

Projections and Hardware Rendering

Brian Curless
CSE 557
Fall 2013

Reading

Required:

- ♦ Shirley, Ch. 7, Sec. 8.2, Ch. 18


Further reading:

- ♦ Foley, et al, Chapter 5.6 and Chapter 6
- ♦ David F. Rogers and J. Alan Adams, *Mathematical Elements for Computer Graphics*, 2nd Ed., McGraw-Hill, New York, 1990, Chapter 2.
- ♦ I. E. Sutherland, R. F. Sproull, and R. A. Schumacker, A characterization of ten hidden surface algorithms, *ACM Computing Surveys* 6(1): 1-55, March 1974.

Going back to the pinhole camera...

Recall that the Trace project used, by default, the pinhole camera model.

If we just consider finding out which surface point is visible at each image pixel, then we are just **ray casting**.




For each pixel center P_{ij}


- ♦ Send ray from eye point (COP), C , through P_{ij} into scene.
- ♦ Intersect ray with each object.
- ♦ Select nearest intersection.

Warping space


A very different approach is to take the imaging setup:



then warp all of space so that all the rays are parallel (and distant objects are smaller than closer objects):




and then just draw everything onto the image plane, keeping track of what is in front:



3D Geometry Pipeline

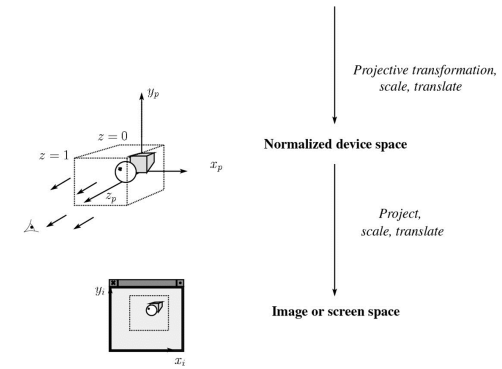
Graphics hardware follows the “warping space” approach.

Before being turned into pixels, a piece of geometry goes through a number of transformations...



5

3D Geometry Pipeline (cont'd)




6

Projections

Projections transform points in n -space to m -space, where $m < n$.

In 3-D, we map points from 3-space to the **projection plane** (PP) (a.k.a., image plane) along **projectors** (a.k.a., viewing rays) emanating from the center of projection (COP):



There are two basic types of projections:

- ◆ Perspective – distance from COP to PP finite
- ◆ Parallel – distance from COP to PP infinite

7

Parallel projections

For parallel projections, we specify a **direction of projection** (DOP) instead of a COP.

There are two types of parallel projections:

- ◆ **Orthographic projection** – DOP perpendicular to PP
- ◆ **Oblique projection** – DOP not perpendicular to PP

We can write orthographic projection onto the $z=0$ plane with a simple matrix.

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

Normally, we do not drop the z value right away. Why not?

8

Z-buffer

The **Z-buffer** or **depth buffer** algorithm [Catmull, 1974] can be used to determine which surface point is visible at each pixel.

Here is pseudocode for the Z-buffer hidden surface algorithm:

```
for each pixel  $(i,j)$  do
  Z-buffer  $[i,j] \leftarrow FAR$ 
  Framebuffer $[i,j] \leftarrow$  <background color>
end for
for each polygon A do
  for each pixel in A do
    Compute depth  $z$  of A at  $(i,j)$ 
    if  $z > Z\text{-buffer}[i,j]$  then
      Z-buffer  $[i,j] \leftarrow z$ 
      Framebuffer $[i,j] \leftarrow$  color of A
    end if
  end for
end for
```


Q: What should FAR be set to?

9

Rasterization

The process of filling in the pixels inside of a polygon is called **rasterization**.

During rasterization, the z value can be computed incrementally (fast!).



Curious fact:

- Described as the “brute-force image space algorithm” by [SSS]
- Mentioned only in Appendix B of [SSS] as a point of comparison for huge memories, but written off as totally impractical.

