

*Computer Graphics
and Visualisation*

Geometry for Computer Graphics

Student Notes

Developed by

F Lin

K Wyrwas

J Irwin

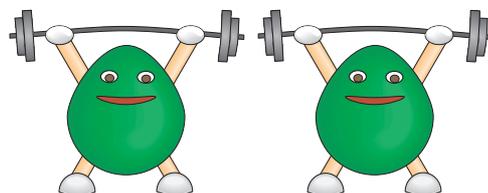
C Lilley

W T Hewitt

T L J Howard

*Computer Graphics Unit
Manchester Computing Centre
University of Manchester*

*Department of Computer Science
University of Manchester*



THE UNIVERSITY
of MANCHESTER

Produced by the ITTI Gravigs Project
Computer Graphics Unit
Manchester Computing Centre
MANCHESTER M13 9PL

First published by UCoSDA in May 1995

© Copyright The University of Manchester 1995

The moral right of Fenqiang Lin, Karen Wyrwas, John Irwin, Christopher Lilley, William Terence Hewitt and Toby Leslie John Howard to be identified as the authors of this work is asserted by them in accordance with the Copyright, Designs and Patents Act (1988).

ISBN 1 85889 059 4

The training materials, software and documentation in this module may be copied within the purchasing institution for the purpose of training their students and staff. For all other purposes, such as the running of courses for profit, please contact the copyright holders.

For further information on this and other modules please contact:
The ITTI Gravigs Project, Computer Graphics Unit, Manchester Computing Centre.
Tel: 0161 275 6095 Email: *gravigs@mcc.ac.uk*

To order further copies of this or other modules please contact:
Mrs Jean Burgan, UCoSDA.
Tel: 0114 272 5248 Email: *j.burgan@sheffield.ac.uk*

These materials have been produced as part of the Information Technology Training Initiative, funded by the Information Systems Committee of the Higher Education Funding Councils.

The authors would like to thank Janet Edwards for her assistance in the preparation of these documents.

Printed by the Reprographics Department, Manchester Computing Centre from PostScript source supplied by the authors.

Table of Contents

1	2D Transformations	1
1.1	Introduction	1
1.2	Types of Transformation	1
1.3	Matrix Representation of Transformations	5
1.4	Concatenation of Transformations	7
1.5	Ordering Transformations	8
1.6	Homogeneous Coordinates	9
1.7	Object and Axis Transformations	10
1.8	The Normalization Transformation in GKS	12
1.9	Summary	13
2	3D Transformations	15
2.1	Introduction	15
2.2	Homogeneous Coordinates	16
2.3	Types of 3D Transformation	16
2.4	Perspective Transformations	19
2.5	Projections	23
2.6	Parallel Projections	26
2.7	Classification of Planar Geometric Projections	26
3	Transformations and Viewing in GKS-3D and PHIGS	29
3.1	The GKS-3D Output Pipeline	29
3.2	The PHIGS Output Pipeline	33
3.3	The Viewing Pipeline	35
3.4	Using the Viewing Model	41
A	References	49

1 2D Transformations

1.1 Introduction

In computer graphics many applications need to alter or manipulate a picture, for example, by changing its size, position or orientation. This can be done by applying a geometric transformation to the coordinate points defining the picture. These notes cover the basic theory of two-dimensional (2D) geometric transformations.

1.2 Types of Transformation

1.2.1 Translation

A common requirement is to move a picture to a new position, as in Figure 1. (The original object is drawn using dotted lines, and the transformed object using solid lines. This convention will be used throughout.)

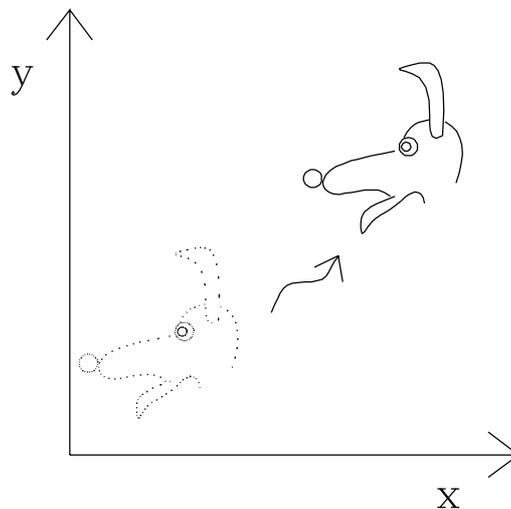


Figure 1: moving a picture

This is achieved by means of a translation or shift transformation. In order to translate a point, constant values are added to the x - and y -coordinates of the point, as shown in Figure 2.

The new coordinates of the point (x', y') are given by

$$\begin{aligned}x' &= x + t_x \\y' &= y + t_y\end{aligned}$$

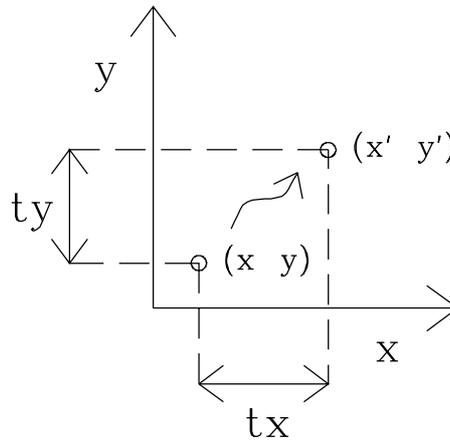


Figure 2: translating a point

1.2.2 Scaling

A scaling transformation is used to change the size of an object. Scaling about the origin (0, 0) is achieved by multiplying the coordinates of a point by x - and y -scale factors:

$$x' = x \cdot s_x$$

$$y' = y \cdot s_y$$

If $|s_x|$ and $|s_y|$ are both >1 , the effect is of increasing the size of an object. In order to reduce the size, $|s_x|$ and $|s_y|$ must be <1 .

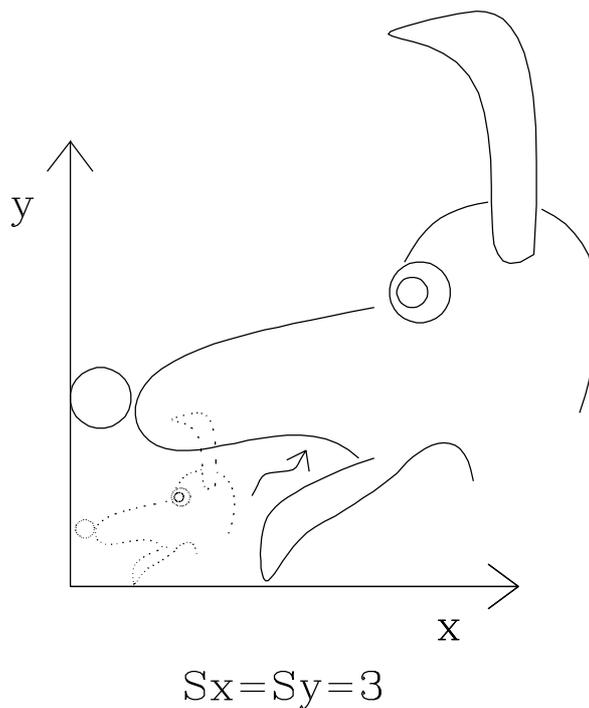


Figure 3: symmetric scaling

Figure 3 illustrates a symmetric or uniform scaling transformation in which the x - and y -scale factors are the same ($s_x = s_y$) so that the object is expanded by the same amount in each axis direction.

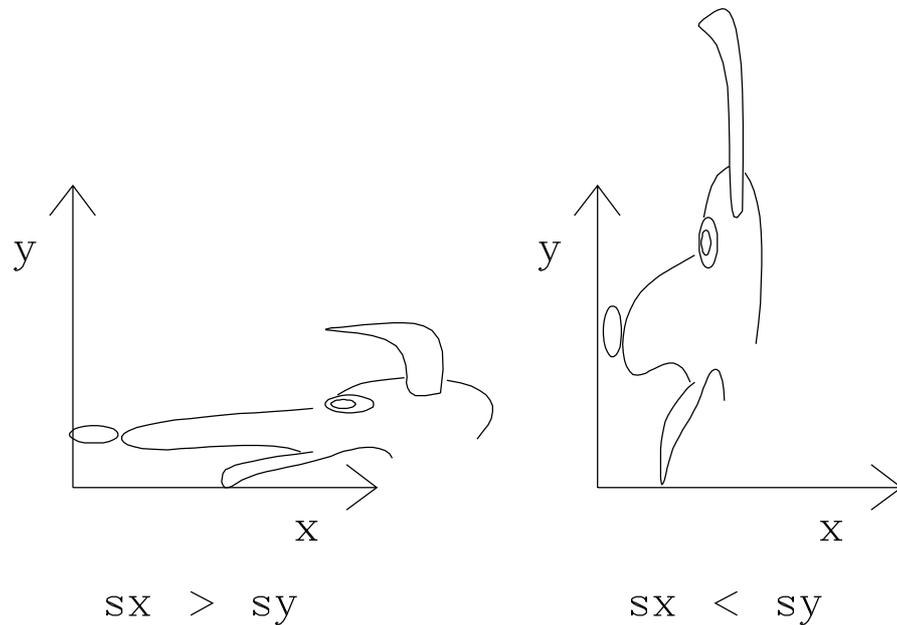


Figure 4: asymmetric scaling

Figure 4 illustrates two asymmetric or non-uniform scaling transformations in which the x - and y -scale factors are not equal ($s_x \neq s_y$). Here the object changes its size by different amounts in the x - and y -axis directions.

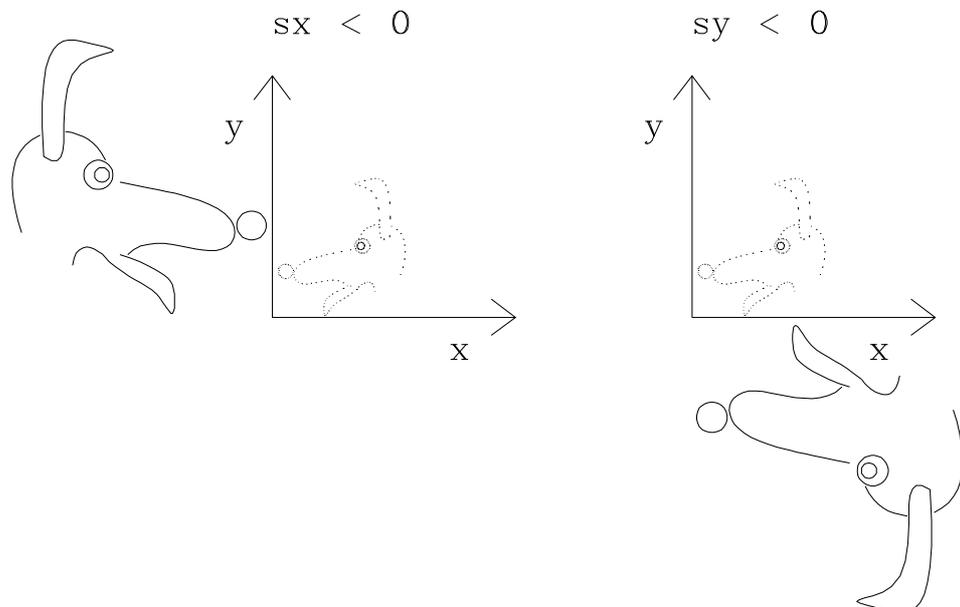


Figure 5: the effect of negative scale factors

If the scale factor in x is negative ($s_x < 0$) then the object is reflected in the y -axis. Similarly, if the scale factor in y is negative ($s_y < 0$) then the object is reflected in the x -axis. These two cases are shown in Figure 5 .

1.2.3 Rotation

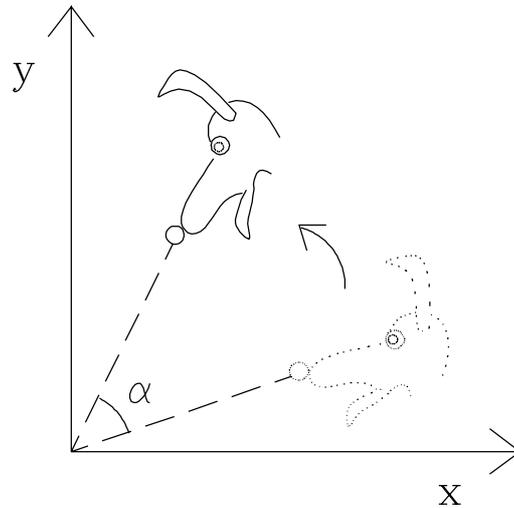


Figure 6: rotating an object about the origin

Another common type of transformation is rotation. This is used to orientate objects. Figure 6 shows an object rotated by an angle α about the origin.

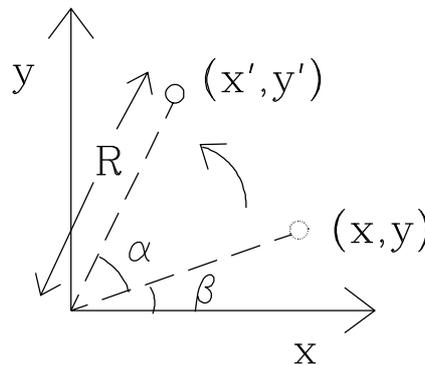


Figure 7: rotating a point about the origin

The rotation of one point in the object is illustrated in Figure 7. A line joining the point with the origin makes an angle β with the x -axis and has length R , hence

$$\begin{aligned}x &= R \cdot \cos \beta \\y &= R \cdot \sin \beta\end{aligned}$$

After rotation the point has coordinates x' and y' with values

$$\begin{aligned}x' &= R \cdot \cos (\alpha + \beta) \\y' &= R \cdot \sin (\alpha + \beta)\end{aligned}$$

Expanding these formulae for $\cos(\alpha+\beta)$ and $\sin(\alpha+\beta)$ and rearranging gives

$$\begin{aligned}x' &= R \cdot \cos \alpha \cdot \cos \beta - R \cdot \sin \alpha \cdot \sin \beta \\y' &= R \cdot \sin \alpha \cdot \cos \beta + R \cdot \sin \beta \cdot \cos \alpha\end{aligned}$$

Finally, substituting for $R \cdot \cos \beta$ and $R \cdot \sin \beta$ gives

$$\begin{aligned}x' &= x \cdot \cos\alpha - y \cdot \sin\alpha \\y' &= x \cdot \sin\alpha + y \cdot \cos\alpha\end{aligned}$$

1.2.4 Shearing

A shear transformation has the effect of distorting the shape of an object.

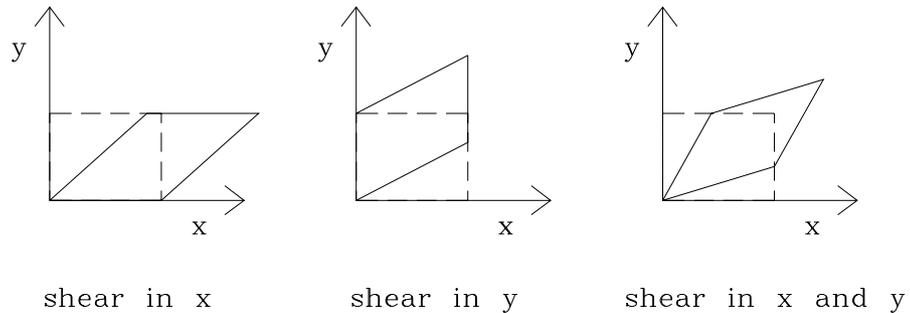


Figure 8: shear transformations

Figure 8 illustrates several different kinds of shear transformation applied to a rectangular object. The first diagram shows a shear in x in which the x -coordinates of points are displaced as a function of their height. The middle diagram shows a shear in y , where the y -coordinates are displaced according to their x -coordinate. Finally, a shear in both x and y is shown. The new x - and y -coordinates of a point after shearing are given by

$$\begin{aligned}x' &= x + y \cdot a \\y' &= y + x \cdot b\end{aligned}$$

If $a \neq 0$ then a shear in x is obtained. Similarly, if $b \neq 0$ then a shear in y is obtained.

