Time allowed: 2 Hours

This paper contains 5 questions, each worth 12 marks.

Candidates should attempt ALL questions.

TOTAL: 60 marks

Please start each question on a separate page.

Exam papers are to be collected with the examination answer booklets.

Calculators with approved sticker are permitted.

Note: 6 questions are given in the sample paper here, but there will be only 5 questions in the actual exam.

Question 1. **(12 marks)**

a) (5 marks) Explain the OpenGL *synthetic camera* model, including how it relates to a *pinhole camera* and how it relates 3D space to 2D images.

b) (3 marks) Describe the shape of the OpenGL default view volume. Describe what you would do if you want to use a view volume other than the default one (**Hint:** think about other shapes of the view volume).

c) (2 marks) What is clipping in the context of OpenGL, and where does it fit in the steps of the OpenGL pipeline architecture?

d) (2 marks) What is rasterization in the context of the OpenGL pipeline architecture, and what are the inputs and outputs to the rasterization step?

Question 2. **(12 marks)**

a) (6 marks) Describe *request mode* and *event mode*. What is the difference between them? Use appropriate diagrams in your explanation.

b) (2 marks) What is the purpose of the *idle callback* in OpenGL? Include an example of what it can be used for.

c) (2 marks) What is *double buffering*? When would it be suitable to use double buffering?

d) (2 marks) Give three common applications of vertex shaders in OpenGL.

Question 3. **(12 marks)**

a) (6 marks) What is the difference between the Phong model of shading and the modified Phong/Blinn-Phong model, and what are their relative advantages?

b) (2 marks) Explain the difference between a directional light and a point source light in the context of OpenGL.

c) (4 marks) Why is a fourth component w added to the usual three dimensional coordinates in computer graphics systems like OpenGL? Give two common uses of the extra entries in the corresponding 4×4 matrices (those in the last row and column).

Question 4. **(12 marks)**

a) (5 marks) Explain clearly what each function call does in the code below:

```
glutInit(&argc, argv);
glutInitDisplayMode(GLUT_DOUBLE | GLUT_RGBA);
glutInitContextVersion( 3, 2 );
glutInitContextProfile( GLUT_CORE_PROFILE );
glutCreateWindow("Example");
glutMouseFunc(mouse);
glutReshapeFunc(reshape);
glutDisplayFunc(display);
init();
glutMainLoop();
```

b) (3 marks) Suppose you want to build a program using OpenGL that draws two kinds of objects - shiny billiard balls, and dull wooden balls. How would you implement these in a way that realistically represents the difference in shininess between the two kinds of objects?

c) (2 marks) Suppose you have a situation where both the viewer and the only light source are always very distant from the objects in a scene. How could you take advantage of this to perform the lighting calculations more efficiently?

d) (2 marks) What are Euler angles? Briefly explain how you would represent a rotation using these angles.

Question 5. **(12 marks)**

a) (3 marks) Consider a horizontal row of fragments that form part of a triangle during rasterization. Suppose the coordinates, normals and texture coordinates of the first and last fragment in the row have already been calculated. How can we quickly calculate the coordinates of the other fragments as we scan along the row from the first fragment to the last fragment?

b) (2 marks) The *painter's algorithm* for hidden surface removal is an alternative to the *z-buffer* algorithm that instead sorts scene objects by distance and then draws them starting from the furthest away, with each overwriting (or "overpainting") what has previously been drawn. Describe two situations where this algorithm will not work well.

c) (2 marks) What is *mipmapping*? What is the advantage of using mipmapping?

d) (2 marks) Describe the two options for texture sampling in OpenGL. Contrast these options in terms of their rendering outputs and computation times.

e) (3 marks) What is aliasing? Describe a way to ease the aliasing problem in texture mapping.

Question 6. **(12 marks)**

a) (3 marks) What kinds of objects are generally suitable for modelling using a subdivision surface technique? Include an example of where using a subdivision surface technique would be appropriate, and another of where it wouldn't.

b) (3 marks) Suppose we apply the Catmull-Clark technique to subdivide surfaces starting with a box shape. What happens to the resulting surface if we add additional vertices near one corner without altering the shape of the box?

c) (3 marks) What are bone weights in the context of animation?

d) (3 marks) In a typical implementation of skinning in OpenGL that takes advantage of shader programs, like that in lectures, say what the main static data is that's passed to the shaders (i.e., not changing between frames) and what the main dynamic data is that's passed to the shaders (i.e., changing each frame).

END OF PAPER