

CHAPTER 3: SOUND SYNTHESIS FOR FLUID-STRUCTURE COUPLING¹

This chapter describes a fast and practical method for simulating the sound of non-empty objects containing fluids. Our sound synthesis algorithm for fluid-structure coupling enhances the rigid-body sound with an added mass operator to account for the amount and type of liquid contained within the object. Although the technique assumes the liquids are inviscid, remain steady, and are not mixed and has some limitations (such as granularity of the solid discretization since modifications to the mass matrix occur at mesh nodes), this method is designed and demonstrated for use in interactive 3D systems, where live sound synthesis is important.

3.1 Introduction

Sound is integral to the level of immersion and sense of presence in virtual and imaginative environments (47). Research has demonstrated that even with eyes closed, we can have a mental imagery response similar to a visual perception, as if we are actually viewing the objects (110). However, in the case of virtual reality, a distraction from any of the senses can cause a ‘break in presence’ (199). Therefore, game engines and VR systems are incorporating physically-based graphics and sound simulation algorithms for interactive and realistic effects to help users remain immersed in the experience.

Given an object and its material properties, we can simulate, in real-time, the resulting sound that it would make when struck (230; 161; 190). To achieve this real-time sound synthesis, we may pre-compute the frequencies and dampings of each element node by performing modal analysis using a generalized eigendecomposition. Then, sound synthesis is based on the results of these pre-processed steps and a given impulse force.

Since both sound and graphics are physics-based, the rendering pipeline can be naturally extended to generate sound based on the same physics. However, 3D objects for graphical display can be filled with

¹ Most of this chapter previously appeared as an article in *The Visual Computer*. The original citation is as follows: Wilson, J., Sterling, A., Rewkowski, N., and Lin, M. C. (2017). Glass half full: sound synthesis for fluid–structure coupling using added mass operator. *The Visual Computer*, 33(6):1039–1048




Figure 3.1: (Top) Simulated liquid xylophone with varying levels of liquid and/or density based on our sound synthesis algorithm for fluid-structure coupling that enhances the rigid-body sound modeled by accounting for both the amount and type of liquid contained within the object. (Bottom) Spectrograms for each goblet, illustrating how the fundamental frequency decreases, resulting in a lower pitch as the amount of liquid or density increases

liquids (and/or other materials) in the scene and we can hear sounds coming from sources that we cannot see. It is, therefore, important to simulate sounds from non-empty objects that may be hollow by design but contain liquid in the virtual world.

In this paper, we present a new and practical method that satisfies both audio and graphics requirements and enhances the existing physically-based sound synthesis model to include non-empty objects. This feature is important, because we interact with filled containers and can distinguish between objects that are either full or empty (193). Therefore, the difference in perceived sounds between empty and non-empty containers should be modeled to increase the realism of the virtual environments.

3.1.1 Contributions

This chapter’s main contributions are:

1. Transforming the problem into a single fluid-structure system using the *added mass operator*;
2. Enhancing the rigid-body sound synthesis pipeline with pre-processing steps for objects containing a liquid;



Figure 3.2: A wooden pot, metallic teapot, and porcelain bowl with fine and coarse subdivision surfaces are a few of the objects simulated with varying volumes of liquid (e.g. water). The far right image is the same porcelain bowl but simulated with a more dense liquid (e.g. milk)

3. Demonstrating the proposed method in interactive 3D VR applications.

3.2 Related Work

Sound Synthesis: Research has modeled rigid body sound using modal analysis to decouple the sound dynamic equations into n independent, damped vibration equations (3; 231). These sound synthesis techniques create sound based on vibration analysis of the object resulting in varying frequency vibration modes. Modal analysis relies on expensive pre-processing to achieve interactive runtime performance. In addition to rigid-body sounds, there are a few other major physically based categories such as fracture (270), fire (34), and liquids (149), to name a few. Please see a recent survey (96) for more details.

Parameter Acquisition: To perform sound synthesis, object specific parameters are required, for example material specific damping coefficients that have traditionally been tuned manually. To automatically determine these material properties, a method to extract parameters from recorded audio was introduced (190). Alternative and more general damping models have also been introduced (219).

