Homework #2

Shading, Ray Tracing, Texture Mapping, Curves

Assigned: Wednesday, November 10^{th}

Due: Tuesday, November 23^{rd}
at the beginning of class

Directions: Please provide short written answers to the following questions on your own paper. Feel free to discuss the problems with classmates, but please ***answer the questions on your own and show your work.***

Please write your name on your assignment!

Problem 1: Short answer (12 points)

Answer True or False to each of the following statements and then *justify your answer*.

a) (4 point) Regardless of the scene geometry, Phong interpolated shading (i.e., shading after interpolating normals) always performs more shading calculations than Gouraud interpolated shading.

b) (4 point) Ray tracing a displacement-mapped sphere takes longer than ray tracing a bump-mapped sphere, but the results are more realistic.

c) (4 point) Anti-aliased ray tracing by using 5x5 non-adaptive supersampling (i.e., a 5x5 pattern of rays for each pixel) and averaging the results requires about 25 times as much computation as doing no anti-aliasing at all.

d) (4 point) Starting with the same number of primary (viewing) rays per pixel, distribution ray tracing eventually ends up casting many more rays than regular (i.e., Whitted) anti-aliased ray tracing in order to simulate glossy reflection.

Problem 2. Blinn-Phong shading (17 Points)

The Blinn-Phong shading model for a scene illuminated by global ambient light and a single directional light can be summarized by the following equation:

$$I_{phong} = k_e + k_a I_a + k_d B\, I_L\, (\mathbf{N}\cdot\mathbf{L})_+ + k_s B\, I_L\, (\mathbf{N}\cdot\mathbf{H})_+^{n_s}$$

Imagine a scene with one white sphere illuminated by white global ambient light and a single white directional light. For sub-problems a) – f), describe – qualitatively, in words – the effect of each step on the shading of the object. At each incremental step, assume that all the preceding steps have been applied first. Assume that the directional light is oriented so that the viewer can see the shading over the surface, including diffuse and specular where appropriate.

a) (2 points) The directional light is off. How does the shading vary over the surface of the object?

b) (2 points) Now turn the directional light on. The specular reflection coefficient k_s of the material is zero, and the diffuse reflection coefficient k_d is non-zero. How does the shading vary over the surface of the object?

c) (2 points) Now translate the sphere straight toward the viewer. What happens to the shading over the object?

d) (2 points) Now increase the specular exponent n_s. What happens?

e) (2 points) Now increase the specular reflection coefficient k_s of the material to be greater than zero. What happens?

f) (2 points) Now decrease the specular exponent n_s. What happens?

g) (2 points) Suppose we assume that the viewing direction **V** is constant regardless of which pixel it passes through. What does this imply about the viewer?

h) (3 points) Assuming that **L** and **V** are constant everywhere, then with a little pre-computation, it is possible to shade faster (i.e., using fewer operations) using the Blinn-Phong model above, than it is to shade using the Phong model, which bases the specular component on $(\mathbf{V}\cdot\mathbf{R})_+^{n_s}$. How is this possible? Explain.

Problem 3. Ray intersection with implicit surfaces (25 points)

There are many ways to represent a surface. One way is to define a function of the form $f(x, y, z) = 0$. Such a function is called an *implicit surface* representation. For example, the equation $f(x, y, z) = x^2 + y^2 + z^2 - r^2 = 0$ defines a sphere of radius *r*. Suppose we wanted to ray trace a so-called "gumdrop torus," described by the equation:

$$4x^4 + 4y^4 + 4z^4 + 17x^2y^2 + 17x^2z^2 + 8y^2z^2 - 20x^2 - 20y^2 - 20z^2 + 17 = 0$$

On the left is a picture of a gumdrop torus, and on the right is a slice through the *x-y* plane.

