

Name _____

GPU Programming
LSU EE 4702-1
Take-Home Final Examination

Wednesday 8 December to Friday 10 December 2021 16:30 CST


Work on this exam alone. Regular class resources, such as notes, papers, solutions, documentation, and code, can be used to find solutions. In addition outside Vulkan and OpenGL references and tutorials, other programming resources, and mathematical references can be consulted. Do not try to seek out references that specifically answer any question here. **Do not discuss this exam with classmates or anyone else**, except questions or concerns about problems should be directed to Dr. Koppelman.

Warning: Unlike homework assignments collaboration is not allowed on exams. Suspected copying will be reported to the dean of students. The kind of copying on a homework assignment that would result in a comment like “See ee4702xx for grading comments” will be reported if it occurs on an exam.

Problem 1	_____	(30 pts)
Problem 2	_____	(30 pts)
Problem 3	_____	(40 pts)
Exam Total	_____	(100 pts)

Good Luck!

Problem 1: [30 pts] An issue faced when approximating a sphere with triangles is the number of triangles to use. If too few are used the sphere looks like a polyhedron, especially around the edges. Using too many can slow things down. One possible solution is to split the triangles in the geometry shader—if necessary. The lower-left image shows a sphere that does not have enough triangles. Some triangles are highlighted (in brown) to make them apparent. Looking along the edge (limb) of the sphere one can see straight lines, which is not what we want to see. The second image shows what happens when each triangle is split into four triangles. (Remember that only some are highlighted in brown shades.) But, looking along the edge we see that this splitting hasn't helped. The reason is that the old triangle and the four new ones all fall in the same plane. The third image shows the way the triangles should be split. The three new vertices are positioned to be on the sphere surface, rather than the plane of the original triangle. The last shows the same sphere as the third, but with a texture applied.



(a) Modify the shader code on the next page so that it splits a triangle into three triangles as described above. The vertices of all triangles must be on the surface of the sphere. The code is provided with some routines that should be helpful. Routine `emit(i)`, to be called in the geometry shader, emits vertex `i`. It is used in the geometry shader main routine, `gs_main()`. Routine `emit(i,j)` emits one vertex, but the vertex attributes are interpolated from geometry shader inputs `i` and `j`. So, `emit(0,1)` will emit a vertex halfway between vertex 0 and 1. That vertex is **not** on the sphere surface.

Solve this problem by first calling `emit` so that the four triangles are emitted, though on the plane of the original triangle. Try to do so using no more than 8 vertices.

Then modify `emit(i,j)` so that the interpolated vertex lies on the sphere surface. Though it is possible to do so using only information currently provided to the geometry shader, it is okay if additional information is passed from the vertex shader, such as the sphere center.

- ☐ Modify `gs_main` so that it emits four triangles as illustrated.
- ☐ Modify `emit(i,j)` so the interpolated vertex lies on the sphere surface.
- ☐ If needed, pass additional information from the vertex shader to the geometry shader.
- ☐ Set `max_vertices` to the number of vertices emitted.

```

void vs_main_instances_sphere()    // Vertex Shader
{
    vec4 pos_rad = sphere_pos_rad[gl_InstanceIndex];
    vec4 center_o = vec4(pos_rad.xyz,1);
    float sphere_radius = pos_rad.w;
    mat4 rot = sphere_rot[gl_InstanceIndex];
    vec4 normr = rot * in_vertex_o;
    vec3 normal_o = normr.xyz;
    normal_e = normalize(gl_NormalMatrix * normal_o );
    vec4 center_e = gl_ModelViewMatrix * center_o;
    vertex_e = center_e + vec4(normal_e * sphere_radius,0);
    tex_coor = in_tex_coor;
    color = sphere_color[gl_InstanceIndex];
}

layout ( location = 0 ) in Data_to_GS {
    vec3 normal_e;
    vec4 vertex_e, color;
    vec2 tex_coor;
} In[];

layout ( triangles ) in;
layout ( triangle_strip, max_vertices = 3 ) out;

void emit(int i0, int i1)
{
    float f = i0 == i1 ? 0 : 0.5;

    normal_e = mix(In[i0].normal_e, In[i1].normal_e, f);

    vertex_e = mix(In[i0].vertex_e, In[i1].vertex_e, f);
    color = mix(In[i0].color, In[i1].color, f);
    tex_coor = mix(In[i0].tex_coor, In[i1].tex_coor, f);
    gl_Position = gl_ProjectionMatrix * vertex_e;
    EmitVertex();
}

void emit(int i) {emit(i,i);}

void gs_main()    // Geometry Shader
{
    emit(0); emit(1); emit(2);
    EndPrimitive();
}

```

(b) The code in the previous part always splits a triangle. But there is no point splitting a triangle if it is already small. Consider a triangle small if an edge spans less than 20 pixels. Also, it is triangles that are near the edge that should be split. Those in the center of our view look good even when relatively large. So it would make sense only to split triangles near the edge.