Acoustic Transfer: However, what we hear is not the modal amplitudes of the vibrating solid but the sound pressure waves radiating from the vibrating surface into the surrounding air. Acoustic transfer techniques couple synthesis and propagation together to generate sound. There are geometric acoustics (GA) techniques (62) and numerical acoustics (NA) methods that solve the wave equation using adaptive Finite Element Method (FEM) (226), Boundary Element Method (BEM) (27), Finite Difference Time Domain (FDTD) (197), spectral methods (24), and Adaptive Rectangular Decomposition (ARD) (179).

Coupled Synthesis-Propagation: In addition to focusing on outgoing waves from the vibrating surface into the air, research also been conducted to evaluate cavity tones for virtual instruments (189) and aerodynamic sound of a swinging sword (54). To simulate the sound propagation into the full 3D environment, various methods have been developed using geometric sound propagation based on ray tracing techniques (194), wave-based algorithms (143), and two types of multipole expansion for radiating sound fields-multipole expansion based on equivalent source methods (97) and single point multipole expansion (194). For single-point multipole expansion (SPME), a single multipole source is placed inside the object while multi-point multipole expansion places a large number at different points inside the object. These representations depict outgoing pressure fields. Source clustering has also been proposed to reduce computation since complexity varies with the number of sound sources (227).

Fluid-Structure Mechanics: To allow for interactions between a solid and a liquid, Müller et al. (150; 151) simulated the interaction of fluids with deformable solids, which estimates the forces between virtual boundary particles of the solid surface and fluid particles. Computational approaches are discussed by Bazilevs et al. (18), where the variational structural mechanics equation in matrix form may be written as:

$$M\ddot{Y} + D\dot{Y} + KY = F \quad (3.1)$$

$$M_{AB} = \int_{\Omega} N_A N_B d\Omega \quad (3.2)$$

where $M = M_{AB}$. For linear elastodynamics, a density term is added to the mass matrix (and damping matrix as well since it is a linear combination of the mass and stiffness matrices). Linear elastodynamics:

$$M = [M_{ij}^{AB}]$$

$$M_{ij}^{AB} = \int_{\Omega} N_A \rho N_B d\Omega \delta_{ij} \quad (3.3)$$

Here, Ω is the material domain of a structure along with the boundary Γ , ρ is the mass density of the structure, and N is associated to its unique mesh nodes.

Added Mass Operator: Viewing the coupled fluid and structure as a single system, added mass or virtual mass is a concept in fluid mechanics that incorporates the surrounding fluid of an accelerating or decelerating object (156). In the marine industry, added mass is often referred to as hydrodynamic added mass and can reach up to 1/3 of the total ship mass. Although less common in aeronautics because of small air density (except lighter-than-air balloons or blimps), the stability and convergence properties were