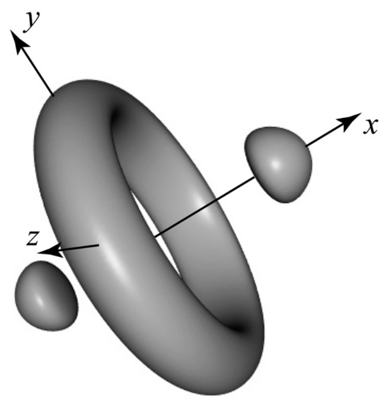

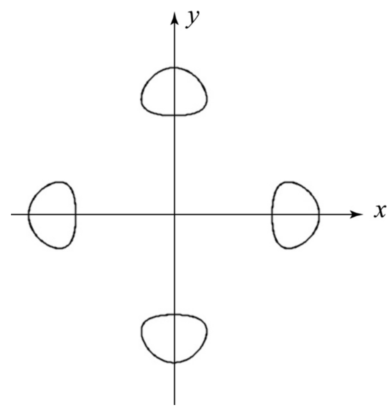

In the next problem steps, you will be asked to solve for and/or discuss ray intersections with this primitive. Performing the ray intersections will amount to solving for the roots of a polynomial, much as it did for sphere intersection. For your answers, you need to keep a few things in mind:

- You will find as many roots as the order (largest exponent) of the polynomial.
- You may find a mixture of real and complex roots. When we say complex here, we mean a number that has a non-zero imaginary component.
- All complex roots occur in complex conjugate pairs. If *A* + i*B* is a root, then so is *A* – i*B*.
- Sometimes a real root will appear more than once, i.e., has multiplicity > 1. Consider the case of sphere intersection, which we solve by computing the roots of a quadratic equation. A ray that intersects the sphere will usually have two distinct roots (each has multiplicity = 1) where the ray enters and leaves the sphere. If we were to take such a ray and translate it away from the center of the sphere, those roots get closer and closer together, until they merge into one root. They merge when the ray is tangent to the sphere. The result is one distinct real root with multiplicity = 2.

a) (10 points) Consider the ray $P + t\mathbf{d}$, where $P = \begin{pmatrix} 0 & 0 & 0 \end{pmatrix}$ and $\mathbf{d} = \begin{pmatrix} 0 & 1 & 0 \end{pmatrix}$.
   - Solve for all values of *t* where the ray intersects the gumdrop torus (**including** any negative values of *t*). Show your work.
   - In the process of solving for t, you should have computed the roots of a polynomial. How many distinct real roots did you find? How many of them have multiplicity > 1? How many complex roots did you find?
   - Which value of *t* represents the intersection we care about for ray tracing?

Problem 3 (cont'd)

b) (15 points) What are all the possible combinations of roots, not counting the one in part (a)? For each combination, describe the 4 roots as in part (a), draw a ray in the *x-y* plane that gives rise to that combination, and place a dot at each intersection point. Assume the origin of the ray is outside of the bounding box of the object. The first diagram below has already been filled in. There are five diagrams that have not been filled in. You may not need all five. (Note: not all conceivable combinations can be achieved on this particular gumdrop torus. For example, there is no ray that will give a root with multiplicity 4.) ***Please write on this page and include it with your homework solutions. You do not need to justify your answers.***

