ADVANCED GRAPHICS – 2015/2016

April 20th – 13.30 – 16.30 - EDUC Beta

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 using the regular questions, 110 including bonus question. Your grade is: `min(10,pts/10)`.

ACCELERATION STRUCTURES

Answer the following questions about the Bounding Volume Hierarchy:

1. We use the Surface Area Heuristic to determine a good position for a split plane during BVH construction.

    a. Write down the cost function used in the Surface Area Heuristic. Note: several variations can be found in literature; any of these will do. *(5 pts)*

E.g. C=Aleft*Nleft+Aright*Nright.

    b. Explain why we use surface area (rather than e.g. bounding box volume) in the cost function. *(5 pts)*

Keyword is ‘probability’: probability of hitting the box with an arbitrary ray is proportional to surface area. Not accepted (by itself): volume may go to 0 when AABB is infinitely thin.

    c. The Surface Area Heuristic is a ‘greedy’ heuristic. What is the meaning of ‘greedy’ in this context? *(5 pts)*

Decisions are made at the current layer, ignoring consequences for the following layers, or anything along those lines. Not: “we do not consider all options / every axis”.

2. Situation: in a real-time application, a BVH is created per mesh in the scene graph. A top-level BVH is used to fuse these BVHs into a single BVH. The BVH for one of the scene graph meshes is *refitted* each frame.

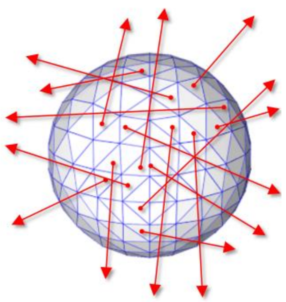

    a. Assuming refitting is the appropriate method to update the BVH for this mesh: what are the characteristics of the animation? *(5 pts)*

Keyword is ‘topology’ or ‘structure’. Two points for ‘waving trees’ and ‘water’.

    b. After a while, one mesh explodes: each triangle starts moving in a random direction (i.e., initially not necessarily away from the center, see figure). If we keep using refitting, what are the consequences for the resulting BVH? What are the consequences for traversal? *(10 pts)*

Massive overlap of the AABBs. But pretty everyone got that. ☺

LITERATURE

1. In the paper “Interactive Rendering with Coherent Ray Tracing” (Wald et al.), seven benefits of ray tracing over rasterization are listed. Write down and explain four. *(10 pts)*

Any selection of four from the paper is correct.

2. Write down the Rendering Equation (variations exist; any correct version will do). Explain the various terms. *(10 pts)*

I know I promised not to do math, but this was pretty much at the core of the theory… Global structure: outgoing radiance is emitted radiance plus reflected radiance, where irradiance is scaled by the BRDF.

3. Whitted-style ray tracing is a point-sampling algorithm. Is this also true for distributed ray tracing (Cook style)? Why / why not? *(10 pts)*

Any discussion / viewpoint showing some insight was accepted. Most of you had no problems with that.

LIGHT TRANSPORT

4. Next Event Estimation is a form of importance sampling. Explain why this statement is true. *(10 pts)*

IS is sampling a function focusing on areas of high contribution. NEE samples those parts of the hemisphere where we expect significant contribution.

5. 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 $\boldsymbol{p}$).

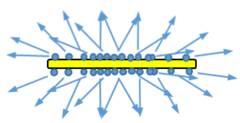

   a. Using a single ray, how do we estimate the visibility of the light source from point $\boldsymbol{p}$? *(5 pts)*
   b. Using a single ray, how do we estimate the visibility of the light source from point $\boldsymbol{p}$? *(5 pts)*

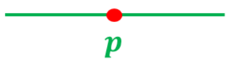

This one was confusing, also because of the missing info in a and b. Intention was to make you describe the difference between approach 1 and 2. Also confusing is the fact that I was asking for *visibility*. Looking at the light, any point can see at most 50% of the light. We can simplify the problem by representing the light by two lights. We now have two options:

- Randomly sample a light: we either sample 0.5 or 0, E = 0.25. Each sample is scaled up by 1/50%, yielding E = 0.5.
- Sample only the facing light: we now always sample 0.5, and don’t scale.

To compensate for the vagueness, any answer mentioning scaling by 1/50% was accepted.

RAY TRACING USING THE GPU

6. Consider the following situation: we are rendering a large scene (~2GB triangles and textures). Most of the rays are very divergent. Explain which GPU hardware characteristic makes the GPU more suitable for this task than a CPU. Assume a GPU with sufficient RAM to store the scene (say, 4GB). *(10 pts)*

Main issue here is that GPUs hide latencies by switching between warps. This approach is effective regardless of the data size. CPUs hide by using caches; this is only effective when the data fits in those caches.

7. The BVH construction algorithm we used on the CPU bears similarities to the Quicksort algorithm. Explain, considering the characteristics of the GPU, why this algorithm does not map well to GPU hardware. *(10 pts)*

Main issue here is that Quicksort starts with 1 task, which is divided in 2, and so on. Early on we thus have very few parallel tasks, which leaves most of the GPU idling.

BONUS QUESTION

8. This exam can be seen as a Monte-Carlo process. Explain why. *(10 pts)*

The integral is your knowledge, exam questions are samples. Any discussion along those lines was accepted (and some other valid interpretations as well).