

Name Solution_____

GPU Programming
EE 4702-1
Take-Home Pre-Final Examination
Due: 30 November 2018 at 16:30

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 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.

Problem 1 _____ (33 pts)

Problem 2 _____ (34 pts)


Problem 3 _____ (33 pts)

Alias Think of something, quick!_____

Exam Total _____ (100 pts)

Good Luck!

Problem 1: [33 pts] Appearing on the next page is the fragment shader code taken from the solution to Homework 3, in which a texture is applied to the triangular spiral in such a way that it looks like the texture image was cut into strips and glued together.



- (a) Modify the shader code so that the text is rotated by 90° , as shown in the screenshot above. The text aspect ratio must be consistent along segments. As with the homework, the text must be readable without skipped or repeated sections. **Do not** assume or make any changes to the host code. In particular, the fragment shader should work with the same primitives as emitted for the Homework 3 code.
- (b) Modify the shader code so that red lines are drawn separating the lines of text.

It is okay to try out your solution on the Homework 3 package, but put the solution in this exam. That is, the solution **will not** be collected by the TA-bot.

See the Homework 3 solution for detailed comments explaining the shader code.

The solution has been checked into the repo with the Homework 3 code. The shader code is in file `pre-fin-shder.cc` and the code can be run by building and executing `pre-fin`.

Use next page for solution.

- Modify shader so text rotated 90°.
- Draw red lines between groups of text.
- Text aspect ratio and size must be consistent on all segments and there must be no gaps or skipped sections.

See next page for solution.

```
void fs_main_hw03() {
    float line_num = floor(tex_coord.x);
    vec2 tc = vec2( tex_coord.x, ( line_num + tex_coord.y ) * tex_ht );
```

```
vec4 color =
    gl_FrontFacing ? gl_FrontMaterial.diffuse : gl_BackMaterial.diffuse;

// No need to modify code below.
vec4 texel = texture(tex_unit_0,tc);
vec3 nne = normalize(normal_e);
float edge_dist = tex_coord.y;
vec3 bnorm = tryout.x ? nne : edge_dist > 0.9
    ? normalize(mix(nne,spiral_normal,2*(edge_dist-0.9))) : edge_dist < 0.1
    ? normalize(mix(-spiral_normal,nne,0.8+2*edge_dist)) : nne;
vec4 lighted_color =
    generic_lighting( vertex_e, color, bnorm, gl_FrontFacing );
gl_FragColor = texel * lighted_color;
gl_FragDepth = gl_FragCoord.z;
}
```

Solution appears below. A common mistake was not continuing a line.

```
void
fs_main_hw03()
{
    /// Pre-Final Problem 1 SOLUTION

    int ncol = 6;    // Number of pieces a line is split into.
    int nlines = 60; // Number of lines per page.

    float s_tex_coord_x = tex_coord.x * tex_ht * nlines;
    float s_line_num = floor(s_tex_coord_x);
    float s_line_offs_01 = fract(s_tex_coord_x);
    float s_line_offs_cols = s_line_offs_01 * ncol;
    float s_col_num = floor(s_line_offs_cols);
    float s_col_offs = fract(s_line_offs_cols);
    float tex_x = ( s_col_num + tex_coord.y ) / ncol;
    float tex_y = ( s_line_num + s_col_offs ) / nlines;

    vec2 tc = vec2(tex_x,tex_y);

    vec4 color =
        tex_x < 0.01 ? vec4(0.1,0.1,0.1,1) :
        s_col_offs < 0.03  s_col_offs > 0.97 ? vec4(1,0,0,1) :
        gl_FrontFacing
        ? gl_FrontMaterial.diffuse : gl_BackMaterial.diffuse;

    /// No need to modify code below.
    vec4 texel = texture(tex_unit_0,tc);
    vec3 nne = normalize(normal_e);
    float edge_dist = tex_coord.y;
    vec3 bnorm = tryout.x ? nne : edge_dist > 0.9
        ? normalize(mix(nne,spiral_normal,2*(edge_dist-0.9))) : edge_dist < 0.1
        ? normalize(mix(-spiral_normal,nne,0.8+2*edge_dist)) : nne;
    vec4 lighted_color =
        generic_lighting( vertex_e, color, bnorm, gl_FrontFacing );
    gl_FragColor = texel * lighted_color;
    gl_FragDepth = gl_FragCoord.z;
}
```

Problem 2: [34 pts] Recall that the true sphere code covered in class emits a square (two triangles) for each sphere to be rendered. The square is chosen so that it perfectly frames the sphere as seen by the viewer. The square is perfect in the sense that if it were made any smaller parts of the sphere would not be visible. But it is still wasteful because the sphere is not visible in the corners of the square and so those fragment shader invocations are wasted.

