

Toward Animating Water with Complex Acoustic Bubbles

Timothy R. Langlois
Cornell University

Changxi Zheng
Columbia University

Doug L. James
Stanford University



Figure 1: Complex acoustic bubbles: Our system is able to capture complex frequency effects due to bubbles' shapes and positions. (Left) Bubbles are colored blue/red if they are lower/higher than the theoretical Minnaert frequency for spherical bubbles, and depicts pitch rise near the surface. (Right) Bubbles are colored (blue/red) based on their (small/large) vibration magnitude.

Abstract

This paper explores methods for synthesizing physics-based bubble sounds directly from two-phase incompressible simulations of bubbly water flows. By tracking fluid-air interface geometry, we identify bubble geometry and topological changes due to splitting, merging and popping. A novel capacitance-based method is proposed that can estimate volume-mode bubble frequency changes due to bubble size, shape, and proximity to solid and air interfaces. Our acoustic transfer model is able to capture cavity resonance effects due to near-field geometry, and we also propose a fast precomputed bubble-plane model for cheap transfer evaluation. In addition, we consider a bubble forcing model that better accounts for bubble entrainment, splitting, and merging events, as well as a Helmholtz resonator model for bubble popping sounds. To overcome frequency bandwidth limitations associated with coarse resolution fluid grids, we simulate micro-bubbles in the audio domain using a power-law model of bubble populations. Finally, we present several detailed examples of audiovisual water simulations and physical experiments to validate our frequency model.

Keywords: acoustic bubbles, sound synthesis, acoustic transfer, fluid animation

Concepts: •Computing methodologies → Physical simulation; Continuous simulation; Massively parallel and high-performance simulations;

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org. © 2016 Copyright held by the owner/author(s). Publication rights licensed to ACM. SIGGRAPH 2016 Technical Paper, July 24-28, 2016, Anaheim, CA

1 Introduction

Liquids, and the sounds they make, are pervasive in our daily lives. Whether it be a dripping faucet, a babbling brook, or your last glass of water, you have most likely heard sounds generated by fluids recently. While there has been significant work on understanding how fluids generate sound using bubbles, and breakthroughs in the *visual* simulation of water, there are no existing methods for computing a realistic *audiovisual* simulation of water with quality anywhere comparable to the purely visual component. For instance, current fluid sound approaches do not even simulate acoustic bubbles realistically by modeling their evolution using two-phase liquid simulations, but instead rely on single-phase flow solvers with *ad hoc* point-like bubble generation techniques [Zheng and James 2009; Moss et al. 2010]. Such approximations are naturally much cheaper to compute, but, unfortunately, they have limited predictive value and ultimately limited realism. In contrast, we seek to understand whether one can simulate bubbly flows with sound from first physical principles, and what modeling challenges and trade-offs must be addressed.

In this paper, we explore a family of methods for sonifying detailed two-phase liquid animations such as the one shown in Figure 1. Given the complexity of liquid sound generation processes, there are many details we consider (see Figure 2 for an overview of our approach). Our approach begins with detailed multi-scale simulation of two-phase incompressible flow to resolve fine bubble geometry needed for higher frequency sounds¹. Accurate modeling of surface tension is needed to resolve bubble pinch-off and subsequent topology changes. Individual bubble geometry is estimated and tracked, and we resolve bubble entrainment, merging, splitting and popping processes at sub-*ms* time scales, and sub-*mm* length scales.

Since the fluid flow is treated as incompressible for performance

ISBN: 978-1-4503-4279-7/16/07

DOI: <http://dx.doi.org/10.1145/2897824.2925904>

¹Recall that Minnaert's frequency model predicts $f \approx 6.52/d_{mm} \text{ kHz}$ (at STP) where d_{mm} is the bubble diameter in mm.

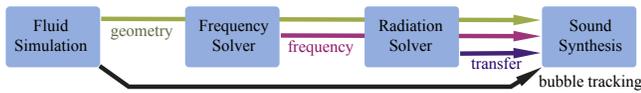


Figure 2: Overview of our system: From the fluid simulator, our method requires fluid surface geometry at each timestep, as well as bubble correspondences between timesteps. With this geometry, we compute each bubble’s frequency and acoustic transfer magnitude, which are used to synthesize the bubble’s sound.

reasons, all acoustic fluctuations of the bubble/fluid/air system are handled using reduced-order vibration models. Toward this end, we propose a method for detailed bubble frequency analysis based on a “bubble capacitance” interpretation. Bubble pitch changes are perceptually important, but previous methods just model them using a parametric “chirp” as the bubble approaches the surface. In contrast, our method can resolve complex pitch shifts due to nonspherical bubble shapes, as well as nearby solid and air interfaces that can result in dramatic pitch decreases and increases, respectively (see Figure 1). We estimate the bubble frequency using a boundary element method, and since we analyze many bubbles, we propose a fast amortized matrix solver which effectively exploits inter-bubble similarities between matrix form factors and dense matrix solves.

To estimate a bubble’s sound pressure at a listener position, we must faithfully estimate the surface-to-air acoustic transfer. To do so, we perform a standard frequency-domain boundary element analysis, with surface vibration data input provided by our bubble-frequency solver. Unlike previous methods, this accounts for the near-field scene geometry to capture container resonance effects, such as the characteristic rising pitch of a container being filled up with water. For cheap, low-accuracy previews, we also propose a greatly simplified bubble-plane transfer model, based on precomputing a lookup table indexed by the size and depth of proxy bubbles.

During the final sound synthesis phase, we simulate the bubble oscillators using the estimated time-varying frequency. Bubble forcing models are devised to account for bubble excitations during surface entrainment, but also subsequent splitting and merging events. While prior works considered only Laplace pressure-jump forcing at entrainment, we leverage improved models of detailed surface-tension modeling (recently proposed by Deane and Czerski [2008; 2010; 2011; 2011]). To model bubble popping sounds, we describe a Helmholtz resonator model.

Since the fluid simulator’s spatial resolution restricts bubble sizes, our synthesized sounds are inherently frequency band-limited. Therefore we propose a bandwidth extension scheme that performs audio-domain simulation of micro-bubbles based on power-law models of micro-bubble populations in breaking waves.

Finally, our results include multiple examples of dripping, pouring, and splashing phenomena, as well as results from laboratory experiments to validate our capacitance-based frequency model.

2 Related Work

Fluid sound can be generated from a variety of sources, including vortex based sounds [Howe 2002], fluid-structure interaction [Howe 1998], shock waves, bubble popping [Deane 2013], drop impact [Howe and Hagen 2011], and, the focus of this paper, harmonic bubble vibrations. Investigation of sounds produced by bubbles dates back almost a century, to the work of Lord Rayleigh [1917] and William Bragg [1920]. Minnaert calculated the frequency of isolated, spherical, harmonically vibrating bubbles [1933]. Strassberg [1953] described how a bubble’s shape and the surrounding

geometry affect its frequency, and by connecting frequency with capacitance went on to compute frequencies for sphere-plane and ellipsoid bubble geometries. Spratt et al. [2015] verified this model for simple nonspherical bubble geometries against a full acoustic scattering computation. We further extend this approach to support frequency computations with general nonspherical bubble systems and arbitrary air and solid interfaces.

Acoustic bubbles have been studied intensively, due to the impact bubbles have on physical processes, and the necessity of passive acoustic sensing of processes which are difficult to observe visually—Leighton’s monumental work [1994] provides a definitive summary. Acoustic bubbles represent a significant portion of ambient ocean noise [Prosperetti 1988a; Prosperetti 1988b], and contribute to gas exchange between the ocean and the air and affect ocean albedo [Deane and Stokes 2010], which in turn affects climate [Bigg et al. 2003; Leighton 1994]. They are important for “accurate quantification of a number of dynamic processes at the air-sea boundary, such as wave energy dissipation, gas exchange rates and the nature of rain” [Klusek and Lisimenka 2013; Longuet-Higgins 1990; Pumphery et al. 1989], such as using passive acoustic remote sensing [Nystuen 1986; Ding and Farmer 1994; Means and Heitmeyer 2002; Wilson and Makris 2008] and improvements (hopefully) in predictive computational models. Acoustic bubbles are even used by humpback whales during bubble net hunting [Leighton et al. 2007]!

Fluid simulation has been a success story of computational physics, with widespread application in computer graphics and animation. However, most fluid simulators in graphics have focussed on single-phase free-surface flow [Stam 1999; Enright et al. 2002; Osher and Fedkiw 2006; Bridson 2008], and have only more recently tackled two-phase flow simulation [Hong and Kim 2005; Losasso et al. 2006; Boyd and Bridson 2012; Ando et al. 2015]. These works almost exclusively aim for visually plausible simulations with large timesteps, and not for acoustic bubble computations. Because of their visual richness, a variety of methods have been developed for simulating air bubbles [Hong and Kim 2003; Hong and Kim 2005], and works range from tiny bubbles [Busaryev et al. 2012] to large volumes of bubbles [Zheng et al. 2009], and complex thin-film interfaces [Da et al. 2015]. These methods lack the ability to track individual acoustic bubbles while preserving their volumes accurately—for example, Kim et al. [2007] artificially inflated bubbles to compensate for their volume change. However, bubble volume contributes significantly toward their estimated frequency, as demonstrated in our work. Consequently, we have built upon *Geris* [Popinet 2003; Popinet 2009], a finite-volume-based multigrid solver used in computational fluid dynamics, to more accurately simulate two-phase flows and track bubbles.

Fluid sound simulation: Despite the importance of bubble sounds, relatively little work has been done on simulating them. Van den Doel [2005] proposed a statistical method to generate bubble sounds. More recently, two works [Zheng and James 2009; Moss et al. 2010] have proposed more physically based methods. However, due to the sheer computational difficulty of predictive modeling of bubble entrainment processes, these methods have relied on single-phase liquid simulators with *ad hoc* stochastic models to estimate point-like bubble creation rates and size distributions, with the unfortunate consequence that bubble creation rates are either unrealistic [Moss et al. 2010] or must be laboriously hand-tuned for examples [Zheng and James 2009]. Furthermore, these single-phase flows solvers can not estimate realistic nonspherical bubbles, and they lack realistic time-varying bubble frequencies. Additionally, Zheng and James [2009] use an acoustic transfer approximation which cannot capture scattering effects from enclosing solid interface, while Moss et al. [2010] ignore acoustic radiation entirely.

