# CS 4620 Final, 11 December 2014

This 150-minute exam has 7 questions worth a total of 100 points. Use the back of the pages if you need more space.

Academic Integrity is expected of all students of Cornell University at all times, whether in the presence or absence of members of the faculty. Understanding this, I declare I shall not give, use, or receive unauthorized aid in this examination.

Signature: _____   Date _____

Trig identities	$-\sin(x) = \cos(x + \frac{\pi}{2}) \; \forall x \in \mathcal{R}$ $\sin^2(x) + \cos^2(x) = 1 \; \forall x \in \mathcal{R}$
Snell's law	$\dfrac{sin\theta_i}{sin\theta_t} = \dfrac{n_1}{n_2}$
Fresnel's formula	$R_s = \left| \dfrac{n_1 cos\theta_i - n_2 cos\theta_t}{n_1 cos\theta_i + n_2 cos\theta_t} \right|^2$ $R_p = \left| \dfrac{n_1 cos\theta_t - n_2 cos\theta_i}{n_1 cos\theta_t + n_2 cos\theta_i} \right|^2$ $R = \dfrac{R_s + R_p}{2}$
Sperical Linear Interpolation	$\cos(\Omega) = \vec{v_1} \cdot \vec{v_2}$ $q(t) = \dfrac{\sin((1-t)\Omega)}{\sin(\Omega)} \vec{v_1} + \dfrac{\sin(t\Omega)}{\sin(\Omega)} \vec{v_2}$

1. [12 points] **Spherical Linear Interpolation**. In the animation process, we were often interested in interpolating unit quaternions over time. We discussed two plausible methods for this task. The first was the same method used to interpolate normal vectors for shading: linear interpolation followed by a normalization step (normalized LERP). The second was spherical linear interpolation (SLERP). In this problem, we will be examining the differences between these two methods by looking at the following example, where we are interpolating between a pair of 2D unit vectors, $\vec{v_1}$ and $\vec{v_2}$.

V2 = (0,1)

V1 = (1,0)

   (a) In 1-2 sentences, explain what the primary advantage of SLERP is for interpolation in this setting over normalized LERP.

   (b) Prove that SLERP has the property you described in part A by differentiating the interpolated vector with respect to t. You may find some of the formulas on the front of the exam to be helpful.

2. [16 points] **Splines: Blending Functions** Consider the following blending functions for a quadratic spline segment, which is controlled by three control points P_0, P_1, and P_2.
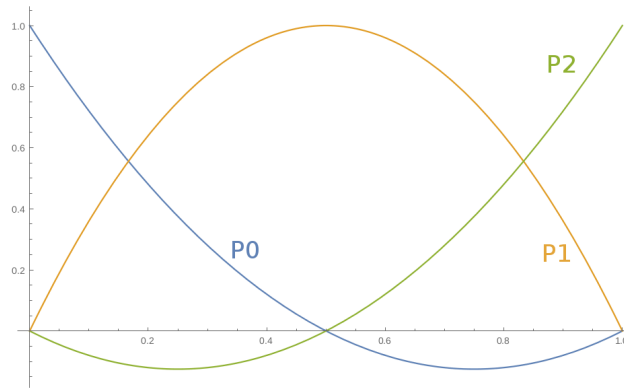
Their equations are given by

$$b_1(u) = 2u^2 - 3u + 1$$
$$b_2(u) = -4u^2 + 4u + 0$$
$$b_2(u) = 2u^2 - u + 0$$

(a) In 1-2 sentences, explain what blending functions are, in general.

(b) If the spline's parametric function is $f(u) = a_0 + a_1 u + a_2 u^2$, give a matrix X such that $(P_0, P_1, P_2)^T = X(a_0, a_1, a_2)^T$. You can answer this question without them, but if your solution includes functions such as matrix determinants/inversions/transposes, you may skip the computations and simply write your answer in terms of determinant/inversion/-transpose notation.

(c) Without doing all the simplification in part C, what features of the graph immediately indicate what the control points P_0, P_1, and P_2 represent?

3. [16 points] **Shader Programming** You will solve a shader programming problem like the ones you had in the Shaders Assignment.