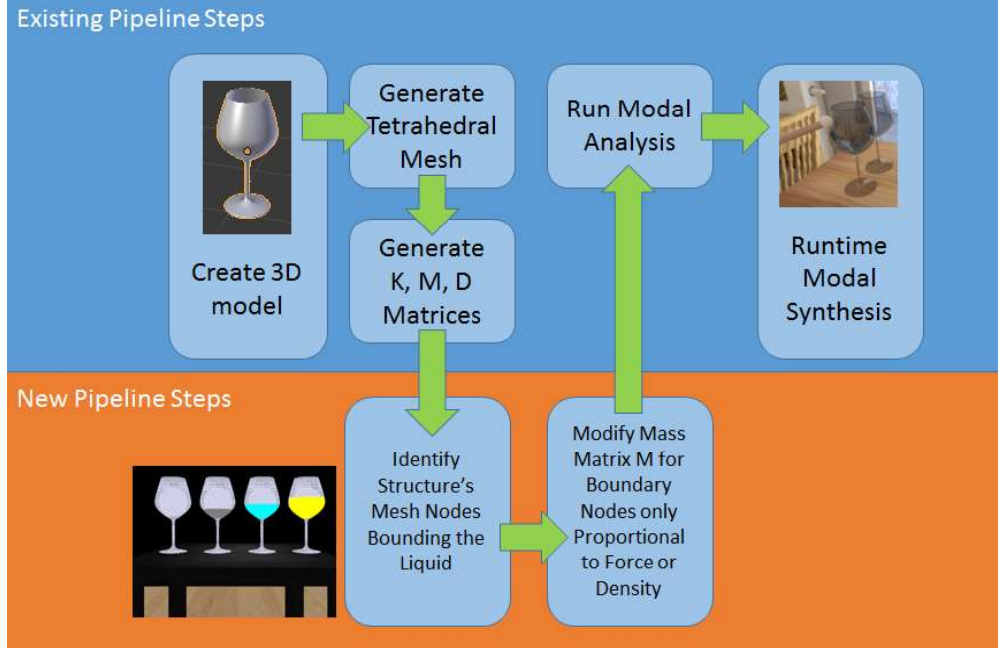


Figure 3.3: Overview of the enhanced rigid body sound pipeline to include new contributed steps to account for different liquid volumes and densities contained within the solid object. These new pre-processing steps come before modal analysis and therefore can be computed before simulation allowing for usage in real-time interactive systems

evaluated based on the ratio of added-mass to the actual structural mass (29). More information on the topic of added mass and fluid inertial forces can be found in a survey by the the Naval Civil Engineering Laboratory (28).

3.3 System Overview

In this paper, we introduce an enhanced sound synthesis model that accounts for the amount of liquid contained within an object by incorporating these added mass and fluid-structure interactions into the existing rigid body sound synthesis model. Our new approach for simulating sound of non-empty objects begins by modifying the mass matrix along the fluid-structure interface prior to performing modal analysis on the system. We present various approaches to calculate the modification required for a specific liquid density and volume.

3.3.1 Rigid-Body Sound Synthesis Pipeline

1. Generate the object's volumetric mesh and stiffness K , mass M , and damping D matrices

2. Apply an added mass operator to modify the mass matrix based on a fluid-structure interface
3. Run generalized eigen-decomposition
4. Construct decoupled modal vibration equations
5. Apply an impulse force
6. Numerically integrate individual modes

3.3.2 Modal Sound Analysis

Modal analysis is the standard linear model for dynamic deformation and physically based sound (207). When an object is struck, it vibrates and deforms. The surrounding air rapidly compresses (compression) and expands (rarefaction or decompression) as the object vibrates outward and inward respectively. As it oscillates periodically, pressure waves are created and air pressure amplitude changes up and down over time (like a sinusoidal wave). Although we may not see the vibrations or deformations, our ears hear the harmonic oscillation and periodic pattern of compression and rarefaction in air pressure as sound. We can simulate the vibration of the solid volume body object in an underdamped response to an impulse using the following equation:

$$M\ddot{u} + D\dot{u} + Ku = f \quad (3.4)$$

where M , D , and K are the mass, damping, and stiffness matrices, respectively; u is the displacement vector and f as the force vector. It is well-established to approximate small levels of damping with *Rayleigh Damping*, i.e. representing the damping matrix as a linear combination of the mass and stiffness matrices.

3.3.3 Modal Sound Synthesis

We can then simulate sound based on a contact position $p = (x, y, z)$, where the object is struck. The impulse direction is usually normal to the contact point but could be tangential to the object. To achieve real-time performance, pre-processing steps are performed for a given object and material.

After solving the generalized eigenvalue problem of Eqn. 3.4, the solution are *modes*, i.e. damped sinusoidal waves where each mode has the form

$$q_i = a_i e^{-d_i t} \sin(2\pi f_i t + \theta_i), \quad (3.5)$$

where f_i is the frequency of the mode, d_i is the damping coefficient, a_i is the excited amplitude, and θ_i is the initial phase. And,

$$\omega_i = 2\pi f_i = \frac{\sqrt{4km - d^2}}{2m} \quad (3.6)$$

We ignore θ_i and safely assume it to be zero since the object is initially at rest. It is also important to note that our approach requires the mass to be normalized ($m = 1$).

3.3.4 Added Mass Operator

In solid and structural mechanics and civil engineering where pressure forces are commonly applied to the structure, the additional force resulting from fluid acting on a structure when formulating the system equation of motion is known as *added mass*. In a physical sense, the added mass of an incompressible flow is the weight added to the system from surrounding fluid that the accelerating or decelerating structural vibration must move with the structure. This force is factored into our sound dynamics equation as:

$$M\ddot{u} + D\dot{u} + Ku = f - m_a\ddot{u} \quad (3.7)$$

where m_a is the added mass. Reordering terms, we arrive at:

$$(M + m_a)\ddot{u} + D\dot{u} + Ku = f \quad (3.8)$$

Note that if mass is also included in the damping matrix, the added mass would automatically be included there as well.

