# CS 184: Problems on Transformations and OpenGL 

Ravi Ramamoorthi

These are some worked out problems that I will go over in the review sessions. They are representative of what you should understand, and may appear on the midterm. Brief solutions are provided in this note. Try to make sure you do understand those. Please also go over the exercises in chapters 6 and 7 of the text.

1. Write the homogeneous $4 \times 4$ matrices for the following transforms:

- Translate by +5 units in the X direction
- Rotate by 30 degrees about the X axis
- The rotation, followed by the translation above, followed by scaling by a factor of 2 .

2. In 3D, consider applying a rotation $R$ followed by a translation $T$. Write the form of the combined transformation in homogeneous coordinates (i.e. supply a $4 \times 4$ matrix) in terms of the elements of $R$ and T. Now, construct the inverse transformation, giving the corresponding $4 \times 4$ matrix in terms of R and T. You should simplify your answer (perhaps writing T as $[\mathrm{Tx}, \mathrm{Ty}, \mathrm{Tz}]$ and using appropriate notation for the 9 elements of the rotation matrix, or using appropriate matrix and vector notation for R and T ). Verify by matrix multiplication that the inverse times the original transform does in fact give the identity.
3. Adapted from the textbook, pages 118 and 119. Consider flatland (without homogeneous coordinates) $2 \times 2$ transformation matrices. Let's say we want to scale by 1.5 (increase length $50 \%$ ) not about the coordinate axes, but about an axis at -45 degrees to the horizontal. What is the resulting transformation matrix?
4. Adapted from the textbook, section 6.1.6. How can any 2D or 3D transformation (without homogeneous coordinates) be written (decomposed) as a combination of rotations and scales?
5. Write the $4 \times 4$ transformation matrix for rotation about an arbitrary point (rather than the origin)?
6. Assume that in OpenGL, your near and far clipping planes are set at a distance of 1 m and 100 m respectively. Further, assume your z-buffer has 9 bits of depth resolution. This means that after the gluPerspective transformation, the remapped z values [ranging from -1 to +1 ] are quantized into 512 discrete depths.

- How far apart are these discrete depth levels close to the near clipping plane? More concretely, what is the z range (i.e. 1 m to ?) of the first discrete depth?
- Now, consider the case where all the interesting geometry lies further than 10 m . How far apart are the discrete depth levels at 10 m ? Compare your answer to the first part and explain the cause for this difference.
- How many discrete depth levels describe the region between 10 m and 100 m ? What is the number of bits required for this number of depth levels? How many bits of precision have been lost? What would you recommend doing to increase precision?

7. Derive the homogeneous $4 \times 4$ matrices for gluLookAt and gluPerspective.
8. Consider the following operations in the standard OpenGL pipeline: Scan conversion or Rasterization, Texture Mapping, Projection Matrix, Transformation of Points and Normals by the ModelView Matrix, Dehomogenization (perspective division), clipping. Briefly explan what each of these operations are, and in what order they are performed and why. Which of these operations are conventionally performed in the vertex shader, fragment shader, or the OpenGL fixed function pipeline?

## Answers

1. Homogeneous Matrices A general representation for 4 x 4 matrices involving rotation and translation is

$$
\left(\begin{array}{cc}
R_{3 \times 3} & T_{3 \times 1}  \tag{1}\\
0_{1 \times 3} & 1_{1 \times 1},
\end{array}\right)
$$

where $R$ is $3 \times 3$ rotation matrix, and $T$ is a $3 \times 1$ translation matrix.
For a translation along the $X$ axis by 3 units $T=(5,0,0)^{t}$, while $R$ is the identity. Hence, we have

$$
\left(\begin{array}{llll}
1 & 0 & 0 & 5  \tag{2}\\
0 & 1 & 0 & 0 \\
0 & 0 & 1 & 0 \\
0 & 0 & 0 & 1
\end{array}\right)
$$