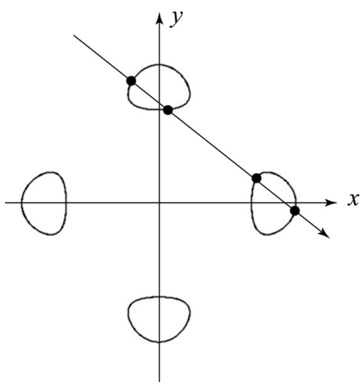

# of distinct real roots: **4**

# of roots w/ multiplicity > 1: **0**

# of complex roots: **0**

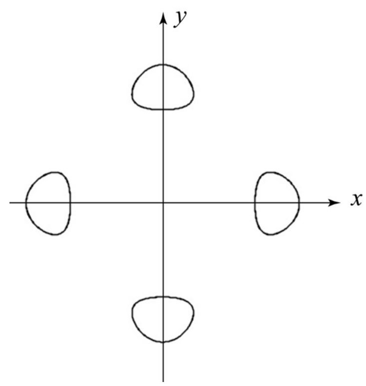

# of distinct real roots:

# of roots w/ multiplicity > 1:

# of complex roots:

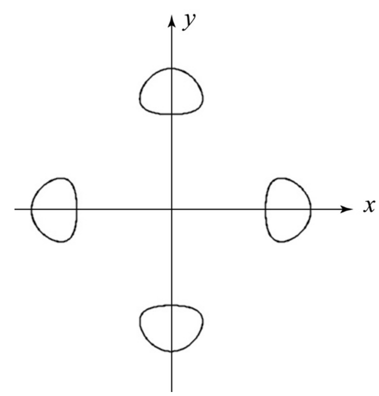

# of distinct real roots:

# of roots w/ multiplicity > 1:

# of complex roots:

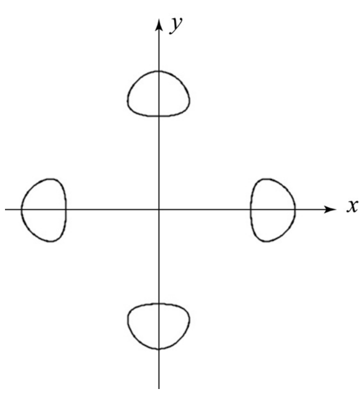

# of distinct real roots:

# of roots w/ multiplicity > 1:

# of complex roots:

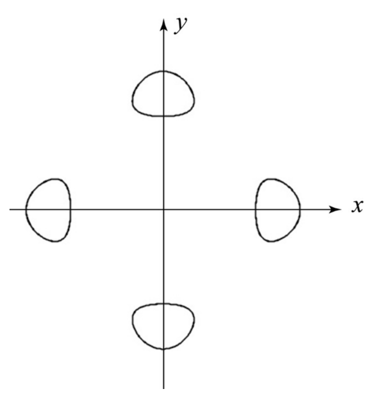

# of distinct real roots:

# of roots w/ multiplicity > 1:

# of complex roots:

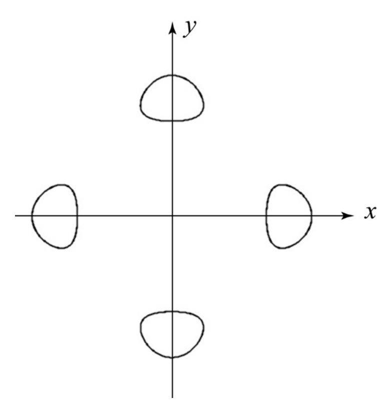

# of distinct real roots:

# of roots w/ multiplicity > 1:

# of complex roots:

Problem 4. Environment mapping (20 points)

One method of environment mapping (reflection mapping) involves using a "gazing ball" to capture an image of the surroundings. The idea is to place a chrome sphere in an actual environment, take a photograph of the sphere, and use the resulting image as an environment map. Let's examine this in two dimensions, using a "gazing circle" to capture the environment around a point.

Below is a diagram of the setup. In order to keep the intersection and angle calculations simple, we will assume that each viewing ray **V** that is cast through the projection plane to the gazing circle is parallel to the z-axis. The circle is of radius 1, centered at the origin.

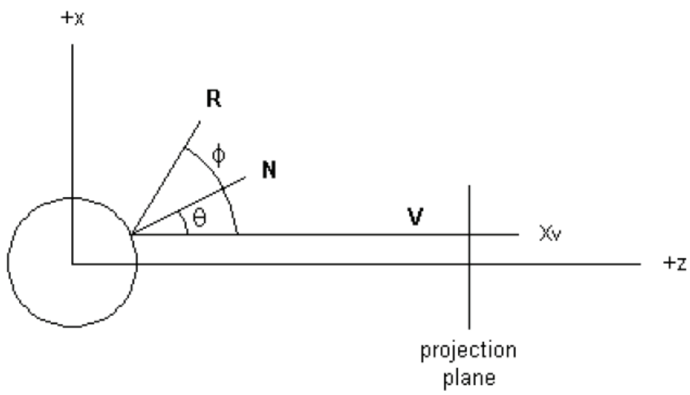

a) (5 points) If the x-coordinate of the view ray is x_v, what are the (x,z) coordinates of the point at which the ray intersects the circle? What is the unit normal vector at this point?

