Physically Based Sound for Computer Animation and Virtual Environments

Course Organizer

Doug L. James
Stanford University

Physically based sound is an important emerging approach for computer synthesis of realistic synchronized sounds for physically based animation and real-time virtual environments. A major challenge for learning and implementing these sound techniques is the wide range of physically based models and sound phenomena involved, as well as the need for optimizations. Furthermore many publications on physically based animation and sound rendering can assume mathematical background that many in the graphics community lack. The result is that learning physics-based sound techniques is unnecessarily difficult for many interested students and practitioners.

The main goal of this course is to address the need to make the principles and methods of physically based sound accessible to a broader computer graphics audience. The course will cover sound models for various animated phenomena, and take a hands-on approach to implementing practical systems. By using supporting software and implementation short-cuts, students will be able to start using physically based sound immediately after completing the course. Students should be familiar with mainstream computer graphics, and related mathematics such as basic vector/matrix math and differential equations. Any knowledge of physically based modeling will be helpful, especially for the more advanced topics. No knowledge of sound modeling will be assumed.

This course covers sound source models for sonifying important phenomena from physics-based animation: rigid bodies, brittle fracture, thin-shells, cloth, deformable collisions and contact, fluids and fire. Material related to rigid-body sound is covered in greater detail in the first half of the course, and addresses fundamental topics such as modeling modal vibrations and sound radiation from surfaces. Further readings are given throughout the notes to survey prior work, and to identify other existing approaches. Several recent SIGGRAPH papers will also serve as supplemental notes for more advanced topics.

Course materials and updates can be found at:

`http://graphics.stanford.edu/courses/sound`

Course Schedule

9:00 am **Introduction** **James**

9:05 am **Modal Vibration** **Zheng**

Simple harmonic oscillator, mass-spring systems; modal vibration of 3D solids; mass and stiffness matrices; meshing and discretization; generalized eigenvalue problem; eigenfrequencies and eigenmodes; damping models; time-stepping modal vibrations; integration with rigid-body dynamics engines.

9:25 am **Acoustic Transfer** **James**

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

10:00 am **Implementing Rigid-Body Sound** **Zheng**

Pipeline implementation; supporting libraries; demonstration.

10:15 am **Break**

10:30 am **Acceleration Noise** **James**

Rigid-body acceleration; Hertz contact and impact time scale; acceleration noise pulses; precomputation; rendering details.

10:45 am **Brittle Fracture** **Zheng**

Modeling background; rigid-body sound approximation; fracture impulses; ellipsoidal sound proxies; simplified implementation using pre-scored fracture.

11:00 am **Thin Shells** **James**

Modeling background; nonlinearly coupled modes; subspace integration.

11:10 am **Liquids** **Langlois**

Acoustic bubble models; particle-based bubble simulation; acoustic transfer; speed-accuracy trade-offs; simplified implementation.

11:35 am **Fire** **James**

Modeling background; combustion noise; low-frequency approximation; sound texture synthesis.

11:50 am **Sound in VR** **Mehra**

Challenges for real-time sound in VR; spatial audio; head-related transfer functions (HRTFs), etc.

12:15 pm **Close**

Course Schedule (Brief)

9:00 am	**Introduction**	**James**
9:05 am	**Modal Vibration**	**Zheng**
9:25 am	**Acoustic Transfer**	**James**
10:00 am	**Implementing Rigid-Body Sound**	**Zheng**
10:15 am	**Break**	**–**
10:30 am	**Acceleration Noise**	**James**
10:45 am	**Brittle Fracture**	**Zheng**
11:00 am	**Thin Shells**	**James**
11:10 am	**Liquids**	**Langlois**
11:35 am	**Fire**	**James**
11:50 am	**Sound in VR**	**Mehra**
12:15 pm	**Close**	**–**

Course Speakers

Doug James is a Full Professor of Computer Science at Stanford University (since June 2015), and was previously an Associate Professor of Computer Science at Cornell University (2006-2015), and an Assistant Professor of CS/Robotics at Carnegie Mellon University (2002-2006). He holds three degrees in applied mathematics, including a Ph.D. in 2001 from the University of British Columbia. Doug is a recipient of a National Science Foundation CAREER award, and a fellow of both the Alfred P. Sloan Foundation and the Guggenheim Foundation. He received a 2012 Technical Achievement Award from The Academy of Motion Picture Arts and Sciences for "Wavelet Turbulence," and the 2013 Katayanagi Emerging Leadership Prize from Carnegie Mellon University and Tokyo University of Technology. He was the Technical Papers Program Chair of ACM SIGGRAPH 2015, and is currently a consulting Senior Research Scientist at Pixar Animation Studios.

Timothy Langlois received a B.S. from the University of Massachusetts Amherst (2009), and a Ph.D. from Cornell University's Computer Science Department (2016). His PhD research is on rigid-body sound models, eigenmode compression, and two-phase liquid sound synthesis. He is a recipient of an NSF graduate research fellowship.

