



20. DECEMBER 2018

INTRODUCTION TO COMPUTER GRAPHICS ASSIGNMENT 7

Submission deadline for the exercises: 10. January 2019

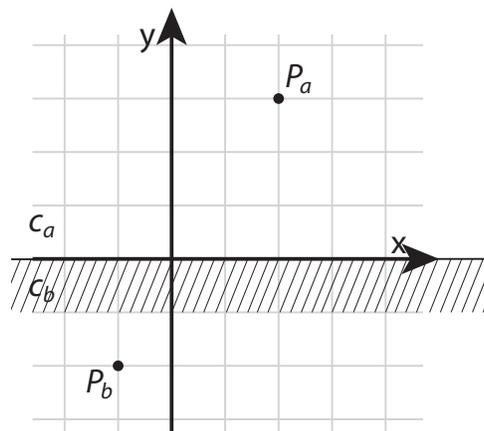
The paper copies for the theoretical parts of the assignments will be collected at the beginning of the lecture on the due date. The programming parts must instead be marked as release before the beginning of the lecture on the due date. The code submitted for the programming part of the assignments is required to reproduce the provided reference images. The submission ought to tag the respective commit in your group's git repository as well as attach the mandatory generated images to the release. The submission should also contain a creative image show-casing all extra-credit features that have been implemented.

The projects are expected to compile and work out of the box. A successful build by Drone CI is a good indicator. If it fails to do so on our Windows or Linux computers, we will try to do so on the CIP-pool students' lab as a "fallback".

To pass the course you need for every assignment at least 50% of the points.

7.1 Snell's Law (5 + 3 + 5 Points)

- Let us consider a 2-dimensional slice through a 2-layer dielectric material such that the half-space of positive y coordinates lies in a medium where light travels at constant speed c_a and the half-space of negative y coordinates in a medium where light travels at constant speed c_b . Assuming that light travels between the 2 points $P_a(2, 3)$ and $P_b(-1, -2)$ by crossing the interface between the two media at some point $P_i(x_i, 0)$, write the expression for the time of travel as a function of x_i , c_a and c_b .



- According to Fermat's principle, light always travels between 2 points along the path with minimal time of travel. Write the equation that x_i must satisfy for the path taken between P_a and P_b to be valid.
- Use the formulations you have computed above to rederive Snell's law, which relates the refractive indices n_a and n_b to the angles of incidence and exitance of the light rays.

7.2 Fourier Transformation (10 Points)

The transformation of a signal f to Fourier space is given by:

$$\mathcal{F}(f) = F : F(k) = \int_{-\infty}^{+\infty} f(x) \cdot e^{-2\pi i k x} dx$$

Show that the Fourier transformation of the box function $B(x)$ is a *sinc* type function. The sinc function is defined as $\text{sinc}(x) = \frac{\sin(\pi x)}{\pi x}$.

$$B(x) = \begin{cases} 0 & \text{for } x \leq -1 \\ 1 & \text{for } -1 < x < 1 \\ 0 & \text{for } 1 \leq x \end{cases}$$

7.3 Duality of Multiplication and Convolution (10 Points)

The convolution of a function $f(t)$ with a second function $g(t)$ is defined as:

$$(f \otimes g)(t) = \int_{-\infty}^{+\infty} f(\tau) \cdot g(t - \tau) d\tau$$

The multiplication of two function is defined as the point-wise multiplication:

$$(f \cdot g)(t) = f(t) \cdot g(t)$$

The transformation of a signal f to Fourier space is given by:

$$\mathcal{F}(f) = F : F(k) = \int_{-\infty}^{+\infty} f(x) \cdot e^{-2\pi i k x} dx$$

We call \mathcal{F} the operator mapping f to Fourier space: $\mathcal{F}f = F$. Show that convolving in signal space is the same as multiplication in Fourier space:

$$\mathcal{F}[f \otimes g] = \mathcal{F}[f] \cdot \mathcal{F}[g]$$

7.4 Super Sampling (10 Points)

In order to reduce the aliasing effects one can shoot multiple, slightly perturbed rays per pixel and average the results out. You will need to:

- Implement `Renderer::setSamples` which should specify how many rays should be shot for each pixel.
- The `Renderer::render` should take that information into account. If more than one ray per pixel is being shot, it should no longer go through the center of the pixel, but anywhere within the pixel area.