Multi-frequency vibration: For small acoustic bubbles, it is usu-

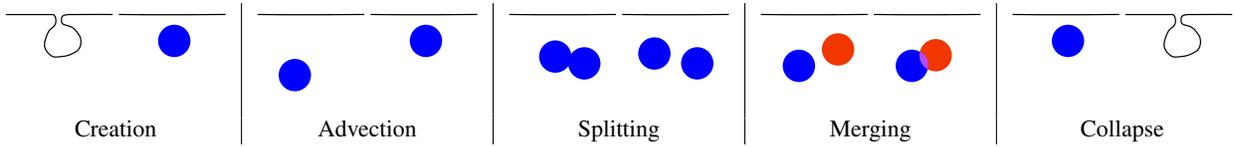


Figure 3: Bubble Tracking: We use colors to denote different bubble ids b_i . There are five types of bubble actions between timesteps. The left figure in each column denotes timestep i , and the right figure denotes timestep $i + 1$. Creation or entrainment: no b_i overlap with b_{i+1} . Advection: b_{i+1} overlaps with one value of b_i . Splitting: one bubble id b_i overlaps with ≥ 2 different bubble ids, b_{i+1} . Merging: two different b_i 's overlap with one b_{i+1} . Collapse: no b_{i+1} overlaps with a b_i .

ally assumed that they vibrate at a single frequency. However, it is possible that other vibration modes can radiate sound at different frequencies. Lamb [1932] investigated higher-order vibration modes using a first-order perturbation analysis, and concluded that they do not radiate efficiently enough to be important. A series of papers by Longuet-Higgins [1989a; 1989b; 1991] showed that in certain situations, resonant coupling can occur between shape modes and the volume mode, and higher vibration modes can radiate as a monopole, causing a single bubble to produce sound at several frequencies. The recent work in graphics by [Moss et al. 2010] proposed using a related multi-frequency zonal-harmonic vibration model. While multi-frequency radiation has been observed in certain specific experimental conditions [Chicharro and Vazquez 2014], it is believed that under real-world forcing conditions these coupling effects are less important [Longuet-Higgins 1991], and arguments have been given that shape-mode coupling accounts for a very small portion of the total air-domain sound radiation of bubbles [Longuet-Higgins 1990]. Longuet-Higgins himself also mentions that his model is not valid for realistic pinch-off scenarios, which would require fully nonlinear equations [Longuet-Higgins 1989b].

Experimental evidence to support multi-frequency bubble sounds is also lacking. Medwin and Beaky [1989] analyzed a large number of cases of single- and multi-bubble events (over 2000) generated in a wave tank. They laboriously classified four types of bubble radiation. The large majority of events show no signs of multiple frequencies. The one type of event (“type D”) that does show multiple frequencies is speculated to be caused by interference from two bubbles. Therefore, while it is possible for a single bubble to radiate at multiple frequencies, it seems to be rare, and we ignore it in the present study.

3 Fluid Simulation

Due to the dependence of a bubble’s frequency on its size, shape, and position, accurate fluid simulation is necessary for realistic sound. We solve the incompressible, variable-density Navier-Stokes equations with surface tension.

$$\rho(\partial_t \mathbf{u} + \mathbf{u} \cdot \nabla \mathbf{u}) = -\nabla p + \nabla \cdot (2\mu \mathbf{D}) + \sigma \kappa \delta_s \mathbf{n} \quad (1)$$

$$\partial_t \rho + \nabla \cdot (\rho \mathbf{u}) = 0 \quad (2)$$

$$\nabla \cdot \mathbf{u} = 0 \quad (3)$$

with \mathbf{u} the fluid velocity, ρ the fluid density, μ the dynamic viscosity, and \mathbf{D} the deformation tensor defined as $D_{ij} = (\partial_i u_j + \partial_j u_i)/2$. The surface tension force is only nonzero on the interface, which the Dirac distribution δ_s signifies. σ is the surface tension coefficient, κ is the surface curvature (not related to our κ in equation (4)), and \mathbf{n} is the surface normal.

For two-phase flows, densities ρ_1, ρ_2 and viscosities μ_1, μ_2 are given for the first and second fluids respectively. The continuous volume fraction field $c(\mathbf{x}, t)$ is used to denote the volume of the first

fluid and defines the density and viscosity as

$$\rho(c) = c\rho_1 + (1 - c)\rho_2$$

$$\mu(c) = c\mu_1 + (1 - c)\mu_2$$

The advection equation for the density (2) is then replaced with an equivalent equation for the volume fraction

$$\partial_t c + \nabla \cdot (c\mathbf{u}) = 0.$$

We use the open-source *Gerris* solver [Popinet 2003; Popinet 2009], which is based on a finite volume method implementation with an adaptive octree data structure, parallel multigrid Poisson solver, and an accurate surface tension model. *Gerris* enables high-fidelity simulation of bubble shapes and motion, with good volume preservation properties, as well as realistic bubble formation from surface entrainment. However, despite its merits, there are significant spatial and temporal resolution requirements in order to resolve bubble entrainment and subsequent topological events, and thin interfaces. In our simulations, we use spatial resolutions on the order of [0.625mm, 5mm], timesteps (determined by CFL conditions) on the order of [30 μ s, 250 μ s], and spatial adaptation near interfaces. Nevertheless, our ability to resolve tiny bubbles below 1 mm is limited. Also, *Gerris* has no way to preserve thin films when bubbles come into contact with each other or the air surface, so merging events are over estimated and surface popping is immediate.

4 Bubble Identification and Tracking

The link between the fluid simulation and our sound pipeline is the interface geometry, consisting of the surface of each bubble and the enclosing fluid volume. We describe here how this is implemented in *Gerris*, but note that any fluid simulation method which can provide this geometry will be compatible with our method.

Our representation of the interface geometry is given by c on the octree grid. Specifically, bubbles are connected components of cells where $c < 1$, which are identified and uniquely numbered with a flood fill algorithm. To track bubbles between time steps, a new variable b_i is defined at timestep i and initialized to the number of each bubble. The b_i variable is advected between time steps. Overlaps between the current b_{i+1} and the previous b_i advected from the last timestep are used to correlate the bubbles between time steps. There are several situations:

1. **New bubble:** No b_i touching b_{i+1}
2. **Advected bubble:** One value of b_i touching one value of b_{i+1}
3. **Split bubble:** Two or more values of b_{i+1} overlap one of b_i
4. **Merged bubbles:** Two or more values of b_i overlap one of b_{i+1}
5. **Collapsed bubble:** No b_{i+1} touching a value of b_i

These are illustrated in Figure 3. The CFL condition ensures that interface fragments do not move more than one grid cell per timestep, and we did not observe any bubbles getting lost during tracking. Marching cubes [Lewiner et al. 2003] is used to generate interface geometry for each bubble, and for the enclosing fluid volume.

5 Bubble Frequency Estimation

After a brief introduction to bubble vibrations (§5.1), we describe a new model for estimating the instantaneous bubble frequency (§5.2, §5.3), and an efficient algorithm that amortizes computation across many bubbles (§5.4).

5.1 Bubble Basics

The equations of spherical bubble vibration were originally proposed by Minnaert [1933] and are described in detail in Chapter 3 of [Leighton 1994]. Briefly, when bubbles are formed, they vibrate, creating pressure waves which travel through the fluid and then pass through the fluid-air interface into the air. The simple harmonic oscillator model is similar to a linear spring-damper-mass system, where the spring forces are provided by internal gas pressure and surface tension, and the mass is due to the surrounding liquid. We use the “volume-pressure” frame, so that the infinitesimal volume pulsation of a bubble $v = V(t) - V_0$ (where V_0 is the average volume) satisfies

$$m\ddot{v} + \alpha\dot{v} + \kappa v = p(t) \quad (4)$$

where $p(t)$ is a forcing term in units of pressure. The equation is usually divided through by m and written as

$$\ddot{v} + 2\beta\dot{v} + \omega^2 v = \frac{p(t)}{m}$$

with $\omega^2 = \frac{\kappa}{m}$. For a bubble of equivalent radius r , $\beta = \omega\delta/\sqrt{\delta^2 + 4}$, with $\delta = \delta(\omega, r) = \delta_{rad} + \delta_{vis} + \delta_{th}$ where

$$\delta_{rad} = \frac{\omega r}{c} \quad \delta_{vis} = \frac{4\mu}{\rho\omega r^2} \quad \delta_{th} = 2 \frac{\sqrt{\psi - 3} - \frac{3\gamma - 1}{3(\gamma - 1)}}{\psi - 4}$$

with $\psi = \frac{16}{9(\gamma - 1)^2} \frac{G_{th}g}{\omega^2 c}$, c is the speed of sound in the fluid, μ is the liquid’s shear viscosity, ρ is the fluid density, γ is the gas’ heat capacity ratio, $G_{th} = \frac{3\gamma p_f}{4\pi\rho D_g}$ is the thermal damping constant at resonance, D_g is the gas’ thermal diffusivity, g is gravitational acceleration, and p_f is the hydrostatic pressure of the liquid. A full derivation is provided by Leighton [1994].

Minnaert computed ω for a spherical bubble. However, nearby fluid-air and fluid-solid surfaces, as well as bubble shape, can affect ω , and are the reason for the familiar pitch shift as bubbles approach the fluid surface. In the following sections, we describe one of our main contributions: a method for accurately computing ω (by computing κ and m) for complex geometries, and thereby providing more realistic sound.

5.2 Frequency Model

Extending Strasberg’s [1953] derivation, we seek to better estimate the bubble frequency ω , by more accurately modeling the scene-specific effective stiffness, κ , and effective mass, m . We consider a bubble’s infinitesimal volume pulsation $v(t)$ in the volume-pressure frame (4), in the undamped case ($\alpha=0$).

