CS535 Fall 2024

Final exam (100 + 10 extra points), December 17, 2024

Name_____

1. Given the normal of a surface at a given point *N* and an incident ray *L*, we need to compute the specular reflection ray *R* at that point to compute its shade. Design an incremental method to compute that vector for points of a triangle. Use the following triangle as an example to illustrate the concept, assuming *R1*, *R2* and *R3* at the three vertices of the triangle are known to us. (10 points)

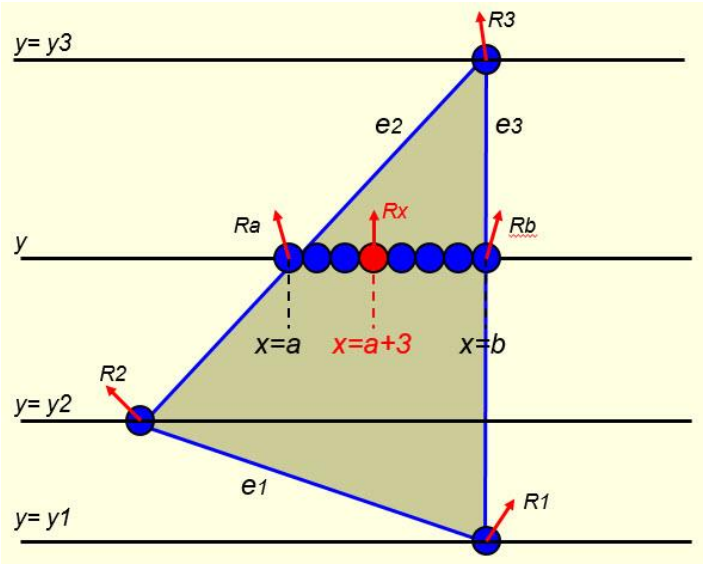

2. Gouraud shading (intensity-interpolation shading) and Phong shading (normal-interpolation shading) can both be used to eliminate intensity discontinuities when rendering a polygonal mesh. However, Gouraud shading would generate the so-called ***Mach band effect*** and Phong shading would not. Can you think of a reason for Gouraud shading to get the Mach band effect? (10 extra points)

3. The *shadow volume* based *'shadow generation*' algorithm can be integrated with the scan-line hidden surface elimination process so that we can do hidden surface elimination and shadow generation at the same time. How are **shadow polygons** used by the scan-line method to determine if a point (pixel) is in shadow? If necessary, draw a figure to illustrate the process. (10 points)

