COMS 4160: Problems and Questions on Rendering

Ravi Ramamoorthi

Questions and Problems

We first give a number of standard questions (some from last year's final). These questions ask basic definitions and concepts that we will be covering in lecture and you should be familiar with. Answers will not be provided here; ask in the review sessions if you are unsure of anything.

1. Briefly explain "flat shading", "gouraud shading" and "phong shading", and describe the differences between them. Which of these modes does OpenGL implement or not implement and why?

2. Explain how to intersect a a ray with a sphere in ray tracing. Show how to do this for a general implicit surface. How can intersection be implemented elegantly in C++?

3. Define the terms Radiance and Irradiance, and give the units for each. Write down the formula (integral) for irradiance at a point in terms of the illumination $L(\omega)$ incident from all directions ω. Write down the local reflectance equation, i.e. express the net reflected radiance in a given direction as an integral over the incident illumination. Prominently label the main terms of the equation such as the BRDF. What are the BRDF formulae for (i) Lambertian surfaces (ii) Mirror surfaces, (iii) Dark glossy materials?

4. What is the rendering equation? Derive one version of it. Explain how ray tracing and radiosity approximate the rendering equation. What is the radiosity equation? Make the appropriate approximations to derive it from the full rendering equation.

Now, we give a few problems. In general, the exam material for this part of the course will focus only on high-level concepts, as in these problems and the material above, and will not be too technical.

1. Match the surface material to the formula (and goniometric diagram shown in class). Also, give an example of a real material that reasonably closely approximates the mathematical description. Not all materials need have a corresponding diagram. The materials are *ideal mirror, dark glossy, ideal diffuse, retroreflective*. The formulae for the BRDF f_r are $k_a(\vec{R} \cdot \vec{V}), k_b(\vec{R} \cdot \vec{V})^4, k_c/(\vec{N} \cdot \vec{V}), k_d\delta(\vec{R}), k_e$.

2. Consider the Cornell Box (as in the radiosity lecture, assume for now that this is essentially a room with only the walls, ceiling and floor. Assume for now, there are no small boxes or other furniture in the room, and that all surfaces are Lambertian. The box also has a small rectangular white light source at the center of the ceiling.) Assume we make careful measurements of the light source intensity and dimensions of the room, as well as the material properties of the walls, floor and ceiling. We then use these as inputs to our simple OpenGL renderer. Assuming we have been completely accurate, will the computer-generated picture be identical to a photograph of the same scene from the same location? If so, why? If not, what will be the differences? Ignore gamma correction and other nonlinear transfer issues. Now, answer this question again with the two small boxes added, i.e. the floor has two smaller

boxes sitting on it. You may assume we have accurately measured geometric and material properties of the smaller boxes also.

3. Consider a simplified skylight model, so the radiance along any direction is given by $A + B \sin \alpha$ where A and B are positive constants, and α is the elevation angle (i.e. the angle to the horizontal, being 0 degrees toward the horizontal and 90 degrees toward the zenith or top of the sky). That is, the radiance is more higher up in the sky. The lighting is isotropic; there is no variation with azimuthal angle (ϕ). Assume for this problem that there is no occlusion by trees, buildings etc., the sky hemisphere is the only source of illumination [no ground lighting, direct sunlight etc.], the surfaces are Lambertian with albedo 1, and the sky can be assumed to be a distant source. What is the irradiance on the ground, assumed to be a horizontal surface?. Now, assume we have a sphere suspended (or on the ground, if that makes things more logical for you). Remember the assumptions, i.e. lighting only from the (distant) sky, no occlusions etc. Which point on the sphere will be brightest? What will be the reflected radiance at this point? Which point will be the dimmest? What will be the reflected radiance at that point? Qualitatively, how will the brightness on the sphere vary as a function of location (parameterized by spherical coordinates for instance)? Extra credit for deriving an analytic quantitative formula for the brightness of the sphere as a function of the surface normal.

4. History: In 1998, for the 25th SIGGRAPH conference, there was a list compiled of seminal graphics papers. Which were the papers included for shading? Obviously very related, which Computer graphics achievement awards have been given for rendering?

Answers

1. Materials An ideal mirror is something like a normal reflective mirror, and the BRDF corresponds to $k_d \delta \left(\vec{R}\right)$. A dark glossy surface is close to glossy plastic with a formula like $k_b \left(\vec{R} \cdot \vec{V}\right)^4$. An ideal diffuse surface is Lambertian, close to wall paint or ideally, spectralon, with BRDF a constant k_e. A retroreflective surface is something like a highway reflector, that reflects light back toward the viewer and would have a formula using $\delta \left(\vec{L}\right)$ instead of $\vec{R}$ as in an ordinary reflector.

2. Cornell Box The missing feature will be global illumination (effects like color bleeding). Furthermore, area light sources are not supported in OpenGL, and approximating them with a point light will lead to other minor differences. After the smaller boxes are added, we will also not get shadowing effects in OpenGL.