The effective stiffness (κ) accounts for internal gas and surface tension effects, and is the simplest of the two values. The stiffness in the volume-pressure frame is the rate of change of bubble pressure with volume,

$$\kappa = -\frac{dp}{dv}. \quad (5)$$

Assuming a polytropic gas law, the reference volume V_0 and pressure P_0 are related to the modified volume $V = V_0 + v$ and pressure $P = P_0 + p$ by

$$P_0 V_0^\gamma = P V^\gamma,$$

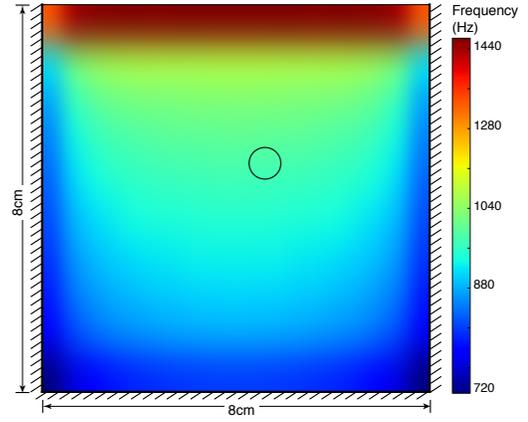


Figure 4: Spatially varying bubble frequency (in Hz) depicted for a spherical bubble of 3mm radius (Minnaert frequency of approximately 1100 Hz). The pitch lowers as the bubble nears rigid walls (left, right, and bottom), and rises sharply as the bubble nears the fluid surface (top). Even in this small 8cm-by-8cm tank, the bubble’s frequency can differ by over 700 Hz depending on position. Note that for this figure, frequencies were only sampled 3.6mm from the boundary (wall or surface), and extrapolated.

where γ is the polytropic index; we use $\gamma = 1.4$ for air. Evaluating the derivative (5) we obtain the bubble stiffness,

$$\kappa = \frac{\gamma P_0}{V_0}. \quad (6)$$

The bubble air pressure P_0 is the sum of the hydrostatic pressure (approximately $P_{atm} + \rho g d$ for a bubble at depth d), and surface tension ($\frac{2\sigma}{R_0}$ for a spherical bubble). To support pressure and surface tension for nonspherical bubbles in complex fluid domains, we obtain a reliable approximation of P_0 from our two-phase fluid simulator. Observe that κ only depends on the internal gas and surface tension, which is consistent with Minnaert’s derivation [Leighton 1994]. Interestingly, note that κ is incapable of producing increasing pitch as a bubble rises (it actually predicts the opposite dependence on d), and thus the mass factor is responsible for this effect.

The effective mass (m) can be derived similar to Strasberg [1953], but with consideration for complex fluid domains. The key idea is to equate the oscillator kinetic energy $\frac{1}{2}m\dot{v}^2$, for a given harmonic bubble volume velocity \dot{v} , with the kinetic-energy volume integral of the surrounding fluid, W_f , to determine m . W_f is also the work that the bubble does to the fluid. In order to compute the fluid kinetic energy, we exploit the fact that bubbles are acoustically compact low-frequency sources², and thus the surrounding fluid’s vibration is well approximated by incompressible, irrotational flow modeled using Laplace’s equation. Specifically, let the surrounding fluid’s particle velocity be $\nabla\phi$, where ϕ is the velocity potential which satisfies

$$\nabla^2 \phi(x) = 0, \quad x \in \Omega,$$

subject to a pressure-like Dirichlet boundary condition on the bubble, $\phi = \phi_b$, and other suitable boundary conditions elsewhere (discussed later). Therefore we can express the effective mass as

$$m = \frac{2W_f}{\dot{v}^2} = \frac{2}{\dot{v}^2} \left(\frac{\rho}{2} \int_{\Omega} (\nabla\phi)^2 d\Omega \right).$$

²Note that for a bubble at 1 atm pressure, we have $fR \approx 3m/s$, or $\omega R \approx 19m/s$, so that $kR = \omega r/c \approx 0.013 \ll 1$. Also note that $\lambda/R \approx 114$.

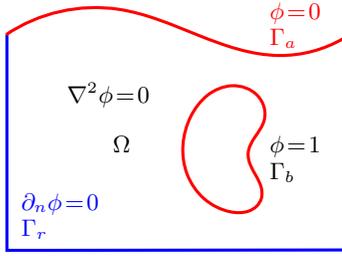


Figure 5: Interior Laplace BVP for bubble capacitance: We use the solution's $\partial_n \phi$ gradient on the bubble boundary Γ_b to compute the bubble capacitance using (9), and on the air boundary Γ_a to evaluate acoustic radiation (in §6).

Converting the volume integral to a boundary integral using Green's first identity we obtain

$$\int_{\Omega} (\nabla \phi)^2 d\Omega + \underbrace{\int_{\Omega} \phi \nabla^2 \phi d\Omega}_{\approx 0} = - \int_{\partial\Omega} \phi \partial_n \phi dS,$$

where the minus sign is due to normals pointing out of the fluid domain Ω (and into air or solid domains). Denoting the bubble, air, and rigid surfaces as Γ_b , Γ_a , and Γ_r , the mass integral becomes

$$\begin{aligned} m &= -\frac{\rho}{\dot{v}^2} \left(\int_{\Gamma_b} \phi \partial_n \phi dS + \int_{\Gamma_a} \phi \partial_n \phi dS + \int_{\Gamma_r} \phi \partial_n \phi dS \right) \\ &= m_b + m_a + m_r. \end{aligned}$$

The terms can be evaluated by making use of the boundary conditions for the generalized problem (see Figure 5). The air-interface m_a contribution is zero since $\phi = 0$ on Γ_a , as the acoustic pressure $p = -\rho \frac{\partial \phi}{\partial t} = -i\omega \rho \phi = 0$ there. The rigid-surface m_r contribution is also zero: the acoustic particle velocity must be zero in the normal direction on the rigid boundary, and thus $\partial_n \phi = 0$ on Γ_r . Therefore, *the only mass contribution arises from the bubble term* $m = m_b$.

5.3 Capacitance Interpretation of Bubble Frequency

Since the velocity potential is constant on the surface of the bubble, $\phi = \phi_b$, it can be taken outside the integral,

$$m = -\frac{\rho}{\dot{v}^2} \int_{\Gamma_b} \phi \partial_n \phi dS \approx -\frac{\rho}{\dot{v}^2} \phi_b \int_{\Gamma_b} \partial_n \phi dS.$$

Then since the volume velocity is $\dot{v} = - \int_{\Gamma_b} \partial_n \phi dS$, it follows that $m = \rho \phi_b / \dot{v}$. The bubble is a uniform-potential conductor-like surface, so Strasberg noticed that treating ϕ as electrostatic potential and \dot{v} as flux, the ratio \dot{v} / ϕ is mathematically equivalent to 4π times the surface's electrostatic capacitance, C . Therefore, the effective mass can be identified as

$$m = \frac{\rho}{4\pi C}. \quad (7)$$

By combining equations (6) and (7), the bubble frequency can be given in terms of its capacitance,

$$\omega^2 = \frac{\kappa}{m} = \frac{4\pi\gamma P_0}{\rho V_0} C. \quad (8)$$

Bubble capacitance BVP: To compute the capacitance we can interpret the bubble as a conductor with a unit potential boundary

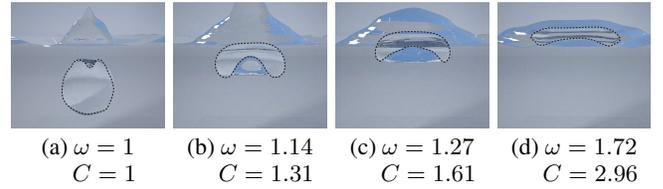


Figure 6: Capacitance-based frequency estimation for a rising bubble: We recover increasing “chirp-like” frequency and capacitance (normalized) as the bubble (initial radius $R=5.8$ mm) nears the surface (a-d). The rising pitch produced as the bubble's water layer (lamella) thins corresponds to a thin-plate capacitor of increasing thinness.

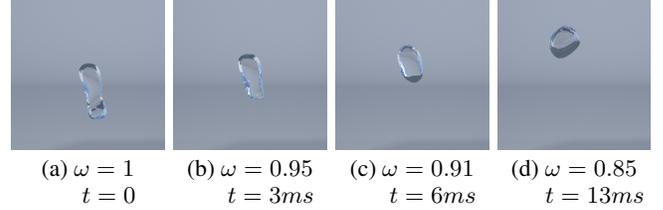


Figure 7: Shape-dependent bubble frequency is demonstrated here for a simulated bubble undergoing natural shape changes. Frequency is normalized. Our bubble frequency estimation method can resolve musical pitch fluctuations occurring on semi-tone magnitudes, on the order of 10ms.

condition ($\phi_b = 1$), the fluid-air surface as a conductor at zero potential ($\phi_a = 0$), and rigid interfaces as insulators ($\partial_n \phi_r = 0$). Our generalized “bubble capacitor” boundary value problem (BVP) is shown in Figure 5.

Given the solution to Laplace's equation $\hat{\phi}$ for this BVP, we compute the capacitance as

$$C = -\frac{1}{4\pi} \int_{\Gamma_b} \partial_n \hat{\phi} dS. \quad (9)$$

Since this formula only requires $\partial_n \hat{\phi}$ on Γ_b , we can solve for each bubble's capacitance using a boundary integral formulation of Laplace's equation. We discuss an optimized bubble capacitance solver in §5.4. Finally, our approach reproduces and generalizes the frequency models of Minnaert and Strasberg, and supports nontrivial frequency estimation in complex fluid geometries (see Figure 6), and for complex bubble shape changes (see Figure 7).

5.4 Fast Amortized Solution of Capacitance BVP

We now describe an efficient method for estimating bubble capacitances that exploits common computations between the bubbles for speed. We can use the boundary element method (BEM) to solve the capacitance BVP using established codes for the interior Laplace problem with mixed boundary conditions (BCs) [Banerjee and Butterfield 1981].

