CS 559: Computer Graphics - Final Exam

December 17, 2004

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.
- Question 10 can be used to make up points lost on Questions 1-9, but you cannot score more than 72 points in total.
- 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:	___/20
Question 2:	___/8
Question 3:	___/6
Question 4:	___/7
Question 5:	___/6
Question 6:	___/3
Question 7:	___/8
Question 8:	___/10
Question 9:	___/4
Base Total:	___/72
Question 10:	___/4
Total:	___/72

Question 1: (20 points)

Answer each of the following questions with one or two sentences. Each question is worth 2 points.

a. What must you do to correctly render a set of partially transparent objects using a z-buffer in OpenGL?

b. Under what circumstances should mipmaps be used?

c. Consider the problem of modeling Christmas gifts of various sizes wrapped in the same patterned wrapping paper. The basis for your model will be a texture mapped box. Would you account for the different box sizes by applying a simple scaling transformation to a generic box, or by using a parametric model for the box? Why?

d. Would a cube or a tetrahedron be a preferable starting shape for the sphere subdivision scheme discussed in class? Why? Both shapes are sketched below.

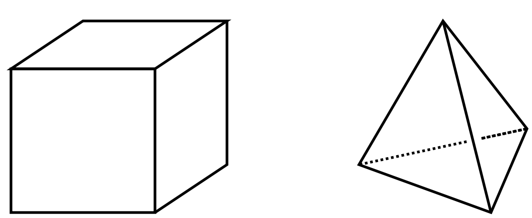

Continued . . .

e. How does the appearance of a specularity change when you increase its specular exponent parameter?

f. You want to construct an open-ended, piecewise cubic, C^2 continuous B-spline curve with **four** consecutive segments (that is, parameter values are in the range $[0, 4)$. How many control points do you need to specify?

g. Describe an object that you would model with hierarchical modeling. Why is a hierarchy the appropriate technique for your object?

h. How would you describe the objects that results from sweeping an annulus (shown below) along a straight path perpendicular to the page?

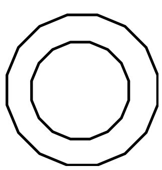

i. Which animation method offers the greatest control over the final appearance of the animation?

j. Describe an animated effect that you would model with a particle system.

Question 2: (8 points)

Consider the Z-buffer and A-buffer hidden surface removal algorithms. You draw the following polygons in the order given:

$z = 0.6$, rgba color (1,0,0,1); $z = 0.2$, rgba color (0,1,0,0.5); $z = 0.4$, rgba color (0,0,1,1)

Increasing depth means further from the viewer, and colors are in the OpenGL rgba format, which is NOT pre-multiplied alpha. Each polygon completely covers the pixel (x, y), and we are interested in what happens at that pixel. When drawing one polygon over another, use the over compositing operator, which in this context is:

$$c_{out} = \alpha_{frag} c_{frag} + (1 - \alpha_{frag}) c_{buf}$$

where c_{out} is the output color, c_{frag} is the incoming color, c_{buf} is the current color in the buffer. For example, (0,0,1,1) on top of (0,1,0,0.5) gives (0,0,1,1) and (0,1,0,0.5) over (1,0,0,1) gives (0.5,0.5,0,1).

a. For the Z-buffer algorithm, fill in the table below indicating the state of the color and depth buffer at (x, y) after each polygon is drawn. The buffers are initialized to black and infinite depth before drawing starts.

Operation	Color	Depth
initialize	(0,0,0)	∞
draw $z = 0.6$, (1,0,0,1)		
draw $z = 0.2$, (0,1,0,0.5)		
draw $z = 0.4$, (0,0,1,1)		

b. For the A-buffer algorithm, fill in the table below with the list associated with pixel (x, y) after each polygon has been processed, before the final pixel value is computed.

Operation	List
initialize	$\rightarrow$ $\boxed{z = \infty, (0,0,0,1)}$
draw $z = 0.6$, (1,0,0,1)	
draw $z = 0.2$, (0,1,0,0.5)	
draw $z = 0.4$, (0,0,1,1)	

c. Again for the A-buffer, what will the final rgba color value be, based on your answer to part (b)?