(a) Modify the true-sphere geometry shader below so that it emits a $2s$ -sided polygon rather than a square. Show the code for computing the polygon coordinates, **do not** use the assumed functions from the next problem. Note that when $s = 2$ the code should emit a bounding square, which is what the original code does.

Use next page for solution.

- Modify code to emit 2s-sided bounding polygon.
 Take advantage of the fact that 2s is an even number.
- See next page for solution.

```

const int s = 8;      // Don't assume that s is always 8!!
layout ( points ) in;
layout ( triangle_strip, max_vertices = 4 ) out;

void gs_main() {
  vertex_id = In[0].vertex_id;
  mat4 mvp = trans_proj * gl_ModelViewMatrix;
  vec4 pos_rad = sphere_pos_rad[vertex_id];
  vec3 ctr_o = pos_rad.xyz;
  float r = pos_rad.w;
  vec3 e_o = vec3(gl_ModelViewMatrixInverse*vec4(0,0,0,1)); // Eye location in object space.

  vec3 ec_o = ctr_o - e_o; // Eye to Center (of sphere).
  float ec_len = length(ec_o);
  vec3 ne = ec_o / ec_len;

  // Find vectors orthogonal to ec_o.
  vec3 atr_o = abs(ec_o);
  int min_idx = atr_o.x < atr_o.y ? ( atr_o.x < atr_o.z ? 0 : 2 )
    : ( atr_o.y < atr_o.z ? 1 : 2 );
  vec3 nx_raw = min_idx == 0 ? vec3(0,-ec_o.z,ec_o.y)
    : min_idx == 1 ? vec3(-ec_o.z,0,ec_o.x) : vec3(-ec_o.y,ec_o.x,0);
  vec3 nx = normalize(nx_raw), ny = cross(ne,nx);

  // Compute center of the limb, lmb_o. (The limb is the sphere outline.)
  float sin_theta = r / ec_len;
  float a = r * sin_theta;
  vec3 lmb_o = ctr_o - a * ne;

  // Compute axes to draw the limb.
  float b = r * sqrt( 1 - sin_theta * sin_theta );
  vec3 ax = b * nx, ay = b * ny;

  // Emit a bounding square for the limb. --- Solution should start here.
  for ( int i = -1; i < 2; i += 2 ) for ( int j = -1; j < 2; j += 2 ) {
    vertex_o = lmb_o + ax * i + ay * j;
    gl_Position = mvp * vec4(vertex_o,1);
    EmitVertex(); }
  EndPrimitive();
}

```

Solution appears below. A single triangle strip is used for the entire polygon taking advantage of the fact that the number of sides is even. First, a bounding circle for the polygon is found. A loop iterates over an angle effectively from $\theta = \pi/(4s)$ to $\theta = \pi(1 - 1/4s)$. Coordinates of points on the circle at both $-\theta$ and θ are computed and emitted. In this way only $2s$ vertices need to be emitted. A brute-force approach would emit $2s$ individual triangles (pie slices) of 3 vertices each.

```

void gs_main()
{
    vertex_id = In[0].vertex_id;
    vec4 pos_rad = sphere_pos_rad[vertex_id];
    vec3 ctr_o = pos_rad.xyz;
    float r = pos_rad.w;

    if ( mirrored ) ctr_o.y = -ctr_o.y;

    // Eye location in object space.
    vec3 e_o = vec3(gl_ModelViewMatrixInverse * vec4(0,0,0,1));

    // Vectors from eye to sphere center.
    vec3 ec_o = ctr_o - e_o; // Eye to Center (of sphere).
    float ec_len = length(ec_o);
    vec3 ne = ec_o / ec_len;

    // Vectors orthogonal to ec_o.
    //
    vec3 atr_o = abs(ec_o);
    int min_idx = atr_o.x < atr_o.y ? ( atr_o.x < atr_o.z ? 0 : 2 )
        : ( atr_o.y < atr_o.z ? 1 : 2 );
    vec3 nx_raw = min_idx == 0 ? vec3(0,-ec_o.z,ec_o.y)
        : min_idx == 1 ? vec3(-ec_o.z,0,ec_o.x) : vec3(-ec_o.y,ec_o.x,0);
    vec3 nx = normalize(nx_raw);
    vec3 ny = cross(ne,nx);

    // Compute center of the limb, lmb_o, the circle formed by the most
    // distant visible parts of the sphere surface.
    float sin_theta = r / ec_len;
    float a = r * sin_theta;
    vec3 lmb_o = ctr_o - a * ne;

    // Compute axes to draw the limb.
    float b = r * sqrt( 1 - sin_theta * sin_theta );
    vec3 ax = b * nx;
    vec3 ay = b * ny;

    mat4 mvp = trans_proj * gl_ModelViewMatrix;

#ifdef 0
    // Emit a bounding square for the limb.
    for ( int i = -1; i < 2; i += 2 )
        for ( int j = -1; j < 2; j += 2 )
            {
                vertex_o = lmb_o + ax * i + ay * j;
            }
#endif
}