1.3 Matrix Representation of Transformations

In the last section we looked at the basic types of transformation and for each derived an expression for the new coordinates of a point after transformation. We can now write down a general formula for the transformation of points

$$\begin{aligned}x' &= a \cdot x + b \cdot y + c \\y' &= d \cdot x + e \cdot y + f\end{aligned}$$

where a , b , c , d , e and f are all constants. The expressions for x' and y' are linear functions of x and y . This can be re-expressed using matrices as

$$\begin{bmatrix} x' \\ y' \end{bmatrix} = \begin{bmatrix} a & b \\ d & e \end{bmatrix} \cdot \begin{bmatrix} x \\ y \end{bmatrix} + \begin{bmatrix} c \\ f \end{bmatrix}$$

Now include all of the constants in one matrix

$$\begin{bmatrix} x' \\ y' \end{bmatrix} = \begin{bmatrix} a & b & c \\ d & e & f \end{bmatrix} \cdot \begin{bmatrix} x \\ y \\ 1 \end{bmatrix}$$

A square matrix is much easier to deal with so the matrix is extended to a 3×3 matrix

$$\begin{bmatrix} x' \\ y' \\ w' \end{bmatrix} = \begin{bmatrix} a & b & c \\ d & e & f \\ g & h & i \end{bmatrix} \cdot \begin{bmatrix} x \\ y \\ 1 \end{bmatrix}$$

The column vectors representing points now have an extra entry. If the bottom row of the matrix is [0 0 1] then w' will be 1 and we can ignore it. The effect of setting the bottom row of the matrix to values other than [0 0 1] is dealt with later (see section 6 - Homogeneous Coordinates).

The formulae for each of the different types of transformation can now be re-written using this matrix notation:

- Translate $\begin{bmatrix} 1 & 0 & t_x \\ 0 & 1 & t_y \\ 0 & 0 & 1 \end{bmatrix}$
- Scale $\begin{bmatrix} s_x & 0 & 0 \\ 0 & s_y & 0 \\ 0 & 0 & 1 \end{bmatrix}$
- Rotate $\begin{bmatrix} \cos\alpha & -\sin\alpha & 0 \\ \sin\alpha & \cos\alpha & 0 \\ 0 & 0 & 1 \end{bmatrix}$
- Shear $\begin{bmatrix} 1 & a & 0 \\ b & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}$

There is a special matrix which leaves the coordinates x' and y' equal to x and y . This is known as the unit or identity matrix:

$$\begin{bmatrix} x' \\ y' \\ w' \end{bmatrix} = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix} \cdot \begin{bmatrix} x \\ y \\ 1 \end{bmatrix}$$

$$x' = x$$

$$y' = y$$

$$w' = 1$$

1.4 Concatenation of Transformations

We have examined the basic types of transformation and derived the corresponding matrices. In this section we will see how these transformations can be combined to perform more complex operations such as rotation or scaling about an arbitrary point.

Consider the transformation to rotate an object about its centre point (x_c, y_c) . This can be broken down into a series of basic transformations as follows:

- Translate the object by $(-x_c, -y_c)$ so that the centre coincides with the origin.
- Perform a rotation about the origin.
- Translate the object by (x_c, y_c) to return it to its original position.

This series of transformations is illustrated in Figure 9.

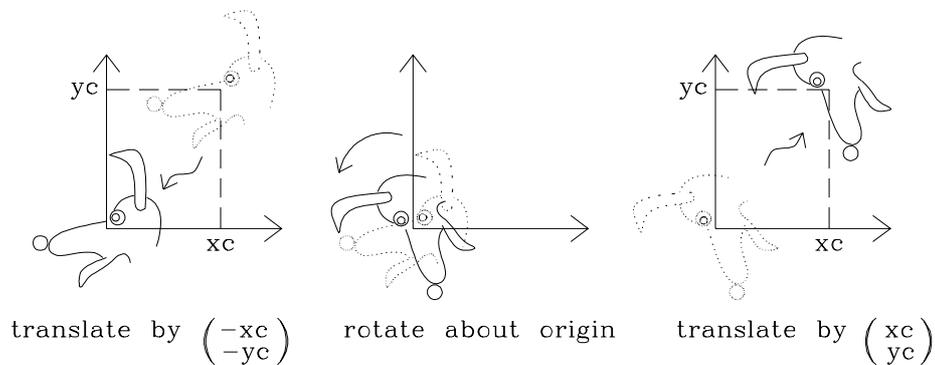


Figure 9: rotating an object about its centre

How can this composite transformation be expressed in terms of matrices? If we apply each of the component transformations separately:

- Translate by $(-x_c, -y_c)$:

$$\begin{bmatrix} x_1 \\ y_1 \\ 1 \end{bmatrix} = \begin{bmatrix} 1 & 0 & -x_c \\ 0 & 1 & -y_c \\ 0 & 0 & 1 \end{bmatrix} \cdot \begin{bmatrix} x \\ y \\ 1 \end{bmatrix}$$

- Rotate by an angle about the origin:

$$\begin{aligned} \begin{bmatrix} x_2 \\ y_2 \\ 1 \end{bmatrix} &= \begin{bmatrix} \cos \alpha & -\sin \alpha & 0 \\ \sin \alpha & \cos \alpha & 0 \\ 0 & 0 & 1 \end{bmatrix} \cdot \begin{bmatrix} x_1 \\ y_1 \\ 1 \end{bmatrix} \\ &= \begin{bmatrix} \cos \alpha & -\sin \alpha & 0 \\ \sin \alpha & \cos \alpha & 0 \\ 0 & 0 & 1 \end{bmatrix} \cdot \begin{bmatrix} 1 & 0 & -x_c \\ 0 & 1 & -y_c \\ 0 & 0 & 1 \end{bmatrix} \cdot \begin{bmatrix} x \\ y \\ 1 \end{bmatrix} \end{aligned}$$

- Translate by (x_c, y_c) :

$$\begin{aligned}
 \begin{bmatrix} x_3 \\ y_3 \\ 1 \end{bmatrix} &= \begin{bmatrix} 1 & 0 & x_c \\ 0 & 1 & y_c \\ 0 & 0 & 1 \end{bmatrix} \cdot \begin{bmatrix} x_2 \\ y_2 \\ 1 \end{bmatrix} \\
 &= \begin{bmatrix} 1 & 0 & x_c \\ 0 & 1 & y_c \\ 0 & 0 & 1 \end{bmatrix} \cdot \begin{bmatrix} \cos \alpha & -\sin \alpha & 0 \\ \sin \alpha & \cos \alpha & 0 \\ 0 & 0 & 1 \end{bmatrix} \cdot \begin{bmatrix} 1 & 0 & -x_c \\ 0 & 1 & -y_c \\ 0 & 0 & 1 \end{bmatrix} \cdot \begin{bmatrix} x \\ y \\ 1 \end{bmatrix} \\
 &= \begin{bmatrix} \cos \alpha & -\sin \alpha & (x_c - \cos \alpha \cdot x_c + \sin \alpha \cdot y_c) \\ \sin \alpha & \cos \alpha & (y_c - \sin \alpha \cdot x_c - \cos \alpha \cdot y_c) \\ 0 & 0 & 1 \end{bmatrix} \cdot \begin{bmatrix} x \\ y \\ 1 \end{bmatrix}
 \end{aligned}$$

The net effect of the transformation is to map the point (x, y) onto the point (x_3, y_3) . This mapping can be expressed as the matrix multiplication of the three basic transformation matrices used. The value of using square matrices to represent transformations can now be seen. Square matrices can be multiplied together to produce another square matrix of the same dimensions. Hence composite transformations can be expressed as a single transformation matrix by multiplying them together. Each point to be transformed is multiplied by this matrix which performs all the component transformations in one step. When there are many points to be transformed, this is considerably more efficient than multiplying the points by each component transformation matrix in turn.

1.5 Ordering Transformations

We have already seen how more than one transformation can be combined by multiplying together the corresponding transformation matrices. Matrix multiplication is not a commutative operation $\mathbf{M}_1 \cdot \mathbf{M}_2 \neq \mathbf{M}_2 \cdot \mathbf{M}_1$. In the same way, the application of transformations is not, in general, commutative and therefore the order in which transformations are combined is important. For example, consider the two transformations illustrated in Figure 10.

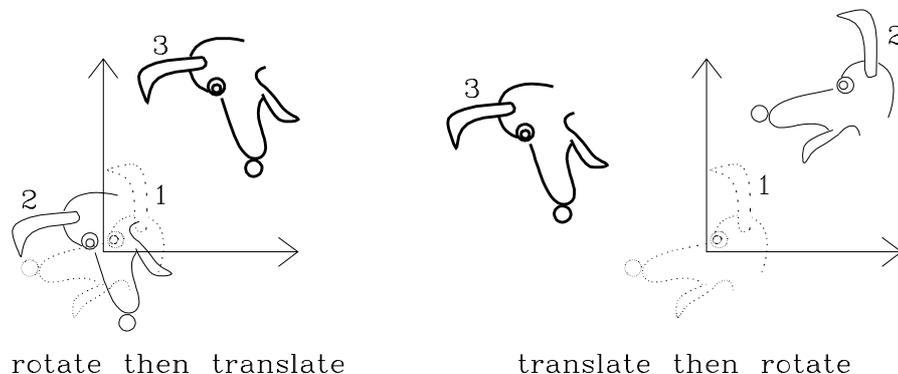


Figure 10: different orders of transformations

The first example shows the effect of rotating and then translating the object. The second example does the same translation and rotation but in a different order, first translating and then rotating the object. The effect in both cases is clearly not the same.

The transformation matrix \mathbf{M}_1 maps the point \mathbf{p} onto the point \mathbf{p}' :

$$\mathbf{p}' = \mathbf{M}_1 \cdot \mathbf{p}$$

If a second transformation \mathbf{M}_2 is to be combined with \mathbf{M}_1 such that \mathbf{M}_1 is applied first followed by \mathbf{M}_2 , then \mathbf{M}_2 is postconcatenated with \mathbf{M}_1 so that

$$\mathbf{p}' = \mathbf{M}_2 \cdot \mathbf{M}_1 \cdot \mathbf{p}$$

Alternatively, \mathbf{M}_2 may be preconcatenated with \mathbf{M}_1 . This will cause \mathbf{M}_2 to be applied first:

$$\mathbf{p}' = \mathbf{M}_1 \cdot \mathbf{M}_2 \cdot \mathbf{p}$$

Note that the order in which transformations are applied can be seen by reading outwards from the vector being transformed. In other words, the transformation which is applied first appears closest to the vector.

Two other terms used for combining transformation matrices are premultiply and postmultiply. Premultiply corresponds to postconcatenate and postmultiply corresponds to preconcatenate. The *pre* and *post* terms in the example transformation program refer to *premultiply* and *postmultiply*.

1.6 Homogeneous Coordinates

To obtain square matrices an additional row was added to the matrix and an additional coordinate, the w -coordinate, was added to the vector for a point. In this way a point in 2D space is expressed in three-dimensional homogeneous coordinates. This technique of representing a point in a space whose dimension is one greater than that of the point is called homogeneous representation. It provides a consistent, uniform way of handling affine transformations.

On converting a 2D point (x, y) to homogeneous coordinates the w -coordinate is set to 1, giving the corresponding homogeneous coordinate point $(x, y, 1)$. This may then be transformed by the 3×3 homogeneous transformation matrix as shown below.

$$\begin{bmatrix} x' \\ y' \\ w' \end{bmatrix} = \begin{bmatrix} a & b & c \\ d & e & f \\ g & h & i \end{bmatrix} \cdot \begin{bmatrix} x \\ y \\ 1 \end{bmatrix}$$

This gives the transformed point (x', y', w') . All of the transformation matrices examined up to now have had $[0 \ 0 \ 1]$ in the bottom row and therefore w' has always been 1. In this case the transformed 2D point is (x', y') . In general, the elements in the bottom row of the matrix, g , h and i , may be set to any value resulting in $w' \neq 1$. The effect of this general transformation matrix is to transform a point $(x, y, 1)$ in the $w = 1$ plane onto the point (x', y', w') in the $w = w'$ plane. The *real-world* coordinate space is the plane $w = 1$ and therefore the transformed point must be mapped back onto the $w = 1$ plane. This is done by projecting the point

(x', y', w') onto the plane $w = 1$, as shown in Figure 11. This process is known as homogeneous division.

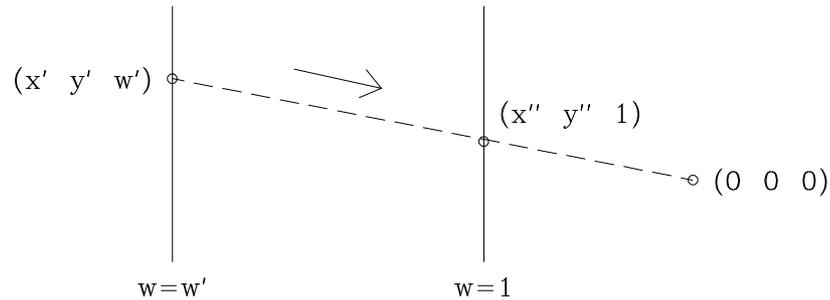


Figure 11: homogeneous division

The 2D *real-world* point is (x'', y'') where x'' and y'' are the x - and y -coordinates of the projected point. The mathematical effect of the projection is that of dividing the x - and y -components by the w -component. Hence

$$x'' = x' / w'$$

$$y'' = y' / w'$$

Now let's look at an example of a transformation matrix with values other than $[0 \ 0 \ 1]$ on the bottom row. Consider the following transformation:

$$\begin{bmatrix} x' \\ y' \\ w' \end{bmatrix} = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 4 \end{bmatrix} \cdot \begin{bmatrix} x \\ y \\ 1 \end{bmatrix}$$

The resulting homogeneous point is $(x, y, 4)$. We obtain the corresponding 2D point by performing the homogeneous division. This gives us the point $(x/4, y/4)$. From this we can see that the bottom right element in the matrix performs overall or uniform scaling.

1.7 Object and Axis Transformations

The types of transformation we have examined up to now are known as object transformations. We think of the object being transformed, while the axes remain fixed. There is another way of looking at transformations - as axis transformations. Here, the object remains fixed while the axes are changed. Figure 12 illustrates the difference between these two types of transformation. In the first example, the sequence of points making up the object are shifted by (d_x, d_y) , the transformed points are plotted relative to the same set of axes. This is an object transformation.

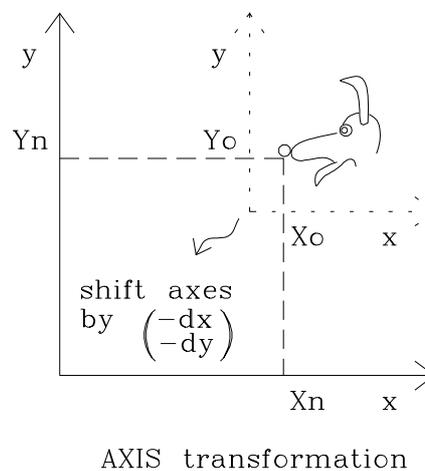
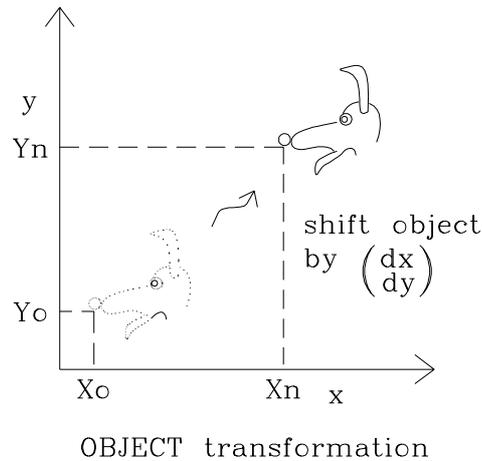


Figure 12: equivalent object and axis transformations

The second example shows an axis transformation in which the axes are shifted by $(-d_x, -d_y)$. This time the points are plotted with respect to the new axes, although they remain fixed in space. After the object transformation which shifts the object by (d_x, d_y) , the new coordinates of the point (x_0, y_0) are given by

$$\begin{aligned}x_n &= x_0 + d_x \\y_n &= y_0 + d_y\end{aligned}$$

This is the same as the new coordinates of the same point after the axis transformation shifting the axes by $(-d_x, -d_y)$. We can deduce from this that an axis translation is equivalent to an equal and opposite object translation. This rule also applies to the other types of transformation. Hence an object transformation which scales an object by (s_x, s_y) is equivalent to an axis transformation which scales the axes by $(1/s_x, 1/s_y)$. Similarly the rotation of an object by α is equivalent to an axis rotation of $-\alpha$. The transformation matrices corresponding to some common object and axis transformations are given below:

- Translate OBJECT by (t_x, t_y)

$$\begin{bmatrix} 1 & 0 & t_x \\ 0 & 1 & t_y \\ 0 & 0 & 1 \end{bmatrix}$$
- Translate AXES by (t_x, t_y)

$$\begin{bmatrix} 1 & 0 & -t_x \\ 0 & 1 & -t_y \\ 0 & 0 & 1 \end{bmatrix}$$
- Scale OBJECT by (s_x, s_y)

$$\begin{bmatrix} s_x & 0 & 0 \\ 0 & s_y & 0 \\ 0 & 0 & 1 \end{bmatrix}$$
- Scale AXES by (s_x, s_y)

$$\begin{bmatrix} 1/s_x & 0 & 0 \\ 0 & 1/s_y & 0 \\ 0 & 0 & 1 \end{bmatrix}$$
- Rotate OBJECT by α

$$\begin{bmatrix} \cos\alpha & -\sin\alpha & 0 \\ \sin\alpha & \cos\alpha & 0 \\ 0 & 0 & 1 \end{bmatrix}$$
- Rotate AXES by $-\alpha$

$$\begin{bmatrix} \cos\alpha & \sin\alpha & 0 \\ -\sin\alpha & \cos\alpha & 0 \\ 0 & 0 & 1 \end{bmatrix}$$

As a general rule, the inverse of an object transformation is the corresponding axis transformation.