In the second case, where we are rotating about the $X$ axis, the translation matrix is just 0 . We need to remember the formula for rotation about an axis, which is (with angle $\theta$ ),

$$
\left(\begin{array}{cccc}
1 & 0 & 0 & 0  \tag{3}\\
0 & \cos \theta & -\sin \theta & 0 \\
0 & \sin \theta & \cos \theta & 0 \\
0 & 0 & 0 & 1
\end{array}\right)=\left(\begin{array}{cccc}
1 & 0 & 0 & 0 \\
0 & \sqrt{3} / 2 & -1 / 2 & 0 \\
0 & 1 / 2 & \sqrt{3} / 2 & 0 \\
0 & 0 & 0 & 1
\end{array}\right) .
$$

Finally, when we are combining these transformations, $S * T * R$, we apply the rotation first, followed by a translation. It is easy to verify by matrix multiplication, that this simply has the same form as equation 1 (but see the next problem for when we have $\mathrm{R} * \mathrm{~T}$ ). The scale just multiplies everything by a factor of 2 , giving

$$
\left(\begin{array}{cccc}
2 & 0 & 0 & 10  \tag{4}\\
0 & \sqrt{3} & -1 & 0 \\
0 & 1 & \sqrt{3} & 0 \\
0 & 0 & 0 & 1
\end{array}\right)
$$

It is also possible to obtain this result by matrix multiplication of $\mathrm{S} * \mathrm{~T} * \mathrm{R}$

$$
\left(\begin{array}{llll}
2 & 0 & 0 & 0  \tag{5}\\
0 & 2 & 0 & 0 \\
0 & 0 & 2 & 0 \\
0 & 0 & 0 & 1
\end{array}\right)\left(\begin{array}{llll}
1 & 0 & 0 & 5 \\
0 & 1 & 0 & 0 \\
0 & 0 & 1 & 0 \\
0 & 0 & 0 & 1
\end{array}\right)\left(\begin{array}{cccc}
1 & 0 & 0 & 0 \\
0 & \sqrt{3} / 2 & -1 / 2 & 0 \\
0 & 1 / 2 & \sqrt{3} / 2 & 0 \\
0 & 0 & 0 & 1
\end{array}\right)=\left(\begin{array}{cccc}
2 & 0 & 0 & 10 \\
0 & \sqrt{3} & -1 & 0 \\
0 & 1 & \sqrt{3} & 0 \\
0 & 0 & 0 & 1
\end{array}\right) .
$$

2. Rotations and translations Having a rotation followed by a translation is simply $\mathrm{T}^{*} \mathrm{R}$, which has the same form as equation 1 . The inverse transform is more interesting. Essentially $(T R)^{-1}=R^{-1} T^{-1}=$ $R^{t} *-T$, which in homogeneous coordinates is

$$
\left(\begin{array}{cc}
R_{3 \times 3}^{t} & 0_{3 \times 1}  \tag{6}\\
0_{1 \times 3} & 1
\end{array}\right)\left(\begin{array}{cc}
I_{3 \times 3} & -T_{3 \times 1} \\
0_{1 \times 3} & 1
\end{array}\right)=\left(\begin{array}{cc}
R_{3 \times 3}^{t} & -R_{3 \times 3}^{t} T_{3 \times 1} \\
0_{1 \times 3} & 1
\end{array}\right) .
$$

Note that this is the same form as equation 1 , using $R^{\prime}$ and $T^{\prime}$ with $R^{\prime}=R^{t}=R^{-1}$ and $T^{\prime}=-R^{t} T$.
Finally, we may verify that the product of the inverse and the original does in fact give the identity.

$$
\left(\begin{array}{cc}
R_{3 \times 3}^{t} & -R_{3 \times 3}^{t} T_{3 \times 1}  \tag{7}\\
0_{1 \times 3} & 1
\end{array}\right)\left(\begin{array}{cc}
R_{3 \times 3} & T_{3 \times 1} \\
0_{1 \times 3} & 1
\end{array}\right)=\left(\begin{array}{cc}
R_{3 \times 3}^{t} R_{3 \times 3} & R_{3 \times 3}^{t} T_{3 \times 1}-R_{3 \times 3}^{t} T_{3 \times 1} \\
0_{1 \times 3} & 1
\end{array}\right)=\left(\begin{array}{cc}
I_{3 \times 3} & 0_{3 \times 1} \\
0_{1 \times 3} & 1
\end{array}\right)
$$

