J. Y., GRINSPUN, E., AND ZORIN, D. 2004. A Discrete Model for Inelastic Deformation of Thin Shells. Tech. rep., Courant Institute of Mathematical Sciences, New York University, Aug.
- GREEN, S., TURKIYYAH, G., AND STORTI, D. 2002. Subdivision-based multilevel methods for large scale engineering of thin shells. In *Proceedings of ACM Solid Modeling*, 265–272.
- GRINSPUN, E., HIRANI, A., DESBRUN, M., AND SCHRODER, P. 2003. Discrete shells. In *ACM SIGGRAPH Symposium on Computer Animation*, 62–67.
- GUMEROV, N., AND DURAISWAMI, R. 2005. *Fast Multipole Methods for the Helmholtz Equation in Three Dimensions*. Elsevier Science.
- HAMMING, R. W. 1983. *Digital Filters*. Prentice-Hall, Englewood Cliffs, NJ.
- 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.
- JAMES, D. L., AND PAI, D. K. 2004. BD-Tree: Output-sensitive collision detection for reduced deformable models. *ACM Transactions on Graphics* 23, 3 (Aug.), 393–398.
- JAMES, D. L., BARBIĆ, J., AND PAI, D. K. 2006. Precomputed Acoustic Transfer: Output-sensitive, accurate sound generation for geometrically complex vibration sources. *ACM Transactions on Graphics* 25, 3 (July), 987–995.
- KRYSL, P., LALL, S., AND MARSDEN, J. E. 2001. Dimensional model reduction in non-linear finite element dynamics of solids and structures. *International Journal for Numerical Methods in Engineering* 51, 479–504.
- LIU, Y. J. 2009. *Fast Multipole Boundary Element Method: Theory and Applications in Engineering*. Cambridge University Press, Cambridge.
- MALATKAR, P. 2003. *Nonlinear vibrations of cantilever beams and plates*. PhD thesis, Virginia Polytechnic Institute and State University, Blacksburg, VA, USA.
- MOUSSAOUI, F., AND BENAMAR, R. 2002. Non-linear vibrations of shell-type structures: a review with bibliography. *Journal of Sound and Vibration* 255, 1, 161–184.
- NAYFEH, A., AND MOOK, D. 1979. *Nonlinear oscillations*. Wiley-Interscience.
- NAYFEH, A., AND NAYFEH, S. 1995. Nonlinear normal modes of a continuous system with quadratic nonlinearities. *Journal of Vibration and Acoustics* 117, 199.
- O'BRIEN, J. F., COOK, P. R., AND ESSL, G. 2001. Synthesizing sounds from physically based motion. In *Proceedings of ACM SIGGRAPH 2001*, Computer Graphics Proceedings, Annual Conference Series, 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.
- PENTLAND, A., AND WILLIAMS, J. 1989. Good vibrations: Modal dynamics for graphics and animation. In *Computer Graphics (Proceedings of SIGGRAPH 89)*, 215–222.
- PIERCE, J., AND VAN DUYN, S. 1997. A passive nonlinear digital filter design which facilitates physics-based sound synthesis of highly nonlinear musical instruments. *The Journal of the Acoustical Society of America* 101, 1120.
- 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*. Springer-Verlag, New York, NY.
- SHABANA, A. 2005. *Dynamics of Multibody Systems*, 3rd ed. Cambridge.
- SHEN, L., AND LIU, Y. J. 2007. An adaptive fast multipole boundary element method for three-dimensional acoustic wave problems based on the Burton-Miller formulation. *Computational Mechanics* 40, 3, 461–472.
- TERZOPOULOS, D., PLATT, J., BARR, A., AND FLEISCHER, K. 1987. Elastically Deformable Models. In *Computer Graphics (Proceedings of SIGGRAPH 87)*, 205–214.
- TOUZÉ, C., THOMAS, O., AND CHAIGNE, A. 2004. Hardening/softening behaviour in non-linear oscillations of structural systems using non-linear normal modes. *Journal of Sound and Vibration* 273, 1-2, 77–101.
- 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*, Computer Graphics Proceedings, Annual Conference Series, 537–544.
- VORLANDER, M. 2007. *Auralization: Fundamentals of Acoustics, Modelling, Simulation, Algorithms and Acoustic Virtual Reality*. Springer Verlag.
- WU, S. 2008. Methods for reconstructing acoustic quantities based on acoustic pressure measurements. *The Journal of the Acoustical Society of America* 124, 5, 2680.