1.8 The Normalization Transformation in GKS

The normalization transformation in GKS maps the contents of a window in world coordinate space into a viewport specified in normalized device coordinate space. This is shown in Figure 13.

This type of window to viewport transformation is very common in graphics systems. The window and viewport are aligned with their respective axes and are therefore defined by their bottom left and top right hand corners. The two corners of the window are given by the points (w_{xmin}, w_{ymin}) and (w_{xmax}, w_{ymax}) . The corresponding corners of the viewport are (v_{xmin}, v_{ymin}) and (v_{xmax}, v_{ymax}) . The stages in the transformation are as follows:

- Apply a translation to map the bottom left hand corner of the window to the origin:

$$\mathbf{P}' = \begin{bmatrix} 1 & 0 & -w_{xmin} \\ 0 & 1 & -w_{ymin} \\ 0 & 0 & 0 \end{bmatrix} \bullet \mathbf{P}$$

- Next apply a scaling to make the size of the window the same as the size of

the viewport. If the x - and y -scale factors, s_x and s_y are

$$s_x = \frac{v_{xmax} - v_{xmin}}{w_{xmax} - w_{xmin}}$$

$$s_y = \frac{v_{ymax} - v_{ymin}}{w_{ymax} - w_{ymin}}$$

the transformation becomes

$$\mathbf{P}' = \begin{bmatrix} s_x & 0 & 0 \\ 0 & s_y & 0 \\ 0 & 0 & 1 \end{bmatrix} \cdot \begin{bmatrix} 1 & 0 & -w_{xmin} \\ 0 & 1 & -w_{ymin} \\ 0 & 0 & 0 \end{bmatrix} \cdot \mathbf{P}$$

- Finally, apply a translation so that a point at the origin is mapped onto the bottom left hand corner of the viewport:

$$\mathbf{P}' = \begin{bmatrix} 1 & 0 & v_{xmin} \\ 0 & 1 & v_{ymin} \\ 0 & 0 & 0 \end{bmatrix} \cdot \begin{bmatrix} s_x & 0 & 0 \\ 0 & s_y & 0 \\ 0 & 0 & 1 \end{bmatrix} \cdot \begin{bmatrix} 1 & 0 & -w_{xmin} \\ 0 & 1 & -w_{ymin} \\ 0 & 0 & 0 \end{bmatrix} \cdot \mathbf{P}$$

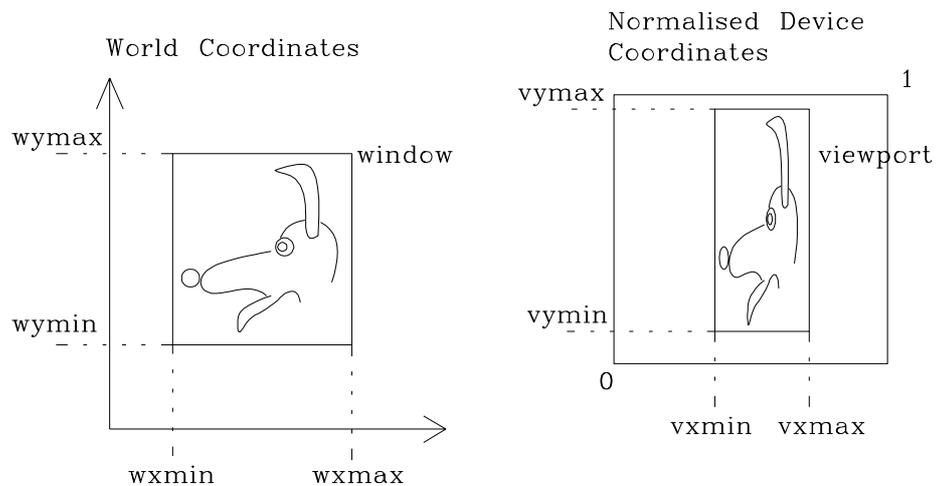


Figure 13: the normalization transformation in GKS

1.9 Summary

We have examined the basic types of transformation: translation, scale, rotation and shear. A transformation may be either an object transformation in which the points of the object are transformed, or an axis transformation in which the coordinate axes are transformed and the object points re-expressed relative to the new axes. All of these transformations can be expressed in a 3×3 matrix which is multiplied with the vector for a point to obtain the coordinates of the transformed point. A 3×3 matrix is used to enable different transformations to be combined by multiplying the matrices together. This means that a 2D point to be transformed must be represented as a three-dimensional homogeneous point $(x, y, 1)$. After

transformation we have the point (x', y', w') . The *real-world* 2D coordinates are obtained by dividing the x - and y -components by the w -component.

2 3D Transformations

2.1 Introduction

We have already looked at the two-dimensional transformations which are used to manipulate pictures. Equivalent transformations are needed to manipulate three-dimensional pictures in three-dimensional space. However, not only are transformations useful as a tool for creating and subsequently altering a picture, they can also help us to visualise the three-dimensional shape of the resulting picture. We would examine an unfamiliar object by picking it up and turning it round to look at it from above, below and from the side, or by holding it at arm's length or standing back from it. In the same way, transformations can be used to rotate, translate or scale a picture to obtain an understanding of its shape. This is particularly important in computer graphics in which the medium for displaying pictures is the two-dimensional display screen on which depth information may not be obvious.

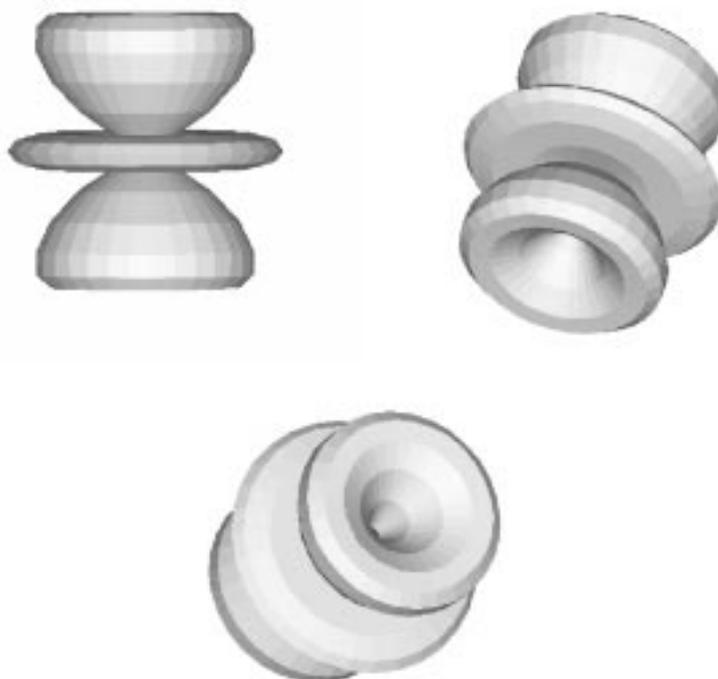
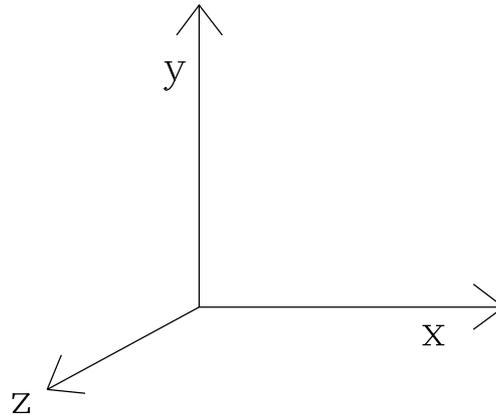


Figure 14: investigating the three-dimensional shape of a surface

In Figure 14 a complex surface is displayed showing the same surface rotated in different ways. Only by looking at a number of these different views can we begin to understand the 3D shape of the surface. These notes will develop techniques for expressing three-dimensional transformations by extending the two-dimensional

techniques already presented. Right-handed coordinate systems will be used, as shown in Figure 15.



z axis points OUT
of the paper

Figure 15: a right-handed coordinate system

2.2 Homogeneous Coordinates

Based on our experience in 2D, we immediately introduce homogeneous coordinates so that a point in 3D space (x, y, z) is represented by a four-dimensional position vector (x, y, z, w) . This point may then be transformed by the following matrix operation:

$$\begin{bmatrix} x' \\ y' \\ z' \\ w' \end{bmatrix} = \begin{bmatrix} a & b & c & d \\ e & f & g & h \\ i & j & k & l \\ m & n & o & p \end{bmatrix} \cdot \begin{bmatrix} x \\ y \\ z \\ w \end{bmatrix}$$

In order to obtain the 3D coordinates from the transformed homogeneous point, we divide the x -, y - and z -components by the w -component:

$$\begin{aligned} x' &= x' / w' \\ y' &= y' / w' \\ z' &= z' / w' \end{aligned}$$

2.3 Types of 3D Transformation

2.3.1 Translation

The matrix to perform 3D translation is shown below

$$\begin{bmatrix} 1 & 0 & 0 & d \\ 0 & 1 & 0 & h \\ 0 & 0 & 1 & l \\ 0 & 0 & 0 & 1 \end{bmatrix}$$

Matrix element d is the displacement added to the x -coordinate, h is the displacement added to the y -coordinate, and l is added to the z -coordinate.

2.3.2 Scaling

3D scaling is performed by the elements on the main diagonal of the matrix:

$$\begin{bmatrix} a & 0 & 0 & 0 \\ 0 & f & 0 & 0 \\ 0 & 0 & k & 0 \\ 0 & 0 & 0 & p \end{bmatrix}$$

The x -, y - and z -scale factors are given by a , f and k respectively. The element p provides overall scaling by a factor of $1/p$.

2.3.3 Rotation

The terms in the upper-left 3×3 component matrix control 3D rotation:

$$\begin{bmatrix} a & b & c & 0 \\ e & f & g & 0 \\ i & j & k & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}$$

The basic 2D rotation was a rotation about the origin in the xy -plane. There are three basic 3D rotations: rotation about the x -axis, rotation about the y -axis and rotation about the z -axis. These are illustrated in Figure 16.



Figure 16: three-dimensional rotations

A rotation about the z -axis is equivalent to the 2D rotation about the origin. Hence we can write down the x and y terms of the matrix straight away. Since we

are rotating about the z -axis, the z -coordinate should not be changed, and so the z -row and column should both be $[0\ 0\ 1\ 0]$ (as in the identity matrix). Similarly for rotation about x and y , the row and column corresponding the axis of rotation are taken from the identity matrix. The cosine and sine terms are then used to fill the remaining elements of the 3×3 component matrix. The matrices for rotation about the three axes are:

- Rotation about the x -axis:
$$\begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & \cos\alpha & -\sin\alpha & 0 \\ 0 & \sin\alpha & \cos\alpha & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}$$

- Rotation about the y -axis:
$$\begin{bmatrix} \cos\alpha & 0 & \sin\alpha & 0 \\ 0 & 1 & 0 & 0 \\ -\sin\alpha & 0 & \cos\alpha & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}$$

- Rotation about the z -axis:
$$\begin{bmatrix} \cos\alpha & -\sin\alpha & 0 & 0 \\ \sin\alpha & \cos\alpha & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}$$

2.3.4 Shearing

The off-diagonal elements in the upper-left 3×3 component matrix produce 3D shearing effects:

$$\begin{bmatrix} 1 & b & c & 0 \\ e & 1 & g & 0 \\ i & j & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}$$

In 3D a shear in x may be obtained as a function of the y - and z -coordinates. This is controlled by matrix elements b and c respectively. Similarly elements e and g control shearing in y as a function of x and z and elements i and j control shearing in z as a function of x and y .

2.3.5 Rotation about an arbitrary axis

Having looked at the basic 3D transformations we will now look at a more complex example involving a combination of these transformations. A common requirement is to rotate an object about an arbitrary axis rather than one of the coordinate axes. In the notes on 2D transformations the transformation for rotating about an arbitrary point was derived. This involved shifting the object and point so that the point coincides with the origin, a rotation about the origin and finally a second shift which is the inverse of the first to place the object back in its original position. Similarly a 3D rotation about an arbitrary axis involves transforming the object and axis of rotation so that the axis coincides with one of the coordinate axes, followed by a rotation about the coordinate axis and finishing

with a transformation which is the inverse of the first. The individual steps are as follows:

- Translate so that axis of rotation passes through the origin.
- Rotate object so that axis of rotation coincides with one of the coordinate axes.
- Perform the specified rotation about appropriate coordinate axis.
- Apply inverse rotations to bring axis of rotation back to original orientation.
- Apply inverse translation to bring rotation axis back to original position.

For more details of this transformation see Hearn and Baker[1]. This derives the rotations required to orientate the axis of rotation so that it is parallel to one of the coordinate axes.

2.4 Perspective Transformations

All of the 3D transformation matrices examined so far have been of the form

$$\begin{bmatrix} a & b & c & d \\ e & f & g & h \\ i & j & k & l \\ 0 & 0 & 0 & p \end{bmatrix}$$

i.e. elements m , n and o are equal to zero. This section looks at the effect achieved when one or more of these values is non-zero. Consider the following transformation

$$\begin{bmatrix} x' \\ y' \\ z' \\ w' \end{bmatrix} = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & \gamma & 1 \end{bmatrix} \cdot \begin{bmatrix} x \\ y \\ z \\ 1 \end{bmatrix} = \begin{bmatrix} x \\ y \\ z \\ \gamma \cdot z + 1 \end{bmatrix}$$

After homogeneous division the real 3D coordinates of the transformed point are (x'', y'', z'') where

$$\begin{aligned} x'' &= x / (\gamma \cdot z + 1) \\ y'' &= y / (\gamma \cdot z + 1) \\ z'' &= z / (\gamma \cdot z + 1) \end{aligned}$$

As the original z -coordinate tends to infinity

$$\begin{aligned} x'' &\rightarrow 0 \\ y'' &\rightarrow 0 \\ z'' &\rightarrow 1/\gamma \end{aligned}$$

Hence, after transformation lines originally parallel to the z -axis will appear to pass through the point $(0, 0, 1/\gamma)$, known as the *vanishing point*. This kind of transformation is known as a *perspective* transformation and is illustrated in Figure 17. A house aligned with the x -, y - and z -axes is shown before and after a perspective transformation.

A perspective transformation has a distorting effect which gives the transformed object a natural appearance, similar to that which would be seen by the eye from the point $(0, 0, -1/\gamma)$. The eye point is often referred to as the *centre of projection*.

