Practical Physically-Based Shading in Film and Game Production

SIGGRAPH 2012 Course

Course Organizers

Stephen Hill
Stephen McAuley
Ubisoft Montreal

Presenters

Brent Burley
Walt Disney Animation Studios

Yoshiharu Gotanda
tri-Ace

Naty Hoffman
Activision Studio Central

Adam Martinez
Sony Pictures Imageworks

Brian Smits
Pixar Animation Studios

Course Description

Physically-based shading is becoming of increasing importance to both film and game production. By adhering to physically-based, energy-conserving shading models, one can easily create high quality, realistic materials that maintain that quality under a variety of lighting environments. Traditional "ad-hoc" models have required extensive tweaking to achieve the same result, thus it is no surprise that physically-based models have increased in popularity in film and game production, particularly as they are often no more difficult to implement or evaluate. After the success of the *Physically-Based Shading Models in Film and Game Production* course at SIGGRAPH 2010, this new course presents two years of advances in the subject. New research in the area will be covered, as well as more production examples from film and game.

The course begins with a brief introduction into the physics and mathematics of shading, before speakers share examples of how physically-based shading models have been used in production. New research is introduced; its practical usage in production explained; then the advantages and disadvantages are discussed. Real-world examples are a particular focus of this year's course, giving attendees a practical grounding in the subject.

LEVEL OF DIFFICULTY: Intermediate

Intended Audience

Practitioners from the videogame, CG animation, and VFX fields, as well as researchers interested in shading models.

Prerequisites

An understanding of shading models and their use in film or game production.

Course Website

All course materials can be found at `http://selfshadow.com/publications/s2012-shading-course`

Contact

Address questions or comments to `s2012course@selfshadow.com`

About the Presenters

BRENT BURLEY is a Principal Software Engineer at Walt Disney Animation Studios working on production rendering software. Recently he led the development of a new physically-based BRDF model now being used in all current productions. Prior to joining Disney in 1996, he worked at Philips Media developing a cross-platform game engine, and also worked on aircraft training simulators at Hughes Training Inc.

YOSHIHARU GOTANDA is the CEO and CTO of tri-Ace, Inc, which is a game development studio in Japan.

STEPHEN HILL is a 3D Technical Lead at Ubisoft Montreal, where his current focus is on physically-based methodologies. He previously held this role on *Splinter Cell Conviction*, where he steered development of the renderer for the entire development period. During that time, he developed systems for dynamic ambient occlusion and visibility.

NATY HOFFMAN is a Technical Director at Activision Studio Central, where he assists Activision's worldwide studios with graphics research and development. Prior to joining Activision in 2008, Naty worked for two years on *God of War III* at SCEA Santa Monica Studio. Naty has also worked at Naughty Dog (where he had an instrumental role in the development of the ICE libraries for first-party PS3 developers), at Westwood Studios (where he was graphics lead on *Earth and Beyond*) and at Intel as a microprocessor architect, assisting in the definition of the SSE and SSE2 instruction set extensions.

ADAM MARTINEZ is a Shader Writer for Sony Pictures Imageworks and a member of the Shading Department, which oversees all aspects of shader writing and production rendering at Imageworks. He is a pipeline developer, look development artist, and technical support liaison for productions at the studio and he is one of the primary architects of Imageworks' rendering strategy behind *2012* and *Alice In Wonderland*.

STEPHEN MCAULEY joined Ubisoft in 2011 after spending 5 years at Bizarre Creations, where he worked on games such as *Blood Stone*, *Blur* and *Project Gotham Racing*, focusing on rendering architecture, physically-based shading and deferred lighting.

BRIAN SMITS has worked at Pixar Animation Studios since 2000. He developed the reflection model used for *WALL-E* and *Up*. He currently works in the RenderMan group. Before Pixar, he was a research professor at the University of Utah working on rendering. He received his PhD in computer science from Cornell University.

Presentation Schedule

09:00–09:05   **Introduction: The Importance of Physically-Based Shading** *(Hill)*

09:05–09:30   **Background: Physically-Based Shading** *(Hoffman)*

09:30–10:00   **Calibrating Lighting and Materials in Far Cry 3** *(McAuley)*

10:00–10:30   **Beyond a Simple Physically-Based Blinn-Phong Model in Real-Time** *(Gotanda)*

10:30–10:45   **Break**

10:45–11:15   **Physical Production Shaders with OSL** *(Martinez)*

11:15–11:45   **Physically-Based Shading at Disney** *(Burley)*

11:45–12:15   **Reflection Model Design for WALL-E and Up** *(Smits)*

Abstracts

Background: Physically-Based Shading

Naty Hoffman

We will go over the fundamentals behind physically-based shading models, starting with a qualitative description of the underlying physics, followed by a quantitative description of the relevant mathematical models, and finally discussing how these mathematical models can be implemented for shading.