BEM BVP matrix structure: After discretizing the direct boundary integral equation formulation of the interior Laplace problem for the velocity potential associated with a single bubble, we arrive at the linear matrix problem,

$$\mathbf{H}\phi = \mathbf{G}v,$$

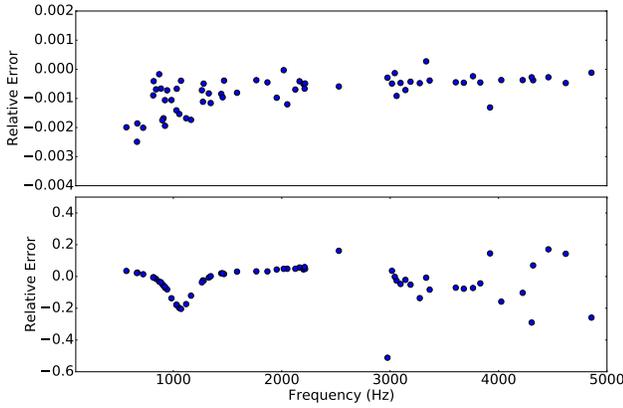


Figure 8: Solver Error: For one timestep of the pouring faucet example, we meshed the domain at a high resolution (5mm maximum edge length), and computed frequency and transfer for all the bubbles. Although we use aggressive mesh simplification, the relative errors for our frequency solver compared to the high resolution results are very small (top). Pressure magnitude errors (bottom) are tolerable in our range of interest, and increase at higher frequencies as expected.

with the following block structure,

$$\begin{bmatrix} H_{bb} & H_{ba} & H_{bs} \\ H_{ab} & H_{aa} & H_{as} \\ H_{sb} & H_{sa} & H_{ss} \end{bmatrix} \begin{pmatrix} \phi_b \\ \phi_a \\ \phi_s \end{pmatrix} = \begin{bmatrix} G_{bb} & G_{ba} & G_{bs} \\ G_{ab} & G_{aa} & G_{as} \\ G_{sb} & G_{sa} & G_{ss} \end{bmatrix} \begin{pmatrix} v_b \\ v_a \\ v_s \end{pmatrix},$$

where ϕ is the BEM vector of potential values and v is the BEM vector of (outward) normal derivative values, $\frac{\partial \phi}{\partial n}$; here the three boundary regions are denoted by b (bubble), a (air), and s (solid). Applying the Capacitance BVP boundary conditions (see Figure 5), we arrive at the linear system $\mathbf{A}\mathbf{x} = \mathbf{b}$ with the block structure,

$$\underbrace{\begin{bmatrix} G_{bb} & G_{ba} & H_{bs} \\ G_{ab} & G_{aa} & H_{as} \\ G_{sb} & G_{sa} & H_{ss} \end{bmatrix}}_{\mathbf{A}} \underbrace{\begin{pmatrix} v_b \\ v_a \\ -\phi_s \end{pmatrix}}_{\mathbf{x}} = \underbrace{\begin{bmatrix} H_{bb} \\ H_{ab} \\ H_{sb} \end{bmatrix}}_{\mathbf{b}} (\phi_b),$$

where the unknowns are $v_b \in R^{n_b}$, $v_a \in R^{n_a}$, and $\phi_s \in R^{n_s}$. Solving this system provides three quantities: v_b which helps approximate $\int_{\Gamma_b} \partial_n \hat{\phi} dS \approx a_b^T v_b$ and thus (9) for the bubble capacitance; v_a which describes the free surface vibration, and will be used for acoustic radiation modeling (in §6); and ϕ_s which is ignored since it is not needed by our application.

In our implementation we use the BEM++ software library [Smigaj et al. 2012] to evaluate the various \mathbf{G} and \mathbf{H} matrix blocks; we used Galerkin discretization (which is more robust to meshing irregularities than collocation schemes) with constant elements for Neumann data (on Γ_s), and piecewise linear elements for Dirichlet data (on Γ_b and Γ_a).

Fast amortized matrix solver: For a given frame, we solve the matrix problem many times for many bubbles. A naive LU-based evaluation of $\mathbf{x} = \mathbf{A}^{-1}\mathbf{b}$ would require $O((n_b + n_a + n_s)^3)$ flops. If we denote the four blocks of \mathbf{A} by

$$\mathbf{A} = \left[\begin{array}{c|c} G_{bb} & B \\ \hline C & D \end{array} \right],$$

we observe that the huge lower-right “domain” submatrix D relating the self-effect form factors for the air and solid boundaries ($\Gamma_a \cup \Gamma_s$) is constant across problems. We can exploit this fact

for fast evaluation of $\mathbf{A}^{-1}\mathbf{b}$ for different bubbles. By exploiting the Sherman-Morrison-Woodbury formula for U [Hager 1989], the matrix inverse can be written as

$$\mathbf{A}^{-1} = \left[\begin{array}{c|c} X & Y \\ \hline Z & U \end{array} \right]$$

where

$$X = (G_{bb} - BD^{-1}C)^{-1}, \quad (10)$$

$$Y = -XBD^{-1}, \quad (11)$$

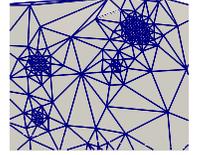
$$U = D^{-1}(I - CY), \text{ and} \quad (12)$$

$$Z = -UCG_{bb}^{-1}. \quad (13)$$

By carefully exploiting common subexpressions and cached LU factorization of the large D block, the product $\mathbf{x} = \mathbf{A}^{-1}\mathbf{b}$ can be evaluated in $O((n_a + n_s)^2)$ flops for small bubbles of size $n_b = O(1)$. Please see Appendix A for explicit algorithmic details. In our examples, we observe that the amortized capacitance solver accelerates bubble frequency estimation by 10-12x for small bubbles, while for very large bubbles the speed-up may only be 3x, but still worthwhile.

5.5 Adaptive Meshing

The generation of meshes used for BEM analysis should be discussed briefly, since adaptive meshing is needed to keep our frequency solve times manageable. Detailed meshes are obtained from the fluid simulator using Marching Cubes, which we adaptively simplify using Quadric Error Simplification [Garland and Heckbert 1997]. However, we must properly resolve interfaces and inter-surface gaps as bubbles approach the fluid surface or container walls in order to compute accurate capacitance values. Therefore we use a sizing function that ensures each triangle’s maximum edge length is less than its distance to the nearest bubble. The inset shows an illustrative example. Although we use aggressive mesh simplification, our errors (shown Figure 8) are still tolerable.



6 Bubble Acoustic Transfer

We now describe two acoustic transfer solvers that estimate the pressure amplitude at the listener’s position(s) from an acoustic bubble vibrating with unit pressure: (1) a BEM-based solver that includes scene geometry (§6.1), and (2) a fast but very approximate solver that uses a proxy bubble-plane transfer model (§6.2). In practice, we sample frequency and bubble-to-ear transfer values at a fixed rate during the lifetime of the bubble; in our implementation these solves are done every 1ms, and interpolated with a cubic spline. These transfer values are used for sound synthesis later in §8.

6.1 BEM-based Acoustic Transfer Solver

We now describe how to approximate realistic sound amplitudes from an harmonically vibrating fluid surface. Since sound scattering and resonance effects from external geometry (such as the walls of a glass container) can introduce perceptually important and pitch-dependent amplitude variations (such as when a glass is filled up with water) we seek to include near-field scene geometry in our transfer solver that has been neglected in prior works [Zheng and James 2009].

Mathematically we approximate solutions to the exterior wave radiation problem specified by the Helmholtz equation

$$(\nabla^2 + k^2)p(x) = 0, \quad x \in \Omega_a,$$

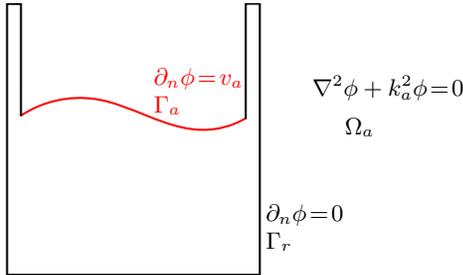


Figure 9: Exterior Helmholtz BVP for acoustic radiation

where $k = \omega/c_a$, and c_a is the speed of sound in air. As boundary conditions, we impose a vibration BC on the fluid-air interface, and a no-vibration BC on the rigid scene geometry (see Figure 9). Following our bubble frequency solve, the vibration of the fluid-air surface is known³ from $\partial_n \hat{\phi}$ on Γ_a , which in turn becomes input boundary data for the air-domain acoustic transfer estimation.

To support arbitrary scene geometry, we use established BEM solvers for the exterior Helmholtz radiation BVP. In our implementation, we use the BEM++ implementation [Smigaj et al. 2012]. A major practical task is generating a suitable scene mesh with BC data at each time step. We first mesh the fluid and external container together to generate detailed geometry for the external radiation problem (see Figure 11). Since these meshes are typically too detailed for efficient BEM analysis, we again use adaptive meshing (as in §5.5) to decimate the mesh. We further restrict the largest edge length so as to sufficiently resolve the smallest wavelength; in our examples, we have used a 3cm limit which can resolve $\lambda_a = 17.2\text{cm}$ at the upper 2kHz range we simulate. Finally the interior capacitance BVP and exterior transfer BVP may have different meshes for the fluid-air interface, and therefore we interpolate the previously obtained velocity-like solution data, v_a (i.e., $\partial_n \hat{\phi}$), from the interior mesh to the mesh of the exterior fluid-air interface Γ_a . The pressure BC on that interface is the same up to a scaling factor, since $\partial_n p_f = \frac{\rho_f}{\rho_a} \partial_n p_a = \frac{\rho_f}{\rho_a} \partial_n \hat{\phi}$.

6.2 Fast Bubble-Plane Proxy Transfer

Transfer computations can be expensive for detailed water surfaces and scenes. For many applications and “fast preview” renderings, we can use a cheap transfer model based on a simplified bubble-plane geometry (see Figure 10).

To do this, we precompute a lookup table of transfer values for spherical bubbles of various radii (r) at various depths (d) beneath a planar water surface. We sample bubble radii from 0.25mm to 1cm in 0.25mm increments, and depths down to 8cm in 0.5mm increments. For each radius and depth, we solve the interior frequency BVP and exterior radiation BVP. Then we fit a multipole expansion [Zheng and James 2010] to the bubble’s exterior radiation data, and only store the multipole coefficients (as opposed to storing all the boundary data). Further, we assume the bubble’s response is cylindrically symmetrical, so we only store the zonal multipole coefficients.