On short time scales, the effect of the fluid on the structure can be represented as an added mass. Ratio of this added mass to the structural mass is critical to convergence and should be less than or equal to 1; else, the system may become unstable (29). In reality, the fluid will be accelerated but for simplicity, it is

modeled as a volume moving with the object as a second-order tensor, relating the fluid acceleration vector to the resulting force vector on the body.

Added mass is analogous to the amount of work needed to change the kinetic energy T associated with the motion of the fluid:

$$T = \frac{\rho}{2} I U^2 = \frac{\rho}{2} \int_V (u_1^2 + u_2^2 + u_3^2) dV \quad (3.9)$$

where u is the fluid velocity in Cartesian coordinates, V is the volume of fluid, ρ is the liquid density, U is the structure's velocity, and I is a number proportional to the liquid's volume. We assume a non-moving domain and a rectilinear velocity, which could be generalized to other motions if required (?). When the solid body accelerates causing changes in U , kinetic energy T increases, supplying additional work by the body on the fluid. The rate of additional work is the rate of change of T with respect to time dT/dt and is considered added drag by the body. A viscous fluid also contributes a drag force; however, in our implementation, we assume inviscid liquids. Then, the added drag, F , can be represented as:

$$F = -\frac{1}{U} \frac{dT}{dt} = -\rho I \frac{dU}{dt} = -\rho \frac{V}{2} \ddot{u} \quad (3.10)$$

$$I = \int_V \frac{u_i}{U} \frac{u_i}{U} dV \quad (3.11)$$

where I is half the volume and F is in the form of $m_a \ddot{u}$ or $\rho_{fluid} \cdot \frac{V_{fluid}}{2} \ddot{u}$ if we approximate our liquids to be spherical, similar to the concept of spherical bounding volume hierarchy (69). Since cylindrical objects result in a full volume rather than half (28), future work is required to analyze the trade-offs between ease of implementation and accuracy in accounting for the geometric complexity of the liquid.

As a result, if we approximate our liquid as a sphere, added mass for each liquid object is approximated by:

$$m_a = \rho_{fluid} \cdot V_{fluid}/2 \quad (3.12)$$

For example, the added mass for a sphere (of radius r) is $\frac{2}{3}\pi r^3 \rho_{fluid}$ which is half the volume of a sphere times fluid density. Therefore, a spherical air bubble rising in water has a mass of $\frac{4}{3}\pi r^3 \rho_{air}$ and added mass of $\frac{2}{3}\pi r^3 \rho_{water}$.

Since the liquid must move with the same phase as the structure's motion, this may be referred to as a rigid *double body* (156). In general, the value of the added mass also depends on the direction of acceleration and is incorporated by projecting the area of the body in the direction of acceleration; for instance, tangential acceleration yields zero added mass. In our case, projection is not necessary since we assume a non-moving domain where the fluid motion is in the same direction as the structural vibration.

3.4 Sound Synthesis for Fluid-Structure Coupling

Added mass from the liquid is distributed to mesh nodes of the structural object along the fluid-structure boundary. In our system, we represent both the solid and liquids as tetrahedral meshes to detect boundary elements although particle-based methods (such as Lagrangian particles) are also possible. We also make the following assumptions:

- Solid is impermeable & the liquid is incompressible
- Fluid is at rest in hydrostatic equilibrium or the flow velocity at each point is constant over time
- Domain is non-moving, i.e. fluid motion coincides with structure motion
- Fluid is inviscid, having no or negligible viscosity

3.4.1 Identifying Nodes at Fluid-Structure Interface

A number of collision detection methods exist (68) and may be used to identify the necessary boundary nodes. Since the liquid is assumed to be at rest, boundary detection of symmetric objects may be simplified by selecting all nodes above and below the min and max nodes that intersect between the structure and fluid (see Fig. 3.4).