Today, Z-buffers are commonly implemented in hardware.

10

Properties of parallel projection


Properties of parallel projection:

- Not realistic looking
- Good for exact measurements
- Are actually a kind of affine transformation
 - Parallel lines remain parallel
 - Ratios are preserved
 - Angles not (in general) preserved
- Most often used in CAD, architectural drawings, etc., where taking exact measurement is important

11

Derivation of perspective projection

Consider the projection of a point onto the projection plane:



By similar triangles, we can compute how much the x and y coordinates are scaled:

12

Homogeneous coordinates revisited

Remember how we said that affine transformations work with the last coordinate always set to one.

What happens if the coordinate is not one?

We divide all the coordinates by w :

$$\begin{bmatrix} x \\ y \\ z \\ w \end{bmatrix} \rightarrow \begin{bmatrix} x/w \\ y/w \\ z/w \\ 1 \end{bmatrix}$$

If $w = 1$, then nothing changes.

Sometimes we call this division step the "perspective divide."

13

Homogeneous coordinates and perspective projection

Now we can re-write the perspective projection as a matrix equation:

$$\begin{bmatrix} x' \\ y' \\ w' \end{bmatrix} = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & -1/d & 0 \end{bmatrix} \begin{bmatrix} x \\ y \\ z \\ 1 \end{bmatrix} = \begin{bmatrix} x \\ y \\ -z/d \end{bmatrix}$$

After division by w , we get:

$$\begin{bmatrix} x' \\ y' \\ 1 \end{bmatrix} = \begin{bmatrix} -\frac{d}{z}x \\ -\frac{d}{z}y \\ 1 \end{bmatrix}$$


Again, projection implies dropping the z coordinate to give a 2D image, but we usually keep it around a little while longer.

14

Projective normalization

After applying the perspective transformation and dividing by w , we are free to do a simple parallel projection to get the 2D image.


What does this imply about the shape of things after the perspective transformation + divide?



15

Viewing angle

An alternative to specifying the distance from COP to PP is to specify a viewing angle:



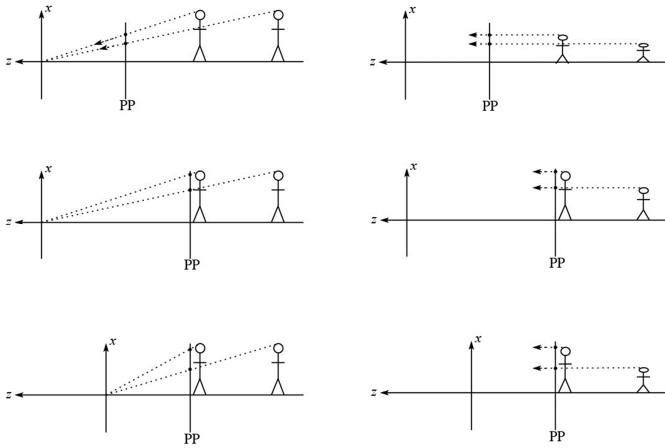
Given the height of the image h and θ , what is d ?

What happens to d as θ increases (while h is constant)?

16

Zoom and dolly


We can make things seem bigger by **zooming** in (increasing d) or **dollying** in (getting closer to the scene).



17

Zoom and dolly

Alternatively, we can keep one plane at the same size by carefully controlling zoom and dolly...




18

Vanishing points

What happens to two parallel lines that are not parallel to the projection plane?

Think of train tracks receding into the horizon...



The equation for a line is:

$$\mathbf{l} = \mathbf{p} + t\mathbf{v} = \begin{bmatrix} p_x \\ p_y \\ p_z \\ 1 \end{bmatrix} + t \begin{bmatrix} v_x \\ v_y \\ v_z \\ 0 \end{bmatrix}$$

After perspective transformation we get:

$$\begin{bmatrix} x' \\ y' \\ w' \end{bmatrix} = \begin{bmatrix} p_x + tv_x \\ p_y + tv_y \\ -(p_z + tv_z)/d \end{bmatrix}$$

19

Vanishing points (cont'd)

Dividing by w :

$$\begin{bmatrix} x' \\ y' \\ w' \end{bmatrix} = \begin{bmatrix} \frac{-p_x + tv_x}{p_z + tv_z} d \\ \frac{-p_y + tv_y}{p_z + tv_z} d \\ 1 \end{bmatrix}$$

Letting t go to infinity:

We get a point!