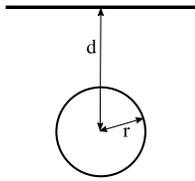


Figure 10: Fast proxy transfer model

³We refer to the pressure gradient as a velocity, but they are merely proportional: $\partial_n p = -i\omega\rho V_n$ for a harmonic vibration of velocity V_n . Since the normal velocity is continuous across the fluid-air interface, $-i\omega V_n = \partial_n p_a / \rho_a = \partial_n p_f / \rho_f$, and so up to a multiplicative constant density factor we can say that $\partial_n p_a$ and $\partial_n p_f$ are equal velocity-like quantities.

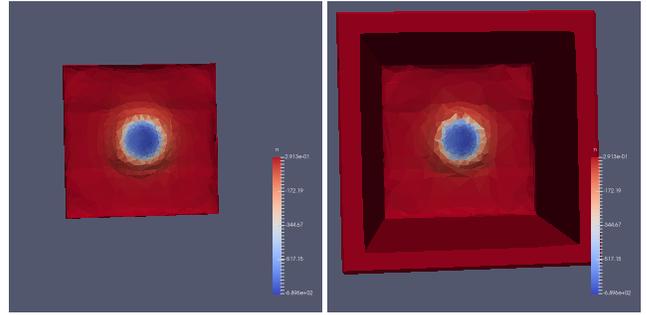


Figure 11: Interpolating velocity BC data for a rising bubble in a square container (top view): (Left) velocity BCs from the interior solve geometry (on Γ_a) are (Right) interpolated onto the mesh (of $\Gamma_a \cup \Gamma_r$) for the exterior Helmholtz BVP.

At runtime, instead of solving the exterior BVP, we compute the equivalent spherical radius of the bubble, and the distance to the fluid surface. These values are used to lookup the closest set of multipole coefficients in our database, which are evaluated at the listening position. When performing lookups in our database, we first find the closest radius, then find the closest depth sample with that radius. It took 2.2 hours to construct the database using 32 cores. Runtime evaluation times for each example are given in Table 1.

7 Bubble Forcing

There are several mechanisms which can drive bubble vibrations. Previous fluid sound work in graphics has used Laplace pressure forcing, which approximates entrainment forcing by a pressure jump due to surface tension. The additional pressure, $p_\sigma = \frac{2\sigma}{R}$ (R is the mean bubble radius), provides an initial impulse to the bubble oscillator (possibly smoothed in time). However, estimates of the Laplace pressure jump, as well as hydrostatic pressure and shape mode coupling effects, show that they represent a minor ($< 10\%$) amount of the total forcing [Pumphrey and Ffowcs Williams 1990]. A recent set of papers from Deane and Czerski [2008; 2010; 2011; 2011] propose a family of models based on neck collapse (for entrainment and splitting events) and neck expansion (for merging events) where surface tension effects account for the majority of forcing (summarized in Figure 12).

Entrainment: The forcing of bubbles released from an underwater tube were analyzed in [Deane and Czerski 2008; Czerski and Deane 2010]. As the bubbles separate from the nozzle, a conical neck forms with very sharp curvature (causing high surface tension) at the tip. At separation, surface tension causes the neck to rapidly shoot into the bubble, quickly decreasing the bubble’s volume, which forces oscillations. A similar situation happens during bubble entrainment: a neck is formed as a bubble separates from

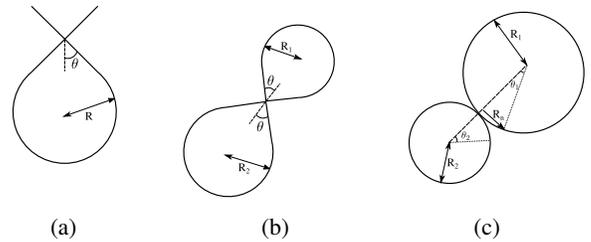


Figure 12: Bubble forcing: There are 3 types of forcing events we model: (a) entrainment, (b) splitting, and (c) merging.

the fluid surface. This geometry is illustrated in Figure 12a. Their time-dependent forcing model based on a conical neck is

$$f(R, t) = -\frac{9\kappa\sigma\eta p_{in,0}\sqrt{1+\eta^2}}{4\rho R^3}t^2, \quad (14)$$

where κ is the gas polytropic index, σ is the surface tension coefficient, $p_{in,0}$ is the equilibrium pressure in the bubble, R is the equilibrium radius, ρ is the fluid density, and η is the slope of the neck cone ($\eta = \tan(\theta)$). We have multiplied their original equation by ρR^2 to transform to the correct units (pressure). We sample the neck angle uniformly between $[20^\circ, 55^\circ]$, which corresponds to $\eta \in [0.36, 1.43]$. It was observed that this forcing function is valid for approximately 300-400 μs for 2mm bubbles (Minnaert frequency = 1.6 kHz, period = 625 μs). For a bubble with period τ , we use this forcing function for the min ($\frac{1}{2}\tau, 600\mu\text{s}$), and set it to 0 afterwards.

Splitting: For bubble splitting events, the same mechanism is proposed in [Czerski and Deane 2011] as the main source of forcing. Both child bubbles are driven by the neck that forms during pinchoff (illustrated in Figure 12b). We again use equation (14), and again assume it is active for min ($\frac{1}{2}\tau, 600\mu\text{s}$) for each bubble. The same sampling method for η is also used.

Merging: For bubble coalescence events, a similar mechanism is proposed in [Czerski 2011]. When two bubbles merge, very strong surface tension is generated at the merge point. This causes the bubble surface to expand, rapidly increasing the volume of the new bubble (see Figure 12c). Their model is

$$f(t, R) = \frac{6\sigma\kappa p_{in,0}}{\rho R^3}t^2, \quad (15)$$

where again we have multiplied their original equation by ρR^2 to transform to pressure. The time that this forcing function is active is labeled t_{lim} , and is defined as the amount of time it takes for the expanding radius to reach a fixed fraction (sampled uniformly between $[0.4, 0.8]$) of the smaller bubble radius. An arbitrary modulation function is also added, resulting in the forcing function

$$f(t, R) = \left[\frac{1}{2} - \frac{1}{\pi} \tan^{-1} \left(3 \frac{t - t_{lim}}{t_{lim}} \right) \right] \frac{6\sigma\kappa p_0}{\rho R^3} t^2.$$

We again limit the forcing to min ($t_{lim}, 600\mu\text{s}$).

8 Sound Synthesis

The last part of our pipeline (after fluid simulation, bubble identification and tracking, frequency estimation, and radiation analysis) is to synthesize the resulting sound at the listening position. While conceptually similar to [Zheng and James 2009], there are several different and important details in our approaches.

8.1 Audio Synthesis Details

Culling silent bubbles: To avoid unnecessary frequency and radiation solves for inaudible bubbles, we cull silent bubbles. Specifically, we do not perform any solves for bubbles that are older than $-\ln(.01)/\beta$, where β is estimated from the bubble's equivalent spherical radius and Minnaert frequency.

Oscillator tracking: To avoid discontinuities in the synthesized sound, we continue oscillators through split and merge events. During split events, the parent bubble's oscillator continues to the largest child bubble. During merge events, the largest parent bubble's oscillator continues to the child bubble. We use the standard RK4

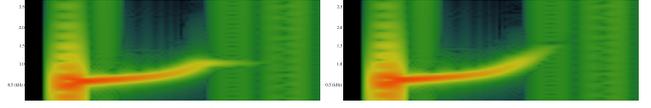


Figure 13: Bubble frequency extension model: (Left) a bubble collapses before it has finished oscillating, resulting in an audible sample-and-hold frequency artifact. (Right) To improve the approximation we extrapolate the frequency using a fitted exponential.

method to integrate the oscillator equations, and did not need any special treatment to handle abrupt frequency changes.

Simultaneous events: A single bubble can undergo multiple events during a single timestep of our simulation. For example, a near-surface bubble could split and one of the daughter bubbles could touch the surface and disappear. Similarly, a bubble could be entrained and merge with another bubble during one time step. To handle these volume changes and missed events correctly, we monitor each bubble's volume during bubble tracking. If there is a sudden increase, we add a new bubble entrainment event and merge it immediately. When there is a sudden volume decrease, we add a split event followed immediately by a collapse event.

8.1.1 Frequency Extension

For various reasons such as resolution limits or the idealized damping model, simulated bubbles may reach the fluid surface and collapse before they are done oscillating. This singularity can limit our ability to resolve the bubble's chirp-like frequency response. To avoid such artifacts, we fit a small exponential model, ae^{ct} , to the frequency samples, and use it to extrapolate the bubble's frequency in time (see Figure 13). Parameters a and c are calculated to ensure C_0 and C_1 continuity.

8.1.2 Bubble Popping Sound Model

Bubbles that "pop" at the surface have their oscillators die out, but in reality there is a characteristic chirp-like popping sound due to the small pressurized cavity that briefly rings like a so-called Helmholtz resonator [Spiel 1992]. These sounds are quiet, but occupy a part of the audio spectrum which is otherwise quiet in our model due to the predominantly lower pitch of bubbles resolvable by our fluid simulator. The physical process of bubble bursting and aerosol generation is actually terribly complicated [Lhuissier and Villermaux 2012]. We use a simplified model for the sound produced by bursting bubbles proposed for remote acoustic sensing of important bubble properties, such as bubble cap film thickness [Deane 2013].

The model gives the time dependent frequency of a bubble pop as (equation 5 in [Deane 2013])

$$f_H(t) = \frac{c}{2\pi} \sqrt{\frac{3\pi^2}{16V} R \sin\left(\frac{ut}{R}\right)}, \quad 0 \leq t \leq t_{max},$$

where u is the velocity of the retracting bubble film, R is the bubble cap radius, and V is the bubble volume. t_{max} is the time it takes for the film to retract fully. The film retraction velocity can be estimated from the length of time the film has been draining for before it nucleates. We define the minimum drain time t_{min} as the time it would take for the film thickness to reach $R/10$.