3.4.2 Modifying the Mass Matrix

We use the added mass equation for a sphere (see Eqn. 3.12) to calculate total added mass for each liquid based on the liquid density and volume. For example, if the liquid added is pure water, then the density is $1,000 \text{ kg/m}^3$; milk, $1,050 \text{ kg/m}^3$; olive oil, about 860 kg/m^3 ; and boiling water, approximately 958 kg/m^3 . This total added mass is then distributed to the mesh nodes of the structural object. An approach

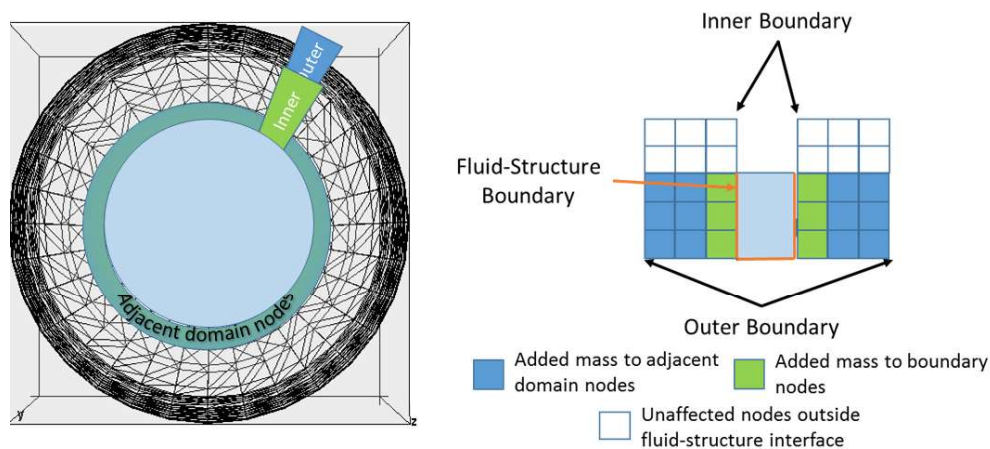


Figure 3.4: Left: top view. Right: cross sectional view. Surface boundary and adjacent domain nodes in between the inner and outer boundary are detected, stored, and modified by an added mass based on the amount and type of liquid in each respective object




Figure 3.5: Added mass should be distributed along the fluid-structure interface. Randomly distributing (left) the added mass over the entire structure also modifies the sound but uniformly (middle) or weighted (right) distributions result in frequencies closest to real-world recordings




Figure 3.6: Free-body diagram for the hydrostatic force along the fluid-structure boundary of a curved surface depends on depth below the surface, liquid density, liquid volume, and gravitational constant g

to randomly distribute this added mass across the entire structure is less accurate than uniformly or force weighted distributions along the fluid-structure boundary.

A uniform distribution calculates the total added mass and then distributes an equal amount to each boundary and adjacent domain mesh node, as illustrated in Fig. 3.4 and Fig. 3.5. Alternatively, rather than uniformly distributing additional mass from the liquid, we may distribute it relative to the local force. For example, mesh nodes with a greater depth will have a greater force and therefore should obtain a greater proportion of the total added weight. Given the boundary nodes and the weight of the liquid, we modify the mass matrix elements for each node based on its hydrostatic force, as shown in Fig. 3.6 and Eqn. (3.18).

3.4.3 Hydrostatic Force on a Curved Surface

In our system, we can also evaluate the force along the boundary of an object to determine the weighted distribution of added mass. This is based on fluid statics where the pressure increases linearly with depth below the water's surface for fluids in a non-moving domain (32). For example, dams are designed to be parabolic for improved stability, allowing the weight of the water to press the dam into the ground.

Newton's 3rd law dictates that $\vec{F}_N = -\vec{F}_R$ and 2nd law that $\sum \vec{F} = m\vec{a} = 0$ since there is no motion. Given these conditions, our goal is to solve for the force along the fluid-structure boundary $\vec{F}_N = F_H\hat{i} + F_V\hat{j}$ where N stands for normal, R for reactive, H for horizontal, and V for vertical, as illustrated in Fig. 3.6. For both the horizontal and vertical components of the normal force, the force magnitude

increases with depth. Therefore, the force along the boundary of the solid is greatest at points farthest away from the surface. In addition to depth, the hydrostatic force equations also demonstrate the linear relationship with density. Since the fluid body is not moving, the net force is zero which means that the horizontal force from the structure is equal to the horizontal pressure force from the liquid.