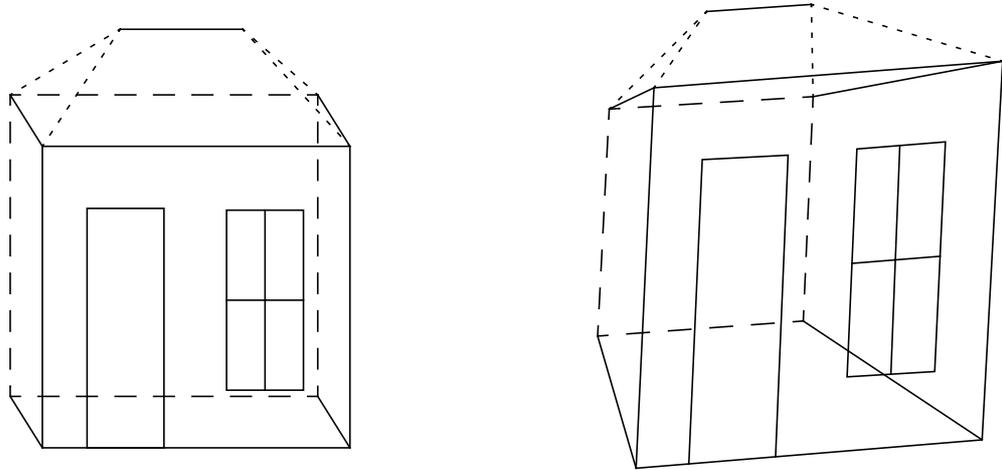


Figure 17: a perspective transformation

Different types of perspective transformation are obtained if the other two elements on the bottom row are set. For example, the matrix

$$\begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ \alpha & 0 & 0 & 1 \end{bmatrix}$$

would create a perspective transformation with a vanishing point for lines originally parallel to the x -axis at $(1/\alpha, 0, 0)$ and a centre of projection at $(-1/\alpha, 0, 0)$. Similarly the matrix

$$\begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & \beta & 0 & 1 \end{bmatrix}$$

would create a perspective transformation with a vanishing point for lines originally parallel to the y -axis at $(0, 1/\beta, 0)$ and a centre of projection at $(0, -1/\beta, 0)$. These two cases are illustrated in Figure 18.

Perspective transformations with only one vanishing point are known as one point perspective transformations. If two or three of the matrix elements are non-zero together, a two or three point perspective transformation is obtained, as shown in Figure 19.

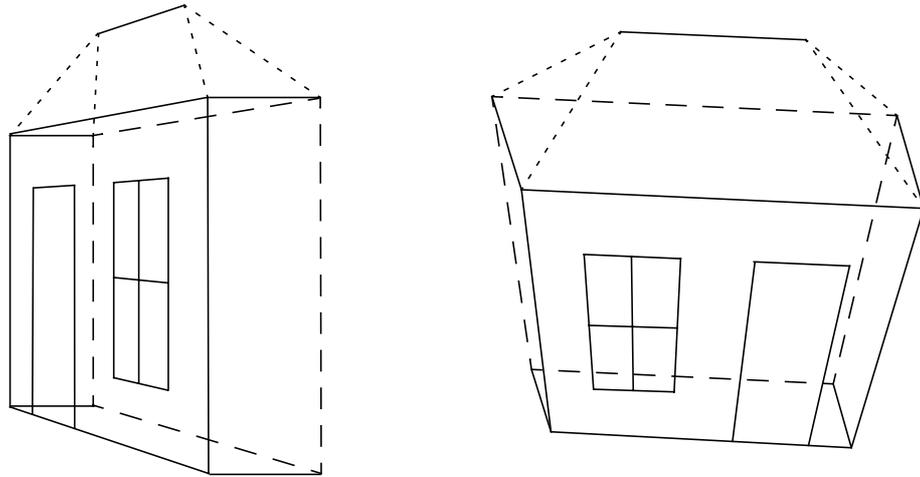


Figure 18: one point perspective transformations

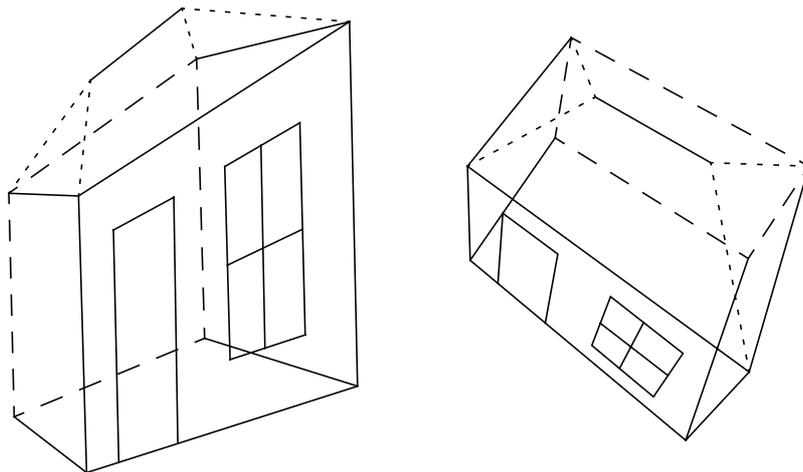


Figure 19: two and three point perspective transformations

2.4.1 Points behind the eye point

The following shows a perspective transformation applied to the point (x, y, z) with the eye point (centre of projection) at $(0, 0, c)$:

$$\begin{bmatrix} x' \\ y' \\ z' \\ w' \end{bmatrix} = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & -1/c & 1 \end{bmatrix} \cdot \begin{bmatrix} x \\ y \\ z \\ 1 \end{bmatrix} = \begin{bmatrix} x \\ y \\ z \\ (c-z)/c \end{bmatrix}$$

The value of w' varies depending on the value of the original z -coordinate as shown below

$$z < c \rightarrow w' > 0$$

$$z = c \rightarrow w' = 0$$

$$z > c \rightarrow w' < 0$$

This is illustrated in Figure 20.

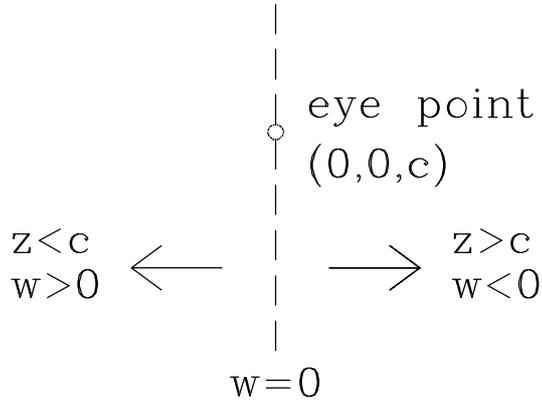


Figure 20: variation of w' with the original z -coordinate in a perspective transformation

Now consider a line which joins the points \mathbf{p}_1 and \mathbf{p}_2 which are positioned on either side of the eye point (\mathbf{p}_1 has $z < c$ and \mathbf{p}_2 has $z > c$). After transformation \mathbf{p}_1 will have a positive w value and \mathbf{p}_2 will have a negative w value. Figure 21 shows the transformed line in homogeneous coordinate space (for simplicity, only the xw -plane is drawn). Homogeneous division then projects the transformed points \mathbf{p}_1' and \mathbf{p}_2' onto the $w = 1$ plane. The resulting line however is not the line joining the projected points \mathbf{p}_1'' and \mathbf{p}_2'' . In Figure 22 other points along the line $\mathbf{p}_1'\mathbf{p}_2'$ are projected onto the $w = 1$ plane. It can be seen from this that the projected line is actually in two parts: \mathbf{p}_1'' to positive infinity and negative infinity to \mathbf{p}_2'' . These are known as *external line segments* and is the correct interpretation of the application of a perspective transformation to a line joining points on either side of the eye point. More details can be found in Blinn and Newell[2].

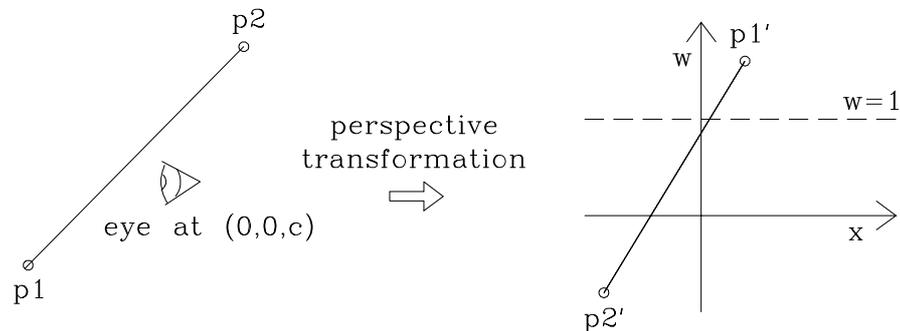


Figure 21: perspective transformation of the line $\mathbf{p}_1\mathbf{p}_2$

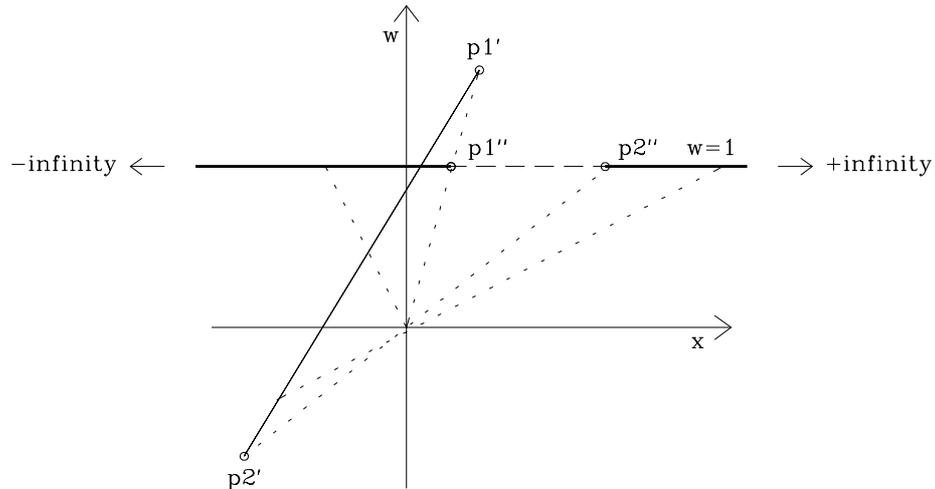


Figure 22: projecting line $\mathbf{p}_1'\mathbf{p}_2'$ onto the $w = 1$ plane

2.4.2 Clipping

In general we want the results of a perspective transformation to simulate what the eye would actually see. Since the eye can only see objects in front of it, items behind the eye should be clipped out. This can be done in one of two ways:

- Clip before the perspective transformation by removing all those parts with $z \geq c$.
- Clip after perspective transformation but *before* homogeneous division, in this case remove parts with $w \leq 0$.

After the homogeneous division it is impossible to distinguish between points which were originally behind and in front of the eye.

2.5 Projections

The 3D transformations examined so far have transformed one 3D object into another 3D object. There is another class of transformations which transform a 3D object into a 2D object by *projecting* it onto a plane. These transformations, known as *projections*, are of particular interest in computer graphics in which the finished picture must always be projected onto the flat viewing surface of a graphics display. Projecting a three dimensional object onto a plane is similar to casting a shadow of the object onto a flat surface. Figure 23 illustrates two cases, one where the light source is a finite distance from the object, and the other where the light source is a long distance from the object, so that the light rays hitting the object are approximately parallel. Similarly there are two types of projection. A perspective projection in which the centre of projection is a finite distance from the object, and a parallel projection in which the centre of projection is a long distance from the object, so that the projectors hitting the object are approximately

parallel. Rather than a light source, a projection is identified by the centre of projection, projectors replace rays of light.

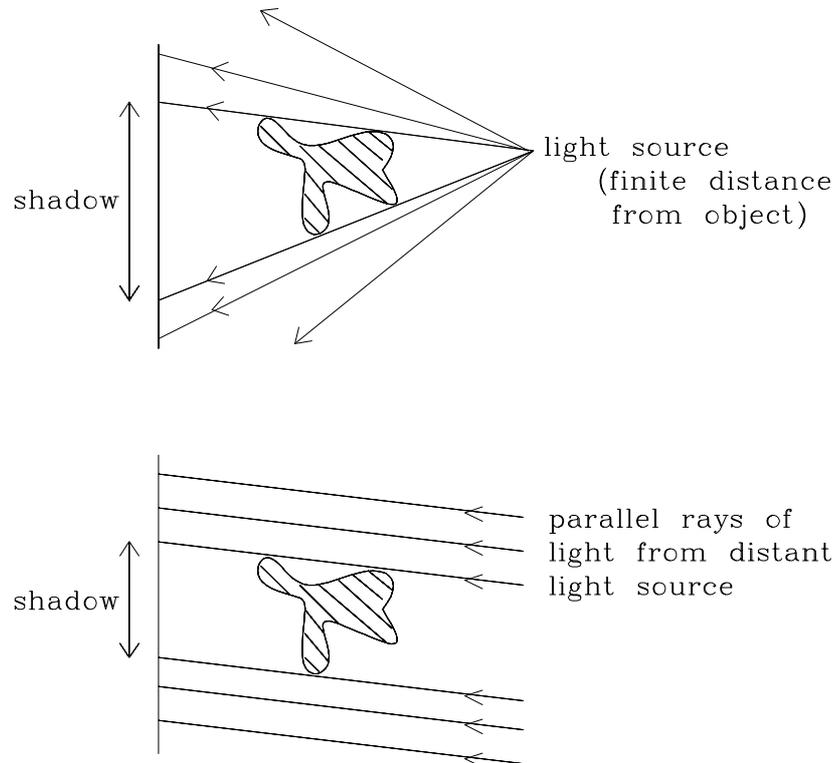


Figure 23: casting shadows

2.5.1 Perspective projections

A perspective projection is shown in Figure 24. It is defined by a centre of projection and a view plane onto which a 2D image of the 3D object is projected. The resulting image is the same as if the eye had been placed at the centre of projection. Figure 25 illustrates a perspective projection in which the centre of projection is the point $(0, 0, c)$ and the view plane is the plane $z = 0$. The point \mathbf{p} is projected onto the point \mathbf{p}' by drawing a projector from the centre of projection, \mathbf{e} , through \mathbf{p} onto the view plane. The point at which the projector intersects the view plane is \mathbf{p}' . From the figure, similar triangles $\mathbf{ep}'\mathbf{b}$ and $\mathbf{pp}'\mathbf{a}$ can be used to find the coordinates of \mathbf{p}' :

$$x'/c = (x' - x)/z$$

so that

$$x' = x/(1 - z/c)$$

Similarly,

$$y' = y/(1 - z/c)$$

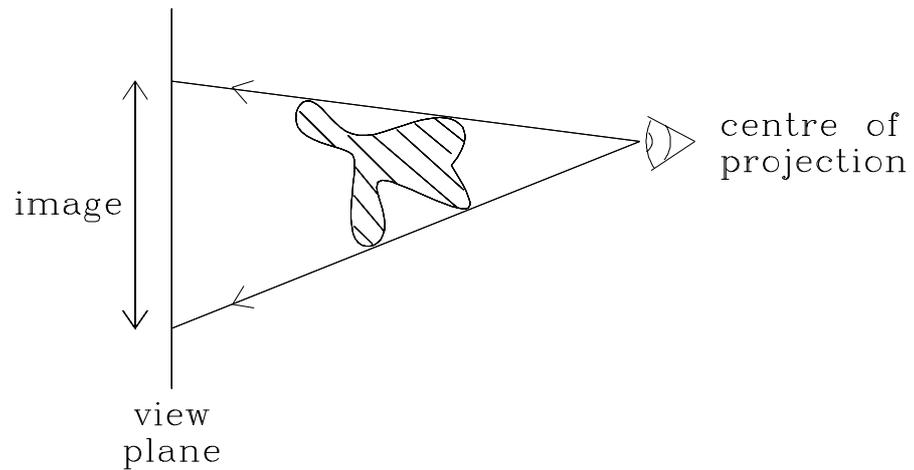


Figure 24: a perspective transformation

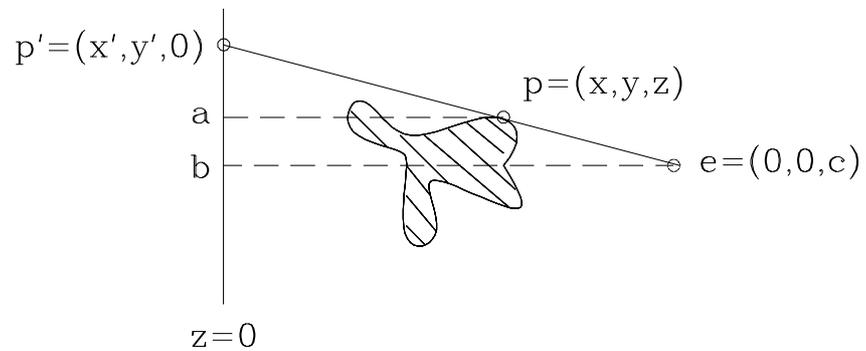


Figure 25: another perspective projection

We can see from this that the effect of a perspective projection on the x - and y -coordinates is the same as that of a perspective transformation. Since the view plane is the plane $z = 0$, all z -coordinates will be mapped to zero.