```

```

        gl_Position = mvp * vec4(vertex_o,1);
        EmitVertex();
    }
    EndPrimitive();
#else

    /// Pre-Final Solution
    //
    const int n_sides = s * 2;

    // Find radius of bounding circle for 2s-sided regular polygon.
    //
    const float pi = 3.1415926535;
    const float delta_theta = pi / s;
    const float delta_theta_h = delta_theta/2;
    const float b_outer = b / cos(delta_theta_h); // Radius of bounding circle.

    // Compute axes for this circle.
    //
    const vec3 bx = b_outer * nx;
    const vec3 by = b_outer * ny;

    for ( int i=0; i<s; i++ )
    {
        const float theta = delta_theta_h + i * delta_theta;

        // Compute vertex coordinates for "left" and "right" half of circle.
        //
        for ( int j=-1; j<2; j+=2 )
        {
            vertex_o = lmb_o + bx * cos(theta) - j * sin(theta) * by;
            gl_Position = mvp * vec4(vertex_o,1);
            EmitVertex();
        }
    }
    EndPrimitive();
#endif
}

```


(b) Let t_f denote the execution time of 1 invocation of the true-sphere fragment shader when the sphere covers the fragment, and let t_n denote the execution time of 1 invocation when the sphere does not cover the fragment. Assume that $t_f > t_n$. Let $t_{g0} + st_{g1}$ denote the execution time of one invocation of the true-sphere geometry shader that emits a $2s$ -sided polygon.

In general, emitting a polygon with more sides reduces the work performed by the fragment shader but increases the work done by the geometry shader, improving overall performance. But there will be a point at which increasing s reduces performance.

Using the symbols defined above **plus other information**, find an expression for the time needed to render a sphere. The other information is needed to compute a sphere rendering time using the timings above. Of course, no credit will be given if that other information is the execution time. Also determine the value of s that will minimize the total time used by the geometry and fragment shaders for rendering a sphere. When computing the s that minimizes time it is okay to apply simplifications based on s being large. (Otherwise a closed-form solution can't easily be found.) *Hint: Yes, the solution requires calculus.*

✓ Indicate the other information that is needed.

The other information is the number of fragments. Let r denote the number of fragments along the limb radius (from the center of the projected sphere to an edge). This value of r is determined by the sphere size, the location of the sphere and user, and the projection matrix.

✓ Show an expression for the time needed to render a sphere using a $2s$ -sided bounding polygon in terms of t_f , etc., and the other information.

The number of fragment within the circle, taking t_f time, is πr^2 . The number within the square bounding box is $4r^2$. The number of fragments within a $2s$ -sided regular polygon with an inner radius of r is $2s \tan \frac{\pi}{2s} r^2$.

When rendering a sphere the geometry shader will be invoked just once and the fragment shader will be invoked $2s \tan \frac{\pi}{2s} r^2$ times, of which πr^2 invocations will take t_f and the remainder t_n . The total time needed, $t(s)$, is then:

$$t(s) = t_{g0} + st_{g1} + \pi r^2 t_f + \left(2sr^2 \tan \frac{\pi}{2s} - \pi r^2 \right) t_n.$$

✓ Find an expression for s that will minimize time in terms of t_f , etc., and the other information.