3. Scaling about an axis For all these cases of non-standard transforms, we first apply an operation, here a rotation to align the coordinate axes, then apply the scale, and then undo the rotation. In particular, we would first rotate by +45 degrees to align the axis with the horizontal, then scale by $(1.5,1)$, and then rotate by -45 degrees to undo the initial rotation. The net transform $M$ is

$$
\begin{equation*}
M=R(-45) S(1.5,1) R(+45)=R S R^{t} \tag{8}
\end{equation*}
$$

where $R$ and $S$ are the rotation matrix for -45 degrees, and scale matrices respectively, and we use that $R^{-1}=R^{t}$. Plugging in numerical values, the net transformation is

$$
M=\left(\begin{array}{rr}
\frac{1}{\sqrt{2}} & \frac{1}{\sqrt{2}}  \tag{9}\\
-\frac{1}{\sqrt{2}} & \frac{1}{\sqrt{2}}
\end{array}\right)\left(\begin{array}{rr}
1.5 & 0 \\
0 & 1
\end{array}\right)\left(\begin{array}{rr}
\frac{1}{\sqrt{2}} & -\frac{1}{\sqrt{2}} \\
\frac{1}{\sqrt{2}} & \frac{1}{\sqrt{2}}
\end{array}\right)=\left(\begin{array}{rr}
\frac{5}{4} & -\frac{1}{4} \\
-\frac{1}{4} & \frac{5}{4}
\end{array}\right)=\left(\begin{array}{rr}
1.25 & -.25 \\
-.25 & 1.25
\end{array}\right)
$$

4. Decomposing Transformations As illustrated by the previous problem, any symmetric matrix can be written (via eigenvalue decomposition) as $R S R^{t}$ where $R$ is an orthogonal (and hence rotation) matrix and $S$ is diagonal (and hence a scale matrix). Let's call the columns of $R$ (the eigenvectors) unit vectors $\mathbf{v}_{1}$ and $\mathbf{v}_{2}$, and the diagonal elements (eigenvalues) of $S$ as $\lambda_{1}$ and $\lambda_{2}$.

Then, any symmetric $2 \times 2$ or $3 \times 3$ transformation matrix can be considered as non-uniform scaling by the eigenvalues $\left(\lambda_{1}, \lambda_{2}\right)$ about the new axes $\left(\mathbf{v}_{1}, \mathbf{v}_{2}\right)$. In other words, we first rotate $\left(\mathbf{v}_{1}, \mathbf{v}_{2}\right)$ to $(x, y)$ using $R^{t}$. Then, we apply the standard scaling along coordinate axes given by $S$, and finally we undo the rotation using $R$, just as in the previous exercise, i.e., rotate $(x, y)$ to $\left(\mathbf{v}_{1}, \mathbf{v}_{2}\right)$.

For general, non-symmetric transformation matrices, we can use a singular-value decomposition $U S V^{t}$. The diagonal entries of $S$ are now called the singular values and denoted by $\left(\sigma_{1}, \sigma_{2}\right)$, and the columns of $U$ and $V$ are now respectively the left and right singular vectors. We still have the same sequence of steps. First, rotate $\mathbf{v}_{1}$ and $\mathbf{v}_{2}$ to the $x$ and $y$ axes using $V^{t}$, then scale in $x$ and $y$ by $\left(\sigma_{1}, \sigma_{2}\right)$ using $S$, and finally rotate the $x$ and $y$ axes to $\left(\mathbf{u}_{1}, \mathbf{u}_{2}\right)$ (the transform by $U$ ). While we have illustrated this in 2D, the same results hold in 3D (or higher dimensions for that matter).
5. Rotation about arbitrary point The same basic idea applies as for problem 3 . We move the point to the origin (a translation), do a standard rotation, and undo the translation. Let us call the center of rotation c and consider a point $\mathbf{p}$. We really want to compute