Modify the geometry shader below so that it assigns `split_tri` to true if the triangle should be split based on the size (in pixels) and whether it is close to the edge.

To help determining the size the window width in height in pixels are provided. What remains is determining the length of an edge, measured in pixels. That can be done using a particular coordinate space.

To determine if something is on the edge use the vertex normal and other information. (It is not necessary to use the exact method of the sphere shader.)

- ☐ Set `is_large` to true iff an edge is larger than 20 pixels.
- ☐ Set `near_edge` to true if the triangle is close to the edge of the sphere as seen by the viewer.



```
void gs_main_b()
{
    vec3 normal0_e = In[0].normal_e;
    vec3 vertex0_e = In[0].vertex_e;
    int w_px = com.win_width;
    int h_px = com.win_height;

    bool is_large = ;

    bool near_edge = ;

    bool split_tri = is_large && near_edge;
    // Assume remainder of geometry shader code is here.
}
```

Problem 2: [30 pts] The code below, once this problem is solved, will show the cards from the pre-final rendered in pairs as shown in the screenshot on the right. Each card shows the entire texture image (once per card). The screenshot on the left shows the original card, and without a texture applied.




(a) First, get the texture right. The code below inserts texture coordinates into the buffer set for the card. Using the diagram show how the texture will appear based on this code. (It won't be blank.) The sampler has been set to repeat the texture.

☐ Using diagram above show texture mapped by code below.

```
World::render_card(vk::CommandBuffer& cb)
{
    auto pC = [&](float x, float y)
    { return pCoord(x,y,0); };
    bset_card.reset( pipe_card );
    bset_card
        << pC(0,7) << pC(1,6) << pC(4,7) << pC(3,6)
        << pC(4,0) << pC(1,4) << pC(0,0) << pC(1,6)
        << pC(0,7);

    for ( pCoord c: bset_cards.pos.vals )
    {
        float x = c.x, y = c.y;
        bset_cards << pTCoord( x, y );
    }

    bset_cards.to_dev();
}
```



(b) Modify the code below so that it shows the entire texture image on the card, as illustrated on the previous page.

☐ Modify code to show texture image as shown on right image.

```
World::render_card(vk::CommandBuffer& cb)
{
    auto pC = [&](float x, float y) { return pCoor(x,y,0); };
    bset_card.reset( pipe_card );
    bset_card
        << pC(0,7) << pC(1,6) << pC(4,7) << pC(3,6)
        << pC(4,0) << pC(1,4) << pC(0,0) << pC(1,6) << pC(0,7);

    for ( pCoor c: bset_cards.pos.vals )
    {
        // This loop iterates through the coordinates inserted above.
        // That is, on the first iteration c=pCoor(0,7,0), on the second
        // iteration c=pCoor(1,6,0), etc.

        float x = c.x, y = c.y;

        bset_cards << pTCoor( x, y );
    }

    bset_cards.to_dev();
}
```

(c) Modify the code below so that it shows the second card placed to the left of each original card and rotate 30 degrees. The right edge of the new card should touch the left edge of the old card. Do this by inserting a new transformation matrix into `buf_card_xform`. The code below starts off a new transform, `t2`. It needs to be finished and inserted into the buffer.

- ☐ Add a transform so that there is a new card to the left of each existing one, as described above.

```
buf_card_xform.clear();
for ( auto& ball: balls ) {
    pMatrix ctform =
        pMatrix_Translate(ball->position)
        * ball->omatrix
        * pMatrix_Translate(pVect(0,ball->radius * 1.5,0))
        * pMatrix_Rotation(pVect(0,1,0),M_PI)
        * pMatrix_Scale( ball->radius / 3.5 );

    buf_card_xform << ctform;

    pMatrix t2 = ctform * pMatrix_Scale(1);

}

buf_card_xform.to_dev();
pipe_card.storage_bind( buf_card_xform, "BIND_XFORM" );
transform.use_global_for(pipe_card);
pipe_card.record_draw_instanced(cb,bset_card,buf_card_xform.size());
```

Problem 3: [40 pts] Answer each question below.