We are providing a `random()` function¹ in `random.h`, which returns a floating point number in the range $[0, 1)$. The function is thread-safe.

¹We are using the Mersenne Twister algorithm. The state of the machine is kept in a thread-local variable, so calls to the function in one thread have no impact on the results in another.

7.5 Area Lights (20 Points)

Materials can emit light on their own, but so far that kind of light remained ignored during lighting: Shadow rays have been shot only to “virtual” lights (`PointLight`, `SpotLight` and `DirectionalLight`). Your task is to bring those two together. Implement an `AreaLight` that accepts any `Solid` as an argument and uses its material to compute the light intensity. When computing the lighting you should:

- Use `Solid::sample()` to sample a random point on the surface of the solid. Implement the said method for `Triangle` and `Quad`. Please note, that we changed the signature in order to return a struct that contains also the normal of the sampled surface point. Please adapt the interface of all derived solids accordingly.
- Use `Material::getEmission()` to get the light intensity. Currently, the arguments to `getEmission()` can be set to dummy values. We will be using only `LambertianMaterial` with constant textures.

Area light sources generate soft shadows when the source is only partially covered by an obstacle. The multisampling technique from the previous exercise can greatly help reducing the noise.

7.6 Distributed Materials (20 + 20 Points)

So far, all materials described their interaction with light in a deterministic way. This, combined with the fact that a material could specify only a single direction of interest through `Material::getSampleReflectance()` is very limiting. In this exercise we allow materials to have multiple directions that are worth sampling and expect it to pick one at random.

Implement:

- `GlassMaterial`, representing a material that reflects and refracts incoming light. A single index of refraction `eta` defines the new medium. If the incoming ray is coming from “below” the solid (the angle between the ray direction and normal is greater than 90°), assume that the ray is leaving the medium and use the inverse of `eta`.
- `FuzzyMirrorMaterial`, which behaves similarly to mirror, but the reflected ray is randomly perturbed. The perturbation is defined by a maximum angle that the ray may deviate from the perfect reflection. You can approximate this deviation by sampling a random point p on an imaginary disc perpendicular to the perfect reflection ray, at a unit distance from the hit point. Then, set the ray direction such that it goes through p .

You may also want to update the `CombineMaterial` such that it supports multiple sampling sub-materials by picking one at random, but we are not going to test it.

7.7 Depth of Field (10 Points)

Implement the depth of field effect, as an extension to the perspective camera (`DOFPerspectiveCamera`). Two new parameters appear:

- `focalDistance` defines the distance from the center of the camera to an imaginary plane at which all the objects should appear sharp.
- `apertureRadius` defines the maximum ray origin perturbation to control the strength of DOF.

7.8 Motion Blur (10 Points)*

Implement the Motion Blur effect.