$$\sum F_x = 0 = F_2 - F_H \quad (3.13)$$

$$F_H = F_2 = \gamma h_c A_{\text{left}} = \gamma \left(d + \frac{h_s}{2} \right) (w_s h_s) \quad (3.14)$$

where $\gamma = \rho \cdot g$, h_c is the vertical distance from the free surface to the centroid of the left planar surface, area of the left planar surface is A_{left} , d is depth, h_s is height from bottom to top of the liquid surface, and w_s is the width of the liquid surface.

$$\sum F_y = 0 = F_V - F_1 - W \quad (3.15)$$

$$W = m_f g = (\rho V_f) g = \gamma V_f = \gamma (d w_s l_s) \quad (3.16)$$

$$F_V = F_1 + W \quad (3.17)$$

The sum of the forces in the vertical direction are also equal to zero. F_1 is the pressure force due to the weight W of the fluid directly above the isolated fluid body and in our case is equal to zero. The weight W is mass m times the gravitational constant g , where mass can be calculated as density ρ times volume V .

$$F_N = \sqrt{F_H^2 + F_V^2} = F_R \quad (3.18)$$

With this approach, we are able to calculate the forces along the boundary. Rather than uniformly distributing additional mass from the liquid, we can distribute relative to the local force. Given the boundary nodes and the weight of the liquid, we modify the mass matrix elements for each node based on its force contribution. If we calculate the forces using a particle approach, then we can distribute the liquid particle forces to the structure mesh nodes using Gaussian quadrature rules (151).

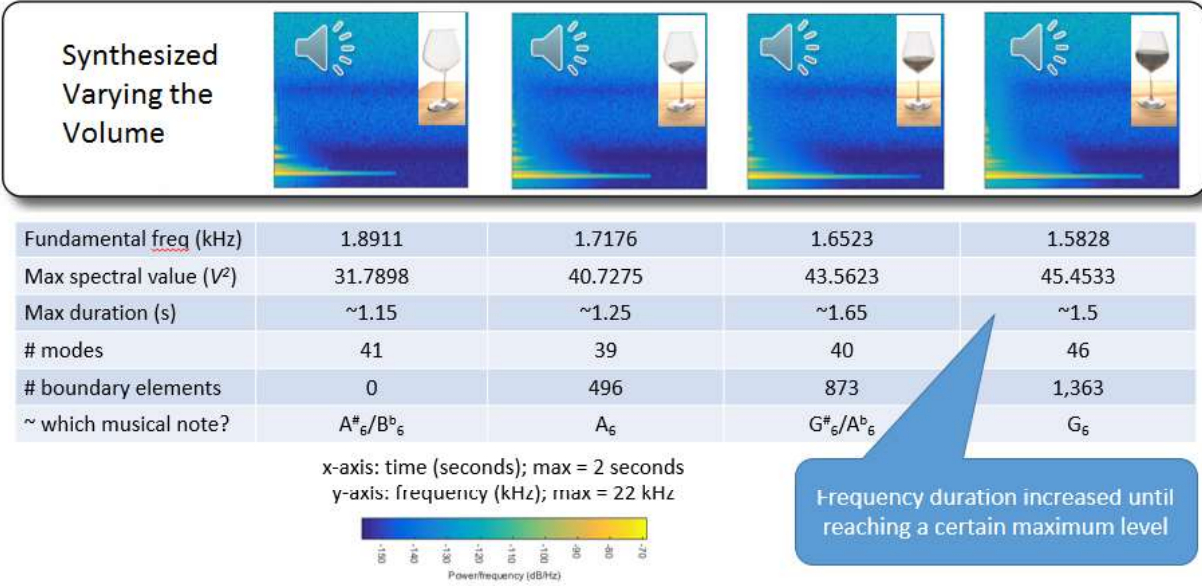


Figure 3.7: This figure displays the results of sound synthesis for our approach with liquid volume increasing from empty to glass half full; notice that the fundamental frequency decreases resulting in a lower pitch as the liquid increases

3.5 Results and Analysis

We have implemented our algorithm in C++, while performing the power spectral density and spectrogram analysis using MATLAB. The virtual reality demo application was created with Unreal Game Engine and viewed with the HTC Vive, as shown in the supplementary video posted at: <http://gamma.cs.unc.edu/GlassHalfFull>

3.5.1 Comparison: Synthesized vs. Recording

To evaluate the effectiveness of our method, we compared synthesized sounds to real-world recordings. The results show that the direction of the change in sound is accurate. For example, the pitch decreases as the volume or density increases. This inverse relationship is referred to as “wet” natural frequencies (with liquid) that are lower than corresponding “dry” natural frequencies (without liquid) (17). Also, although the synthesized and actual frequencies differ, the similarity in sound between the two are imperceptible (206). Our synthesized sounds would provide the expected auditory difference among different liquids, as shown in Table 3.1.