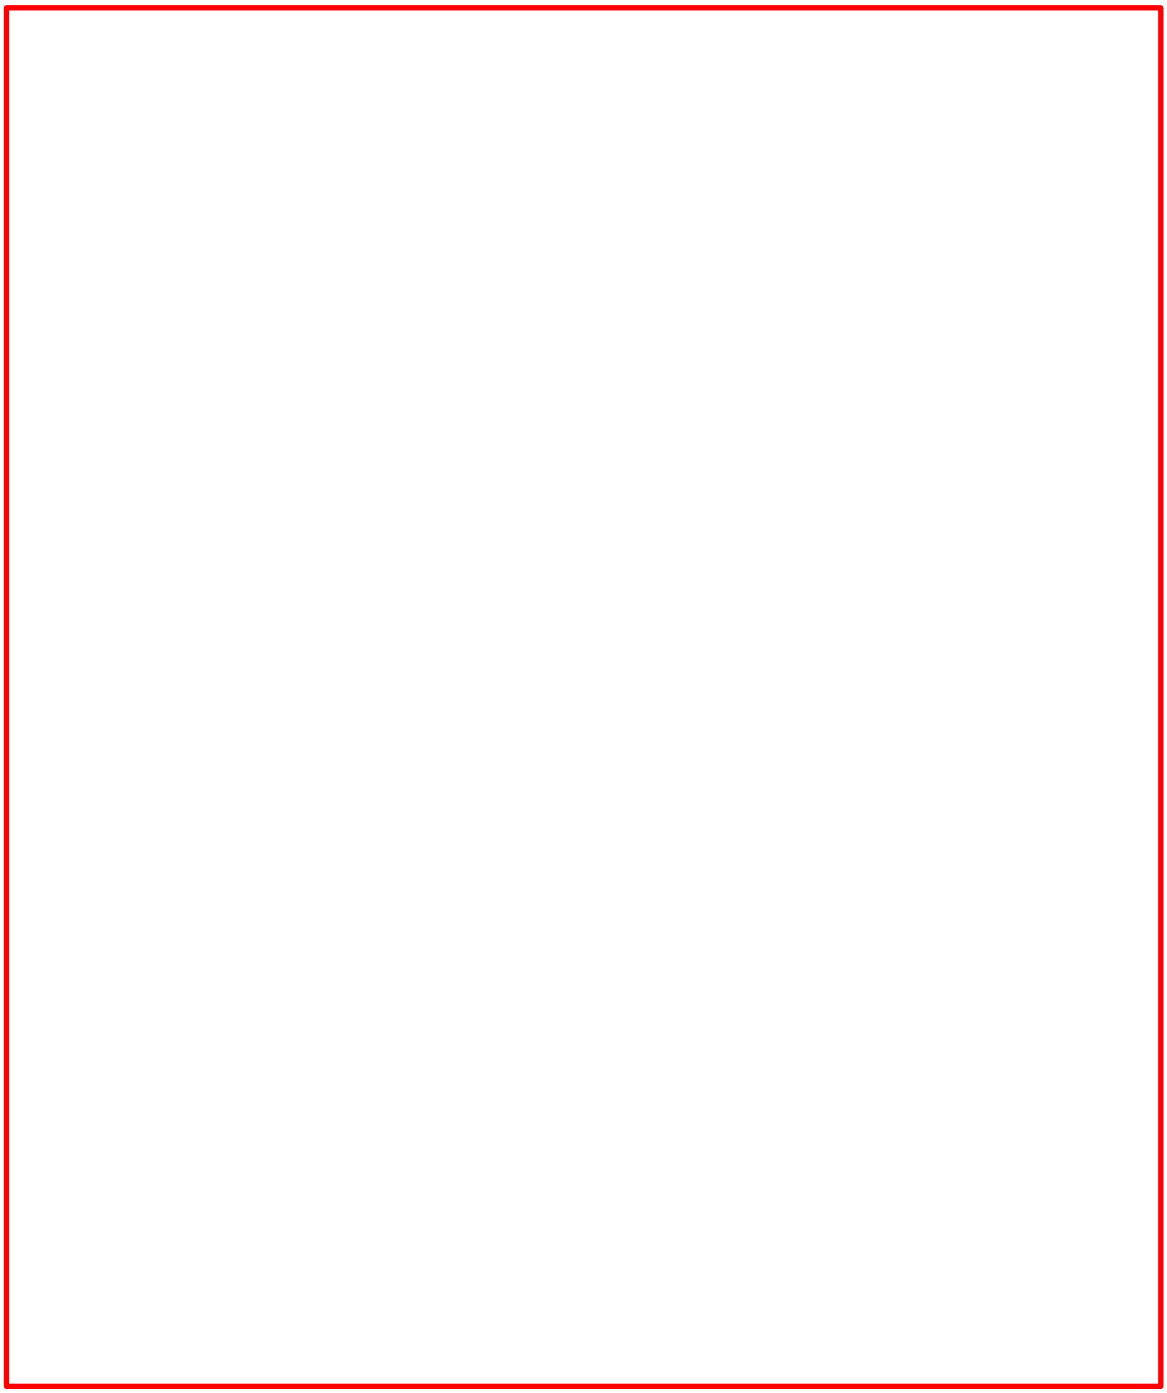

4. The ***shadow map based*** 'shadow generation' algorithm is easy to implement. But it has a potential problem. What is it? What is the reason for getting this potential problem? Is there a way to overcome this potential problem? (10 points)

5. When ray trace an instance of an object transformed by a matrix M, we usually perform the ray tracing process in the space of the original object/primitive. What is the advantage of doing the tracing this way? (10 points)

6. The CSG (Constructive Solid Geometry) tree representation technique introduced in Section 10.10 is not unique, i.e., there are usually more than one CSG representation for a CSG object. Are there occasions that the CSG representation for a CSG object is unique? Either way, justify your answer. (10 points)

7. Can ray tracing reproduce texture of a surface? Justify your answer. (10 points)

8. Why ***perspective correctio***n is necessary when doing texture mapping? How should it be done? Your answer should address two issues here:
   (1) Theoretically, how can it be done?
   (2) Practically, how should it be done efficiently?
   The term "efficiently" means the process is so efficient that it can be integrated with the triangle rasterization process without changing its performance much at all. (10 points)

