

## Homework 2

Prepared by: Jon Borchardt, Eugene Hsu and Zoran Popović  
Assigned: Friday, May 4, 2001  
Due: Monday, May 21, 2001

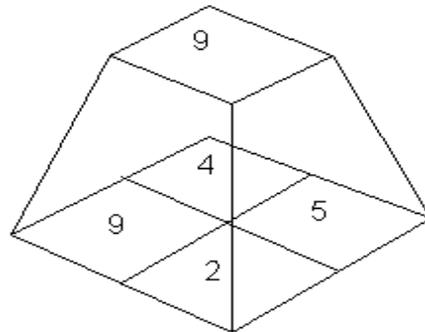
### DIRECTIONS

Please provide short written answers to the questions in the space provided. If you require extra space, you may staple additional pages to the back of your assignment. Feel free to talk over the problems with classmates, but please answer the questions on your own.

NAME: \_\_\_\_\_

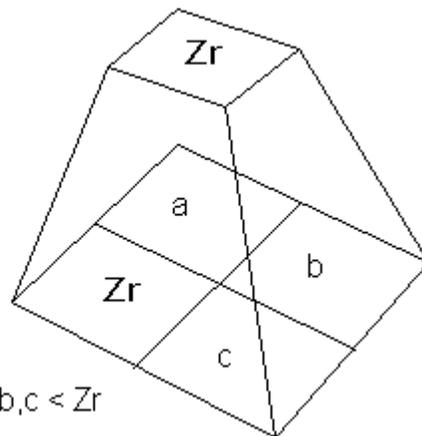
**Problem 1 (14 points)**

The Z-buffer algorithm can be improved by using an image space “Z-pyramid.” The basic idea of the Z-pyramid is to use the original Z-buffer as the finest level in the pyramid, and then combine four Z-values at each level into one Z-value at the next coarser level by choosing the farthest (largest) Z from the observer. Every entry in the pyramid therefore represents the farthest (largest) Z for a square area of the Z-buffer. A Z-pyramid for a single 2x2 image is shown below:



- a. (3 Points) At the coarsest level of the pyramid there is just a single Z value. What does that Z value represent?

Suppose we wish to test the visibility of a polygon **P**. Let **Z<sub>p</sub>** be the nearest (smallest) Z value of polygon **P**. Let **R** be the smallest region in the Z-pyramid that *completely* covers polygon **P**, and let **Z<sub>r</sub>** be the Z value that is associated with region **R** in the Z-pyramid.



where  $a, b, c < Z_r$

### Problem 1 - continued

b. (3 Points) What can we conclude if  $Z_r < Z_p$ ?

c. (3 Points) What can we conclude if  $Z_p < Z_r$ ?

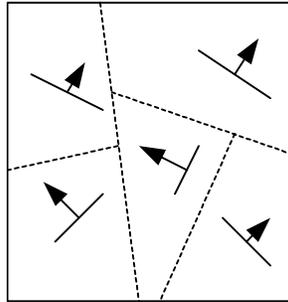
If the visibility test is inconclusive, then the algorithm applies the same test recursively: it goes to the next finer level of the pyramid, where the region  $\mathbf{R}$  is divided into four quadrants, and attempts to prove that polygon  $\mathbf{P}$  is hidden in each of the quadrants  $\mathbf{R}$  of that  $\mathbf{P}$  intersects. Since it is expensive to compute the closest  $Z$  value of  $\mathbf{P}$  within each quadrant, the algorithm just uses the same  $Z_p$  (the nearest  $Z$  of the *entire* polygon) in making the comparison in every quadrant. If at the bottom of the pyramid the test is still inconclusive, the algorithm resorts to ordinary  $Z$ -buffered scan conversion to resolve visibility.

d. (5 Points) Suppose that, instead of using the above algorithm, we decided to go to the expense of computing the closest  $Z$  value of  $\mathbf{P}$  within each quadrant. Would it then be possible to always make a definitive conclusion about the visibility  $\mathbf{P}$  of within each pixel, without resorting to scan conversion? Why or why not?

## Problem 2 (15 points)

BSP trees are widely used in computer graphics. There are many variations that can be used to increase performance. The following questions deal with some of these variations.