To minimize $t(s)$ solve $\frac{d}{ds}t(s) = 0$ for s .


$$\frac{d}{ds}t(s) = t_{g1} + 2t_n r^2 \tan \frac{\pi}{2s} - \frac{\pi t_n r^2}{s \cos^2 \frac{\pi}{2s}} = 0$$

Assuming that s is large reduces the equation above to $st_{g1} = 0$, which isn't useful. *Note: I made an error taking the derivative when constructing the problem, with my error it would be possible to solve for s . Any solution attempting to differentiate the expression for time, setting it to zero, and solving for s would get full credit.*

A better solution approach is to approximate the number of non-sphere fragments using a bounding circle for the polygon, one of radius $r / \cos(\pi/2s)$.

Problem 3: [33 pts] The screenshot below is based on the code from Homework 4. The code is in *lenses* mode, meaning that the sphere is visible only at certain points. Unlike Homework 4, the lenses are much further apart from each other, and therefore much less of the sphere is visible. If the code in the Homework 4 solution were used then most invocations of the fragment shader would be wasted because they end up discarding the fragment.

There would be much less waste if a separate primitive were emitted for each lens, rather than one primitive for the entire sphere.



Suppose that each sphere has the same number of lenses, call that number h , and refer to the lenses as lens 0, lens 1, etc. Assume that a library of functions is available for computing positions for each lens. Among these is a function `lens_get_ctr_o(i, sphere_info)` which returns the object-space coordinates of the center of lens i . Variable `sphere_info` holds the value of the sphere center, radius, and orientation. Assume that `sphere_info` has already been prepared. You may assume that related functions exist such as `lens_get_ctr_e(i, sphere_info)` for eye-space coordinates. Also convenient for this problem is function `lens_get_bb_o(i, j, n, sphere_info)` which returns the object-space coordinates of point j of an n -sided bounding polygon for lens i . The bounding polygon is like the bounding square used in the true-sphere code: it tightly fits around the lens as seen by the viewer (or equivalently, when projected on to the frame buffer). The points wrap around the lens in a clockwise direction with increasing j .

Show changes needed to the geometry and fragment shaders on the following pages to render the lenses efficiently for cases in which h is relatively small, say on the order of 10.

Problem 3, continued:

(a) Appearing below is a shortened version of the geometry shader from the Homework 4 solution. The comments have been removed, but feel free to look at comments in the posted solution. Modify the geometry shader below so that it emits a primitive for each lens and provides information needed by the fragment shader (see next problem). Use the functions described above to find coordinates of bounding polygons for the lenses.

- Modify so that it emits a primitive for each lens.
- If necessary, change layout and output declarations.
- Provide information needed by fragment shader (see next subproblem).

The solution appears after the unmodified code.

```
out Data_to_FS // Data to fragment shader
{
    flat int vertex_id;
    vec3 vertex_o;

};

const int h = 8; // Number of holes. Don't assume it is always 8.
layout ( points ) in;
layout ( triangle_strip, max_vertices = 4 ) out;

void gs_main() {
    vertex_id = In[0].vertex_id;
    vec4 pos_rad = sphere_pos_rad[vertex_id];
    vec3 ctr_o = pos_rad.xyz;
    float r = pos_rad.w;
    Sphere_Info sphere_info = make_sphere_info(pos_rad,sphere_rot[vertex_id]);

    // Eye location in object space.
    vec3 e_o = vec3(gl_ModelViewMatrixInverse * vec4(0,0,0,1));

    // Vectors from eye to sphere center.
    vec3 ec_o = ctr_o - e_o; // Eye to Center (of sphere).
    float ec_len = length(ec_o);
    vec3 ne = ec_o / ec_len;

    // Vectors orthogonal to ec_o.
    vec3 atr_o = abs(ec_o);
    int min_idx = atr_o.x < atr_o.y ? ( atr_o.x < atr_o.z ? 0 : 2 )
        : ( atr_o.y < atr_o.z ? 1 : 2 );
    vec3 nx_raw = min_idx == 0 ? vec3(0,-ec_o.z,ec_o.y)
        : min_idx == 1 ? vec3(-ec_o.z,0,ec_o.x) : vec3(-ec_o.y,ec_o.x,0);
    vec3 nx = normalize(nx_raw);
    vec3 ny = cross(ne,nx);
```

```

// Compute center of the limb, lmb_o, the circle formed by the most
// distant visible parts of the sphere surface.
float sin_theta = r / ec_len;
float a = r * sin_theta;
vec3 lmb_o = ctr_o - a * ne;

// Compute axes to draw the limb.
float b = r * sqrt( 1 - sin_theta * sin_theta );
vec3 ax = b * nx, ay = b * ny;
mat4.mvp = trans_proj * gl_ModelViewMatrix;

// Emit a bounding square for the limb.
for ( int i = -1; i < 2; i += 2 ) for ( int j = -1; j < 2; j += 2 ) {
    vertex_o = lmb_o + ax * i + ay * j;
    gl_Position =.mvp * vec4(vertex_o,1);
    EmitVertex();
}
EndPrimitive();