3. Irradiance The irradiance is given by

$$E = \int_{\theta=0}^{\pi/2} \int_{\phi=0}^{2\pi} L(\theta, \phi) \cos\theta \sin\theta \, d\theta, \tag{1}$$

where θ stands for the angle with respect to the zenith, or normal to the ground plane. In this coordinate frame, the elevation angle = $\pi/2 - \theta$, and so the incident radiance $L = A + B\cos\theta$. Removing the azimuthal dependence in the above integral, we obtain

$$E = 2\pi \int_{0}^{\pi/2} (A + B\cos\theta) \cos\theta \sin\theta \, d\theta. \tag{2}$$

To evaluate this integral, we put $u = \cos\theta$, to obtain

$$E = 2\pi \int_{0}^{1} Au + Bu^2 \, du = \pi \left(A + \frac{2B}{3}\right). \tag{3}$$

The brightest point is clearly the top of the sphere, or the point facing directly upward, which receives the same irradiance as the ground. The dimmest point will be at the bottom of the sphere, which will be completely dark (since it sees only the lower hemisphere that is dark, since we assume not interreflection from the ground). The radiance will decrease from top to bottom, remaining azimuthally symmetric.

Finding an analytic formula for the radiance as a function of angle is an advanced topic, well outside the scope of this course. One can try to use the spherical harmonic irradiance formula in my thesis.

4. History The relevant papers are

- Continuous Shading of Curved Surfaces by H. Gouraud. This paper introduces what we now know as Gouraud shading.

- Illumination for Computer Generated Pictures by B. Phong. This paper introduces Phong illumination and shading. One side note is that Bui Tong Phong does not figure in the list of achievement award winners discussed later. He was briefly a professor at Stanford University, after completing his doctorate at Utah, when he was tragically killed in a car crash in the mid 70s.

- Models of Light Reflection for Computer Synthesized Pictures by J. Blinn. Introduces some of the reflection models commonly in use like Blinn-Phong, and talks about Torrance-Sparrow microfacet reflection. Jim Blinn won SIGGRAPH's inaugaral computer graphics achievement awardee in 1983, and became SIGGRAPH's first double awardee when he won the Coons award in 1999.

- An Improved Illumination Model for Shaded Display by T. Whitted. Introduces recursive ray tracing, often called Whitted ray tracing. Turner Whitted (currently the manager of the Microsoft Research graphics group) was awarded the 1986 Computer Graphics Achievement Award.

- Shade Trees by R. Cook. Introduces the idea of composing illumination calculations in a tree-like fashion, for instance, combining diffuse and specular components, an idea commonly used in shading calculations and in products like Pixar's Renderman. Rob Cook was honored with the 1987 Computer Graphics Achievement award, and (along with Ed Catmull and Loren Carpenter) most recently won an oscar in 2000 for technical achievement for "their significant advancements to the field of motion picture rendering, as exemplified in Pixar's 'Renderman"'. While plaques and certificates for technical excellence are given at the Academy Awards every year, this was the first (and only) time that the familiar statuette was awarded for the development of computer software.

- Modeling the Interaction of Light Between Diffuse Surfaces by C. Goral, K. Torrance, D. Greenberg and B. Battaile. The original radiosity paper. Don Greenberg has pioneered much of realistic rendering and global illumination in computer graphics over the last 3 decades at Cornell University, and has been honored with the 1987 Coons award. Ken Torrance wrote the original Torrance-Sparrow paper and was given the 1994 Computer graphics achievement award for his work on rendering, including reflection models and radiosity.

- An Image Synthesizer by K. Perlin. Introduces what is now known as "Perlin Noise" for effects like procedural textures (marble), turbulence and so on. Ken Perlin was awarded the Technical Achievement Award from the Academy of Motion Picture Arts and Sciences in 1997 for the development of Perlin Noise, "a technique used to produce natural appearing textures on computer generated surfaces for motion picture visual effects".

- The Rendering Equation by J. Kajiya. Provides the unifying framework for global illumination. Jim Kajiya (then a professor at Caltech, now director of research at Microsoft) was awarded the Computer Graphics Achievement Award in 1991 for the rendering equation, and the technical achievement award from the Academy of Motion Picture Arts and Sciences (with Tim Kay) for his work on rendering hair and fur.

- Ray Tracing JELL-O(tm) Brand Gelatin by P. Heckbert. This paper is a joke, providing a satire of rendering research, the only such paper ever to appear at SIGGRAPH.
- A Progressive Refinement Approach to Fast Radiosity Image Generation by M. Cohen, S. Chen, J. Wallace and D. Greenberg. Introduces hemicubes, essential for effective radiosity. Michael Cohen was given the 1998 Computer Graphics Achievement Award for the development of radiosity.

Since much of computer graphics research involves rendering, most awards have been made in the rendering area. The awardees are

- James F. Blinn (1983; Coons Award 1999) for his work on many of the basic illumination and shading techniques like bump and texture mapping, reflection mapping, and reflection models.
- Loren Carpenter (1985; Oscar for Renderman 2000) for a number of image synthesis algorithms, including fractals.
- Turner Whitted (1986) for the development of ray tracing.
- Rob Cook (1987) perhaps best known for shade trees, the Cook-Torrance reflection model and distribution ray tracing.
- James T. Kajiya (1991) for the rendering equation.
- Pat Hanrahan (1993; Coons Award 2003; Academy Award for Scientific and Technical Achievement for Renderman 1992) for rendering systems (in particular, Renderman) and algorithms (volume rendering, hierarchical and wavelet radiosity, and subsurface scattering).
- Kenneth E. Torrance (1984) for reflection models and radiosity.
- Marc Levoy (1996) for volume rendering.
- Michael F. Cohen (1998) for radiosity.

In addition, Coons awards have been given to Donald P. Greenberg (1987), Ed Catmull (1993) and Lance Williams (2001).