You will implement a toon shader that renders objects in the cartoon-like style shown in Figure 1. This is achieved by rendering with quantized colors (to imitate cel shading). The quantized shader uses a modified version of Lambertian shading in which the cosine of the angle between the surface normal and the light vector is quantized before using it to compute the final color. There is only one directional light source in the scene. To make sure the boundaries between the color regions form sharp, smooth curves, perform this quantization calculation for each pixel rather than interpolating across triangles. We ask you to quantize the cosine value according to the following pseudocode:

```
intensity = floor(4 * cos_theta) / 4.0;
```

You will have to define two files, one for the vertex shader and one for the fragment shader. We don't need you to write syntactically perfect GLSL code; pseudocode or GLSL-like code is fine, but it needs to be unambiguously clear what your shader is doing.

Assume your Java code has access to vectors containing:

- the camera position in world coordinates,
- the direction toward the light source in world coordinates,
- the diffuse surface color,

as well as 4x4 matrices containing the usual transformations:

- projection,
- viewing (aka. camera),
- modeling (aka. object-to-world).

Your meshes have the following data available per vertex:

- position,
- normal,
- texture coordinate.

Here are some questions to get you started, you don't need to answer them, just include the needed variables in you shader code. Remember to unambiguously specify what is stored in the variable with comments, referring specifically to the data we've said is available. For instance, you might say that a particular variable contains the inverse of the modeling transformation.

1. Which transformations should be passed to the vertex shader? Create uniform variables for them in your program and use them! Write comments that explain exactly what is to be stored into each one.
2. Which vertex attributes should be passed to the vertex shader? Create vertex attribute variables in your program and use them! Write comments that explain exactly what data is passed into each one.
3. Which variables should be passed between the vertex and fragment shader? Create varying variables for them in your program and use them! Write comments to indicate the function of each one.

Vertex shader:

```
// Put uniform variables here
// uniform mat4 apple;




// Put vertex attributes here
// attribute vec3 peach;




// Put varying variables here
// varying vec3 pear;




void main() {
    // Write vertex shader code here







    gl_Position = ...
}
```

Fragment shader:

```
// Put uniform variables here
// uniform float box;




// Put varying variables here
// varying vec3 pear;




void main() {
    // Write fragment shader code here








    gl_FragColor = ...
}
```

(a)   (b)

Figure 1: Building up the toon shader. Image (a) shows the teapot under diffuse shading, (b) is the final toon result.

4. [14 points] **Texture Mapping** This problem is about reflection and normal mapping.

   (a) You have a scene with a camera at (0, 0, 1), looking in the (0, 0, –1) direction, and a disc centered at the origin, with radius $r = 2$ and normal vector (0, 0, 1). The lighting comes from an environment map, and we use reflection mapping on the disc. There is a bright point in the environment map at the direction $(0, \frac{\sqrt{2}}{2}, \frac{\sqrt{2}}{2})$. What is the coordinate of the point on the surface of the disk where we see this bright point reflected?

   (b) Now we want to make our disk fancier, so we add normal mapping to it. The normal map is determined in a way that the normal of the disk is the same as if we had a hemisphere with the same radius and center and projected it to the disk keeping the hemisphere's normals. Therefore a point at distance $r \sin \theta$ from the center has a surface normal that makes an angle θ with the disk's geometric normal. In this case we place the camera very far away, giving it the viewing direction $(0, \sin 15, -\cos 15)$, i.e. the camera is looking up by 15° from the original (0, 0, –1) direction. Now what is the coordinate of the point on the surface of the disk where we see this bright point reflected?