Figure 3.8: Given equal volumes of liquid, the fundamental frequency of the sounding object decreases as the contained liquid density increases

Keeping the volume constant and changing density (e.g. milk versus water) also changes the frequency. For example, in Figure 3.8, we can see that frequency is inversely proportional to liquid density; that is, as density increases, frequency decreases. The difference in sound by varying the type may be harder to distinguish unless there is a more significant difference in density (e.g. olive oil versus milk).

Table 3.1: Generalizations for different liquid densities. Syn Freq and Actual Freq are the synthesized and actual fundamental frequencies respectively, in kHz, of a wineglass with 1/4 liquid volume

| Liquid | Syn Freq (kHz) | Act Freq (kHz) | Rel Error |
|-----------|----------------|----------------|-----------|
| Milk | 1.7037 | 1.6823 | 1.27% |
| Water | 1.7176 | 1.7607 | 2.45% |
| Hot Water | 1.7297 | 1.7771 | 2.67% |
| Olive Oil | 1.7597 | 1.7824 | 1.28% |

Similar to density, volume is also inversely proportional to pitch. Compared to real-world recordings, the change in frequency direction is aligned with expectations and the magnitude of the frequencies are reasonable within single-digits, (see Fig. 3.9 and Table 3.2).

Table 3.2: Generalizations for different liquid volumes where Syn Freq and Actual Freq are the synthesized and actual fundamental frequencies, in kHz, of a porcelain bowl with water

| Vol | Syn Freq(kHz) | Act Freq(kHz) | Rel Error |
|-------|---------------|---------------|-----------|
| Empty | 2.9419 | 2.9709 | 0.98% |
| 1/4 | 2.9389 | 2.9597 | 0.70% |
| 1/3 | 2.8816 | 2.9453 | 2.16% |
| 1/2 | 2.7810 | 2.8759 | 3.30% |




Figure 3.9: Fundamental frequency of the object decreases as the amount of liquid increases

3.5.2 Spectrogram Analysis

We used power spectrograms to analyze the synthesized sounds in a time-varying frequency representation. This allows us to view the fundamental frequency as well as the duration of each frequency. As expected, the fundamental frequency decreases, resulting in a lower pitch, as the liquid increases. Figures 3.10 and 3.11 show that the empty wineglass has a fundamental frequency of 1.89 kHz for a duration of about 1.15 seconds while the half full glass has 1.58 kHz for approximately 1.5 seconds.

The fundamental frequency duration increases as the volume rises; however, at about the volume halfway point, it begins to decrease. This requires additional analysis in future work but we believe that this may be due to the convergence constraint in the added mass operator that the added mass be below some maximum percentage of total mass.

Performance: the wineglass 3D model contained about 17,000 vertices, taking less than 30 seconds of precomputation to run our sound synthesis algorithm for fluid-structure coupling on a 2.40 GHz Lenovo Intel Core i7-4700MQ machine. Since our method enhances an existing pipeline by pre-processing steps only, we refer to related work for a computational cost analysis on the interactivity of sound synthesis techniques (178).

3.6 Applications

We integrated our prototype implementations with two applications: simulated liquid xylophone and interaction with virtual environments.

