CS 559: Computer Graphics - Final Exam

Dec 17, 2002

Name: _____

ID: _____

Login: _____

- You have 2 hours to complete the exam.
- Before beginning, write your name, ID number and login on the front page, and your login on every page.
- On your desk you may have something to write with, one double-sided piece of paper with anything on it, an optional ruler, and nothing else.
- Do all your work on the pages provided, going to the back side if necessary. If you do use the back, indicate on the front side that there is something on the back.
- If you need to make assumptions in order to answer a question, say what they are. However, all the questions should be unambiguous.

Question 1: ____/12

Question 2: ____/6

Question 3: ____/8

Question 4: ____/4

Question 5: ____/6

Question 6: ____/4

Question 7: ____/3

Question 8: ____/6

Question 9: ____/5

Question 10: ____/3

Question 11: ____/10

Total: ____/67

Question 1: (12 points)

Briefly answer each of the following:

a. How many pixels does Bresenham's algorithm fill when drawing a line from (2,4) to (5,10)?

b. When drawing a convex polygon with a scanline algorithm, at most how many edges will be in the active edge list at any time?

c. Give one reason why rational splines (such as NURBS) would be preferred over B-splines.

d. Is a mipmap useful when a textured object is nearby, or far away?

e. Give one advantage of a subdivision surface over a B-spline surface.

f. If it takes 1 minute to render a 400×400 image with a basic raytracer, approximately how long will it take to render a 800×800 image of the same scene?

Question 2: (6 points) Consider the partially built BSP tree for the scene below. Arrows on the edges point to the "inside."

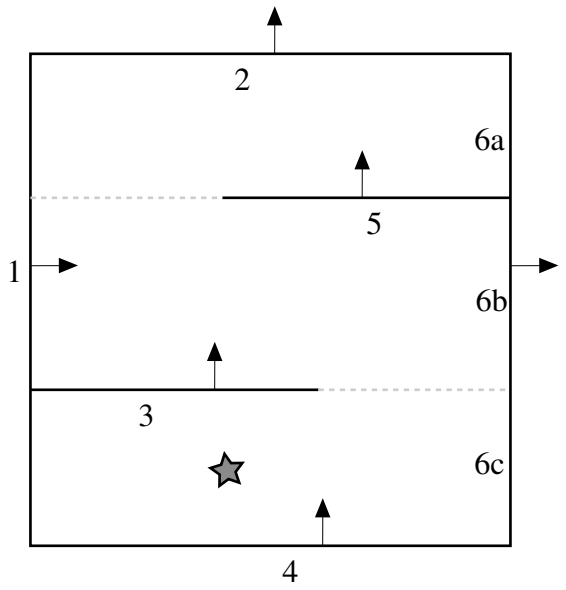

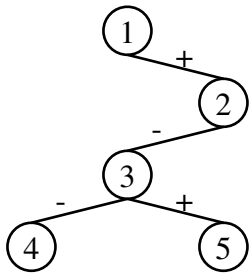

a. Complete the tree by adding nodes for 6a, 6b and 6c.

b. Give the rendering order for back-to-front rendering, derived from the BSP tree, if the viewer is located at the star on the image.

Question 3: (8 points)

Consider the set of polygons shown below, along with a point light source and a viewer location. The normal vectors are specified per-vertex, and Phong shading interpolation is used in lighting the scene.