5. [14 points] **Reflection and cube maps**. Imagine that we would like to render a glazed diffuse sphere under environment lighting. The sphere (diffuse reflectance 0.2, index of refraction 1.5) is centered at $(0, 0, 0)$ and its radius is 1. The cube map is a cross-like texture (just like the ones in Ray2 assignment), with width 600 and height 800 (See the figure below). A directional light is from direction $(0, 0, 1)$.
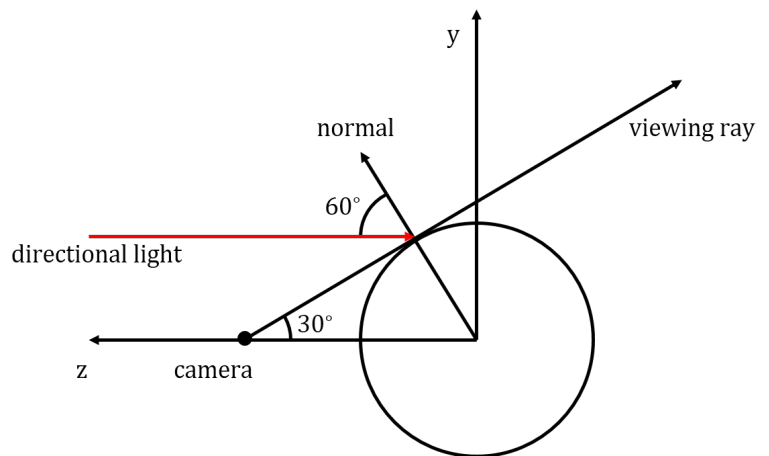
(a) A perspective camera is positioned at $(0, 0, \sqrt{2})$, viewing at $(0, 0, 0)$. Which part of the cube map will be reflected by the sphere? Draw the region on the cross in the next page and mark your sketch with enough positions and dimensions (in texel units) so that it's clear exactly what region is shaded. The left part in the above figure is an orthogonal projection of the scene onto the y-z plane and it shows the viewing ray tangent to the sphere.

(b) If the camera is positioned at $(0, 0, 2)$ and the cube map has a constant value 1.0, how bright should the directional light be, in order to make the border and center of the sphere look equally bright in the rendered image? You don't need to scale the diffuse term by $(1 - R)$ here; the shading is simply the sum of the Fresnel reflection from the surface and a Lambertian diffuse component. The right part of the figure in the previous page is an orthogonal projection of the scene onto the y-z plane. It shows the viewing ray, directional light and normal of a point on the border.

6. [12 points] **Bounding Volume Hierarchies**.

   Suppose we have a number of objects in a 2D scene that we want to ray trace, as shown on
   the left sides of Figures 2 and 3 (the scenes are the same in both figures). We can set up
   an acceleration structure using axis-aligned bounding boxes to try to avoid costly intersection
   operations. Suppose that surface A has a glazed shader, and we want to calculate the first
   intersection of the reflecting rays shown (assuming the rays start some small distance epsilon
   away from the surface). The right sides of Figures 2 and 3 give two possible configurations for
   a bounding volume hierarchy containing the surfaces. For each ray and for each tree, please list
   the objects (including bounding volumes) on which intersection operations are performed (the
   objects can be listed in any order—no need to take tree traversal order into account). How
   many total intersection operations are performed in each case? What about in the naïve case,
   i.e. without any acceleration structure?

   Ray 1, Tree 1:

   Ray 2, Tree 1:

   Ray 1, Tree 2:

   Ray 2, Tree 2:

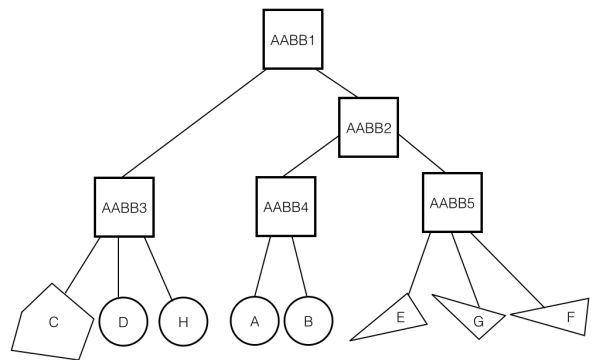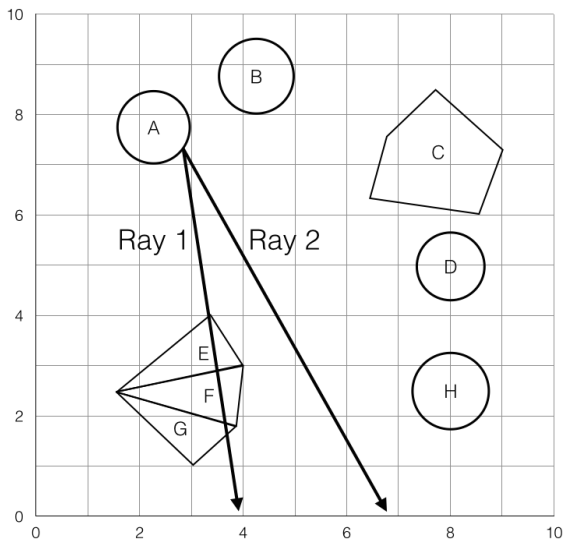Figure 2: The 2D scene to ray trace, along with BVH tree 1.