```

```

}

```

The solution appears below. The geometry shader emits h primitives using the convenience functions to find the primitive limits. For each lens the shader also sends the object-space coordinates of the lens center and whether the lens faces the viewer. Note that it is now the geometry shader that computes the lens location, not the fragment shader. That's possible because one primitive is emitted for each lens.

/// SOLUTION

```

out Data_to_FS
{
    flat int vertex_id;
    vec3 vertex_o;

    /// SOLUTION – Add variables for lens center and face.
    //
    flat vec3 lens_ctr_o;
    flat bool lens_front;

};

// Type of primitive at geometry shader input.
//
layout ( points ) in;

// Type of primitives emitted geometry shader output.
//
/// SOLUTION – Increase number of vertices.
layout ( triangle_strip, max_vertices = h * 4 ) out;

void
gs_main()
{
    vertex_id = In[0].vertex_id;
    vec4 pos_rad = sphere_pos_rad[vertex_id];
    vec3 ctr_o = pos_rad.xyz;
    float r = pos_rad.w;

    if ( mirrored ) ctr_o.y = -ctr_o.y;

    // Eye location in object space.
    vec3 e_o = vec3(gl_ModelViewMatrixInverse * vec4(0,0,0,1));

    // Vectors from eye to sphere center.
    vec3 ec_o = ctr_o - e_o; // Eye to Center (of sphere).
    float ec_len = length(ec_o);
    vec3 ne = ec_o / ec_len;

    // Vectors orthogonal to ec_o.
    //
    vec3 atr_o = abs(ec_o);
    int min_idx = atr_o.x < atr_o.y ? ( atr_o.x < atr_o.z ? 0 : 2 )

```

```

    : ( atr_o.y < atr_o.z ? 1 : 2 );
vec3 nx_raw = min_idx == 0 ? vec3(0,-ec_o.z,ec_o.y)
    : min_idx == 1 ? vec3(-ec_o.z,0,ec_o.x) : vec3(-ec_o.y,ec_o.x,0);
vec3 nx = normalize(nx_raw);
vec3 ny = cross(ne,nx);

// Compute center of the limb, lmb_o, the circle formed by the most
// distant visible parts of the sphere surface.
float sin_theta = r / ec_len;
float a = r * sin_theta;
vec3 lmb_o = ctr_o - a * ne;

// Compute axes to draw the limb.
float b = r * sqrt( 1 - sin_theta * sin_theta );
vec3 ax = b * nx;
vec3 ay = b * ny;

mat4 mvp = trans_proj * gl_ModelViewMatrix;

/// SOLUTION – Remove original code emitting primitive.
#if 0
// Emit a bounding square for the limb.
for ( int i = -1; i < 2; i += 2 )
    for ( int j = -1; j < 2; j += 2 )
        {
            vertex_o = lmb_o + ax * i + ay * j;
            gl_Position = mvp * vec4(vertex_o,1);
            EmitVertex();
        }
    EndPrimitive();
#endif

/// SOLUTION – Use convenience functions to find bounding square ..
/// .. of each lens.
//
// Provide lens center and face to fragment shader.

for ( int l=0; l<h; l++ )
    {
        lens_ctr_o = lens_get_ctr_o(l,sphere_info);
        vec3 norm_o = lens_ctr_o - ctr_o;
        lens_front = dot(ec_o,norm_o) < 0;
        // Emit a bounding square for the limb.
        int indices[] = {0, 1, 3, 2};
        for ( int s=0; s<4; s++ )
            {
                vertex_o = lens_get_bb_o(l,indices[s],4,sphere_info);
                gl_Position = mvp * vec4(vertex_o,1);
                EmitVertex();
            }
        EndPrimitive();
    }

```

```

    }
}

#endif

}

```

(b) Appearing below is the fragment shader from the Homework 4 solution. The comments have been removed for brevity, feel free to look at the full Homework 4 solution. Show the changes needed to the fragment shader. The shader should be substantially simpler.