2.6 Parallel Projections

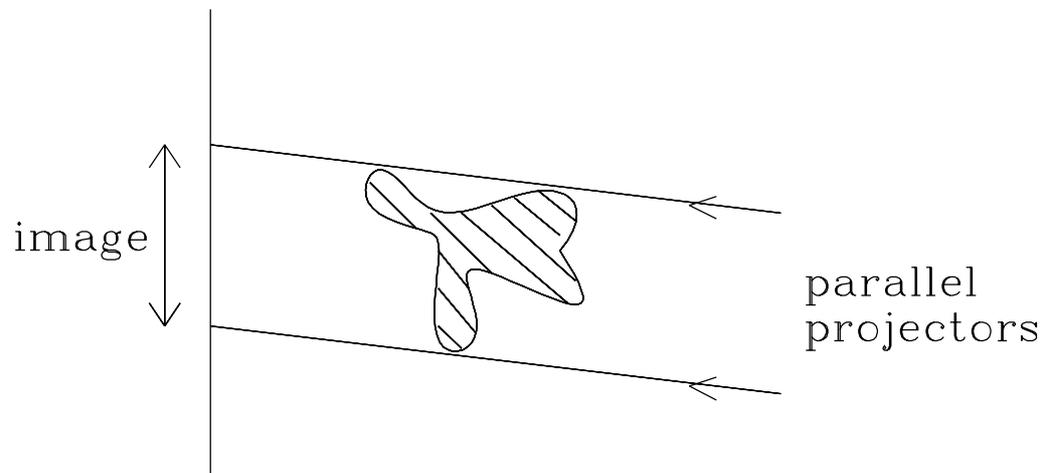


Figure 26: a parallel projection

A parallel projection is shown in Figure 26. This can be described as a perspective projection in which the centre of projection has been moved to infinity. In this case, rather than a centre of projection, the direction of the parallel projectors defines the projection together with the view plane. Although parallel projections tend to produce less realistic views of an object, they are useful as they usually maintain some of the original size information about the original object and can therefore be used for measurements.

2.7 Classification of Planar Geometric Projections

There are many different types of parallel and perspective projections. These are classified in Figure 27. Each type can be obtained by choosing a suitable view plane and centre of projection. This section gives a brief description of the different types of projection, how they are obtained and where they might be used. A thorough review can be found in Carlbom and Paciorek[3].

- **Parallel** Obtained when the centre of projection is at infinity, the projectors are always parallel. Parallel projections are divided into:
 - **Orthographic** The projectors are perpendicular to the view plane. There are two different types of parallel orthographic projection:
 - **Multiview orthographic** The view plane is parallel to one of the principal planes of the object. Usually a number of views are shown together. The exact shape of one face in the object is retained but it is usually hard to visualise the three dimensional shape from the projection. This type of projection is typically used in engineering drawings.
 - **Axonometric** The view plane is chosen so that the projection will illustrate the general three dimensional shape of the object. Parallel lines are equally foreshortened and remain parallel. Uniform foreshortening occurs along principal axes enabling measurements to be taken to scale along these axes. Axonometric projections are classified according to the angles made between the principal axes and the view plane:

- **Trimetric** The angles between the principal axes and the view plane are all different. This means that the foreshortening along each axis will be different.
- **Dimetric** Two of the angles between the principal axes and the view plane are equal and therefore foreshortening along two axes will be the same.
- **Isometric** All of the angles between the principal axes and the view plane are equal and therefore foreshortening along all three axes will be the same.
- **Oblique** The projectors are NOT perpendicular to the view plane. Usually the view plane is positioned parallel to the principal face of the object so that this face is projected without distortion. This allows direct measurements to be made of this face. The angle between the projectors and the view plane is chosen so as to best illustrate the third dimension, two examples are:
 - **Cavalier** The angle is 45° - faces parallel and perpendicular to the view plane are projected at their true size.
 - **Cabinet** The angle is $\text{arccot}(0.5)$ (approximately 64°) - faces parallel to the view plane are projected at true size, and faces perpendicular to the view plane are projected at half size.
- **Perspective** Obtained when the centre of projection is a finite distance from the object, the projectors emanate from the centre of projection. Perspective projections are classified according to the number of vanishing points in the projection.
 - **One point** When only one of the principal axes intersects the view plane.
 - **Two point** When exactly two of the principal axes intersect the view plane.
 - **Three point** When all three principal axes intersect the view plane.

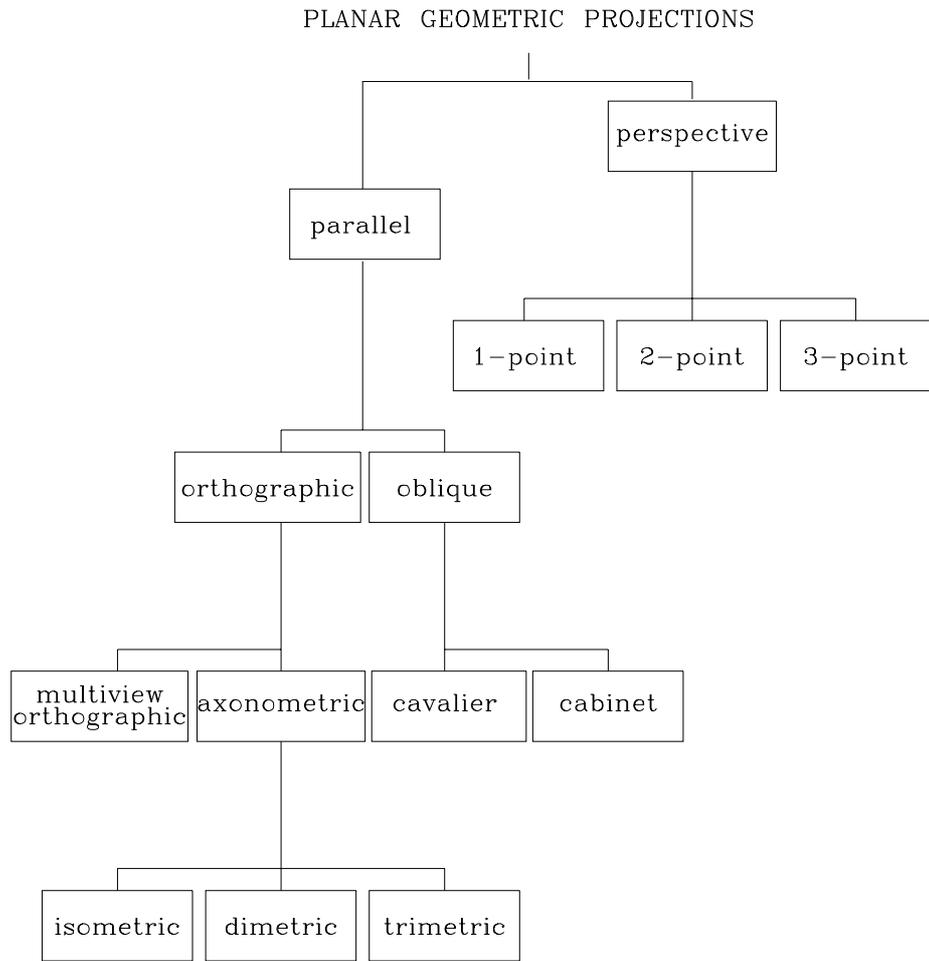


Figure 27: classification of the planar geometric projections

3 Transformations and Viewing in GKS-3D and PHIGS

We have already studied the theory of three-dimensional transformations. These notes will look at the practical aspects of three-dimensional transformations, in particular, how transformations are used in GKS-3D[5] and PHIGS[6].

In GKS, two-dimensional graphical objects defined by the application programmer undergo a series of transformations to convert them to device coordinates for display. These transformations are the normalization transformation, segment transformation and workstation transformation, which together with the clip are collectively known as the GKS output pipeline. In the same way GKS-3D and PHIGS each define an output pipeline which maps the user's 3D picture onto a physical graphics device. These two pipelines however are more complex than the GKS output pipeline because they must also include transformations to produce different views of 3D scenes.

In these notes we will examine the output pipelines in GKS-3D and PHIGS, concentrating particularly on the part of the pipeline which generates the view, known as the viewing pipeline.

3.1 The GKS-3D Output Pipeline

The GKS-3D output pipeline is illustrated in Figure 28.

Figure 29 lists the functions which allow the application programmer to control the effects of the pipeline.

In GKS-3D, output primitives are specified by the user in World Coordinates ($WC3$) and these are mapped to Normalized Device Coordinates ($NDC3$) by the *normalization transformation*. This mapping is similar to the normalization transformation in GKS, the difference being that the window and viewport are cuboids rather than rectangles. The functions SET WINDOW 3 and SET VIEWPORT 3 are used to define the respective volumes in $WC3$ and $NDC3$ for a given normalization transformation. SELECT NORMALIZATION TRANSFORMATION sets the current normalization transformation.

Figure 30 shows the normalization transformation in action.

Depending upon the level of GKS-3D in use, a *segment transformation* may be applied. This is a mapping from $NDC3$ to $NDC3$ which is specified as a 3×4 matrix and therefore may translate, scale, shear and rotate primitives contained in a segment. The segment transformation is defined using the function SET SEGMENT TRANSFORMATION 3. Animation effects can be achieved by repeatedly changing the segment transformation.

The *normalization clip* is optional and is only performed on primitives with an associated clipping indicator value of CLIP. Primitives are clipped against the cuboid in *NDC3* defined by the current normalization viewport.

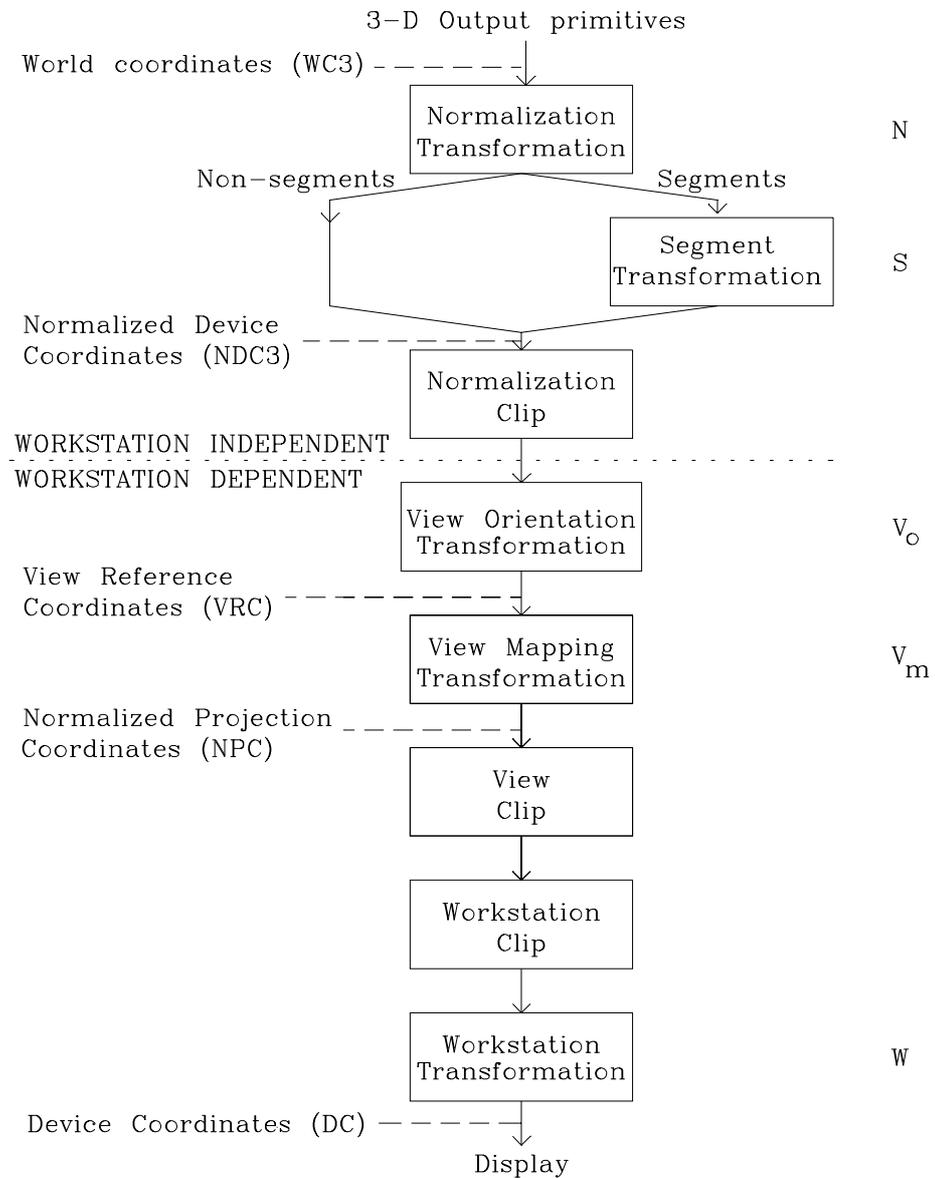


Figure 28: GKS-3D output pipeline

Figure 31 illustrates the effect of changing the segment transformation with the normalization clip enabled. As in GKS, the clipping boundary is not affected by the transformation and this may lead to unexpected results for the user.

The action of the pipeline up to this point is workstation independent. This means that the normalization and segment transformations and the normalization clip will have the same effect on all workstations. Following the normalization clip, the workstation takes charge and therefore subsequent stages may vary between workstations.

	<i>GKS-3D</i>	<i>PHIGS</i>
<i>Normalisation Transformation</i>	SET WINDOW SET VIEWPORT SELECT NORMALIZATION TRANSFORMATION	_____
<i>Segment Transformation</i>	SET SEGMENT TRANSFORMATION	_____
<i>Modelling Transformation</i>	_____	SET LOCAL TRANSFORMATION SET GLOBAL TRANSFORMATION
<i>Viewing Pipeline</i>		SET VIEW INDEX
<i>1. Utility Functions</i>		EVALUATE VIEW ORIENTATION MATRIX EVALUATE VIEW MAPPING MATRIX
<i>2. View Representation</i>		SET VIEW REPRESENTATION
<i>Workstation Transformation</i>		SET WORKSTATION WINDOW SET WORKSTATION VIEWPORT