$$
\begin{equation*}
\mathbf{p}^{\prime}=\mathbf{c}+R(\mathbf{p}-\mathbf{c})=R \mathbf{p}+(\mathbf{c}-R \mathbf{c}), \tag{10}
\end{equation*}
$$

with the net transformation matrix written as

$$
\begin{equation*}
M=T(\mathbf{c}) R T(-\mathbf{c}) . \tag{11}
\end{equation*}
$$

Multiplying this out for rotation and translation matrices (see problems 1 and 2),

$$
M=\left(\begin{array}{cc}
I_{3 \times 3} & \mathbf{c}_{3 \times 1}  \tag{12}\\
0_{1 \times 3} & 1
\end{array}\right)\left(\begin{array}{cc}
R_{3 \times 3} & 0_{3 \times 1} \\
0_{1 \times 3} & 1
\end{array}\right)\left(\begin{array}{cc}
I_{3 \times 3} & -\mathbf{c}_{3 \times 1} \\
0_{1 \times 3} & 1
\end{array}\right)=\left(\begin{array}{cc}
R_{3 \times 3} & \mathbf{c}_{3 \times 1}-R \mathbf{c}_{3 \times 1} \\
0_{1 \times 3} & 1
\end{array}\right),
$$

which agrees with equation 10 .
7. gluLookAt and gluPerspective I wrote this answer earlier to conform in notation to the Unix Man Pages. Some of you might find it easier to just understand this from the lecture slides in the transformation lectures than this derivation and may want to skip over this section if you already understand the concepts.
gluLookat defines the viewing transformation and is given by gluLookAt(eyex, eyey, eyez, centerx, centery, centerz, upx, upy, upz), corresponds to a camera at eye looking at center with up direction $u p$. First, we define the normalized viewing direction. The symbols used here are chosen to correspond to the definitions in the man page.

$$
F=\left(\begin{array}{c}
C_{x}-E_{x}  \tag{13}\\
C_{y}-E_{y} \\
C_{z}-E_{z}
\end{array}\right) \quad \mathbf{f}=F /\|F\| .
$$

This direction $\mathbf{f}$ will correspond to the $-Z$ direction, since the eye is mapped to the origin, and the lookat point or center to the negative z axis. What remains now is to define the $X$ and $Y$ directions. The $Y$ direction corresponds to the up vector. First, we define $U P^{\prime}=U P /\|U P\|$ to normalize. However, this may not be perpendicular to the $Z$ axis, so we use vector cross products to define $X=-Z \times Y$ and $Y=X \times-Z$. In our notation, this defines auxiliary vectors,

$$
\begin{equation*}
\mathbf{s}=\frac{\mathbf{f} \times U P^{\prime}}{\left\|\mathbf{f} \times U P^{\prime}\right\|} \quad \mathbf{u}=\frac{\mathbf{s} \times \mathbf{f}}{\|\mathbf{s} \times \mathbf{f}\|} \tag{14}
\end{equation*}
$$

Note that this requires the UP vector not to be parallel to the view direction. We now have a set of directions $\mathbf{s}, \mathbf{u},-\mathbf{f}$ corresponding to $X, Y, Z$ axes. We can therefore define a rotation matrix,

$$
M=\left(\begin{array}{cccc}
\mathbf{s}_{x} & \mathbf{s}_{y} & \mathbf{s}_{z} & 0  \tag{15}\\
\mathbf{u}_{x} & \mathbf{u}_{y} & \mathbf{u}_{z} & 0 \\
-\mathbf{f}_{x} & -\mathbf{f}_{y} & -\mathbf{f}_{z} & 0 \\
0 & 0 & 0 & 1
\end{array}\right),
$$