- Changes needed, based on geometry shader from previous part.
- The fragment shader should render lenses only, not full spheres.
- Simplify based on the fact that geometry shader is providing lens information.

The solution appears after the original code.

```

void fs_main()
{
    vec4 pos_rad = sphere_pos_rad[vertex_id];
    vec3 ctr_o = pos_rad.xyz;
    float rsq = pos_rad.w * pos_rad.w;
    vec3 e_o = gl_ModelViewMatrixInverse[3].xyz;
    vec3 ef = vertex_o - e_o;    // Eye to fragment.
    vec3 ce = e_o - ctr_o;
    float qfa = dot(ef,ef), qfb = 2 * dot(ef,ce), qfc = dot(ce,ce) - rsq;
    float discr = qfb*qfb - 4 * qfa * qfc;

    if ( discr < 0 ) discard;

    const float pi = 3.14159265359;
    const float two_pi = 2 * pi;
    const bool holes = opt_holes == OHO_Holes;    // If true sphere has holes
    const bool lenses = opt_holes == OHO_Lenses;
    const bool solid = !holes && !lenses;        // If true show complete sphere.

    vec4 sur_o; // Global surface cordinates (front or back) of sphere.
    vec3 normal_o, sur_l;
    bool front; // If true, cordinates are of front of sphere.
    float theta, eta; // Local angular cordinates of sphere.
    bool found_hole = false;

    for ( float dir = -1; dir <= 1; dir+=2 ) {
        front = dir == -1;
        float t = ( -qfb + dir * sqrt( discr ) ) / ( 2 * qfa );

        sur_o = vec4(e_o + t * ef, 1);
        normal_o = normalize(sur_o.xyz - ctr_o);
    }
}

```

```

if ( !front ) normal_o = -normal_o;

sur_l = normalize(mat3(sphere_rot[vertex_id]) * normal_o);
theta = atan(sur_l.x,sur_l.z);
eta = acos(sur_l.y);

if ( solid ) break;

const float hole_frac = 0.2;
const float radians_per_hole_eq = two_pi / opt_n_holes_eqt;
const float hole_radius = 0.5 * hole_frac * radians_per_hole_eq;

float eta_hole = round( eta / radians_per_hole_eq ) * radians_per_hole_eq;
float r = sin(eta_hole);
const float n_holes = floor( two_pi * r / radians_per_hole_eq );
if ( n_holes < -1 )
{
    if ( holes ) break;
    if ( dir == 1 ) discard;
}

const float radians_per_hole = two_pi / n_holes;
float theta_hole = round( theta / radians_per_hole ) * radians_per_hole;
vec3 hole_dir_l =
    vec3( r * sin(theta_hole), cos(eta_hole), r * cos(theta_hole) );

float dist = distance(hole_dir_l,sur_l);
found_hole = dist < hole_radius;
if ( lenses && found_hole holes && !found_hole ) break;
if ( !front ) discard;
}

vec3 normal_ee = normalize(gl_NormalMatrix * normal_o);
vec4 sur_ee = gl_ModelViewMatrix * sur_o;

// Compute clip-space depth. Take care so that compiler avoids full
// matrix / vector multiplication.
vec4 sur_e = gl_ModelViewMatrix * sur_o;
vec4 sur_c = trans_proj * vec4(0,0,2*sur_e.z,1);
gl_FragDepth = sur_c.z / sur_c.w;

vec2 tcoord = vec2( ( 1.5 * pi + theta ) / two_pi, eta / pi );
vec4 texel = front ? texture(tex_unit_0,tcoord) : vec4(1);
vec4 color2 = front ? sphere_color[vertex_id] : vec4(0.3,0,0,1);
gl_FragColor = texel * generic_lighting(sur_ee, color2, normal_ee);
}