For the version of BSP trees that we learned about in class, polygons in the scene (or more precisely, their supporting planes) were used to do the scene splitting. However, it is not necessary to use existing polygons – one can choose arbitrary planes to split the scene:



- a. (3 Points) What is one advantage of being able to pick the plane used to divide the scene at each step? What is one disadvantage of not just using existing polygons?

Recall that when using a BSP tree as described in class, we must draw *all* the polygons in the tree. This is very inefficient, since many of these polygons will be completely outside of the view frustum. However, it is possible to store information at the internal nodes in a BSP tree that will allow us to easily determine if any of the polygons below that node will be visible. If none of the polygons in that sub-tree will be visible, we can completely ignore that branch of the tree.

- b. (4 Points) Explain what extra information should be stored at the internal nodes to allow this, and how it would be used to do this “pruning” of the BSP tree.

## Problem 2 - continued

- c. (4 Points) In class, we talked about doing a “back to front” traversal of a BSP tree. But it is sometimes preferable to do a “front to back” traversal of the tree, in which we draw polygons closer to the viewer before we draw the polygons farther away. (See part (d) for one reason why this is useful) How should the tree traversal order be changed in order to do a front to back traversal?

When we traverse a BSP tree in back to front order, we may draw over the same pixel location many times, which is inefficient since we would do a lot of “useless” shading computations. Assume we instead traverse the tree in front to back order. As we scan convert each polygon, we would like to be able to know whether or not each pixel of it will be visible in the final scene (and thus whether we need to compute shading information for that point).

- d. (4 Points) What simple information about the screen do we need to maintain in order to know if each pixel in the next polygon we draw will be visible or not?

### Problem 3 (12 points)

The Phong shading model can be summarized by the following equation:

$$I_{\text{phong}} = k_e + k_a I_a + \sum_i \left[ I_i \left[ k_d (\mathbf{N} \cdot \mathbf{L}_i)_+ + k_s (\mathbf{V} \cdot \mathbf{R}_i)_+^{n_s} \right] \min \left\{ 1, \frac{1}{a_0 + a_1 d_i + a_2 d_i^2} \right\} \right]$$

where the summation  $i$  is taken over all light sources. The variables used in the Phong shading equation are summarized below:

$$I \quad a_0 \quad a_1 \quad a_2 \quad d_i \quad k_e \quad k_a \quad k_d \quad k_s \quad n_s \quad I_a \quad I_i \quad \mathbf{L}_i \quad \mathbf{R}_i \quad \mathbf{N} \quad \mathbf{V}$$

a. (3 Points) Which of the quantities above are affected if...

- ...the viewing direction changes?
- ...the position of the  $i^{\text{th}}$  light changes?
- ...the orientation of the surface changes? Assume that the change in orientation is about the intersection point.

Blinn and Newell have suggested that, when  $\mathbf{V}$  and  $\mathbf{L}$  are assumed to be constants, the computation of  $\mathbf{V} \cdot \mathbf{R}$  can be simplified by associating with each light source a fictitious light source that will generate specular reflections. This second light source is located in a direction  $\mathbf{H}$  halfway between  $\mathbf{L}$  and  $\mathbf{V}$ . The specular component is then computed from  $(\mathbf{N} \cdot \mathbf{H})_+^{n_s}$  instead of from  $(\mathbf{V} \cdot \mathbf{R})_+^{n_s}$ .

b. (1 Points) Under what circumstances might  $\mathbf{L}$  and  $\mathbf{V}$  be assumed to be constant?

c. (2 Points) How does the new equation using  $\mathbf{H}$  simplify shading equations?



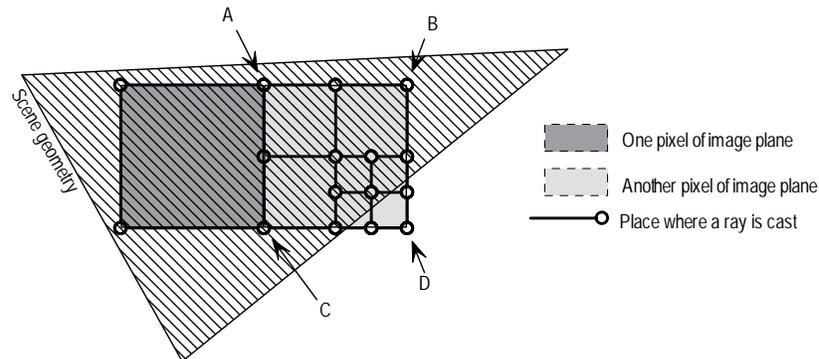
### Problem 4 (10 points)

