



# ADVANCED GRAPHICS – 2017/2018

February 1<sup>st</sup> – 17.00 – 19.00 – EDUC-MEGARON

**PROBLEM 1.** Consider the following scene:

An infinite floor plane with a diffuse BRDF  $f_r(\cdot) = 1/\pi$ , illuminated by two spherical light sources, and surrounded by a skydome.

The skydome has a uniform color, which is (2,2,2).


On the floor plane we found a point  $p$  using a primary ray from the camera.

- Show that we do not need to know the radius of the skydome when we estimate the light transport from a point on the skydome via point  $p$  to the camera.
- A path tracer is used to render the scene. Next Event Estimation (NEE) is used to sample the illumination from the two spherical light sources. The ray resulting from a random bounce at point  $p$  ends up on the skydome. Should this ray return (2,2,2) or (0,0,0)? Why?
- We remove the two spherical lights, and replace the uniform color of the skydome by a (non-uniform) HDR texture, which is now the only source of illumination for the scene. Describe a method to importance sample the illumination from this skydome. Note: this was not discussed during the course, so I am asking you to be creative.

**PROBLEM 2.** When traversing a BVH, doing so in front-to-back order is faster than a random order or back-to-front. We discussed three methods to obtain such an ordering in a (binary) BVH, which rely on different information:

- the first one obtains the distance to each child node on the fly;
- the second one uses information stored with the node about the used split axis;
- the third one guesses the split axis based on the centroids of the child nodes.

With that in mind, we have the following interesting situation:



Nodes A and B are the result of a split along a vertical split plane (the black dotted line).

- In which order will the nodes be traversed, when applying each of the three methods? Describe for each method briefly how you obtain the answer.
- Describe a situation (if one exists) where traversing node A before node B (with the red ray) is advantageous.
- Describe an alternative method to determine which node to visit first. Your method must not rely on calculating the distance to both child nodes. It must visit node A and B in the optimal order in the illustrated case. And finally, it should work well in the general case.

*Some smaller things to ponder:*

**PROBLEM 3.**

The image on the right shows a tree generated by a program called *SpeedTree*. Explain why raytracing or path tracing such an object is relatively expensive.



**PROBLEM 4.**

In the paper "Understanding the Efficiency of Ray Traversal on GPUs", the authors describe several ray traversal kernels in a high-level way. Two of them are called "while-while" and "if-if".

- a) Describe these two traversal loops with some pseudo code. Make sure the difference between the two is clear.
- b) Explain which approach is faster, and why.

**PROBLEM 5.**

Explain why traversal of an MBVH/QBVH (i.e., a BVH in which each interior node can have more than two child nodes) is generally faster than traversal of a regular BVH (with exactly two child nodes per interior node).

**PROBLEM 6.**

Whitted-style ray tracing is a point-sampling algorithm. Is this also true for distributed ray tracing (Cook style)? Explain why / why not.

*That's it! Please note that this is a Caracal →*

*Scoring: each (sub)problem is 1 point, for a maximum total of 11 points; grade will be clamped to [1..10].*



*Good luck, thanks for participating! See you next time,*

Jacco.