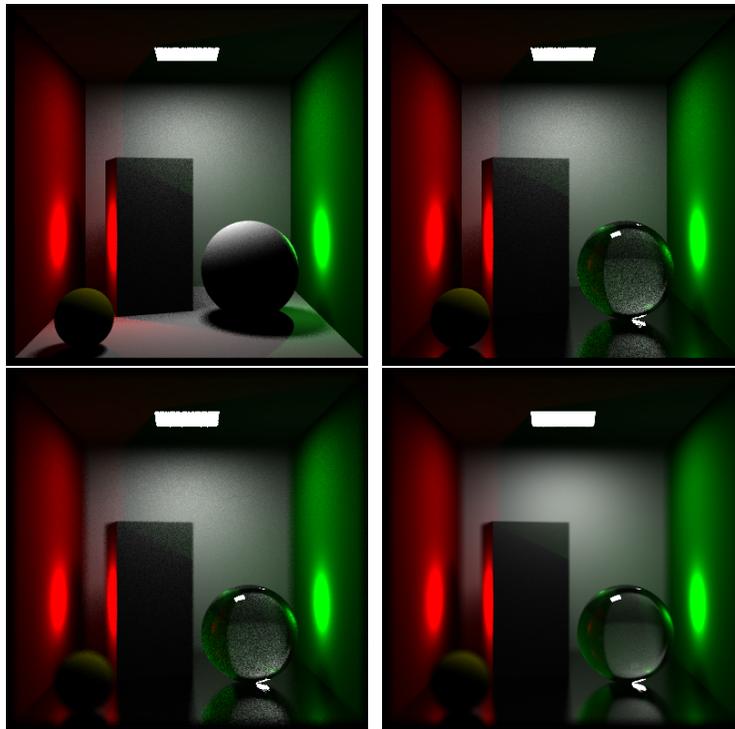


Figure 1: Distributed raytracing. (a) Super sampling and area light source (b) Glass and fuzzy mirror (c) depth of field (d) Higher-quality rendering with 1000 rays shot per pixel

7.9 Smooth Edges (5 Points)*

Most models are described by a collection of flat surfaces. This becomes very apparent when doing lighting as each surface has a single, constant normal.

A common technique for keeping the polygon count low but making it less apparent, is to perturb the normal vector depending on the location where the ray hit the flat surface. For example, a triangle can be given three fake normals on its vertices; when a ray hits the primitive, the normals are interpolated for the given hit point.

Your task is to implement a `SmoothTriangle` solid, which should perform this normal interpolation. The triangle intersection routine should remain unchanged with the exception for the returned normal within the `Intersection` structure.

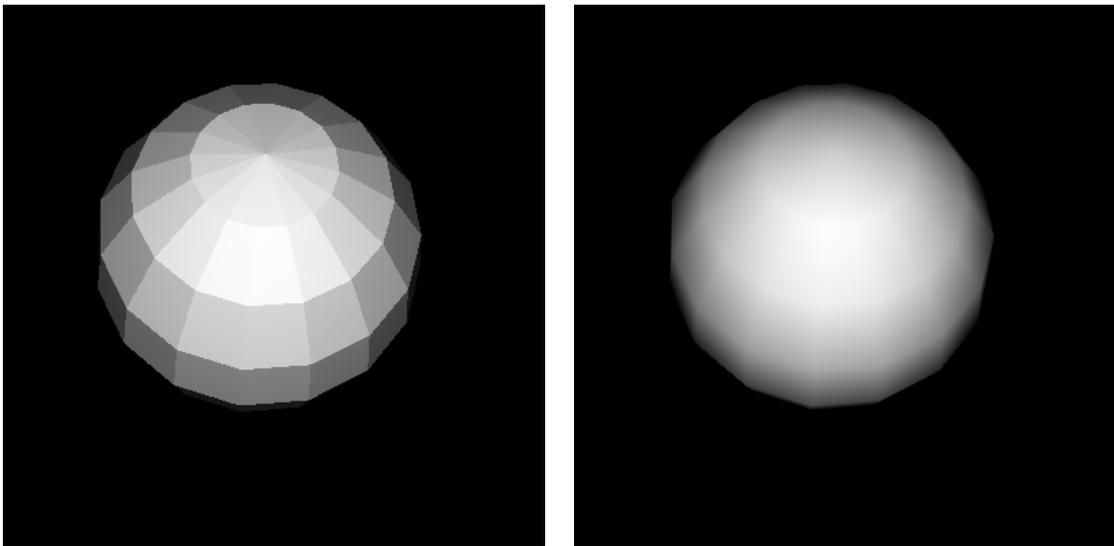


Figure 2: A tessellated sphere, without and with interpolated normals.