```


The solution appears below. The code for computing the hole center coordinates has been removed. Instead the fragment shader uses `lens_ctr_o` provided by the geometry shader. Also, the `dir` loop has been removed. It is not needed because the geometry shader has determined whether the front or back of a lens is visible.

```

void
fs_main()
{
    /// SOLUTION - Pre-final Problem 3b

    vec4 pos_rad = sphere_pos_rad[vertex_id];

    // Center of sphere in original object-space coordinates.
    vec3 ctr_o = pos_rad.xyz;
    float rsq = pos_rad.w * pos_rad.w;

    // Eye location in object-space coordinates.
    vec3 e_o = gl_ModelViewMatrixInverse[3].xyz;

    // Prepare to compute intersection of ray from eye to through fragment with
    // sphere. That intersection is the point on the sphere corresponding
    // to this fragment.
    //
    vec3 ef = vertex_o - e_o;    // Eye to fragment.
    vec3 ce = e_o - ctr_o;
    float qfa = dot(ef, ef);
    float qfb = 2 * dot(ef, ce);
    float qfc = dot(ce, ce) - rsq;
    float discr = qfb*qfb - 4 * qfa * qfc;

    // If outside the limb, return.
    if ( discr < 0 ) discard;

    const float pi = 3.14159265359;
    const float two_pi = 2 * pi;

    vec4 sur_o; // Global surface coordinates (front or back) of sphere.
    vec3 normal_o, sur_l;
    float theta, eta; // Local angular coordinates of sphere.

    /// SOLUTION – No changes made to code above this point ..
    // .. except for deleting unneeded vars.

    /// SOLUTION - Geometry shader determines which side faces eye.
    //
    bool front = lens_front;

    /// SOLUTION: Remove loop. Value of dir based on lens_front from geo shader.
    {
        float dir = front ? -1 : 1;
    }
}

```

```

// Finish computing coordinate of point for this fragment.
//
float t = ( -qfb + dir * sqrt( discr ) ) / ( 2 * qfa );

// Compute true sphere surface coordinate.
//
sur_o = vec4(e_o + t * ef, 1);

// Compute possibly reflected sphere surface coordinate.
//
normal_o = normalize(sur_o.xyz - ctr_o);

// Reverse normal if this is the back surface of the sphere.
//
if ( !front ) normal_o = -normal_o;

// Use sphere-local coordinates to compute texture coordinates.
//
sur_l = normalize(mat3(sphere_rot[vertex_id]) * normal_o);
theta = atan(sur_l.x,sur_l.z);
eta = acos(sur_l.y);

const float hole_frac = 0.4;
const float radians_per_hole_eq = pos_rad.w * two_pi / h;
const float hole_radius = 0.5 * hole_frac * radians_per_hole_eq;

/// SOLUTION – Removed code computing hole location.
// Instead use lens_ctr_o provided by geo shader.
//
// Find distance from center of lens to sphere surface point ..
//
float dist = distance(lens_ctr_o,sur_o.xyz);
//
// .. and check whether surface point is in lens.
//
bool found_lens = dist < hole_radius;
//
// If this part of the sphere is not visible, we're done.
//
if ( !found_lens ) discard;
}

/// SOLUTION – No changes made to code below this point.

vec3 normal_ee = normalize(gl_NormalMatrix * normal_o);
vec4 sur_ee = gl_ModelViewMatrix * sur_o;

// Compute clip-space depth. Take care so that compiler avoids full
// matrix / vector multiplication.
vec4 sur_e = gl_ModelViewMatrix * sur_o;
vec4 sur_c = trans_proj * vec4(0,0,2*sur_e.z,1);

```

```

gl_FragDepth = sur_c.z / sur_c.w;

vec2 tcoord = vec2( ( 1.5 * pi + theta ) / two_pi, eta / pi );

// Get filtered texel.
//
vec4 texel = front ? texture(tex_unit_0,tcoord) : vec4(1);

vec4 color = sphere_color[vertex_id];
vec4 color2 = front ? color : vec4(0.3,0,0,1);

// Multiply filtered texel color with lighted color of fragment.
//
gl_FragColor = texel * generic_lighting(sur_ee, color2, normal_ee);
}

```