Figure 3: The 2D scene to ray trace, along with BHV Tree 2.

7. [16 points] **Rasterization, antialiasing, and the z-buffer**

    (a) Suppose we have the screen-space coordinates for the three vertices that make up a triangle, where $v_0 = (0.2, 0.2, 1)$, $v_1 = (4, 2.75, 7)$, $v_2 = (1, 4.25, 7)$ and the origin is the center of bottom-left pixel in the grid. Draw and label the vertices for this triangle in the left grid below.

    (b) Use the right-hand side grid to shade-in the fragments that will end up being emitted. What rule did you follow to rasterize the triangle?

    (c) The boundaries of the triangle above will have a fairly jaggy appearance. We wish to antialias this using regular supersampling with 4 samples per pixel. Use the following grid to compute the fractional intensities for the rasterized pixels at the positions given. You can assume that all of the previously drawn fragments had an intensity of 1.
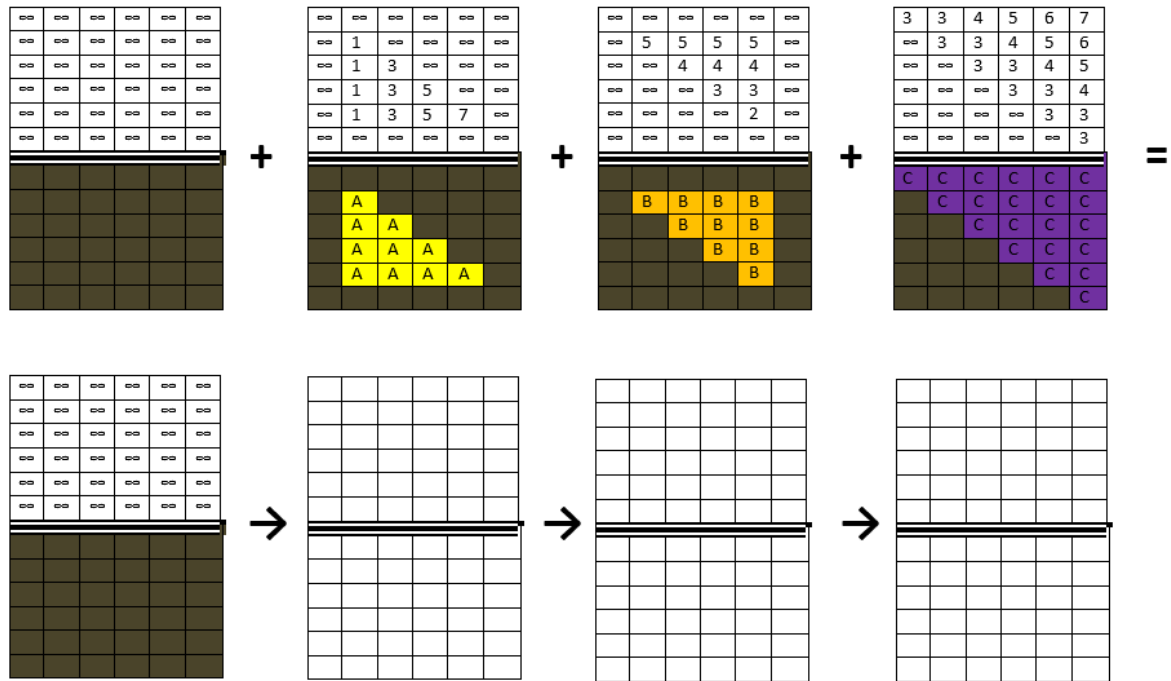
(0,0): _____

(1,2): _____

(2,2): _____

(2,3): _____

(3,1): _____

(3,3): _____

(d) The figure below illustrates the depth values and emitted fragments for three separate polygons in the top row. Complete the bottom row to show how the zbuffer is updated as the rasterized polygons are considered in the given sequence.

(e) Briefly describe, in one or two sentences, the difference between a fragment and a pixel?

Question:	1	2	3	4	5	6	7	Total
Points:	12	16	16	14	14	12	16	100
Score:								