When a bubble reaches the surface in our simulation, we uniformly sample the drain time between $[t_{min}, 20ms]$ to define u and synthesize a cosine chirp with frequency $f_H(t)$. We modulate the pop sound with an ad-hoc function chosen to match experimental data

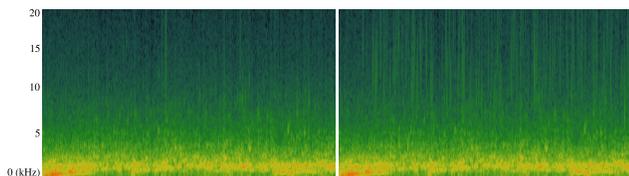


Figure 14: Bubble popping sounds add additional high-frequency content as shown here by comparing spectrograms of sounds produced (Left) without and (Right) with the popping sound model based on [Deane 2013]. The simulation example is the **pouring faucet**.

from [Deane 2013], namely,

$$mod(t) = \exp\left(\frac{\ln(.00001)}{t_{max}}t\right) \frac{2}{\pi} \operatorname{atan}\left(\frac{t}{2t_{max}}\right).$$

The amplitude of the pop is not well defined. Given the maximum absolute value of the corresponding bubble sound s_{max} , we choose to scale the pop sound so that it has a maximum magnitude of $a s_{max}$, where a is uniformly sampled between $[.001, .03]$. The effect of the popping sound model is shown in Figure 14.

8.1.3 Bubble Bandwidth-Extension Scheme

Given the resolution limits of our fluid simulator, bubbles below a certain length scale can not be resolved correctly, resulting in a band-limited frequency response of the bubble oscillator model. Experimental studies of bubble populations in breaking waves have established various bubble size statistics, and have shown that the number of tiny bubbles tends to follow power-law models [Deane and Stokes 2002]. To artificially extend the frequency response of our renderings, we optionally seed audio-domain bubble events from a power-law distribution as follows.

We sample tiny bubbles in the audio domain, based on simulated larger bubbles. Specifically, for each entrained simulation bubble with radius $r_{parent} \geq 2mm$, we assume that the impact which created this bubble also generated other smaller bubbles with radii $r_{tiny} \in [0.1mm, 1mm]$. The number of artificial bubbles generated for each simulation bubble is uniformly sampled between $[0, 3000 * r_{parent}]$. Given the simulation bubble’s creation time t , the start times for each of the artificial bubbles are uniformly sampled from $[t - 0.1, t]$ (because the tiny bubbles are created during the impact). The sizes of the tiny bubbles are sampled from a $-3/2$ power law, consistent with observations of the distribution of bubbles below the Hinze scale. Finally, since we have no geometry or positions for these artificial bubbles, we base their amplitude on the parent bubble. Given the parent bubble’s transfer magnitude p_{parent} , we set the transfer value for a tiny bubble to $50 p_{parent} (r_{tiny})^{1/3}$.

8.2 Sound Synthesis Summary

In summary, our sound synthesis pipeline is very similar to previous work, with the important details that we need to track oscillators through split and merge events to avoid discontinuities, and simultaneous events need to be treated correctly. We proposed several audio domain methods to help add missing detail from our simulations, including frequency extension, a bubble popping model, and a microbubble model. The three latter models are all optional.

9 Results and Discussion

Please see our accompanying video for visual and audio results.

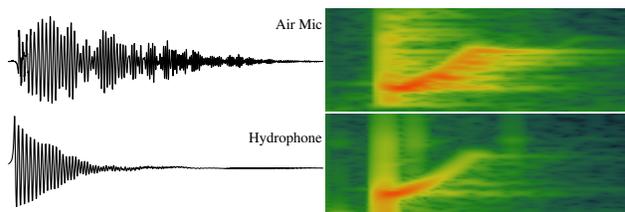


Figure 15: Container effects: Container effects can be strong in the air (top), while in a simultaneous hydrophone recording (bottom) the waveform of the same entrained bubble is much cleaner. Resonances of the container can be seen as lines in the spectrum, which continue after the bubble has popped.

Results were computed on a heterogeneous cluster of 31 nodes, where each node’s core count ranged from 8 to 64 (Intel Xeon X5355 and Xeon X7560 processors), as well as on the NSF cluster SuperMIC (two Intel Xeon E5-2680 processors per node). We report runtimes and simulation statistics in Table 1.

9.1 Discussion of Tank Effects

Container effects such as reverberance can be significant. An example is shown in Figure 15. When a single bubble is entrained in a glass fish tank, strong echo can be seen in the air microphone signal. However, the sound recorded by a hydrophone of the same event looks much closer to a theoretical damped harmonic oscillator.

9.2 Validation

We performed several experiments to validate our frequency model.

Single bubble entrainment: We recorded several bubbles entrained by droplets, and simulated a similar entrainment case. We are able to capture the characteristic frequency chirp of the bubble as it rises (see Figure 16).

Underwater bubble creation: Using a syringe and plastic tubing, we also released underwater bubbles and recorded their emissions. As the bubble moves away from the tube (rigid surface), there is a slight pitch increase, but we do not see the characteristic chirp because the bubble finishes vibrating before it reaches the surface.

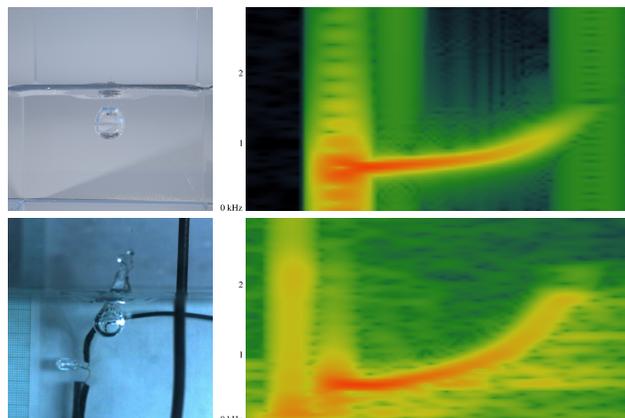


Figure 16: Single Bubble Entrainment: A simulated bubble entrainment event (top) produces a similar spectrum to a recorded entrainment event (bottom).

Example	Domain size (cm)	Length (s)	Simulation time (hours / cores)	# of bubbles / # of solves	Frequency solve time (hours)	Amortized speedup	Radiation solve time (hours)	Proxy evaluation time (hours)	% culled
Dripping Faucet	8 x 8 x 24	9.0	97 / 32	153 / 965	.005	0.81	.07	0.09	6.4
Pouring Faucet	8 x 8 x 24	8.5	402 / 64	331521 / 585311	23	7.35	52	3.8	71.2
Water Step	8 x 24 x 24	4.5	1000 / 96	420134 / 483654	20	4.85	44	20.0	71.8
Dam Break	16 x 16 x 32	2.64	394 / 64	114471 / 121646	2.9	4.72	15	0.66	72.4
Armadillo Drop	16 x 16 x 32	4.0	293 / 64	13981 / 11653	.245	3.82	3.4	0.58	87.9

Table 1: Results. Our fluid simulations used different numbers of cores, which are reported above. The frequency and radiation solves are massively parallel, and were computed using 680 cores. Proxy transfer evaluation was done on a single core, but could be parallelized easily.

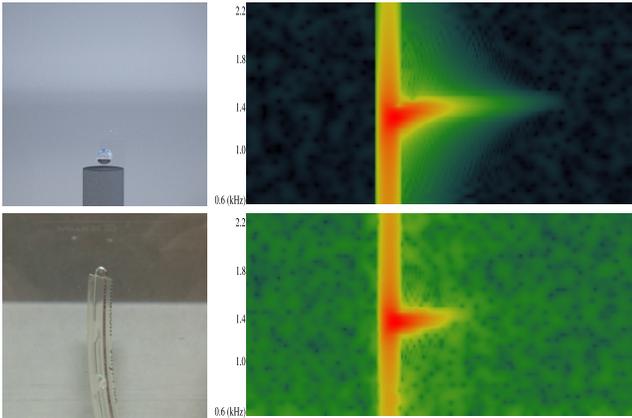


Figure 17: Underwater bubble release: A bubble released from an underwater tube shows a slight frequency rise as it moves away from the rigid tube. The simulation (top) matches well with experiment (bottom).

A simulated scenario of a bubble moving away from an underwater tube produces a similar effect.

9.3 Large Results

Dripping faucet: Our dripping faucet example has a relatively small number of bubbles. It clearly demonstrates the pitch rise as bubbles approach the surface, as well as the benefit of our frequency extension model.

Dam break: There is a loud, low frequency sound produced by the large tubular bubble in the dam break example. The sloshing sounds in this example are convincing. The dam break clearly demonstrates the benefit of our microbubble model, which adds higher frequency texture.

Water step: The water step has the most bubbles of all our simulations. While it has a fairly constant sound spectrum, the differences of our multiple models can be seen. This example also demonstrates the importance of transfer, as it sounds very dissonant when rendered without transfer.

Pouring faucet: We use the same domain as the dripping faucet, but this time use a constant 1cm radius stream to fill the container. This example most clearly shows the importance of transfer, allowing us to capture the characteristic pitch shift as the container fills.