that rotates a point to the new coordinate frame.
However, gluLookAt requires applying this rotation matrix about the eye position, not the origin. It is equivalent to glMultMatrixf( $M$ ) ; glTranslateD(-eyex, -eyey, -eyez) ; This corresponds to a translation $T$ followed by a rotation $R$. We know (using equation 6 as a guideline for instance), that this is the same as the rotation $R$ followed by a modified translation $R_{3 \times 3} T_{3 \times 1}$. Written out in full, the matrix will then be

$$
G=\left(\begin{array}{cccc}
\mathbf{s}_{x} & \mathbf{s}_{y} & \mathbf{s}_{z} & -\mathbf{s}_{x} \mathbf{e}_{x}-\mathbf{s}_{y} \mathbf{e}_{y}-\mathbf{s}_{z} \mathbf{e}_{z}  \tag{16}\\
\mathbf{u}_{x} & \mathbf{u}_{y} & \mathbf{u}_{z} & -\mathbf{u}_{x} \mathbf{e}_{x}-\mathbf{u}_{y} \mathbf{e}_{y}-\mathbf{u}_{z} \mathbf{e}_{z} \\
-\mathbf{f}_{x} & -\mathbf{f}_{y} & -\mathbf{f}_{z} & \mathbf{f}_{x} \mathbf{e}_{x}+\mathbf{f}_{y} \mathbf{e}_{y}+\mathbf{f}_{z} \mathbf{e}_{z} \\
0 & 0 & 0 & 1
\end{array}\right) .
$$

gluPerspective defines a perspective transformation used to map 3D objects to the 2D screen and is defined by gluPerspective(fovy, aspect, zNear, zFar) where fovy specifies the field of view angle, in degrees,
in the $y$ direction, and aspect specifies the aspect ratio that determines the field of view in the x direction. The aspect ratio is the ratio of x (width) to y (height). $z$ Near and $z F a r$ represent the distance from the viewer to the near and far clipping planes, and must always be positive.

First, we define $f=\cot (f o v y / 2)$ as corresponding to the focal length or focal distance. A 1 unit height in $Y$ at $Z=f$ should correspond to $y=1$. This means we must multiply $Y$ by $f$ and corresponding $X$ by $f /$ aspect. The matrix has the form

$$
M=\left(\begin{array}{cccc}
\frac{f}{\text { aspect }} & 0 & 0 & 0  \tag{17}\\
0 & f & 0 & 0 \\
0 & 0 & A & B \\
0 & 0 & -1 & 0
\end{array}\right)
$$

We will explain the form of the above matrix. The form of terms $f /$ aspect and $f$ has already been explained. The term -1 in the last line is needed to divide by the distance $Z$ as required in perspective, and the negative sign is because OpenGL conventions require us to look down the $-Z$ axis.

It remains to find $A$ and $B$. Those are chosen so the near and far clipping planes are taken to -1 and +1 respectively. Indeed, the entire viewing volume or frustum is mapped to a cube between -1 and +1 along all axes. Using the matrix, we can easily formulate that the remapped depth is given by

$$
\begin{equation*}
z^{\prime}=\frac{A z+B}{-z}=-A-\frac{B}{z}, \tag{18}
\end{equation*}
$$

where one must remember that points in front of the viewer have negative $z$ as per OpenGL conventions. Now, the required conditions $z=-z$ Near $\Rightarrow z^{\prime}=-1$ and $z=-z F a r \Rightarrow z^{\prime}=1$ have,

$$
\begin{equation*}
-A+\frac{B}{z N e a r}=-1 \quad-A+\frac{B}{z F a r}=+1 \tag{19}
\end{equation*}
$$

Solving this system gives

$$
\begin{equation*}
A=\frac{z F a r+z \text { Near }}{z \text { Near }-z \text { Far }} \quad B=\frac{2 \cdot z F a r \cdot z \text { Near }}{z \text { Near }-z \text { Far }}, \tag{20}
\end{equation*}
$$

and the final matrix

$$
G=\left(\begin{array}{cccc}
\frac{f}{\text { aspect }} & 0 & 0 & 0  \tag{21}\\
0 & f & 0 & 0 \\
0 & 0 & \frac{z F a r+\text { Near }}{z N e a r-z F a r} & \frac{2 \cdot z \text { Far.zNear }}{z \text { Near-zFar }} \\
0 & 0 & -1 & 0
\end{array}\right) .
$$

6. Z-buffer in OpenGL The purpose of this question is to show how depth resolution degrades as one moves further away from the near clipping plane, since the remapped depth is nonlinear (and reciprocal) in the original depth values.