What happens to the line $\mathbf{l} = \mathbf{q} + t\mathbf{v}$?

Each set of parallel lines intersect at a **vanishing point** on the PP.


Q: How many vanishing points are there?

20

Clipping and the viewing frustum

The center of projection and the portion of the projection plane that map to the final image form an infinite pyramid. The sides of the pyramid are **clipping planes**.

Frequently, additional clipping planes are inserted to restrict the range of depths. These clipping planes are called the **near** and **far** or the **hither** and **yon** clipping planes.



All of the clipping planes bound the the **viewing frustum**.


21

Properties of perspective projections

The perspective projection is an example of a **projective transformation**.

Here are some properties of projective transformations:

- ◆ Lines map to lines
- ◆ Parallel lines do not necessarily remain parallel
- ◆ Ratios are not preserved



One of the advantages of perspective projection is that size varies inversely with distance – looks realistic.


A disadvantage is that we can't judge distances as exactly as we can with parallel projections.

22

Rasterization with color

Recall that the z-buffer works by interpolating z-values across a triangle that has been projected into image space, a process called rasterization.

During rasterization, colors can be smeared across a triangle as well:




23

Gouraud interpolation

Recall from the shading lecture, rendering with per triangle normals leads to faceted. An improvement is to compute per-vertex normals and use graphics hardware to do **Gouraud interpolation**:

1. Compute normals at the vertices.
2. Shade only the vertices.
3. Interpolate the resulting vertex colors.




24

Gouraud interpolation artifacts

Gouraud interpolation has significant limitations.

1. If the polygonal approximation is too coarse, we can miss specular highlights.



2. We will encounter **Mach banding** (derivative discontinuity enhanced by human eye).

This is what graphics hardware does by default.

A substantial improvement is to do...


25

Phong interpolation

To get an even smoother result with fewer artifacts, we can perform **Phong interpolation**.


Here's how it works:

1. Compute normals at the vertices.
2. Interpolate normals and normalize.
3. Shade using the interpolated normals.



26


Default pipeline: Gouraud interpolation



Default vertex processing:

$L \leftarrow$ determine lighting direction
 $V \leftarrow$ determine viewing direction
 $N \leftarrow$ normalize(n_e)
 $c_{\text{blinn-phong}} \leftarrow$ shade with L, V, N, k_d, k_s, n_s
 attach $c_{\text{blinn-phong}}$ to vertex as "varying"
 $v_i \leftarrow$ project v to image

$v_i^1, v_i^2, v_i^3 \rightarrow$ triangle




Default fragment processing:

color $\leftarrow c_{\text{blinn-phong}}^p$

27


Programmable pipeline: Phong-interpolated normals!



Vertex shader:

attach n_e to vertex as "varying"
 attach v_e to vertex as "varying"
 $v_i \leftarrow$ project v to image

$v_i^1, v_i^2, v_i^3 \rightarrow$ triangle



Fragment shader:

$L \leftarrow$ determine lighting direction
 $V \leftarrow$ determine viewing direction
 $N \leftarrow$ normalize(n_e^p)
 color \leftarrow shade with $L, V, N, k_d^p, k_s^p, n_s^p$


28

Texture mapping and the z-buffer

Texture-mapping can also be handled in z-buffer algorithms.

Method:

- ◆ Scan conversion is done in screen space, as usual
- ◆ Each pixel is colored according to the texture
- ◆ Texture coordinates are found by Gouraud-style interpolation



Note: Mapping is more complicated to handle perspective correctly!