The company you work for has just bought rights to a raytracing engine. Unfortunately, you don't have the source code, just a compiled library. You have been asked to determine how rays are terminated. So, you call the authors you find out even they don't remember for sure. All they can tell you is this: *The termination criteria for tracing rays is either (a) rays are traced to a maximum recursion depth of 5, or (b) rays are adaptively terminated based on their contribution to a pixel color.*

- a. (4 Points) Describe a scene that can be used to determine which method is used. Be specific about all relevant aspects of the scene and what you would look for in the resulting image to determine which termination method is used.

One of the features included in the raytracing engine your company bought is a brand new algorithm for antialiasing by adaptive supersampling.

The normal implementation is to sample rays at the corner of every pixel, compare the colors of each sample, and if the difference between neighboring sample colors is too great, subdivide that region recursively and sample more times. (See the diagram below, or Foley, et al., 15.10.4)



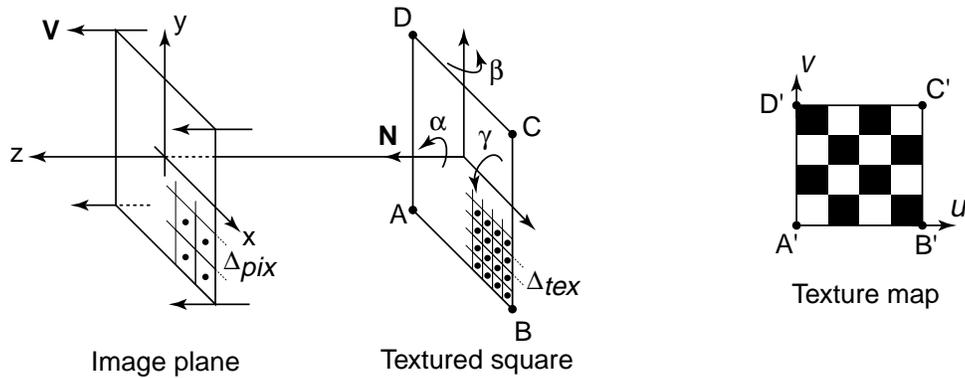
#### **Problem 4 – continued**

However, in this new algorithm, we subdivide and supersample if neighboring rays *intersect different objects*. In other words, note the light-grey pixel above. Three of the four corner samples (a, b, and c) intersect the scene geometry. The fourth corner (d), misses the geometry completely. So we choose to supersample this pixel without ever comparing colors.

b. (6 Points) In what ways is this better than the traditional way? In what ways is it worse?

**Problem 5 (15 points)**

In class, we discussed how brute-force sampling, mip maps, and summed area tables can be employed to anti-alias textures. The latter two techniques average over a region of the texture image very quickly with varying degrees of accuracy, which we consider further in this problem. Consider the scene below: an *orthographic* viewer looking down the  $-z$ -axis views a textured square. The image size and square size are the same and they are initially aligned to one another as shown. The pixel spacing on the image plane and the texel spacing on the square are  $\Delta_{pix}$  and  $\Delta_{tex}$ , respectively.



- a. (4 Points) Assuming  $\Delta_{pix} > \Delta_{tex}$ , how must these sample spacings be related in order for mip mapping to yield the correct values without interpolating among mip map levels?

### Problem 5 – continued

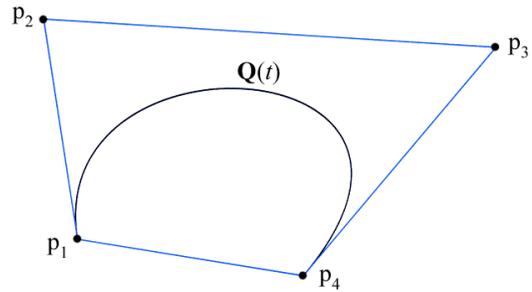
- b. (4 Points) Consider the coordinate system of the square shown in terms of the normal  $\mathbf{N}$  and the two axes aligned with the x and y axes in the figure. Assume that we have the freedom to rotate the square about any *one* of those local axes, as indicated by rotation angles  $\alpha$ ,  $\beta$ , and  $\gamma$ . What restriction do we have on rotation about any one of these axes in order for mip mapping to return the correct average texture values? [For example, you could decide that  $\alpha$ ,  $\beta$ , and  $\gamma$  must all be zero degrees, or you could decide that some of them can vary freely, or you can decide that some can take on a set of specific values. Do not focus on rotations that cause the square to be back-facing.]
- c. (4 Points) Now assume we start again with the unrotated geometry and that we're using summed area tables. If linear interpolation within the summed area table causes no significant degradation, what restriction, if any, should we place on the relative pixel and texel spacings to get correct texture averaging?
- d. (3 Points) As in (b), what restriction must we place on rotation about any one of the given axes in order for summed area tables to return the correct average texture values?