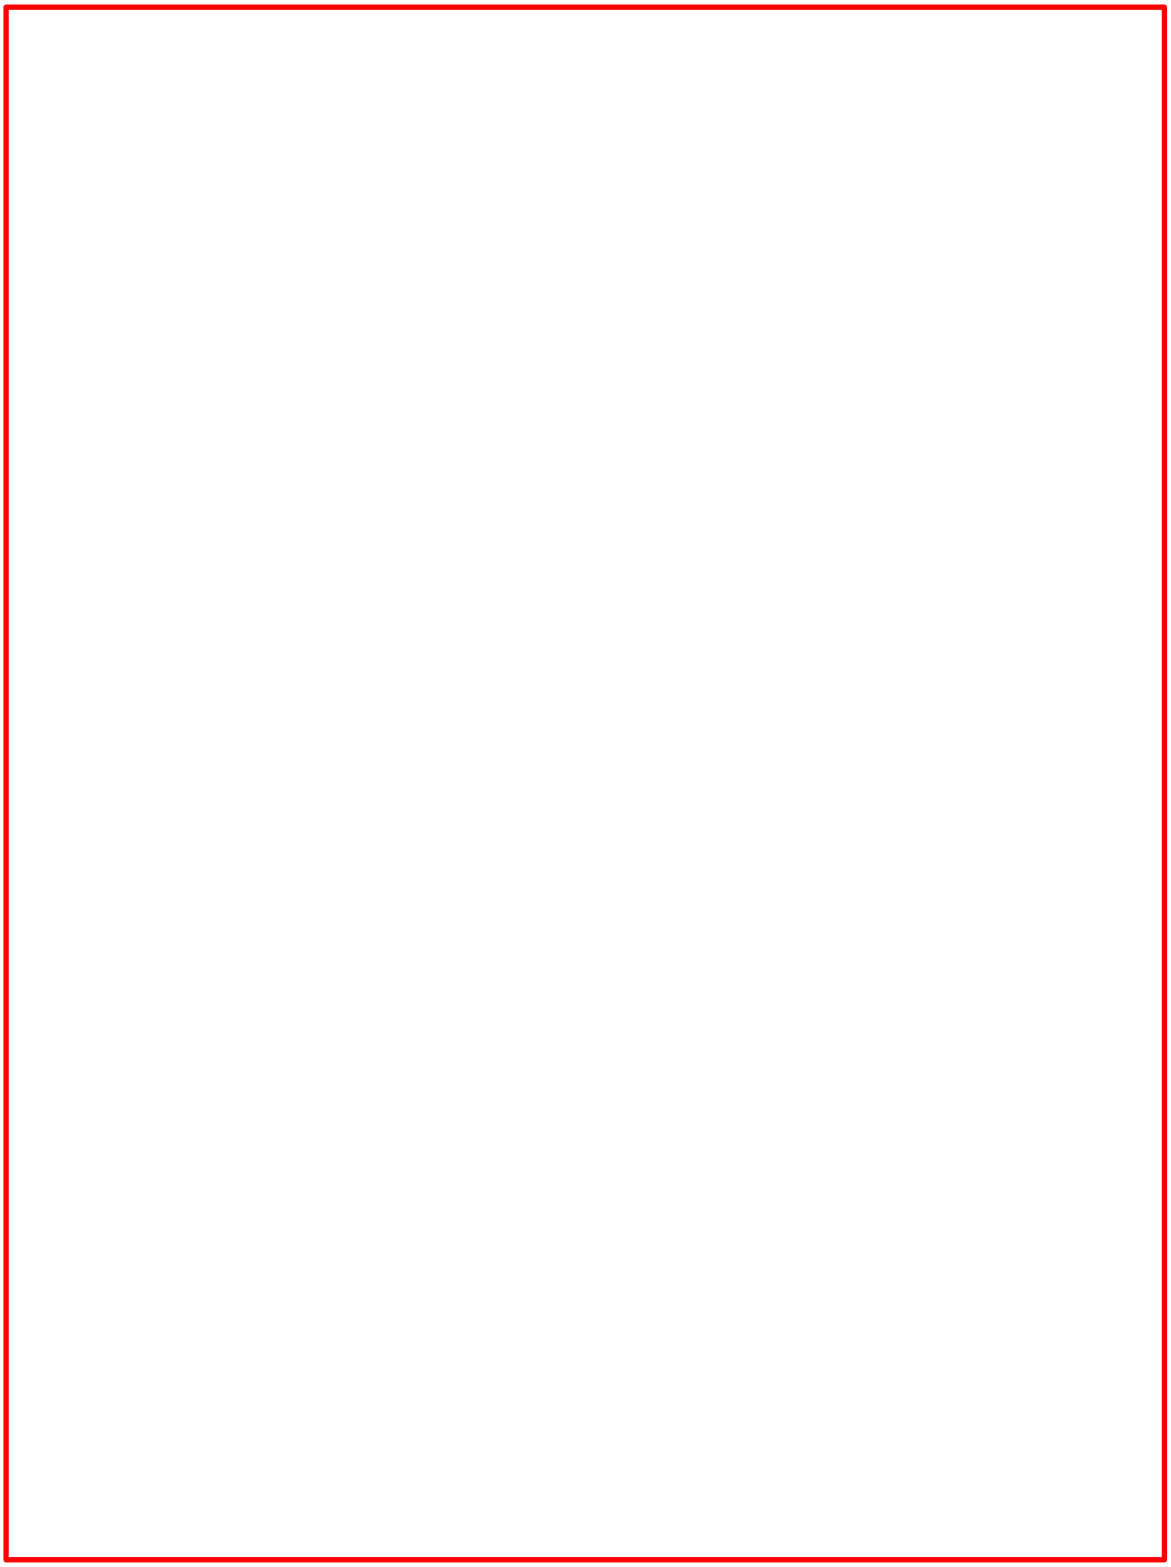

9. **Clipping** is not necessary for the ray tracing process? **Why?** (10 points)

10. A modern CPU can have 4 or 8 cores, but a modern GPU can have thousands. These GPU cores can be used for computationally intensive tasks through the use of ***compute shaders***. Compute shaders are programmed in GLSL and run independently. A ***compute shader*** can perform parallel computing in the following sense: if a ***compute shader*** is required to perform a task on *n* different data sets, one can first creates n copies of the compute shader (invokes the compute shader n times) and then assign each copy of the compute shader a different data set (assign a different task ID (invocationID)). These copies of the ***compute shader*** then run in parallel to perform the task on assigned data sets. In the following box, explain how these two things are implemented in a ***compute shader*** program, especially the GLSL commands/variables needed for these two steps. (10 points)