Armadillo drop: For our last example, we dropped a water shaped armadillo into a pool of water. While this example is fun, it also highlights some of the deficiencies of our system. Even with our microbubble model, there is not much of an impact sound during the initial impact of the armadillo. A microbubble model based on

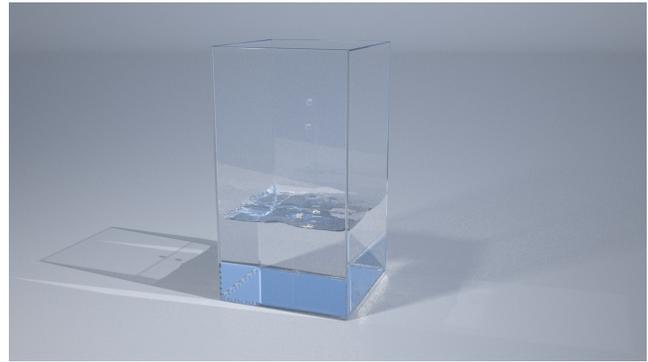


Figure 18: Dripping Faucet

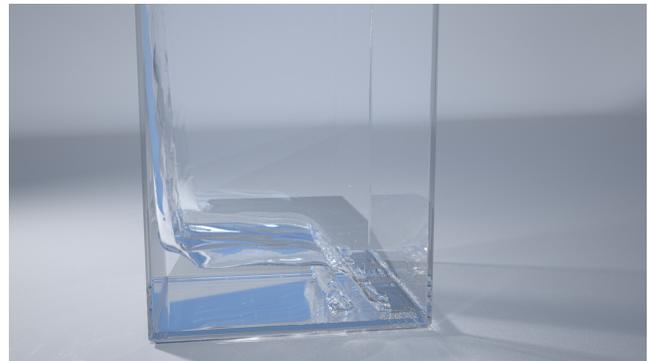


Figure 19: Dam Break



Figure 20: Water Step

droplet impacts, instead of on larger bubbles entrainment times, may help.

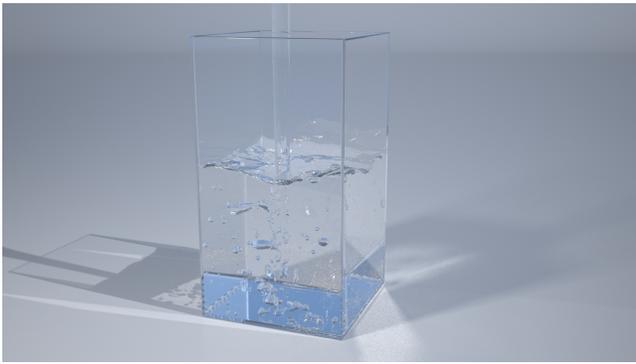


Figure 21: Pouring Faucet

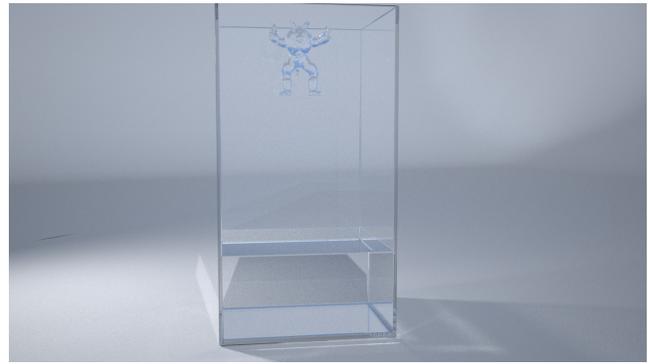


Figure 22: Armadillo

10 Conclusion

We have explored many stages of audiovisual fluid simulation, and identified the need for, and proposed, numerous sound simulation models. We believe that increased resolution of liquid sound generation mechanisms will also lead to improved visual fidelity of fluids in computer animation.

One surprising finding of our bubble frequency model was the large variations predicted based on spatial proximity to boundaries (see Figure 4), whereas the nonspherical pitch variations were comparatively more modest and less perceptually important for sound rendering. As Figure 7 demonstrates, there were some stong pitch variations due to shape, but these seemed rare. Perhaps the most important aspect of nonspherical bubbles was their ability to conform to the fluid-air interface and produce large pitch increases for rising bubbles due to the thin-plate capacitor effect (see Figure 6).

Our frequency model provides a way to capture the complex frequency effects that acoustic bubbles exhibit, and we hope that it will lead to better audiovisual simulation of water in the future, as well as more accurate methods for passive acoustic sensing. Our frequency model can be solved efficiently, and is surprisingly robust to coarse meshes (see Figure 8 top).

Limitations and future work: Far from solving the problem outright, the current study identifies many challenges, limitations, and opportunities for future work in realistic audiovisual fluid simulation. Our formulations calculate frequency and radiation based on an independent bubble assumption; however, bubbles can affect each other to produce frequency-coupled vibrations and in turn change their acoustic emissions. Our frequency-domain transfer model can capture resonance effects in containers and thereby improve the sound quality over previous work [Zheng and James 2009; Moss et al. 2010]. But it misses perceptually significant time-domain reverberation effects. In addition, the container sound is missing, which can be an important source of fluid-solid coupling sounds. Also, we only model the dependence of ω on local geometry, but bubble damping β may also be affected. Acoustic radiation, while perceptually crucial, is computationally expensive. Our bubble-plane proxy is one attempt at performance improvement, but other techniques such as fast-multipole methods, FDTD with PML on the GPU, or geometric acoustics are interesting future work directions.

Our sound model is inherently band-limited due to resolution-limited fluid simulation. Post-processing effects such as popping and microbubble models can help. But it is also desired to incorporate more advanced multi-scale models in our system. The bubble forcing models we use are relatively simple, typically only captur-

ing asymptotic forcing information, but assume idealized geometry and involve undetermined parameters. It is therefore interesting to explore data-driven approaches for automatically extracting forcing models and their parameters. Laboratory experiments reveal a wide-range of surface related sound events which do not radiate effectively underwater (as measured by a hydrophone), but that produce significant high-frequency sound contributions. Methods for approximating these processes, either in the geometric domain or audio domain, are necessary to improve sound realism. Finally, there are numerous important potential applications of fluid-sound synthesis in other fields of science and engineering that should be explored.

Acknowledgements

We would like to acknowledge the help of Arthur Sams on early BEM computations, Andrew Wolfers for laboratory measurement of fluid sounds, Hurf Sheldon for Cornell Measurement Lab support, Joel Gross and Pu Zhang for fluid sound discussions, Landon Boyd and Robert Bridson for sharing *MultiFLIP* software, the developers of the *Gerris* open source software, academic licensing of *Maxwell Renderer*, and early collaborations with Phaedon Sinis, Katherine Breeden, Lindsey Williams (who first suggested *Gerris*) and Pat Hanrahan. We thank the anonymous reviewers for their constructive feedback. We acknowledge funding and support from the National Science Foundation (HCC-0905506, CAREER-1453101, DGE-1144153), Intel, donations from Pixar, Adobe, and XSEDE compute resources. Initial phases of this research were conducted in conjunction with the Intel Science and Technology Center–Visual Computing. 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.

References

- ANDO, R., THUREY, N., AND WOJTAN, C. 2015. A stream function solver for liquid simulations. *ACM Transactions on Graphics (TOG)* 34, 4, 53.
- BANERJEE, P. K., AND BUTTERFIELD, R. 1981. *Boundary element methods in engineering science*, vol. 17. McGraw-Hill London.
- BIGG, G. R., JICKELLS, T. D., LISS, P. S., AND OSBORN, T. J. 2003. The role of the oceans in climate. *International Journal of Climatology* 23, 10, 1127–1159.

- BOYD, L., AND BRIDSON, R. 2012. MultiFLIP for energetic two-phase fluid simulation. *ACM Transactions on Graphics (TOG)* 31, 2, 16.
- BRAGG, S. W. H. 1920. *The World of Sound*. G. Bell and Sons Ltd.
- BRIDSON, R. 2008. *Fluid simulation for computer graphics*. CRC Press.
- BUSARYEV, O., DEY, T. K., WANG, H., AND REN, Z. 2012. Animating bubble interactions in a liquid foam. *ACM Transactions on Graphics (TOG)* 31, 4, 63.
- CHICHARRO, R., AND VAZQUEZ, A. 2014. The acoustic signature of gas bubbles generated in a liquid cross-flow. *Experimental Thermal and Fluid Science* 55, 221–227.
- CZERSKI, H., AND DEANE, G. B. 2010. Contributions to the acoustic excitation of bubbles released from a nozzle. *The Journal of the Acoustical Society of America* 128, 5, 2625–2634.
- CZERSKI, H., AND DEANE, G. B. 2011. The effect of coupling on bubble fragmentation acoustics. *The Journal of the Acoustical Society of America* 129, 1, 74–84.
- CZERSKI, H. 2011. A candidate mechanism for exciting sound during bubble coalescence. *The Journal of the Acoustical Society of America* 129, 3, EL83–EL88.
- DA, F., BATTY, C., WOJTAN, C., AND GRINSPUN, E. 2015. Double bubbles sans toil and trouble: discrete circulation-preserving vortex sheets for soap films and foams. *ACM Transactions on Graphics (TOG)* 34, 4, 149.
- DEANE, G. B., AND CZERSKI, H. 2008. A mechanism stimulating sound production from air bubbles released from a nozzle. *The Journal of the Acoustical Society of America* 123, 6, EL126–EL132.
- DEANE, G. B., AND STOKES, M. D. 2002. Scale dependence of bubble creation mechanisms in breaking waves. *Nature* 418, 6900, 839–844.
- DEANE, G. B., AND STOKES, M. D. 2010. Model calculations of the underwater noise of breaking waves and comparison with experiment. *The Journal of the Acoustical Society of America* 127, 6, 3394–3410.
- DEANE, G. B. 2013. Determining the bubble cap film thickness of bursting bubbles from their acoustic emissions. *The Journal of the Acoustical Society of America* 133, February 2013, EL69–75.
- DING, L., AND FARMER, D. M. 1994. Observations of breaking surface wave statistics. *Journal of Physical Oceanography* 24, 6, 1368–1387.
- ENRIGHT, D., MARSCHNER, S., AND FEDKIW, R. 2002. Animation and rendering of complex water surfaces. *ACM Transactions on Graphics (TOG)* 21, 3, 736–744.
- GARLAND, M., AND HECKBERT, P. S. 1997. Surface simplification using quadric error metrics. In *Proceedings of the 24th Annual Conference on Computer Graphics and Interactive Techniques*, ACM Press/Addison-Wesley Publishing Co., New York, NY, USA, SIGGRAPH '97, 209–216.
- HAGER, W. W. 1989. Updating the inverse of a matrix. *SIAM review* 31, 2, 221–239.
- HONG, J.-M., AND KIM, C.-H. 2003. Animation of bubbles in liquid. In *Computer Graphics Forum*, vol. 22, Wiley Online Library, 253–262.
- HONG, J.-M., AND KIM, C.-H. 2005. Discontinuous fluids. *ACM Transactions on Graphics (TOG)* 24, 3, 915–920.
- HOWE, M. S., AND HAGEN, N. A. A. 2011. On the impact noise of a drop falling on water. *Journal of Sound and Vibration* 330, 4, 625–635.
- HOWE, M. S. 1998. *Acoustics of Fluid-Structure Interactions*. Cambridge Press.
- HOWE, M. S. 2002. *Theory of Vortex Sound*. Cambridge Press.
- KIM, B., LIU, Y., LLAMAS, I., JIAO, X., AND ROSSIGNAC, J. 2007. Simulation of bubbles in foam with the volume control method. In *ACM Transactions on Graphics (TOG)*, vol. 26, ACM, 98.
- KLUSEK, Z., AND LISIMENKA, A. 2013. Acoustic noise generation under plunging breaking waves. *Oceanologia* 55, 4, 809–836.
- LAMB, H. 1932. *Hydrodynamics*. Cambridge University Press.
- LEIGHTON, T., FINFER, D., GROVER, E., AND WHITE, P. 2007. An acoustical hypothesis for the spiral bubble nets of humpback whales and the implications for whale feeding. *Acoustic Bulletin* 22, 1, 17–21.
- LEIGHTON, T. G. 1994. *The Acoustic Bubble*. Academic Press.
- 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 (december), 1–15.
- LHUISSIER, H., AND VILLERMAUX, E. 2012. Bursting bubble aerosols. *Journal of Fluid Mechanics* 696 (4), 5–44.
- LONGUET-HIGGINS, M. S. 1989. Monopole emission of sound by asymmetric bubble oscillations. Part 1. Normal modes. *Journal of Fluid Mechanics* 201, 525–541.
- LONGUET-HIGGINS, M. S. 1989. Monopole emission of sound by asymmetric bubble oscillations. Part 2. An initial-value problem. *Journal of Fluid Mechanics* 201, 543–565.
- LONGUET-HIGGINS, M. S. 1990. An analytical model of sound production by raindrops. *Journal of Fluid Mechanics* 214, 395–410.
- LONGUET-HIGGINS, M. S. 1991. Resonance in nonlinear bubble oscillations. *Journal of Fluid Mechanics* 224, 531.
- LOSASSO, F., SHINAR, T., SELLE, A., AND FEDKIW, R. 2006. Multiple interacting liquids. In *ACM Transactions on Graphics (TOG)*, vol. 25, ACM, 812–819.
- MEANS, S. L., AND HEITMEYER, R. M. 2002. Surf-generated noise signatures: a comparison of plunging and spilling breakers. *The Journal of the Acoustical Society of America* 112, 2, 481–488.
- MEDWIN, H., AND BEAKY, M. M. 1989. Bubble sources of the Knudsen sea noise spectra. *The Journal of the Acoustical Society of America* 86, 3, 1124–1130.
- MINNAERT, M. 1933. On musical air-bubbles and the sounds of running water. *Philosophical Magazine* 16, 104, 235–248.
- MOSS, W., YEH, H., HONG, J.-M., LIN, M. C., AND MANOCHA, D. 2010. Sounding liquids: Automatic sound synthesis from fluid simulation. *ACM Trans. Graph.* 29, 3 (July), 21:1–21:13.