Calibrating Lighting and Materials in Far Cry 3

Stephen McAuley

Far Cry 3 is a first-person shooter due to be released in 2012 on Xbox 360, PlayStation 3 and PC. With indoor and outdoor areas, and a real-time day-night cycle, the game engine must support a huge variety of environments. Materials must be made to react correctly under all possible lighting conditions and all possible camera angles. By using physically-based shading, we were able to ensure that every environment would react as expected under any lighting setup, leading to not only a better quality end result, but an improved workflow for artists who were no longer trying to overcome lighting problems.

For materials, we observed that surface albedo is a physically-based property and developed a colour correction tool that could calculate albedo from photographs, using a Macbeth ColorChecker as a reference object. Textures were then built using this colour-corrected photographic reference. Not only did this help with balancing materials under all lighting conditions, but it enabled us to have the vivid colours required of a tropical island. To calibrate lighting, we generated second-order spherical harmonics from the sky dome to use for the ambient light, then estimated the sun and sky ratio using photographic data and our calibration tool. However, a specific sky post-process was needed to achieve the final look, simulating the effects of a polarising filter.

Finally, we implemented a physically-based shading model in order to accurately use the data we provided. We used a micro-facet BRDF, but found a fast approximation to the geometric term that could be combined for free with the normalisation factor of the NDF. To correctly filter specular highlights, we used Toksvig maps, which could optionally be averaged to a single value to avoid compression issues.

Beyond a Simple Physically-Based Blinn-Phong Model in Real-Time

Yoshiharu Gotanda

Physically-based shading is becoming standard for video games, though it is typically still based around simple models such as Blinn-Phong-Lambert. Our modified Blinn-Phong model not only has physically-based characteristics close to Cook-Torrance, but is also computationally reasonable for lower-end GPUs, such as those in existing game consoles. However, this is still not enough to accurately represent complicated materials that exist in the real world.

In this presentation, sophisticated real-time shading models are discussed that are capable of greater material fidelity and range than a simple physically-based Blinn-Phong model, yet remain feasible even for current-generation consoles.

Physical Production Shaders with OSL

Adam Martinez

Open Shading Language (OSL), being targeted at advanced ray-tracers, is built on a foundation of physically-based shading principles. Concepts such as energy conservation, reciprocity, and physical plausibility are built in to the system. This is not to say that these principles cannot be intentionally violated when necessary, but the BSDF closure is a core component, and first class citizen in the library. The use of BSDF closures means that shaders can't help but be energy conserving and reciprocal. And because of the way that the aggregate closure is sampled by the renderer, you can't make things implausible by composing the BSDFs, either.

From a shader writing perspective, OSL strongly encourages view independence. One can use the eye vector, absolutely, but you have to bend over backwards to do it. Ultimately, this means that the same shaders ought to have the same meaning for traditional *backward* ray (or path) tracing, photon mapping, or bidirectional ray tracing. Certainly, no renderer will support all of these methods, but the rendering method is no longer partially encoded in the shader.

This talk will focus on the nuts and bolts of building practical, physically-based, production-worthy shaders in OSL. Starting with what a BSDF closure in the OSL library looks like, I will take a shader from concept to completion. I will cover the use of OSL's physically based shading features, and how to reconcile those features with typical non-physical needs of feature film production. We will also bid farewell to the *light loop*, and talk about why we don't want to do this in shaders for advanced ray tracers. Finally we will discuss the implications of writing shaders that compute radiance closures, and not final colors.

Physically-Based Shading at Disney

Brent Burley

After attending the 2010 version of this course, several of us at Disney Animation Studios were inspired to investigate general physically-based material models. We had recently had success integrating physically-based hair shading for Tangled and we were now interested in looking at a broader range of materials.

To start our investigation we developed a new BRDF viewer to compare existing physical models with measured BRDF data. We discovered an intuitive way to view measured data and we found interesting features in the measured materials that aren't well-represented by any known model.

During this course we will share observations of different types of measured materials along with insights we've gleaned about which models fit the measured data and where they fall short. We will then present a new model we developed which is now being used on all current productions. We will also talk about our production experience of adopting this new model and how we were able to add the right level of artistic control while preserving simplicitly and robustness.

Reflection Model Design for WALL-E and Up

Brian Smits

I will present the reflection model used for *WALL-E* and *Up*. Before that, Pixar had used a slowly evolving model that was first developed for *A Bug's Life*. We wanted a new one that could quickly generate a broad range of looks, was efficient when used in shaders composed of many layers of materials, would fit accurately into our hardware relighting tool, and minimized the amount of time lighters had to spend adjusting surface shader properties.

We chose to make energy conservation the default behavior. With that constraint and the goal of making it easy to drive the reflection model with patterns, we came up with a small set of parameters that controlled most of the model. Using the same ideas, we developed an approach to layering that gave a lot of flexibility and kept the cost of the reflection model independent of the number of layers. This meant that the data needed for the relighting tool was fixed and thus the graphics hardware could give the same results from the model as RenderMan did. The results required minimal surface shader adjustments from the lighting department.