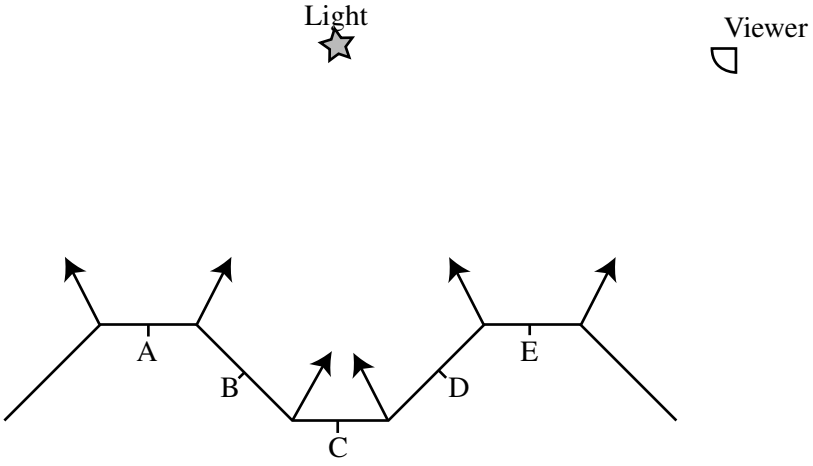

a. The surface is modeled with diffuse reflection only. Which of the points A-E is the brightest, or are you unable to determine which? There may be more than one point with equal maximal brightness, in which case you should give them all.

b. For the same diffuse lighting, which point or points are the darkest, or can you not determine it?

c. The surface is now modeled with a diffuse coefficient of 0.5, a specular coefficient of 0.5, and a specular power of 3. Which point or points are brightest, or can you not determine it?

d. For the same specular lighting, which point or points are the darkest, or can you not determine it?

Question 4: (4 points)

The left image below shows a texture mapped square, labeled with its texture coordinates. The empty square on the right is for you to draw in.

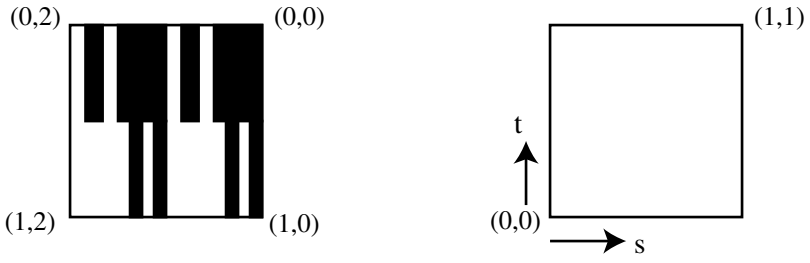

a. Is the texture clamped or repeated in s, or can you not determine it from the image?

b. Is the texture clamped or repeated in t, or can you not determine it from the image?

c. Draw the texture in the space on the right above.

Question 5: (6 points)

A student is planning a polygon mesh data structure in which each vertex stores the normal associated with it, like this:

```
class Vertex {
        float   px, py, pz; // The vertex location.
        float   nx, ny, nz; // The vertex normal.
};
```

a. Is this a convenient way to represent normals if used with flat shading? Why or why not?

b. Suggest an object that this is a good representation for, if used with Gouraud shading. Why?

c. Suggest an object that this is a bad representation for, if used with Grand shading? Why?

Question 6: (4 points)

Consider a Bezier curve with the following control vertices, given in order:

$$\begin{aligned} C_0 &: (-4, -4) \\ C_1 &: (-4, 0) \\ C_2 &: (4, 0) \\ C_3 &: (4, 4) \end{aligned}$$

a. Draw the control polygon on the grid below. Label each vertex with it's coordinates.

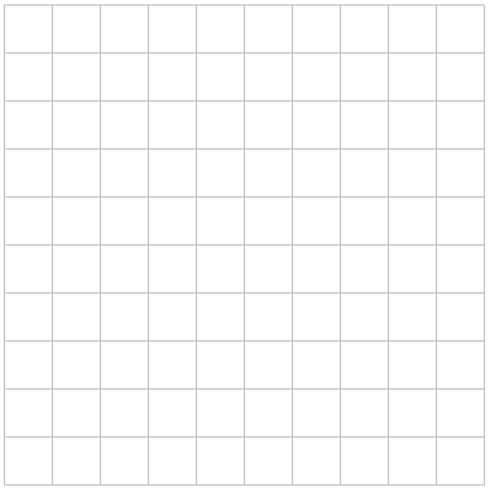

b. The curve is to be subdivided into two sections, one covering the parameter values $0 \leq t \leq 0.5$ and the other the values $0.5 \leq t \leq 1$. What are the control points for the first new piece of curve? Show your working, using the grid below to help you.