11. Given a virtual object represented as a CSG (Constructive Solid Geometry) tree, one can use ray casting or even ray tracing technique to render this virtual object on screen. To use ray casting technique to render a CSG object, we need to find the intersection points of each ray with the object. For instance, for the CSG object given below (left figure), for the given ray, we need to find the parameters of the intersection points of the ray with the object (right figure).

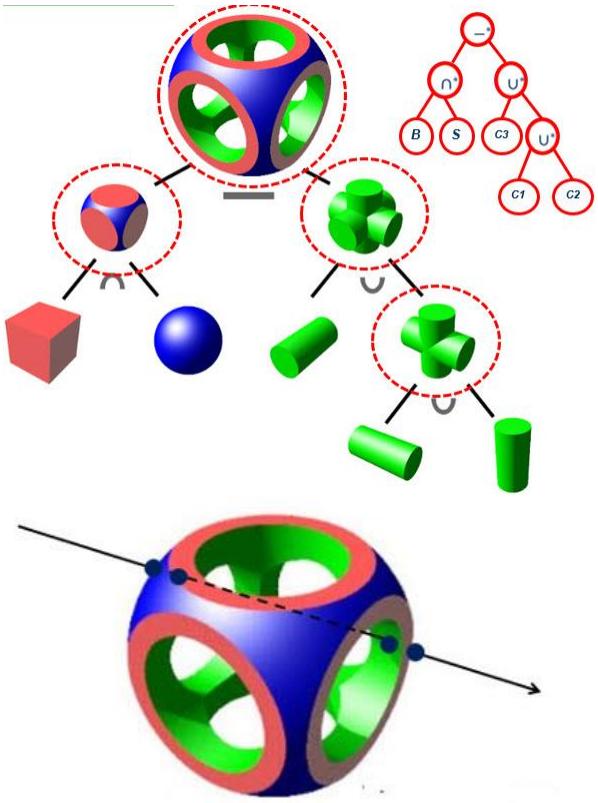

In the following, use the two given cases (intersection of a 2D sphere and a 2D cube, such as the intersection of B and S in the above CSG representation) to explain which two parameters should be reported for each case and why. (10 points)