Question 3: (6 points)

The diagram shows 4 3D polygons. Polygons 1, 2 and 3 are rectangles. The positive side of each polygon is the side you can see. Note that, while you perceive this as a 3D object, the back facing polygons do not exist and should not be considered.

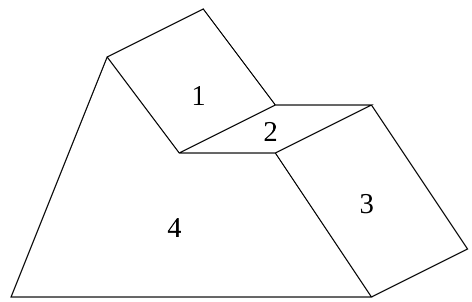

a. Construct the BSP tree for the polygons. Process the polygons in the order 1, 2, 3, 4. If polygons must be split, indicate on the figure above the split line and your labeling of the pieces.

b. Assume the viewer is where you are, looking at the page. What order will your BSP tree render the polygons if they are to be drawn from back to front?

Question 4: (7 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.

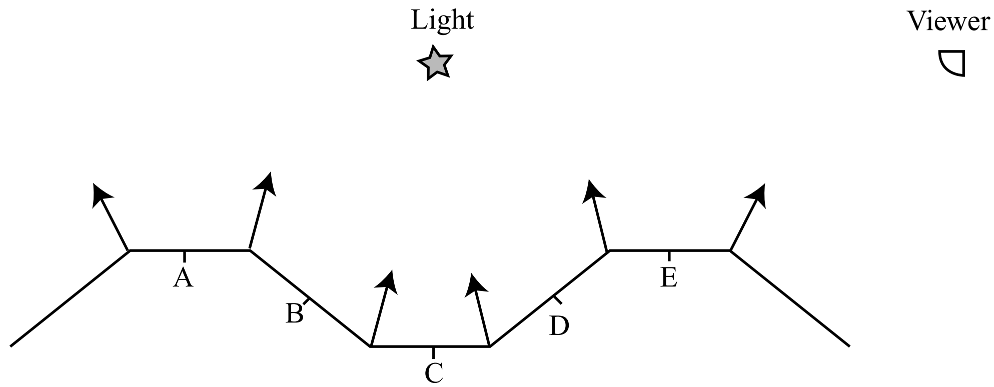

a. Draw, on the figure, the normal vectors that will be used for surface points A-E.

b. 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.

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

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

Question 5: (6 points)

The figure below plots the accurately computed brightness of each point on a flat polygonal surface, as might be computed using the standard lighting model. The vertices of the polygonal surface are shown as small circles, indicating that there are three faces involving four vertices. The brightness at each point is due to some light sources and a particular viewer position, none of which are shown.

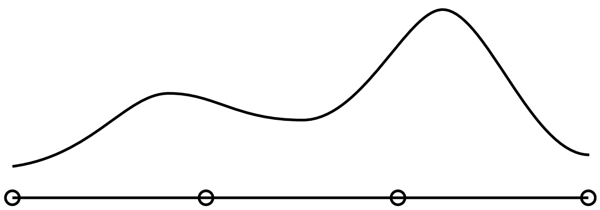

a. On the figure below, plot the intensity that the viewer would see at each point if a flat shading model was used. Assume that the reference point for the flat shading is the leftmost vertex of each face. The accurate values are lightly shown to help you.

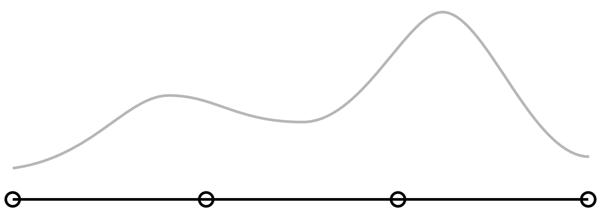

b. On the figure below, plot the intensity that the viewer would see if Gouraud shading interpolation was used.

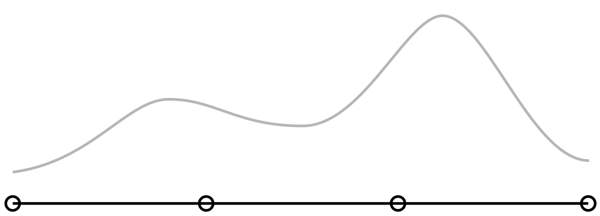

c. On the figure below, plot the intensity that the viewer would see if Phong shading interpolation was used.

