ADVANCED GRAPHICS – 2016/2017

February 2nd – 17.00 – 19.00 – RUPPERT-BLAUW

Please write clearly. Please do not ask for clarification during the exam. If you find a question unclear or ambiguous: write down how you interpret the question, then answer it. You can score up to 100 points. Your grade is: `max(1,pts/10)`*.*

IMPORTANCE

1. Consider a scene consisting of three objects (see figure):
   - an infinite horizontal plane with normal $\vec{N}$;
   - a hemispherical non-reflective skydome with radius r, casting uniform illumination inwards;
   - a single planar occluder with normal $\vec{N}_o$.

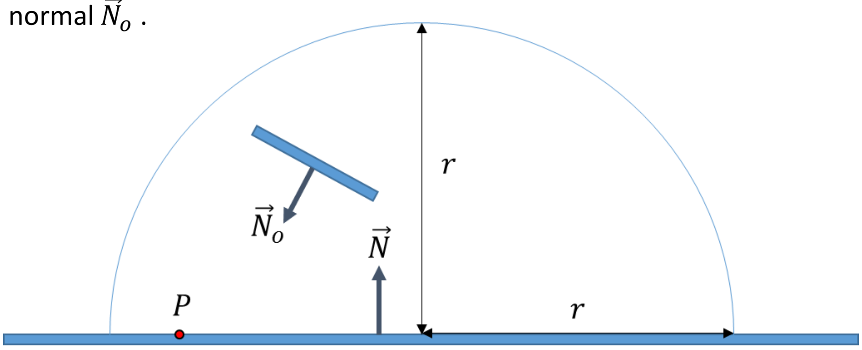

Answer the following questions about this situation. Do not consider indirect light.

a) Do we need to know r to calculate the irradiance arriving at a point P on the plane inside the dome? Why / why not? *(5 pts)*

The 100W radiance is 'spread out' over the area of the dome. Make it larger, and the flux per unit area will decrease. Note that its solid angle remains 2π regardless of r. Not correct: "no", and answers based on distance attenuation.

b) How would you importance sample this setup? Take the occluder into account. *(10 pts)*

We can calculate the irradiance arriving at P without sending any rays (see c). All we have to do now is *deduct* the irradiance blocked by the occluder. To estimate this, we sample the skydome using directions obtained by selecting random points on the occluder. The domain of this integral is the solid angle of the occluder.
Some points were awarded for a somewhat detailed explanation of pdf construction for this situation. Full points (in one case) were awarded to the person that conluded that the occluder is the only object eligible for a random bounce. ☺

c) We now remove the occluder. The dome emits a total of 100W. What is the irradiance arriving at P? Explain your answer. *(5 pts)*

Simplest answer: all energy must arrive at the floor, and is thus spread out over the available area. Divide 100 by this area to get irradiance. Crucial part of this answer is the realization that irradiance is per unit area.

2. Explain why the following techniques can be seen as importance sampling techniques:

   a) Next Event Estimation *(5 pts)*

We expect direct light to be the main part of the energy arriving at a point. NEE explicitly targets this part of the integral, and thus focuses on the assumed importance of lights. Two points for a description of NEE, five points for actually pointing out why this is a form of IS.

   b) Russian Roulette *(5 pts)*

We expect long paths to return less energy than short paths. Same for paths via bright surfaces. We thus focus on short paths and paths via bright surfaces. Again, 2 points for just describing RR.

   c) For dielectrics: basing the probability of generating a reflection on the Fresnel term *(5 pts)*

For Whitted, we would trace transmitted and reflected paths, but they typically do not bring the same amount of energy. In a PT, we sample the path with more expected energy, so we base our choice on Fresnel.

ACCELERATION STRUCTURES

3. Regarding the Surface Area Heuristic:

   a) Explain why the SAH produces better BVHs than midpoint splitting. *(10 pts)*

Keyword here is *probability:* SAH estimates the probability that a ray hits the node. Even better: SAH estimates the cost of the node by multiplying primitive count by surface area, where surface area represents the chance that we actually need to intersect those primitives. Most of you came up with this in one form or another.