Figure 29: GKS-3D and PHIGS output pipeline functions

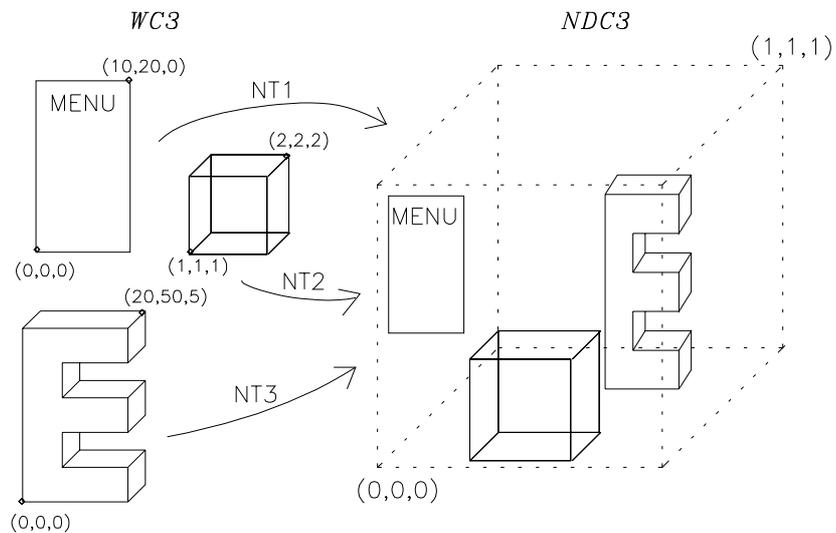


Figure 30: combining different world coordinate systems into NDC3

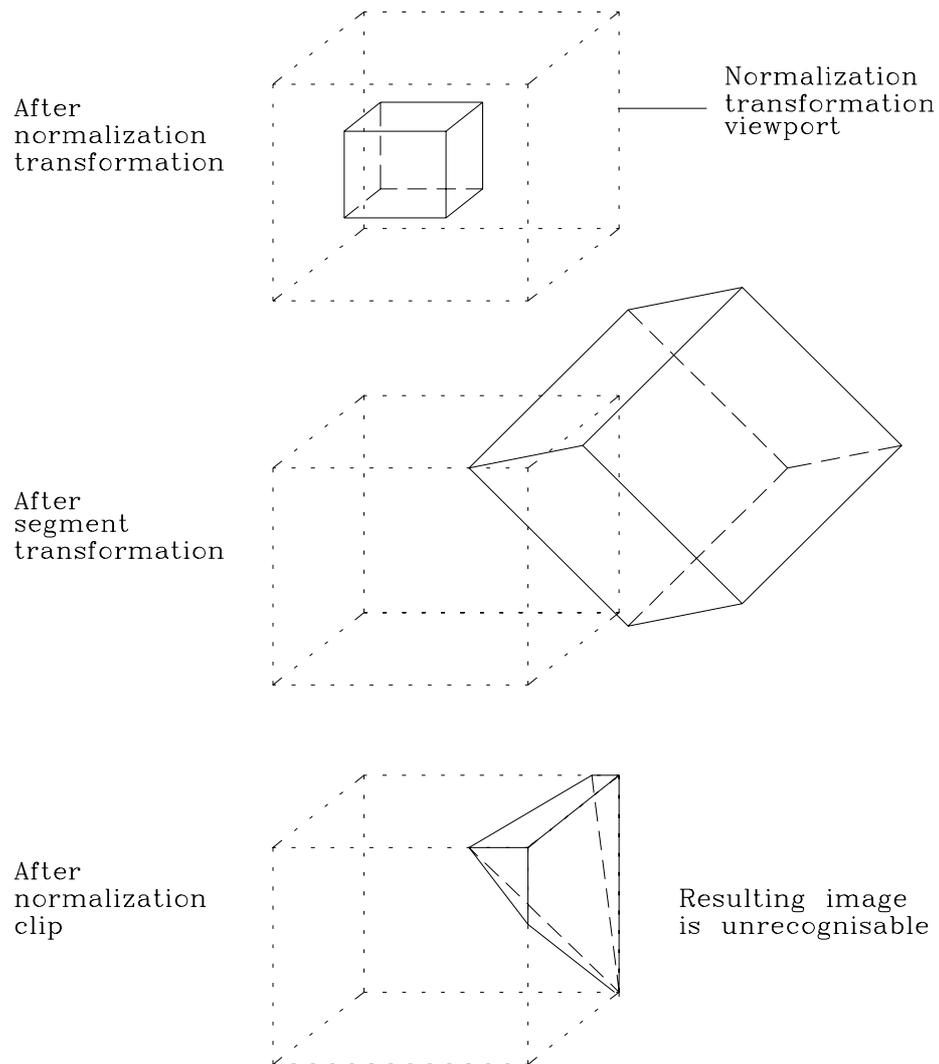


Figure 31: primitives clipped after the segment transformation

The next three stages of the pipeline are concerned with generating views and form the *viewing pipeline*. When primitives are created, they have an associated *view index*. This is defined using the function `SET VIEW INDEX`. The view index selects a particular *view representation* from the workstation's *view table*. The view representation defines two transformations which determine how the primitive will be viewed. It also defines a clipping volume to which primitives may be clipped after viewing. View representations are set on a workstation using the function `SET VIEW REPRESENTATION`. The first viewing transformation defined in the view representation is the *view orientation* transformation. This maps from $NDC3$ to the View Reference Coordinate system (VRC). VRC is an intermediate coordinate system introduced in order to simplify specification of the view required. It is effectively a rotated, translated version of $NDC3$. The second transformation defined in the view representation is the *view mapping transformation*. This creates the view required by applying a parallel or perspective mapping and then converts from VRC to Normalized Projection Coordinates (NPC). Primitives are then clipped to the *view clipping limits* in NPC . The view clipping limits define a cuboid aligned with the NPC axes. Clipping indicators in the view representation determine separately whether or not clipping is performed at the front, back and sides of the volume.

Optional hidden-line or hidden-surface removal (HLHSR) is performed in *NPC* using an implementation defined method.

The last stage in the pipeline is where a portion of *NPC* space is selected for display on the workstation. Before this, the *workstation clip* is performed. Unlike the other clipping operations described so far, the workstation clip is mandatory. Primitives are clipped against a cuboid defined in *NPC* space known as the workstation window which is defined for a given workstation using the function SET WORKSTATION WINDOW 3. The *workstation transformation* is then used to map *NPC* to device coordinates (*DC3*). The contents of the workstation window are mapped onto a cuboid defined in the workstation's physical display space, known as the workstation viewport. The workstation viewport is specified by the function SET WORKSTATION VIEWPORT 3. As in GKS, the workstation transformation preserves aspect ratio in *x* and *y*. However, the *z* extent of the workstation window is always mapped to the entire *z* extent of the workstation viewport.

The total transformation effect of the output pipeline in GKS-3D can be summarized as follows:

$$\mathbf{p}_d = \mathbf{W} \cdot \mathbf{V}_m \cdot \mathbf{V}_o \cdot \mathbf{S} \cdot \mathbf{N} \cdot \mathbf{p}_w$$

where the point \mathbf{p}_w in *WC3* is mapped onto the point \mathbf{p}_d in *DC3* by the transformation matrices: \mathbf{N} = matrix to perform normalization transformation, \mathbf{S} = matrix to perform segment transformation, \mathbf{V}_o = view orientation matrix, \mathbf{V}_m = view mapping matrix and \mathbf{W} = matrix to perform workstation transformation.

Logically there are several places in the pipeline where primitives are clipped. It is possible however, for an implementation to transform all the clipping limits to the end of the pipeline and combine them all together. In this way, the output pipeline can be implemented as a single transformation (as shown above) followed by a single clip. The workstation and viewing clipping volumes are aligned in *NPC* space and can be combined quickly and easily. The normalization clipping volume is not usually aligned with the workstation and viewing clipping volumes, so if the normalization clip is active, then clipping to a convex volume rather than a cuboid must take place.

3.2 The PHIGS Output Pipeline

The PHIGS output pipeline is illustrated in Figure 32. Figure 29 lists the functions provided for the application to control the effects of the pipeline.

In PHIGS, structure elements which create output primitives are specified in Modelling Coordinates (*MC*), which are converted by modelling transformations to World Coordinates (*WC*). There are two types of modelling transformation, the *local* and *global* modelling transformations. These are defined using the functions SET LOCAL MODELLING TRANSFORMATION and SET GLOBAL MODELLING TRANSFORMATION. The local modelling transformation is applied first, followed by the global modelling transformation. Both are specified by the user as 4×4 homogeneous transformation matrices and could be any affine transformation.

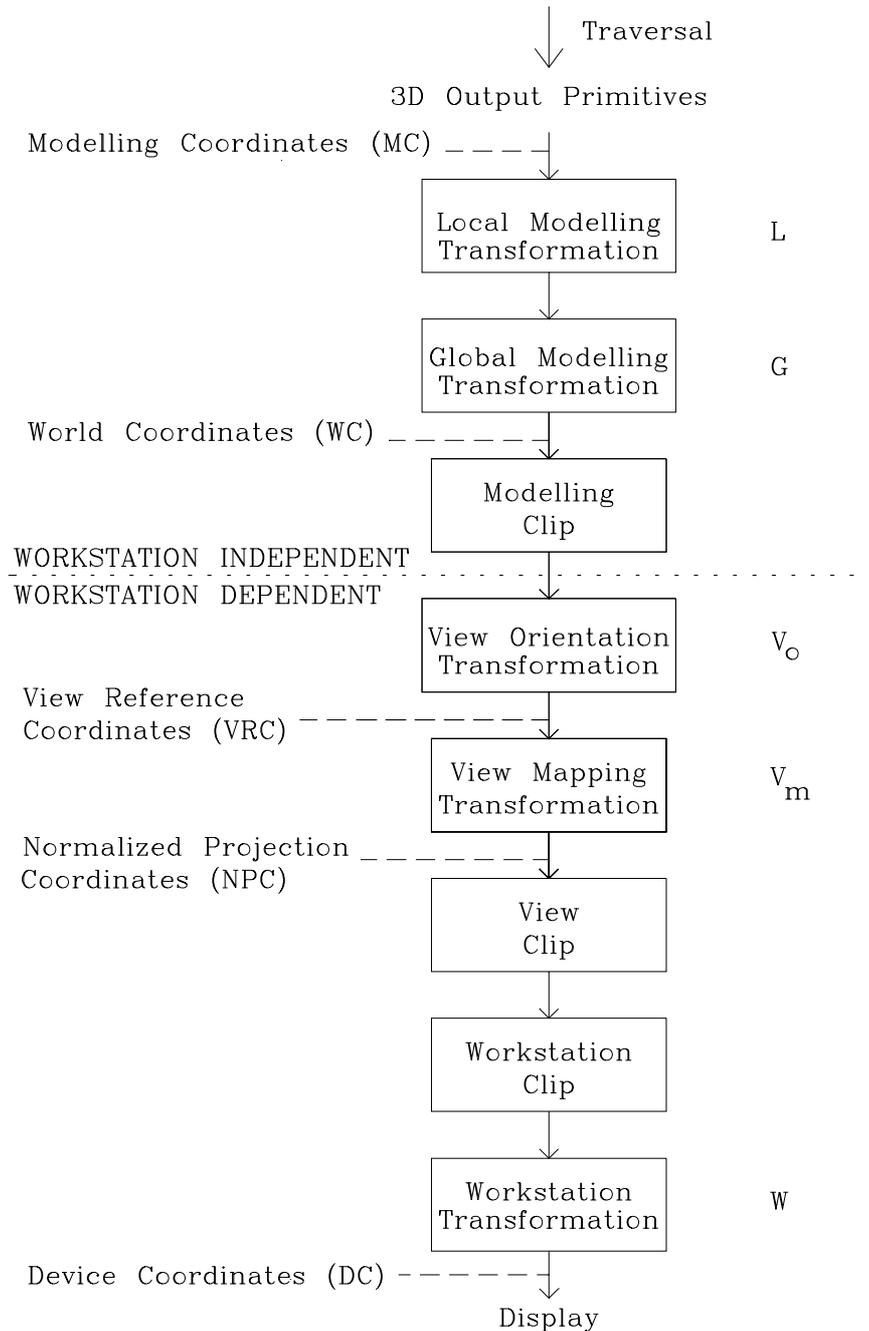


Figure 32: PHIGS output pipeline

Following the modelling transformations, an optional modelling clip to a volume in *WC* is performed. All the clipping operations described so far have involved clipping against a cuboid aligned with the coordinate axes. The modelling clip is more flexible. The user specifies a clipping region by defining an arbitrary number of half planes which are intersected to create a convex space as the clipping region. The half planes are specified in modelling coordinates and are therefore subject to the modelling transformations. This means that the same part of the picture is always visible after clipping, even when the whole picture is transformed. This is illustrated in Figure 33.

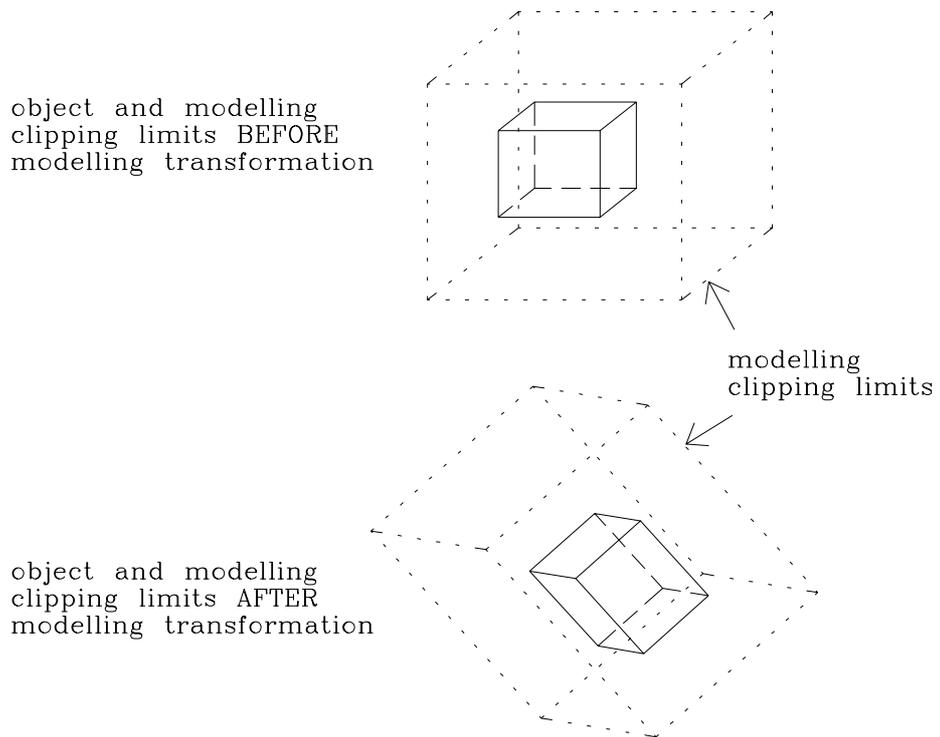


Figure 33: modelling clipping limits subject to modelling transformation

Following the modelling clip, the remaining parts of the pipeline are the same as those in GKS-3D. Once again we can summarize the total transformation effect of the pipeline:

$$\mathbf{p}_d = \mathbf{W} \cdot \mathbf{V}_m \cdot \mathbf{V}_o \cdot \mathbf{G} \cdot \mathbf{L} \cdot \mathbf{p}_m$$

which maps the point \mathbf{p}_m in *MC* onto the point \mathbf{p}_d in *DC3* and where \mathbf{L} = matrix to perform local modelling transformation, \mathbf{G} = matrix to perform global modelling transformation, \mathbf{V}_o = view orientation matrix, \mathbf{V}_m = view mapping matrix and \mathbf{W} = matrix to perform workstation transformation.

As in GKS-3D, by transforming the modelling, view and workstation clipping limits to the end of the pipeline, the PHIGS output pipeline can be implemented as a single transformation followed by a clip against a convex volume.

In the following, wherever *NDC3* is referenced, this should be read as *WC* in the context of PHIGS.

3.3 The Viewing Pipeline

This section looks at the viewing stages of the pipeline in more detail.

In GKS-3D, a view index is bound to each graphical output primitive as it is created. This index specifies an entry in the workstation table of view representations. View representations are defined separately for each workstation and therefore the view of a primitive on one workstation may be different to the view displayed on another workstation. Since the view index is bound to a primitive

when it is created, it is not possible to display different views of the same primitive together on a single workstation.

In PHIGS, the current view index is bound to graphical output primitives at traversal time. A structure may be traversed more than once if it is instanced by other structures. If the current view index is different at each traversal, different views of the structure could be displayed at a single workstation.

We have already looked at the three viewing stages in the pipeline. They are

- View orientation transformation - maps from *NDC3* to *VRC*.
- View mapping transformation - creates the view and maps to *NPC*.
- View clip - optional clip against planes of a cuboid in *NPC*.

The view orientation and view mapping transformations are both specified as 4×4 homogeneous transformation matrices by the user. This gives the user maximum flexibility to set up the view he requires. However, many users will not know how to manipulate a matrix to achieve the required viewing effects and so GKS-3D and PHIGS provide two utility routines, EVALUATE VIEW ORIENTATION MATRIX and EVALUATE VIEW MAPPING MATRIX, which compute the two viewing matrices from a set of viewing parameters. These viewing parameters allow the user to set up a view based on the viewing model defined by GKS-3D and PHIGS.

3.3.1 Model for view orientation

The view orientation transformation maps from *NDC3* to *VRC*. *VRC* is an intermediate coordinate system introduced to simplify the specification of the view mapping. As we will see later, the view mapping is very similar to a projection in that it is defined by a view plane and a centre of projection. In order to obtain different types of views, it must be possible to position the view plane arbitrarily in space. However, the specification of a plane can be complex and therefore the intermediate *VRC* system is introduced in which the view plane is defined as parallel to the *xy*-plane. In this way, the view plane can be described as a distance along the *VRC z*-axis. The view orientation transformation is therefore an *axis* transformation. Objects remain fixed in space but their coordinates are re-expressed relative to the *VRC* system. This is shown in Figure 34. The *x*-, *y*- and *z*-coordinate axes in *VRC* are, by convention, referred to as the *u*-, *v*- and *n*-axes respectively.

The function EVALUATE VIEW ORIENTATION MATRIX computes the view orientation matrix from a set of parameters which define the position and orientation of the *VRC* system with respect to *NDC3*. These parameters are:

- View Reference Point (**VRP**). This is the origin of the view reference coordinate system. It is usually chosen to be a point on or near the object to be viewed.
- View Plane Normal (**VPN**). **VPN** is a vector defining the *n*- (or *z*-) axis of the view reference coordinate system.
- View Up Vector (**VUV**). **VUV** is a vector which is used to compute the direction of the *v*- (or *y*-) axis in the view reference coordinate system. The *v*-axis is defined as an orthogonal projection of **VUV** onto the plane through **VRP** with plane normal **VPN**, as shown in Figure 35. Lines parallel to the projected **VUV** in *NDC3* will appear vertical on the screen.