Question 7: (3 points)

The cubic B-spline loop below has some repeated control points. Label each control point with the number of times it is repeated.

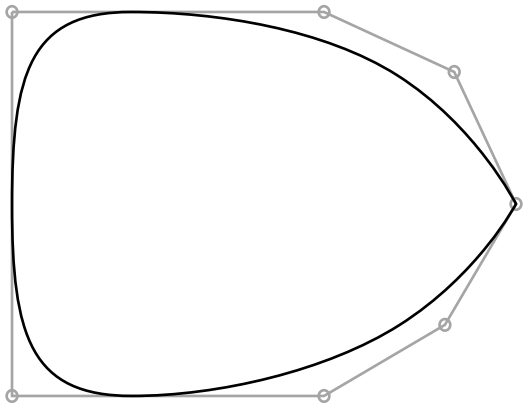

Question 8: (6 points)

Blend surfaces are used to smooth the transition from one surface to another. In the figure below, you are required to design the blend surface between a vertical patch and a horizontal one.

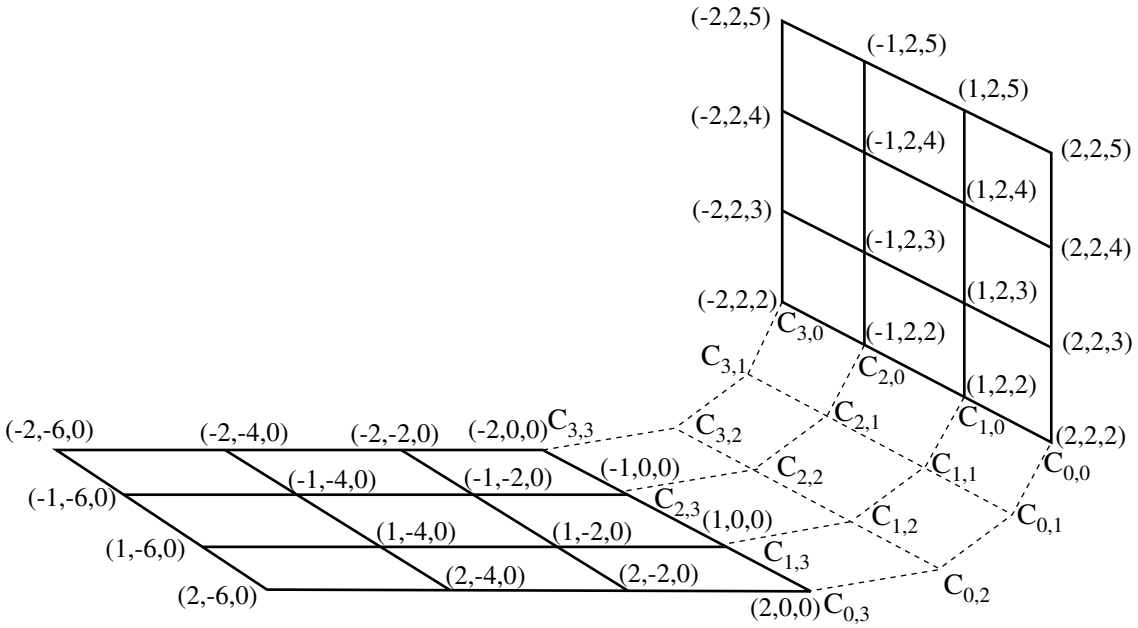

a. Assume that C^1 continuity is required where the blend and vertical patches join. Where should the vertex $C_{0,1}$ be located? Be as precise as you can.

b. Again assuming C^1 continuity, where should the vertex $C_{2,1}$ be located?

c. Now assume a G^1 join between the blend and horizontal patch. Where should the vertex $C_{0,2}$ be located?

d. Again with a G^1 join, where should the vertex $C_{2,2}$ be located?

Question 9: (5 points)

The piece-wise linear curve below is parameterized such that for the segment from V_i to V_{i+1} the parameter varies from i to $i+1$. In other words, $0 \leq t \leq 1$ for the leftmost segment, $1 \leq t \leq 2$ for the next segment, and so on. The parameterization is linear within each segment.