NYSTUEN, J. A. 1986. Rainfall measurements using underwater ambient noise. *The Journal of the Acoustical Society of America* 79, 4, 972–982.

OSHER, S., AND FEDKIW, R. 2006. *Level set methods and dynamic implicit surfaces*, vol. 153. Springer Science & Business Media.

POPINET, S. 2003. Gerris: a tree-based adaptive solver for the incompressible euler equations in complex geometries. *Journal of Computational Physics* 190, 2, 572–600.

POPINET, S. 2009. An accurate adaptive solver for surface-tension-driven interfacial flows. *Journal of Computational Physics* 228, 16, 5838–5866.

PROSPERETTI, A. 1988. Bubble dynamics in oceanic ambient noise. In *Sea Surface Sound*, B. Kerman, Ed., vol. 238 of *NATO ASI Series*. Springer Netherlands, 151–171.

PROSPERETTI, A. 1988. Bubble-related ambient noise in the ocean. *The Journal of the Acoustical Society of America* 84, 3, 1042–1054.

PUMPHERY, H., CRUM, L., AND BJØRNØ, L. 1989. Underwater sound produced by individual drop impacts and rainfall. *J. Acoust. Soc. Am.* 85, 1518–1526.

PUMPHREY, H., AND FLOWCS WILLIAMS, J. 1990. Bubbles as sources of ambient noise. *IEEE Journal of Oceanic Engineering* 15, 4, 268–274.

RAYLEIGH, L. 1917. VIII. On the pressure developed in a liquid during the collapse of a spherical cavity. *Philosophical Magazine Series 6* 34, 200, 94–98.

SMIGAJ, W., ARRIDGE, S., BETCKE, T., PHILLIPS, J., AND SCHWEIGER, M. 2012. Solving boundary integral problems with bem++. *ACM Transactions on Mathematical Software*.

SPIEL, D. E. 1992. Acoustical measurements of air bubbles bursting at a water surface: Bursting bubbles as helmholtz resonators. *Journal of Geophysical Research: Oceans (1978–2012)* 97, C7, 11443–11452.

SPRATT, K. S., LEE, K. M., WILSON, P. S., AND WOCHNER, M. S. 2015. On the resonance frequency of an ideal arbitrarily-shaped bubble. *Proceedings of Meetings on Acoustics* 20, 1.

STAM, J. 1999. Stable fluids. In *Proceedings of the 26th annual conference on Computer graphics and interactive techniques*, ACM Press/Addison-Wesley Publishing Co., 121–128.

STRASBERG, M. 1953. The pulsation frequency of nonspherical gas bubbles in liquids. *The Journal of the Acoustical Society of America* 25, 3, 536–537.

VAN DEN DOEL, K. 2005. Physically based models for liquid sounds. *ACM Transactions on Applied Perception* 2, 4 (Oct.), 534–546.

WILSON, J. D., AND MAKRIS, N. C. 2008. Quantifying hurricane destructive power, wind speed, and air-sea material exchange with natural undersea sound. *Geophysical Research Letters* 35, 10.

ZHENG, C., AND JAMES, D. L. 2009. Harmonic fluids. *ACM Trans. Graph.* 28, 3 (Aug.), 37:1–37:12.

ZHENG, C., AND JAMES, D. L. 2010. Rigid-Body Fracture Sound with Precomputed Soundbanks. *ACM Transactions on Graphics* 29, 4 (July), 69:1–69:13.

ZHENG, W., YONG, J.-H., AND PAUL, J.-C. 2009. Simulation of bubbles. *Graphical Models* 71, 6, 229–239.

A Fast Amortized BEM Solver Details

Here we provide details (from §5.4) on how to efficiently evaluate $\mathbf{x} = \mathbf{A}^{-1}\mathbf{b}$ in our frequency solver for the common case of multiple bubbles (with $n_b \ll n_a + n_b$). First, to clarify, we avoided the expensive inverse in the U block of \mathbf{A}^{-1} expression (12) using the Sherman-Morrison-Woodbury low-rank update formula [Hager 1989] to obtain

$$U = (D - CG_{bb}^{-1}B)^{-1} \quad (16)$$

$$= D^{-1} + D^{-1}C \underbrace{(G_{bb} - BD^{-1}C)^{-1}}_X BD^{-1} \quad (17)$$

$$= D^{-1}(I + C \underbrace{XBD^{-1}}_{-Y}) = D^{-1}(I - CY). \quad (18)$$

The bubble-independent solver setup constructs the air-solid domain form-factor matrix D , and the LU factorization for D^{-1} . Then given a specific bubble geometry, we construct the remaining blocks of A , and then compute

$$G_{bb}^{-1} \leftarrow \text{LUSolver}(G_{bb}) \quad (19)$$

$$T_1 \leftarrow D^{-1}C \quad (n_b \text{ LU solves for } C \text{ RHS}) \quad (20)$$

$$X \leftarrow \text{LUSolver}(G_{bb} - BT_1). \quad (\text{then discard } T_1) \quad (21)$$

Then given the RHS vector (in block form) $\mathbf{b} = \begin{pmatrix} b_b \\ b_{as} \end{pmatrix}$ we can form

$\mathbf{x} = \mathbf{A}^{-1}\mathbf{b} = \begin{pmatrix} x_b \\ x_{as} \end{pmatrix}$ as follows. Since

$$x_b = Xb_b + Yb_{as} = X(b_b - BD^{-1}b_{as}), \quad (22)$$

we can form x_b via (here t_* are temporary variables)

$$t_1 \leftarrow D^{-1}b_{as} \quad (\text{LU solve}) \quad (23)$$

$$t_2 \leftarrow b_b - Bt_1 \quad (24)$$

$$\boxed{x_b} \leftarrow Xt_2. \quad (\text{LU solve}) \quad (25)$$

Second, observe that

$$x_{as} = Zb_b + Ub_{as} = U(-CG_{bb}^{-1}b_b + b_{as}) \quad (26)$$

so we can form $x_{as} = Ut_2 = D^{-1}(I - CY)t_2$ as

$$t_1 \leftarrow G_{bb}^{-1}b_b \quad (\text{LU solve}) \quad (27)$$

$$t_2 \leftarrow b_{as} - Ct_1 \quad (28)$$

$$t_3 \leftarrow Yt_2 = -X(B(D^{-1}t_2)) \quad (2 \text{ LU solves}) \quad (29)$$

$$t_4 \leftarrow t_2 - Ct_3 \quad (30)$$

$$\boxed{x_{as}} \leftarrow D^{-1}t_4 \quad (\text{LU solve}) \quad (31)$$

Assuming (for simplicity) that $n_b = O(1)$ and $N = n_a + n_b$, then it involves only $O(N)$ new form factor computations per bubble (instead of $O(N^2)$), and it costs $O(N^2)$ to construct this A^{-1} transformation (instead of $O(N^3)$), and $O(N^2)$ to apply it to form $A^{-1}\mathbf{b}$. The main per-bubble cost is $n_b + 3$ applications of the D^{-1} LU solver.