Many of you just described the effects of using SAH. This yielded (only) 2 pts.

   b) The Surface Area Heuristic is a simplified cost model. Explain what factors are not taken into account. *(Note: you may have encountered this in a paper, but this is not intended as a knowledge probe; I'm asking you to reason about this.)* *(10 pts)*

Two points for 'lazy evaluation' or something along those lines, but that was not really the question. Correct options include taking into account occlusion (e.g. 6 quads form a box encapsulating 50 triangles) and ray distribution (e.g. nodes close to a point light which will really only ever be traversed by rays in 1 or 2 directions) and the fact that an AABB may be a poorly defined hull. Incorrect: SAH doesn't take into account different prims (it totally can do that).

LIGHT TRANSPORT

4. A scene is illuminated by a single <u>double-sided</u> square light source. Two algorithms are used to sample the light source: the first picks a random point on a random side of the light source, while the second algorithm only picks random points on the side of the light source facing the point we want to shade (point $\boldsymbol{p}$).

   Show, using the mathematical formulation of Monte Carlo integration, that both methods yield the same result when using an infinite number of samples. *(15 pts)*

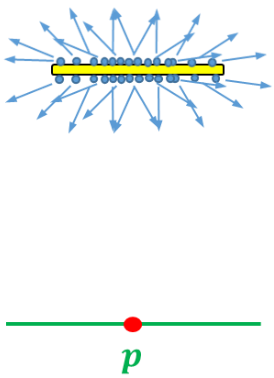

E.g.: let $f(\bar{x})$ denote a sample at the bottom of the light, $g(\bar{x})$ a sample at the top and $h(\bar{x})$ a sample at a random side. Then $f(\bar{x}) = 1$ and $g(\bar{x}) = 0$; $h(\bar{x})$ will be 0.5 on average. case 1: $E \approx \frac{1}{N}\sum_{i=0}^{N} 2\, g(\bar{x}_i) = 1$; case 2: $E \approx \frac{1}{N}\sum_{i=0}^{N} f(\bar{x}_i) = 1$.

GPU RAY TRACING

5. What problem does Wavefront Path Tracing (as described in "Megakernels Considered Harmfull") aim to solve? *(10 pts)*

Two possible answers: register pressure and occupancy. The real problem is of course occupancy. Wrong answer: "the scheduler". You are mixing up the wavefront paper with the "understanding efficiency" paper.

6. Explain how a kD-tree can be traversed without using a stack, without adding data to the nodes (so, no ropes, no short stack). *(10 pts)*

Always go to the near child. After processing a leaf (and not finding an intersection), set tmin to the exit distance for the leaf and restart at the root.

BRDFS

7. The Phong illumination model, without ambient factor, can be formulated as follows:

$$I_x = \sum_{m \in lights} \left(k \left(\underline{\vec{N} \cdot \vec{L}_m} \right) + k \left(\underline{\vec{R}_m \cdot \vec{V}} \right)^{exponent} \right) I_m$$

where k is the material color, $\vec{L}_m$ is a unit vector to light m, $\vec{R}_m$ is vector $\vec{L}_m$ reflected in the surface normal $\vec{N}$ and I_m is the luminance of light m. Underlining of dot products denotes clamping to zero.

a) The model does not ensure energy preservation. This is far less of an issue in Whitted-style ray tracing than in a path tracer. Why? *(5 pts)*

Whitted does not bounce, at least not via diffuse and glossy surfaces. The excessive amount of 'reflected' energy is thus limited to a manageable single bounce, which we can compensate for by tweaking light intensities. In a PT, the variable number of bounces via these surfaces will amplify the effect, and it's no longer fixable via light intensity.

The Modified Phong BRDF is based on the Phong illumination model. It is defined as:

$$f_r(x, \theta_i, \theta_o) = k\frac{1}{\pi} + k\frac{exponent + 2}{2\pi} cos^{exponent} \varphi$$

b) This BRDF does not obey the Helmholtz reciprocity. Why not? *(5 pts)*

The modified Phong BRDF will generate directions into the floor (and then rejects them), but does not reflect directions coming in through the floor. We can thus not always swap incoming and outgoing. Any answer describing this (including floor intersecting lobes) was accepted.

May the Light be with you!