Ravish Mehra is a Research Scientist at the Oculus Research lab where he specializes in sound rendering technology for VR. He completed his PhD in Computer Science at the University of North Carolina at Chapel Hill under the supervision of Prof. Dinesh Manocha and Prof. Ming Lin. He received his M.S. degree from UNC Chapel Hill in 2011 and Bachelor's degree in Computer Science and Engineering from Indian Institute of Technology (IIT) Delhi in 2008.

Changxi Zheng is an Assistant Professor in the Computer Science Department of Columbia University where he co-directs Columbia's Computer Graphics Group (C2G2) in the Columbia Vision and Graphics Center (CVGC). He received a Ph.D. from Cornell University's Computer Science Department in 2012; his Ph.D. thesis was on "Physics-Based Sound Rendering for Computer Animation". He has been serving as the Associated Editor of ACM Transactions on Graphics, and won the NSF CAREER Award, the Hot Paper Award at HotWireless 2015, the Cornell Computer Science Best Dissertation award in 2012, and was named as Forbes "30 under 30" in science and healthcare in 2013.

Contents

I — Course Notes

II — Slides

Physically Based Sound for Computer Animation and Virtual Environments
Introduction

Doug L. James
Stanford University

Introduction

Doug L. James
Stanford University

The past few decades have seen a dramatic rise in the use of physically based modeling techniques in computer animation to simulate everything from rigid and deformable bodies to fracturing solids and splashing fluids. The goal of this course is to explain how physics-based techniques can also be used to synthesize accompanying sounds for these otherwise silent animations. There are several reasons for wanting to do so. For example in future interactive virtual environments, computers will simulate physics with increasing realism in unique and unpredictable manners, and it is desirable to automatically compute synchronized sound for these visual effects to provide realistic and immersive virtual experiences. While traditional computer-animated motion picture production has relied on sound post-production to add sounds and sound effects to films, the rapid increase in 3-D content creation makes it desirable to support automatic computer-generated rendering of visual and auditory stimuli, i.e., just hit the "render" button. Finally enabling sound rendering for physics-based models can enable wide range of unexpected applications, analogous to the creative uses found for computer graphics rendering.

In this course we take a physics-based approach to modeling sounds produced by computer-animated phenomena. We are primarily concerned with *sound source modeling*, namely the process by which some animated phenomena generate sound waves that a virtual listener can hear. The full sound rendering process is more complicated, and involves not only *sound generation* by our animation, but also *sound propagation* through a 3-D environment, as well as modeling of the *listening* process. The latter two stages have been widely considered by prior work, and form the basis for listening to sounds in virtual environments, e.g., a recorded sound emanating from a point in a room. In this course we will focus primarily on the sound generation process, and we refer readers interested in sound propagation and listening to previous materials on those topics.

This course will consider several aspects of sound generation, beginning with fundamentals of vibration and sound, and then later addressing more advanced animated phenomenon. We begin by first considering vibrations of solids, and the important special case of linear modal vibrations which describes small vibrations of stiff objects. Next we consider the process by which sound pressure waves are generated in a surrounding ambient fluid via coupling to the surface accelerations of the solid. We consider two important cases associated with periodic time-harmonic vibrations which can be modeled in the frequency domain, as well as transient and nonperiodic surface accelerations which are more easily modeled in the time domain. Practical details of how to compute solid vibrations and the associated soundwaves generated will be considered in detail. The important special case of rigid-body dynamics, which makes extensive use of linear motor vibration and precomputed sound generation models, will be thoroughly discussed. Various demonstrations and computer codes for exploring these phenomena will be explored.

More advanced phenomena build upon these vibration and sound wave modeling fundamentals. Brittle fracture sounds will be considered using rigid-body sound models. Sounds produced by thin shells tend to have a noisy character and will be explored using nonlinearly coupled vibration modes. Sounds produced by splashing fluids can be modeled by considering the vibrations of the fluid surface produced by tiny acoustic bubbles. We will also consider sounds produced by fire by considering the tiny pressure fluctuations produced by combustion. These more advanced topics will be supplemented by recent SIGGRAPH papers on the same topics.

A major challenge in offering this course is balancing the complexity of the various animated phenomena with ease of understanding. Our goal is that students will be able to understand and implement basic material on solid vibration and sound wave generation so as to be able to build a practical rigid body sound pipeline, and get started immediately using physics-based sound models in computer animation. Additionally we hope to expose students to sound models for more advanced animated phenomena. Shortcuts for simplified practical implementations will be provided whenever possible to help students "get things working" as soon as possible.

Real-time implementations can be explored in many cases, provided that sufficient trade-offs in the physics-based model are made, or more sophisticated preprocessing techniques are used. However the main focus of the course will be on building physics-based sound models which provide realistic results, along with describing practical implementations of these models.