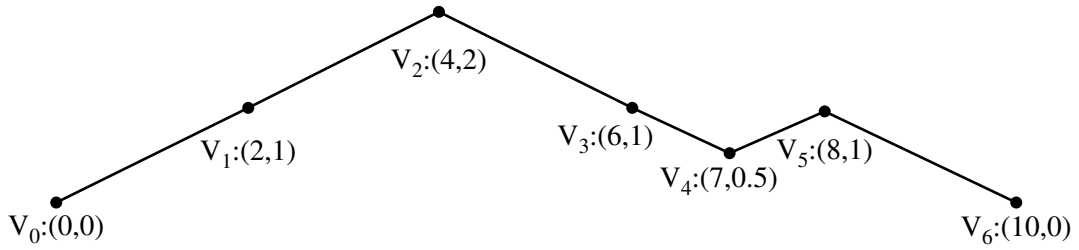

a. Label each join point with its maximum continuity, where C^1 is more continuous than G^1 which is more continuous than C^0.

b. Is it possible for a curve to be C^2 but not C^1? If so, give a situation in which it can occur. The curve above is a hint.

Question 10: (4 points)

These question concern the operation of a basic raytracer.

a. There are several factors that might determine the appropriate number of shadow rays cast toward an area light source for soft shadows. What are two of them?

b. You have a scene containing several diffuse surfaces, one surface that is partly diffuse and partly a mirror, and two point light sources. None of the other surfaces are reflective. What is the maximum number of rays required to compute the brightness of any one pixel?

Question 11: (10 points)

The scenes below contain a range of surfaces, a light, and a viewer. The surfaces are each marked as specular or diffuse. Recall that light paths start at the light and end at the eye.

a. On the figure below draw a light path of the type captured by a raytracer, but not the standard OpenGL lighting model.

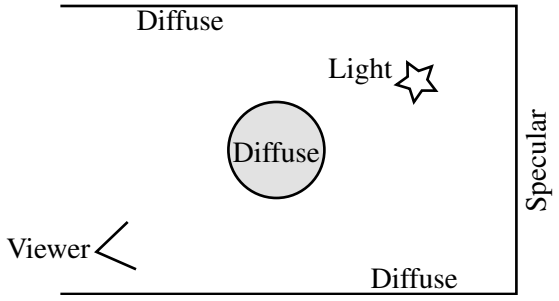

b. On the figure below draw a light path of the type captured by a radiosity algorithm, but not a raytracer or bi-directional raytracer.

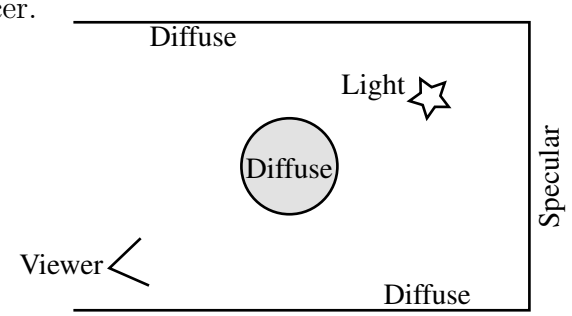

c. On the figure below draw a light path of the type captured by a bi-directional raytracer, but not a regular raytracer or radiosity solver.

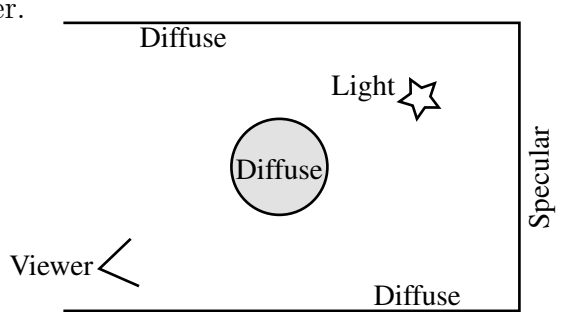

d. Can bi-directional raytracing capture all of the paths caught by regular raytracing?

e. Give two things you could look for in an image that would indicate that a radiosity algorithm had been used to generate it.