**Problem 6 (15 points)**

Suppose the equations relating the Hermite geometry to the Bezier geometry were of the form  $R_1 = \beta(P_2 - P_1)$ ,  $R_4 = \beta(P_4 - P_3)$ . Consider the four equally spaced Bezier control points  $P_1 = (0,0)$ ,  $P_2 = (1,0)$ ,  $P_3 = (2,0)$ ,  $P_4 = (3,0)$ . Show that, for the parametric curve  $Q(t)$  to have constant velocity from  $P_1$  to  $P_4$  the coefficient  $\beta$  must be equal to 3.

**Problem 7 (15 points)**

A nice property of Bezier curves is that the curve itself will always remain within the *convex hull* of its control points. The convex hull of a set of points is defined as the smallest convex polygon containing all those points. Intuitively, you might imagine the convex hull of a set of points in two dimensional space to be the polygon defined by wrapping a rubber band around those points. In three dimensional space, imagining using a rubber sheet instead.



An intuitively true property about convex hulls is as follows. Suppose we are given  $n$  points; call these  $\mathbf{p}_1, \mathbf{p}_2, \dots, \mathbf{p}_n$ . Now suppose we are given  $n$  real numbers,  $w_1, w_2, \dots, w_n$ . If  $0 \leq w_i \leq 1$  for all  $1 \leq i \leq n$  and  $w_1 + w_2 + \dots + w_n = 1$ , then  $\mathbf{q} = w_1\mathbf{p}_1 + w_2\mathbf{p}_2 + \dots + w_n\mathbf{p}_n$  lies within the convex hull of the points  $\mathbf{p}_1, \mathbf{p}_2, \dots, \mathbf{p}_n$ . In other words, taking a weighted average of a set of points necessarily gives a point within the convex hull of those points.

a. (3 Points) A point on a cubic Bezier curve can be defined by the function

$$\mathbf{Q}(t) = \begin{bmatrix} t^3 & t^2 & t & 1 \end{bmatrix} \begin{bmatrix} -1 & 3 & -3 & 1 \\ 3 & -6 & 3 & 0 \\ -3 & 3 & 0 & 0 \\ 1 & 0 & 0 & 0 \end{bmatrix} \begin{bmatrix} \mathbf{p}_1 \\ \mathbf{p}_2 \\ \mathbf{p}_3 \\ \mathbf{p}_4 \end{bmatrix}$$

where  $\mathbf{p}_1, \mathbf{p}_2, \mathbf{p}_3, \mathbf{p}_4$  are the control points of the curve and  $0 \leq t \leq 1$ . Write out the Bezier basis functions  $f_1(t), f_2(t), f_3(t), f_4(t)$  such that

$$\mathbf{Q}(t) = f_1(t)\mathbf{p}_1 + f_2(t)\mathbf{p}_2 + f_3(t)\mathbf{p}_3 + f_4(t)\mathbf{p}_4.$$

**Problem 7 - continued**

b. (4 Points) Show that  $f_1(t) \geq 0$ ,  $f_2(t) \geq 0$ ,  $f_3(t) \geq 0$ ,  $f_4(t) \geq 0$  for all  $0 \leq t \leq 1$ .

c. (3 Points) Show that  $f_1(t) + f_2(t) + f_3(t) + f_4(t) = 1$  for all  $0 \leq t \leq 1$ .

d. (2 Point) Using the property about convex hulls stated previously, argue that any Bezier curve must lie within the convex hull of its control points. (Make sure you use the convex hull property exactly as it is stated)

e. (3 Points) Give an example of a situation in which the convex hull property of Bezier curves might be useful.