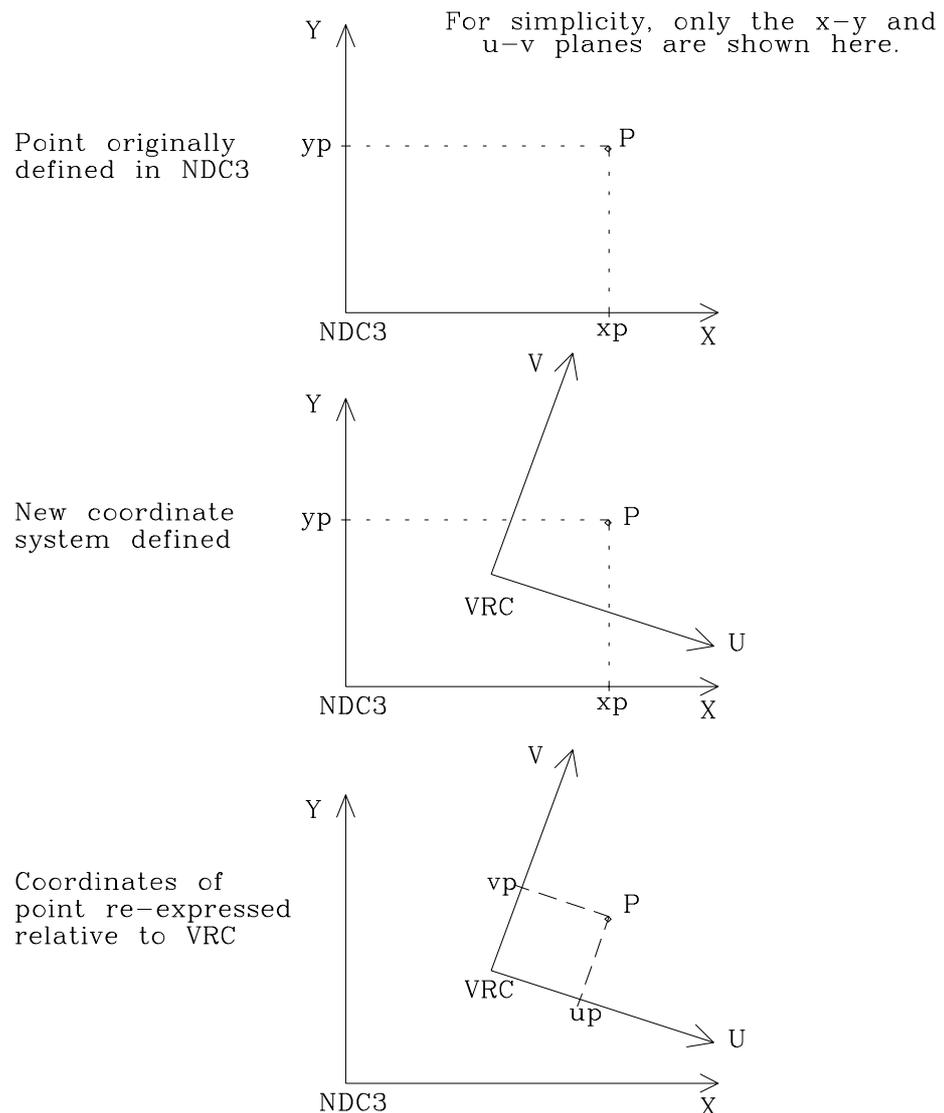


Figure 34: the view orientation transformation

A vector defining the u - (or x -) axis can easily be derived since it must be perpendicular to both \mathbf{VPN} and \mathbf{VUV} . (It is in fact the vector cross product of these two.)

In the lecture on two-dimensional transformations, we saw that the inverse of an axis transformation is the same as an equivalent object transformation. The simplest way to compute the view orientation matrix is by first deriving its inverse. The inverse view orientation transformation is an object transformation which maps unit vectors along the $NDC3$ axes to be coincident with the corresponding VRC axes. This transformation can be divided into two steps:

- Rotate the unit vectors about the origin so that they are aligned with the VRC axes.
- Translate by the view reference point, \mathbf{VRP} , so that a point at the $NDC3$ origin is mapped onto the VRC origin.

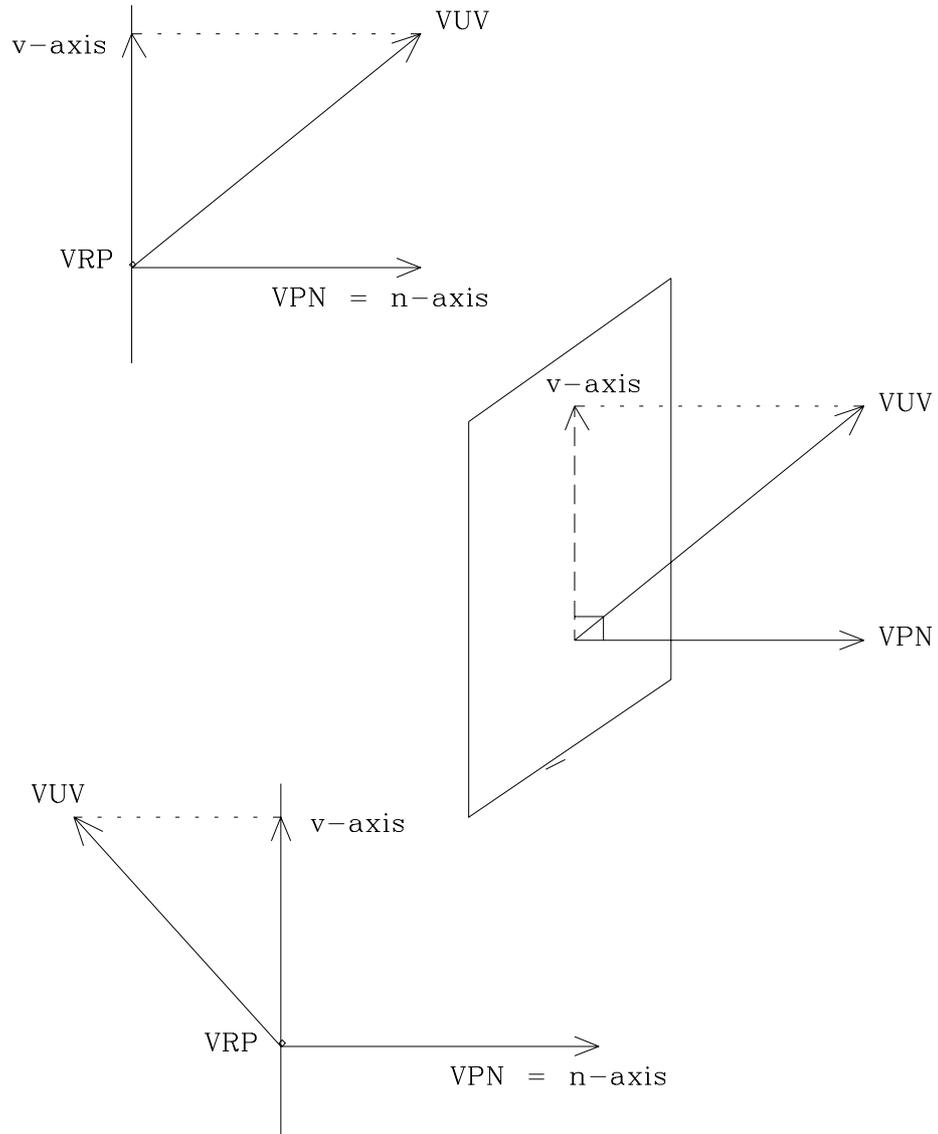


Figure 35: generating the VRC v-axis from VUV

A matrix representing the translation part can be written down immediately:

$$\mathbf{T} = \begin{bmatrix} 1 & 0 & 0 & VPR_x \\ 0 & 1 & 0 & VPR_y \\ 0 & 0 & 1 & VPR_z \\ 0 & 0 & 0 & 1 \end{bmatrix}$$

where VPR_x , VPR_y and VPR_z are the x -, y - and z -components of the view reference point. The rotation maps unit vectors along the x -, y - and z -axes in $NDC3$, $(1, 0, 0)$, $(0, 1, 0)$ and $(0, 0, 1)$ respectively, onto unit vectors along the VRC axes. If \mathbf{u} , \mathbf{v} and \mathbf{n} are the VRC unit vectors with world-coordinates (u_x, u_y, u_z) , (v_x, v_y, v_z) and (n_x, n_y, n_z) respectively, then it can be shown[7] that the matrix \mathbf{R} is given by

$$\mathbf{R} = \begin{bmatrix} u_x & v_x & n_x & 0 \\ u_y & v_y & n_y & 0 \\ u_z & v_z & n_z & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}$$

The inverse view orientation matrix is $\mathbf{T} \bullet \mathbf{R}$. The view orientation matrix is therefore equivalent to $\mathbf{T}^{-1} \bullet \mathbf{R}^{-1}$. The matrix \mathbf{R} has only a rotational effect and thus \mathbf{R}^{-1} is the transpose of \mathbf{R} , while \mathbf{T} describes a translation and therefore its inverse is a translation in the opposite direction. Therefore the view orientation matrix is the following:

$$\begin{bmatrix} u_x & u_y & u_z & 0 \\ v_x & v_y & v_z & 0 \\ n_x & n_y & n_z & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} \bullet \begin{bmatrix} 1 & 0 & 0 & -VPR_x \\ 0 & 1 & 0 & -VPR_y \\ 0 & 0 & 1 & -VPR_z \\ 0 & 0 & 0 & 1 \end{bmatrix}$$

3.3.2 Model for view mapping

The view mapping transformation has two functions. It first creates the view required and then maps the contents of the *view volume* to the *projection viewport* which is a cuboid defined in *NPC* space.

We will first look at how a view is created. The type of transformation used to create a view is similar to the projections described in the three-dimensional transformation lecture whereby 3D objects are mapped onto a view plane. However, the view mapping transformation must retain the third dimension so that hidden-line and hidden-surface removal can be performed further down the pipeline. The utility function EVALUATE VIEW MAPPING MATRIX computes the view mapping matrix from a set of parameters which describe a view volume in *VRC* space and a projection viewport in *NPC* space. The parameters which define the view volume are as follows:

- View plane distance (VPD) - the position of the view plane along the *VRC* *n*-axis. The view plane is always parallel to the *uv*-plane in *VRC*.
- View window limits (U_{min} , U_{max} , V_{min} , V_{max}) - define a rectangle on the view plane which is aligned with the *VRC* *u* and *v*-axes. Projectors passing through the four corners of the view window delimit the view volume in *x* and *y*.
- Projection reference point (PRP) - the centre of projection. If the projection type is PARALLEL, the vector joining PRP to the centre of the view window defines the direction of the projectors. In this case the view volume is a parallelepiped. If the projection type is PERSPECTIVE, all the projectors will pass through PRP, and so the view volume is a truncated pyramid.
- Projection type - either PARALLEL or PERSPECTIVE.
- Front plane distance and back plane distance (FPD and BPD) - describe two planes parallel to the view plane which delimit the view volume in *z*. The front plane is defined as the plane nearest positive infinity.

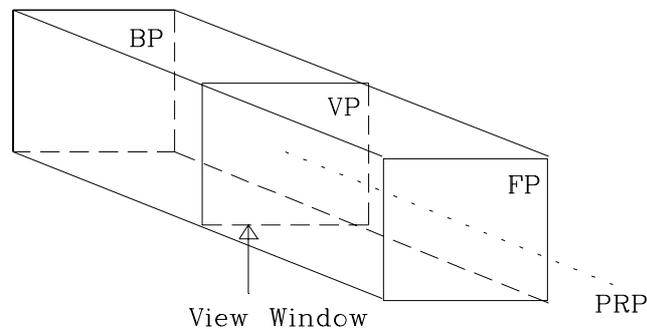
The view volumes for both parallel and perspective projection types are shown in Figure 36.

The view is created by performing a parallel or perspective projection-like transformation, according to the projection type specified. We do not want to perform a true projection since this would lose the relative z -coordinate information which tells us which point is in front of another for subsequent hidden-line and hidden-surface removal. Hence we maintain the relative z information and only apply the projection transformation to the x - and y -coordinates.

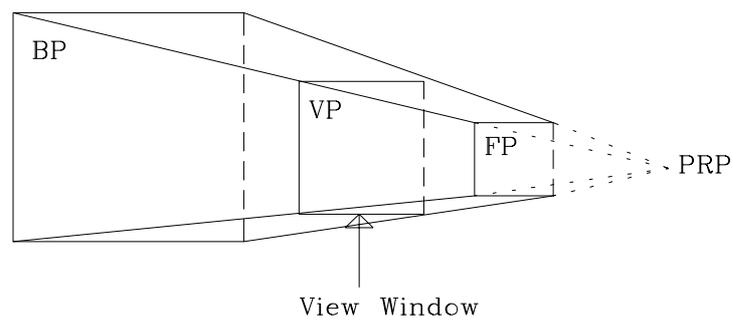
Having created the view using a projection transformation which retains the z information, the *viewed* contents of the view volume must be mapped into the projection viewport. The effect of the projection transformation described above is to transform the view volume to a cuboid. This means that the view volume to projection viewport mapping is simply a window to viewport transformation which behaves in the same way as the normalization transformation in GKS-3D.

Singleton[4] describes the derivation of the view mapping matrix in detail.

View volume for parallel projection types



View volume for perspective projection types



PRP – projection reference point
 FP – front plane
 VP – view plane
 BP – back plane

} The area in which the projectors intersect each infinite plane is illustrated.

Figure 36: view volumes

3.4 Using the Viewing Model

We have examined the viewing parameters used by the two utility routines `EVALUATE VIEW ORIENTATION MATRIX` and `EVALUATE VIEW MAPPING MATRIX` to compute the viewing matrices according to the viewing model defined in GKS-3D and PHIGS. This section looks at how the parameters might be used in practise by deriving the parameter values required to simulate the standard set of planar geometric projections as shown in Figure 37.

It is assumed throughout that the principal faces of the object under consideration are aligned with the principal *NDC3* axes. Those parameters not mentioned do not affect the resulting projection type.

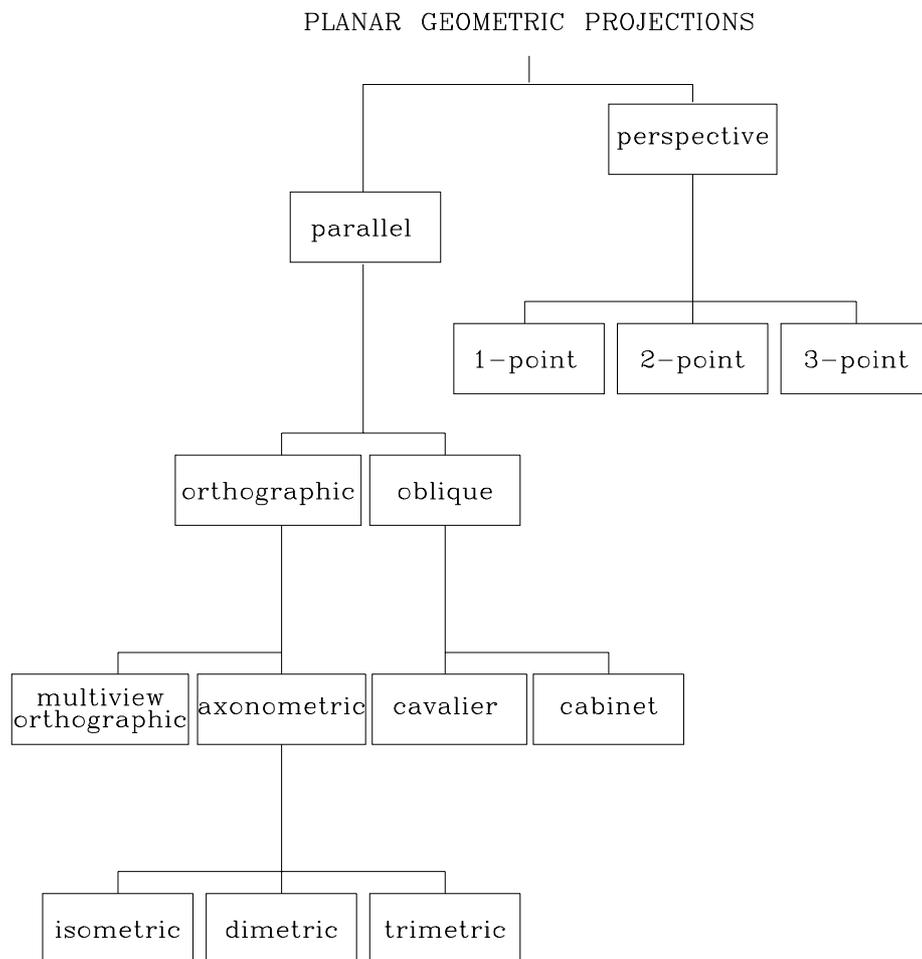


Figure 37: classification of the planar geometric projections

3.4.1 Parallel projections

The following projections require the Projection Type parameter to be set as `PARALLEL`.

erence point (**VRP**). Figure 38 shows how they could be chosen in the most simple case with **VRP** at the bottom left hand corner of the cuboid face representing the current view plane.