Equation 18 gives us the formula for remapping $z$. We just need to find $A$ and $B$, which we can do by solving, or plugging directly into equation 20 using $z N e a r=1$ and $z F a r=100$. We obtain $A=-101 / 99 \approx-1.02$ and $B=-200 / 99 \approx-2.02$. The remapped value is then given by

$$
\begin{equation*}
z^{\prime}=1.02-\frac{2.02}{|z|} \tag{22}
\end{equation*}
$$

Note that for mathematical simplicity, you might imagine the far plane at infinity, so we don't need the .02 .

For the remaining parts of the question, it is probably simplest to just use differential techniques. We can obtain

$$
\begin{equation*}
d z^{\prime} \approx \frac{2}{|z|^{2}} d|z| \Rightarrow|d z|=\frac{|z|^{2}}{2}\left|d z^{\prime}\right| \tag{23}
\end{equation*}
$$

To consider one depth bucket, we simply need to set $\left|d z^{\prime}\right|=1 / 256 \approx 0.004$. Now, using the equation above, setting $|z|=1$, we get $|d z| \approx 0.002$. In other words, the first depth bucket ranges from a depth of $1 m$ to a depth of approximately 1.0019 m , and we can resolve depths 2 mm apart.

Now, consider $|z|=10$, and plug in above. We know that $|d z| \sim|z|^{2}$, so $|d z| \approx 0.2$, and the depth buckets around $z=10 \mathrm{~m}$ are in 20 cm increments and we lose resolving power quadratically, with a danger that many different objects may go into the same depth bucket. This brings us to the fundamental point of this problem that depth levels are closer near the near plane and depth resolution decreases far away.

Finally, we consider the depth levels between 10 m and 100 m . Using equation $22,10 \mathrm{~m}$ transforms to $1.02-2.02 / 10 \approx 0.82$. Thus, only a range of 0.18 remains. Hence, we only have $0.18 * 256=46$ depth buckets, or less than 6 bits of precision. We have lost more than 3 bits of precision. To increase precision, we should move the near clipping plane further out if interesting geometry is in the $10 \mathrm{~m}-100 \mathrm{~m}$ range.

Order of OpenGL operations While the pipeline is programmable, a standard rendering program will still perform operations in this order.

1. Modelview Matrix: Each vertex's spatial coordinates are transformed by the modelview matrix $M$ as $\mathbf{x}^{\prime}=M \mathbf{x}$. Simultaneously, normals are transformed by the inverse transpose, $\mathbf{n}^{\prime}=M^{-t} \mathbf{n}$ and renormalized if specified. This operation is usually conducted in the vertex shader.
2. Projection Matrix: The projection matrix is then applied to project objects into the image plane, also in the vertex shader (sometimes, the combined operation is applied).
3. Clipping: Clipping is then done in the homogeneous coordinates against the standard viewing planes $x= \pm w, y= \pm w, z= \pm w$. Clipping is done before dehomogenization to avoid the perspective divide for vertices that are clipped anyway. This is usually a fixed function OpenGL functionality, though of course some over-riding in the shaders is permitted.
4. Dehomogenization or Perspective divide: Perspective divide by the fourth or homogeneous coordinate $w$ then occurs. Again, this happens in the fixed function OpenGL part between the shaders.
5. Scan conversion or Rasterization: Finally, the primitive is converted to fragments by scan conversion or rasterization. This is the core fixed piece still left in openGL.
6. Texture mapping: Texture mapping is application of image textures to geometry. It is a fragment-level operation (like depth testing) which happens in the fragment shader.