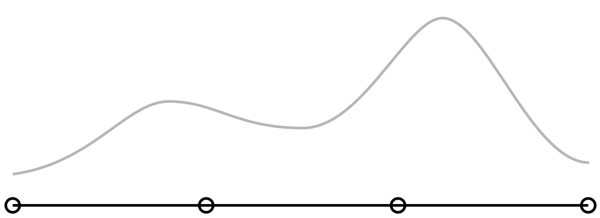

Question 6: (3 points)

Apply the texture image on the left to the polygon on the right. The texture coordinates, (s, t) are shown for each polygon vertex. The texture should be **repeated** in s and **clamped** in t.

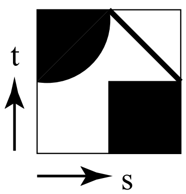

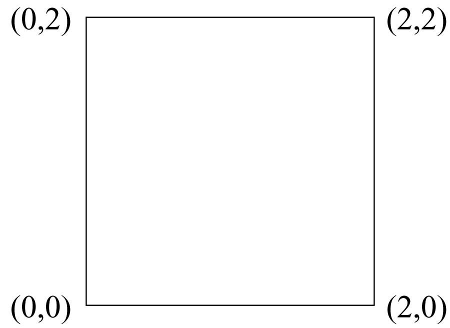

Question 7: (6 points)

You plan to subdivide the mesh below using the modified butterfly scheme. On the left is the 3D shape (a triangulated cube). Center are the vertex locations corresponding to the identifiers shown on the meshes. On the right is the flattened out mesh with the same labeling and connectivity as the cube. The vertex labels and weights for use in the butterfly scheme are given on the final page of the exam, which you can tear off for easy reference.

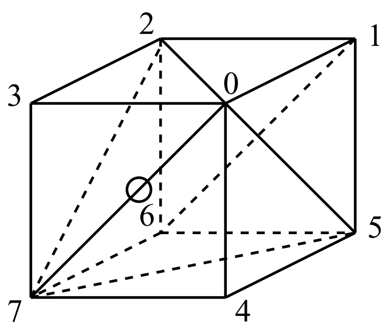


```
0: ( 16, 16, 16)
1: (-16, 16, 16)
2: (-16,-16, 16)
3: ( 16,-16, 16)
4: ( 16, 16,-16)
5: (-16, 16,-16)
6: (-16,-16,-16)
7: ( 16,-16,-16)
```

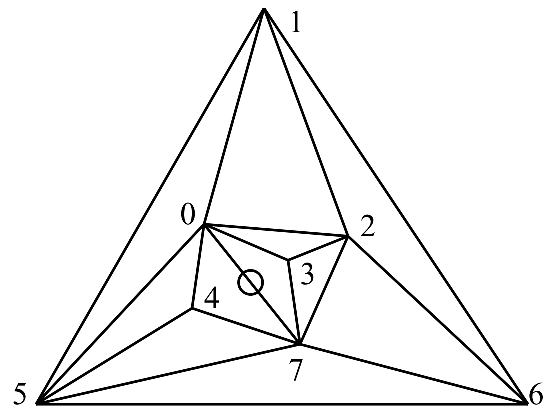

a. Assume we are finding the new vertex between existing vertices 0 and 7, as marked with the small circle. For each of the vertex labels in the mask on the back page, what are the corresponding vertex numbers? Two are given for you, and note that in a mesh like this, some vertices are used for more than one label.

a_0: 0 a_1: 7

b_0: ______ b_1: ______

c_0: ______ c_1: ______ c_2: ______ c_3: ______

d_0: ______ d_1: ______

b. Assume $w = 0$. Give each vertex's contribution to the new point (that is, the weight times the point location). Just give the resulting vector.

a_0: $\frac{1}{2}(16, 16, 16) = (8, 8, 8)$ a_1: $\frac{1}{2}(16, -16, -16) = (8, -8, -8)$

b_0: ____________ b_1: ____________

c_0: ____________ c_1: ____________ c_2: ____________ c_3: ____________

d_0: ____________ d_1: ____________

c. If $w = 0$, what is the location of the new point?