- Axonometric Projections.

In an axonometric projection, the object is positioned so that none of its principle faces is parallel to the view plane. These projections are classified according to the orientation of the view plane which is determined by the view transformation.

- Isometric Projections.

If the angles between the view plane and the *NDC3* coordinate axes are all equal the projection is isometric. The view plane normal must therefore be parallel to one of the four lines $\pm x = \pm y = \pm z$ in *NDC3*. This is true if the absolute values of the components of the view plane normal are all equal. In an isometric projection the coordinate axes are equally foreshortened and the angles between the projected axes are all equal.

- Dimetric Projections.

If exactly two angles between the view plane and the coordinate axes are equal the projection is dimetric. The view plane normal must therefore be parallel to one of the planes $x = \pm y$, $x = \pm z$ or $y = \pm z$ in *NDC3*. In this case two of the absolute values of the components of the view plane normal will be equal. In a dimetric projection two of the coordinate axes are equally foreshortened and two of the angles between the projected axes are equal.

- Trimetric Projections.

If all the angles between the view plane and coordinate axes are different the projection is trimetric. This produces different foreshortening of the three coordinate axes and different angles between the projected axes.

Carlbon and Paciorek[3] show that the view plane normal can be calculated given the required properties of the axonometric projection in terms of either the foreshortening ratios of the three axes or the angles between the three projected coordinate axes.

View reference point	-0.7,	-0.2,	0
View plane normal	0.86,	0,	0.86
View up vector	0,	1,	0
Projection type	PARALLEL		
Projection reference point	0.5,	0.5,	1
View window limits	0,	1	
	0,	1	
View plane distance	0		
Front plane distance:	0.9		
Back plane distance:	0		
Projection viewport limits:	0,	1	
	0,	1	
	0,	1	

Figure 39: a dimetric axonometric view

For example, the foreshortening ratio of an axis is $\cos\theta$ where θ is the angle of intersection between the axis and the view plane. The angles between the axes and view plane can therefore be obtained from the required foreshortening ratios of the principal axes. Given the angle between an axis and the view plane, the angle between the axis and the view plane normal can be calculated. The cosine of the angle between a principal axis and the view plane normal is the corresponding view plane normal direction cosine. If a dimetric axonometric projection is required with foreshortening ratios of 0.5, 0.5 and 1 along the x -, y - and z -axes respectively, the corresponding direction cosines of the view plane normal will be $\sin(\cos^{-1} 0.5)$, $\sin(\cos^{-1} 0.5)$ and $\sin(\cos^{-1} 1.0)$. This projection is illustrated in Figure 39.

3.4.1.2 Oblique projections

The view plane in an oblique projection makes an oblique angle with the projectors. In general, the view plane normal is set perpendicular to a face of the object to project the given face without distortion. An oblique projection is determined by the angle between the projectors and the view plane and the orientation of the projectors with respect to the view plane normal. Carlbom and Paciorek [8] derive formulae for the projector direction, **PD**, given one of the following sets of parameters:

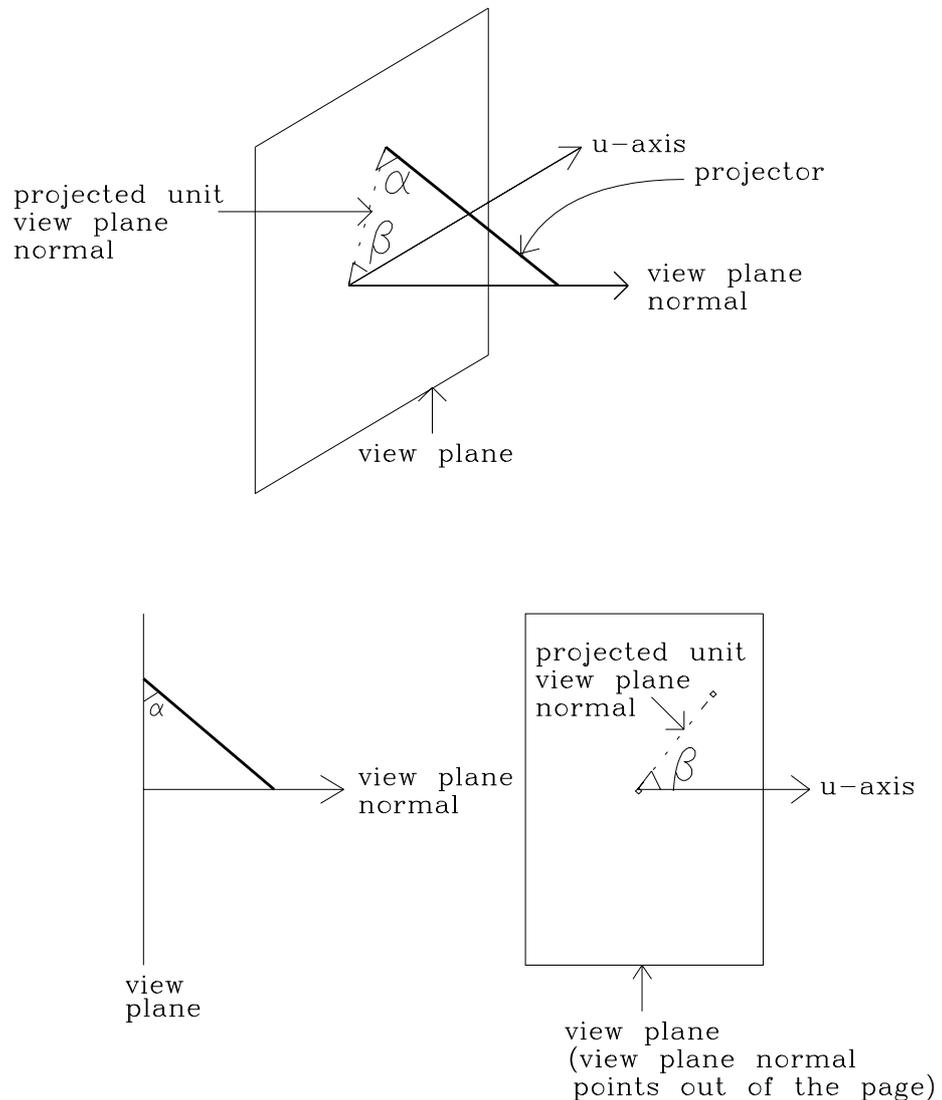


Figure 40: parameters α and β used to define an oblique view

- α : the angle of intersection of the projector with the view plane and β , the angle of rotation of the projector about the view plane normal relative to the x -axis, as shown in Figure 40. Values for the projector direction are

$$PD_x = a \cdot \cos\beta \cdot \cos\alpha$$

$$PD_y = a \cdot \sin\beta \cdot \cos\alpha$$

$$PD_z = -a \cdot \sin\alpha$$

where a is an arbitrary non-zero constant.

- γ : the foreshortening ratio of the projected view plane normal and δ , the angle at which this projected normal intersects a coordinate axis in the view plane. These are illustrated in Figure 41. In this case

$$PD_x = a \cdot \gamma \cdot \cos \delta$$

$$PD_y = a \cdot \gamma \cdot \sin \delta$$

$$PD_z = -a$$

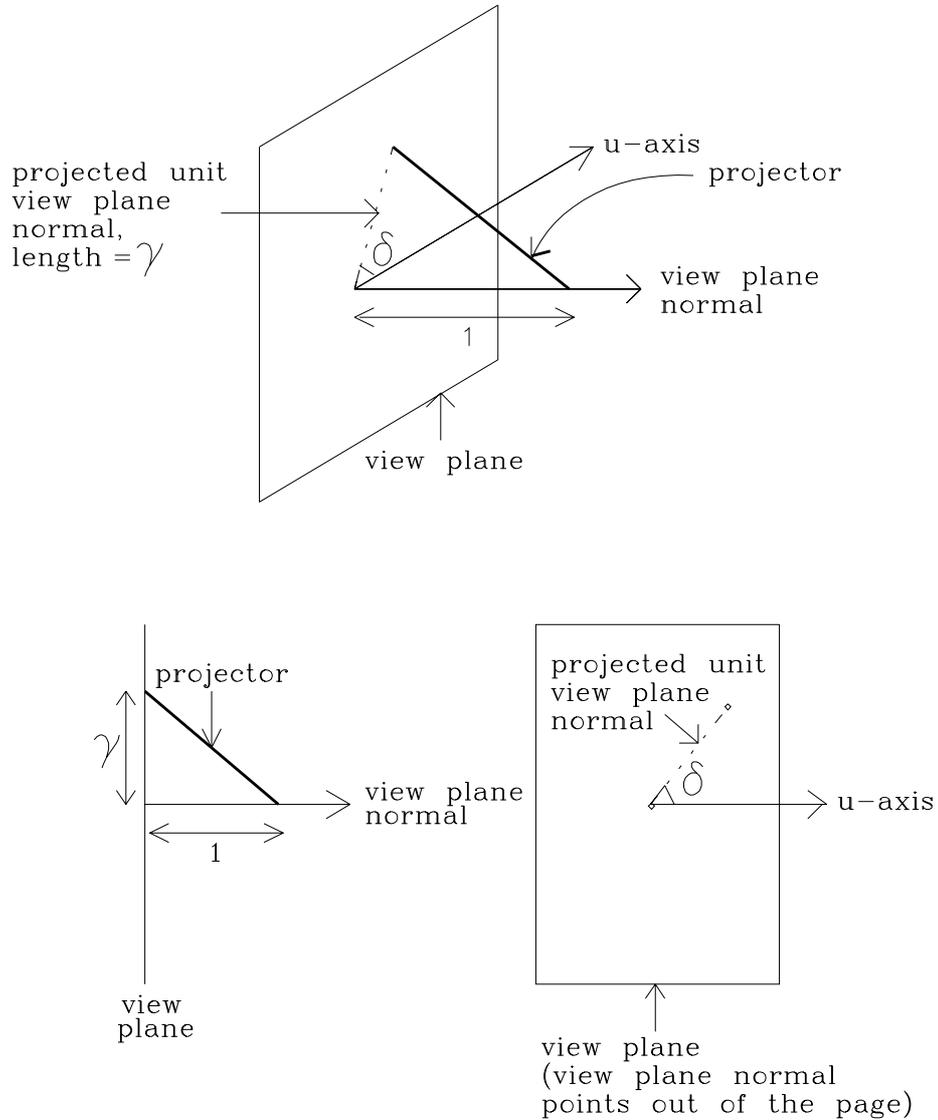


Figure 41: parameters γ and δ used to define an oblique view

The projector direction is determined by the vector joining the projection reference point, **PRP**, and the centre of the view window **C**. Hence

$$\mathbf{PD} = \mathbf{PRP} - \mathbf{C}$$

The coordinates of the projection reference point, given the view window limits and the view plane distance *VPD*, are therefore

$$PRP_x = PD_x + (U_{min} + U_{max})/2$$

$$PRP_y = PD_y + (V_{min} + V_{max})/2$$

$$PRP_z = PD_z + VPD$$

One type of oblique parallel projection is the cavalier projection where the angle between the projectors and the projection plane is 45° . For a cavalier projection, where the angle of rotation of the projectors about the view plane normal is 30° , the projector direction vector would be

$$PD_x = a \cdot \cos 30^\circ \cdot \cos 45^\circ$$

$$PD_y = a \cdot \sin 30^\circ \cdot \cos 45^\circ$$

$$PD_z = -a \cdot \sin 45^\circ$$

A projection of this type can be obtained by assigning values to a , **PRP**, VPD and the view window limits which satisfy the equations given below.

$$PRP_x = a \cdot \cos 30^\circ \cdot \cos 45^\circ + (U_{min} + U_{max})/2$$

$$PRP_y = a \cdot \sin 30^\circ \cdot \cos 45^\circ + (V_{min} + V_{max})/2$$

$$PRP_z = VPD - a \cdot \sin 45^\circ$$

A cavalier projection is illustrated in Figure 42.

View reference point	-0.3,	-0.4,	0
View plane normal	0,	0,	1
View up vector	0,	1,	0
Projection type	PARALLEL		
Projection reference point	1.1,	0.9,	0.7
View window limits	0,	1	
	0,	1	
View plane distance	0		
Front plane distance:	0.9		
Back plane distance:	0		
Projection viewport limits:	0,	1	
	0,	1	
	0,	1	

Figure 42: a cavalier oblique view

3.4.2 Perspective projections

The following projections described require the Projection Type parameter to be set as PERSPECTIVE. In a perspective projection only lines parallel to the view plane remain parallel. Parallel lines that are not parallel to the view plane converge to a single point, called a vanishing point. A *principal vanishing point* is the vanishing point of a principal axis. Perspective projections are classified according to the number of principal vanishing points. This is equivalent to the number of principal *NDC3* coordinate axes that intersect, but do not lie within, the view plane.

- One-Point Perspective Projections.

A one-point perspective projection has one principal vanishing point and

hence only one principal coordinate axis intersects the view plane. The view plane normal vector, **VPN**, must therefore be parallel to the intersecting axis. This vector can be represented by (a, b, c) where exactly one of a, b and c is non-zero. For example, the value of **VPN** when the view plane is required to intersect the z -axis is $(0, 0, c)$ for any non-zero value of c .

- **Two-Point Perspective Projections.**

A two-point perspective projection has two principal vanishing points, and thus the view plane must intersect two of the principal coordinate axes. This is achieved by placing the view plane parallel to one axis but not parallel to any coordinate plane. Hence, **VPN** is perpendicular to exactly one principal axis and can be represented by (a, b, c) where exactly one of a, b and c is zero. The value of **VPN** when the view plane is placed parallel to the x -axis is $(0, b, c)$ for any non-zero values of b and c . A two-point perspective view is illustrated in Figure 43.

- **Three-Point Perspective Projections.**

A three-point perspective projection has three principal vanishing points. This means that the view plane must intersect all three of the principal coordinate axes and thus cannot be parallel to any axis. **VPN** is chosen accordingly as the vector (a, b, c) where a, b and c all have non-zero values.

View reference point	-0.2,	-0.2,	0
View plane normal	-0.4,	0,	1
View up vector	0,	1,	0
Projection type	PERSPECTIVE		
Projection reference point	0.5,	0.5,	1
View window limits	0,	1	
	0,	1	
View plane distance	0.3		
Front plane distance:	0.9		
Back plane distance:	0		
Projection viewport limits:	0,	1	
	0,	1	
	0,	1	

Figure 43: a two point perspective view

A References

1. Hearn D. and Baker M.P., *Computer Graphics*, Prentice-Hall 1986.
2. Blinn J.F. and Newell M.E., *Clipping Using Homogeneous Coordinates*, Proceedings of SIG-GRAPH '78, pp. 245-251.
3. Carlbom I.B. and Paciorek J., *Planar Geometric Projections and Viewing Transformations*, Computing Surveys, Vol. 10 (December 1978), pp. 465–502.
4. Singleton K.M., *An Implementation of the GKS-3D/PHIGS Viewing Pipeline*, Proceedings of Eurographics 1986, North-Holland, Amsterdam.
5. Information Processing Systems - Computer Graphics - Graphical Kernel System for Three Dimensions (GKS-3D) Functional Description ISO DIS 8805.
6. Information Processing Systems - Computer Graphics - Programmer's Hierarchical Interactive Graphics System (PHIGS) ISO DIS 9592.
7. Foley J.D., van Dam A., Feiner S.K. and Hughes J.F., *Computer Graphics: Principles and Practice* (2nd Ed.), Addison Wesley 1990.



These materials have been produced as part of the Information Technology Training Initiative, funded by the Information Systems Committee of the Higher Education Funding Councils.

For further information on this and other modules please contact:

The ITTI Gravigs Project, Computer Graphics Unit, Manchester Computing Centre.
Tel: 0161 275 6095 Email: gravigs@mcc.ac.uk

To order additional copies of this or other modules please contact:

Mrs Jean Burgan, UCoSDA.
Tel: 0114 272 5248 Email: j.burgan@sheffield.ac.uk

ISBN 1 85889 059 4