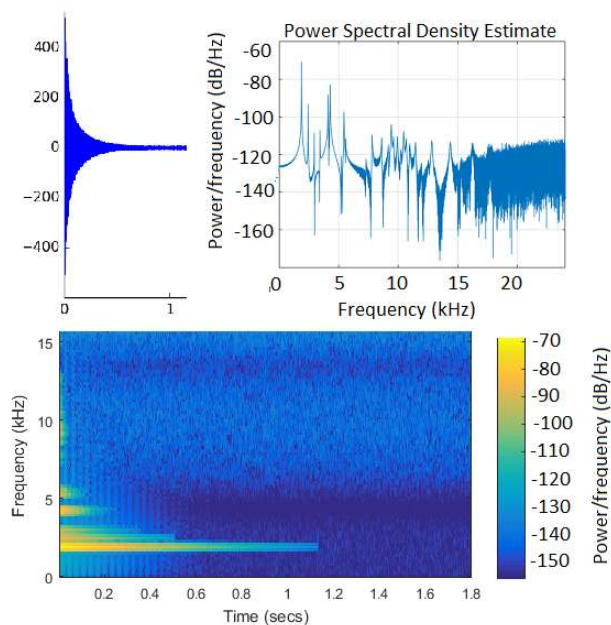


Figure 3.10: Empty wineglass. Top left: time domain signal. Top right: power spectral density estimation graph. Bottom: spectrogram for empty wineglass; higher frequency, higher pitch, shorter duration




Figure 3.11: Half full wineglass. Top left: time domain signal. Top right: power spectral density estimation graph. Bottom: spectrogram for half full wineglass; lower frequency, lower pitch, longer duration

Table 3.3: Solid and liquid are represented as tetrahedral meshes to detect boundary nodes and modify their mass elements although particle-based methods may also be used

| Descriptor | Fluid | Structure |
|--------------------|-------------|-------------|
| Object | Inviscid | Impermeable |
| Mesh | Tetrahedral | Tetrahedral |
| No. Vertices | ~2-5k | ~17k |
| No. Boundary Nodes | ~500-1.5k | ~500-1.5k |



Figure 3.12: User strikes each glass interactively in a virtual environment using a mouse click or virtual reality controller to play an octave of a simulated liquid xylophone

3.6.1 Simulated Liquid Xylophone

Modal analysis is performed before the simulation so the modes are pre-computed, allowing a user to interactively play the simulated liquid xylophone and generate sounds based on object material, liquid type, and liquid volume. Future additions could include making the simulated xylophone even more interactive by allowing the user to modify the type of liquids and/or the amount at run-time.



Figure 3.13: Users can interactively hear the sounds of various objects due to varying amounts and types of liquids in a virtual kitchen scene

3.6.2 Integration With Virtual Environments

To demonstrate the plethora of non-empty objects that we interact with, we applied our method to a virtual kitchen. Next to each empty object is a filled object for sound comparison. Please see the supplementary video to hear the added sound effects.

Table 3.4: Generalizations for different solids and liquid volumes where Syn Freq is the synthesized frequency, in kHz

| Solid | Liquid | Vol | $\rho(kg/m^3)$ | Syn Freq |
|----------------|--------|-------|----------------|----------|
| Porcelain Bowl | Water | Empty | 1,000 | 2.9419 |
| Porcelain Bowl | Water | 1/2 | 1,000 | 2.7810 |
| Metal Teapot | Water | Empty | 1,000 | 10.0070 |
| Metal Teapot | Water | 1/2 | 1,000 | 7.5150 |
| Wineglass | Water | Empty | 1,000 | 1.8911 |
| Wineglass | Water | 1/2 | 1,000 | 1.5828 |

3.7 Conclusion and Future Work

We have presented a novel method that extends the rigid-body sound synthesis pipeline to account for the change in sound frequency and duration resulting from various types and amounts of liquids that an object contains. Although our experiments show results in the general direction of how the frequency changes and produce similar sounds, our work assumes that liquids are inviscid, remain steady, and are not mixed. Our method should be extensible to handle mixed fluids. Liquids are also approximated as spheres to calculate added mass. The granularity of the solid mesh discretization also influences the results since the modifications to the mass matrix occur at the level of the mesh nodes. Other future research directions may include investigation of acoustic transfer and harmonic pressure (97; 271), as well as user evaluation on auditory perception of these added sound effects.

In summary, our sound synthesis algorithm for fluid-structure coupling adds pre-processing steps to the rigid body sound pipeline to simulate sound based on the volume and type of liquid contained within an object. The pre-computed steps are to identify and modify the mass matrix elements of the structural mesh nodes, along the domain of the fluid-structure boundary with an added mass operator proportional to the liquid’s density and volume. Since our technique adds mass from the liquid to the structural object before the simulation, the fluid(s) can be occluded as often is the case in graphical rendering and may be used in interactive 3D graphics and VR systems. Finally, we have demonstrated the effectiveness of our method for simulated musical instruments and composition, as well as enhanced realism of “glass-half-full” sounds in virtual environments.