(a) A piece of information that we often compute and send down the rendering pipeline is the eye-space normal, often called `normal_e`. What is that used for, and how would the image be affected if it were set to always point up.

☐ In the pipelines used in class the eye-space normal is used for:

☐ How might an image, say a sphere, appear if `normal_e` always pointed up.

(b) In ray tracing, the ray generation shader casts a ray from the eye to each pixel. For each ray the frame buffer is written for the geometry item closest to the eye using the color provided (typically) by a closest-hit shader. In rasterization rendering (the material before ray tracing) how do we insure that the item closest to the eye is written to the frame buffer?

☐ Rasterization rendering ensures the frame buffer is written with an item closest to the eye by:

(c) In class we prepared a buffer set holding the vertices of a sphere, and used that for rendering multiple spheres in a scene. We did that for the entire sphere. (This problem is about the tessellated sphere, not the true sphere shader.) That sounds wasteful because the whole sphere is never visible. Why not just store vertices for half of a sphere, say the front half ($-z$). Explain the flaw in the following argument:

One cannot save space by storing vertices for half of a sphere because the eye can be moved to any position around the sphere, and so the eye could move to the part in which the vertices were omitted.

☐ The flaw in the argument above is:

(d) How does the shape of the view volume in a rasterization rendering compare to the shape of the view volume using the following ray-generation shader:

```
void main() {
    const vec2 pixel_p = vec2(gl_LaunchIDNV.xy) + vec2(0.5);
    vec2 inUV = pixel_p/vec2(gl_LaunchSizeNV.xy);
    inUV.y = 1 - inUV.y;
    vec2 pixel_c = inUV * 2.0 - 1.0;

    vec4 eye_g = ut.mvi * vec4(0,0,0,1);
    vec4 pixel_e = ut.pi * vec4(pixel_c, 1, 1) ;
    vec4 eye_to_pixel_g = ut.mvi * vec4(normalize(pixel_e.xyz), 0) ;

    uint rayFlags = gl_RayFlagsOpaqueNV, cullMask = 0xff;
    float tmin = 0.001, tmax = 10000.0;

    traceNV
    ( topLevelAS, rayFlags, cullMask,
      0 /* intersect/hit sbtRecordOffset*/,
      0 /* intersect/hit sbtRecordStride*/,
      0 /*missIndex*/,
      eye_g.xyz, tmin, eye_to_pixel_g.xyz, tmax,
      0 /*payload location */);

    imageStore(fb_image, ivec2(gl_LaunchIDNV.xy), vec4(rp_color, 0.0));
}
```

☐ How do the shapes compare. Be specific, use a sketch.

(e) CPU code is compiled ahead of time. In a production setting, it is compiled by the developers and shipped in binary form to the customer. But shader code is typically compiled when the program is run. Why?

☐ Why is it important for shader code to be compiled at run time on the user's system while its okay for CPU code to be compiled in advance by the developer?

(f) Our true sphere shader enabled us to render a perfect sphere. For each sphere one vertex was sent down the pipeline. The geometry shader emitted a rectangle framing the sphere, and the fragment shader computed which point on the sphere is projected on to that fragment. A true cube shader could be designed along the same lines. Would it improve things over the conventional method of rendering cubes? (The conventional method is 12 triangles, perhaps using 6 triangle strips.)

☐ Would a true cube shader make a better-looking cube than the conventional rendering methods? ☐ Explain.

☐ Would a true cube shader reduce computation over the conventional rendering method? ☐ Explain.

(g) A triangle strip is rendered with different colors, perhaps our card:

```
bset_card.reset( pipe_card );
bset_card
    << pC(0,7) << pC(1,6) << pC(4,7) << pC(3,6)
    << pC(4,0) << pC(1,4) << pC(0,0) << pC(1,6) << pC(0,7);

bset_card
    << color_red << color_white << color_green << color_orange
    << color_blue << color_lsu_spirit_purple << color_lsu_spirit_gold
    << color_turquoise << color_gray;
```

The interface block for the fragment shader input uses the `flat` interpolation qualifier for color:

```
layout ( location = 0 ) out Data_to_FS {
    vec3 normal_e;
    vec4 vertex_e;
    vec2 tex_coor;
    flat vec4 color;
};
```

☐ How would the appearance differ if the `flat` qualifier were omitted? Illustrate with a portion, not the whole thing.