b) (3 points) What is the angle between the view ray **V** and the normal **N** as a function of x_v?

c) (5 points) Note that the angle φ between the view ray **V** and the reflection direction **R** is equal to 2θ, where θ is the angle between **V** and the normal **N**. Plot φ versus x_v. In what regions of the image do small changes in the x_v coordinate result in large changes in the reflection direction?

d) (4 points) We can now treat the photograph of the chrome circle as an environment map (for a 2D world). If we were to ray-trace a synthetic, shiny object and index into the environment map according to each reflection direction, would we expect to get the same rendering as if we had placed the object into the original environment we photographed? Why or why not? In answering the question, you can neglect viewing rays that do not hit the object, assume that the new object is not itself a chrome sphere, and assume that the original environment is some finite distance from the chrome sphere that was originally photographed.

e) (3 points) Suppose you lightly sanded the chrome circle before photographing it, so that the surface was just a little rough.

- What would the photograph of the circle look like now, compared to how it looked before roughening its surface?
- If you used this image as an environment map around an object, what kind of material would the object seem to made of?
- If you did not want to actually roughen the object, what kind of image filter might you apply to the image of the original chrome circle to approximate this effect?

Problem 5. Bezier splines (26 points)

Consider a Bezier curve segment defined by three control points V_0, V_1, and V_2.

a) (4 points) What is the polynomial form of this curve, when written out in the form $Q(u) = A_n u^n + A_{n-1} u^{n-1} + \ldots + A_0$, where n is determined by the number of control points. The coefficients $A_0, \ldots, A_n$ should be substituted in the polynomial equation with expressions that depend on the control points V_0, V_1, and V_2. You may start with recursive subdivision or with the summation over Bernstein polynomials provided in lecture. Either way, show your work.

b) (3 points) What is the first derivative of $Q(u)$ evaluated at $u = 0$ and at $u = 1$ (i.e., what are $Q'(0)$ and $Q'(1)$)? Show your work.

c) (3 points) What is the second derivative of $Q(u)$ evaluated at $u = 0$ and at $u = 1$ (i.e., what are $Q''(0)$ and $Q''(1)$)? Show your work.

d) (5 points) To create a spline curve, we can stitch together consecutive Bezier curves. In this problem, we can add control points W_0, W_1, and W_2. What constraints must be placed on W_0, W_1, and/or W_2 so that, when combined with V_0, V_1, and V_2, the resulting spline curve is C^1 continuous at the joint between the Bezier segments? Write out equations for W_0, W_1, and/or W_2 in terms of V_0, V_1, and/or V_2. (It may be that not all of the W control points are constrained, in which case you would have fewer than three equations.) Show your work. Draw a copy of the control polygon below and place all constrained vertices exactly, and unconstrained vertices wherever you like, and then sketch the spline curve.

e) (5 points) Suppose we wanted to make the spline curve C^2 continuous at the joint between the Bezier segments. Now what constraints must be placed on W_0, W_1, and W_2? Write out equations for W_0, W_1, and/or W_2 in terms of V_0, V_1, and/or V_2. (It may be that not all of the W control points are constrained, in which case you would have fewer than three equations.) Show your work. Draw a copy of the control polygon below and place all constrained vertices exactly, and unconstrained vertices wherever you like, and then sketch the spline curve.

f) (3 points) Is it possible to achieve C^3 continuity with this spline? Explain.

g) (3 points) Suppose that all the control points are points in three dimensions, so that we can create a spline curve in 3-space. Each Bezier curve segment (i.e., the one corresponding to V_0, V_1, and V_2 and the one corresponding to W_0, W_1, and W_2) will lie in a plane, though not necessarily the same plane for both segments. Why? Would this still be the case if the Bezier curve segments were instead defined by four control points each? Explain.