Question 8: (10 points)

In the figure below, a designer is joining two cubic Bezier curve segments in order to form a longer curve. The left segment uses the control points **A**, **B**, **C** and **D** in that order. The right segment will use control points **E**, **F**, **G** and **H**, which are not shown. The coordinate system is indicated, and in the following questions be as precise as possible. For instance, if you know exactly where a point should be, give its coordinates.

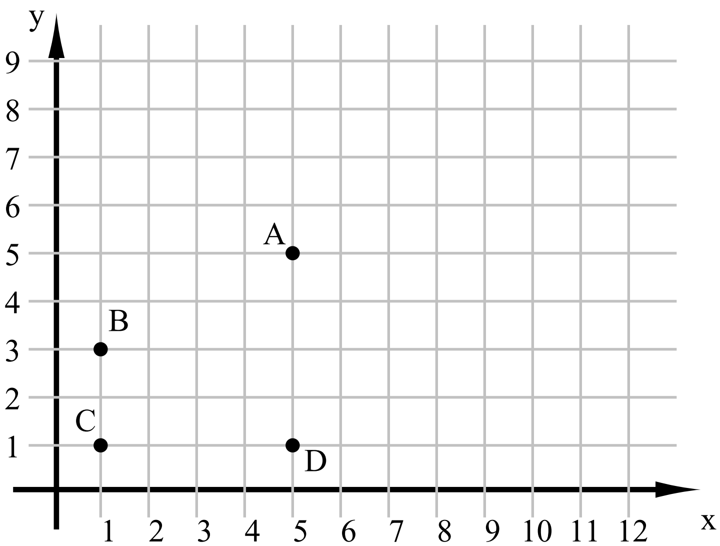

a. Sketch the left segment on the figure. Be sure to respect the properties of Bezier curves.

b. For the combined curve to have C^0 continuity, what, if any, restrictions are there on:
   (i) Point **E**?

   (ii) Point **F**?

c. For the combined curve to have C^1 continuity, what, if any, restrictions are there on point **F**?

d. For the combined curve to have G^1 continuity, what, if any, restrictions are there on point **F**?

e. Which points on the right curve are constrained if C^2 continuity is desired at the join?

Question 9: (4 points)

The cubic B-spline loop below has some repeated control points.

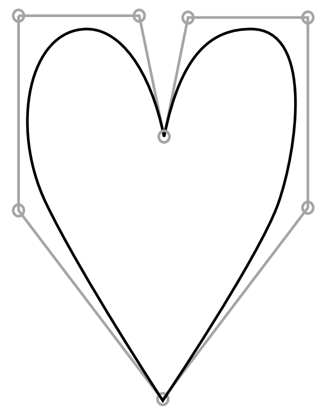

a. Label each control point with the number of times it is repeated.

b. What is one advantage of explicitly representing the above curve as a loop, rather than a linear segment that just happens to meet itself?

Question 10: (Bonus question, 4 points)

These question concerns 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 is one 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. Draw the ray-tree for a ray that leaves the eye and hits the partially mirrored surface first.

Modified Butterfly Scheme

This page is not a question - you can remove it from the exam.

Below is the vertex mask used for computing a new vertex in the modified butterfly subdivision scheme, as presenting in class. Also given is the weight applied to each label, which depends on the parameter w.

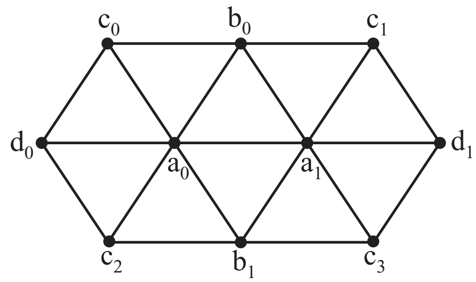

$$
\begin{aligned}
\mathrm{a}_i &: \frac{1}{2} - w \\
\mathrm{b}_i &: \frac{1}{8} + 2w \\
\mathrm{c}_i &: -\frac{1}{16} - w \\
\mathrm{d}_i &: w
\end{aligned}
$$