

CS380: Introduction to Computer Graphics
 Rasterization
 Chapter 12

Min H. Kim
 KAIST School of Computing


Min H. Kim (KAIST) Foundations of 3D Computer Graphics, S. Gortler, MIT Press, 2012



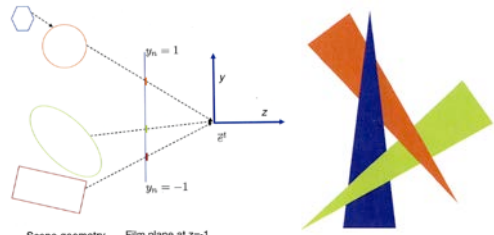
Depth

SUMMARY

Min H. Kim (KAIST) Foundations of 3D Computer Graphics, S. Gortler, MIT Press, 2012




Visibility



Scene geometry Film plane at $z=1$

- We could explicitly store everything hit along a ray and then compute the closest.
 - Make sense in a **ray tracing** setting, where we are working **one pixel per ray at a time**, but not for OpenGL, where we are working **one triangle at a time**.

Min H. Kim (KAIST) Foundations of 3D Computer Graphics, S. Gortler, MIT Press, 2012



Z-buffer

- We will use z-buffer
- Triangle are drawn in any order
- Each pixel in frame buffer stores 'depth' value of closest geometry observed so far.
- When a new triangle tries to set the color of a pixel, we first compare its depth to the value stored in the z-buffer.
- Only if the observed point in this triangle is closer, we overwrite the color and depth values of this pixel.

Min H. Kim (KAIST) Foundations of 3D Computer Graphics, S. Gortler, MIT Press, 2012

Proj. Trans.: Eye coord. → NDC

- Camera projection transformation

$$\begin{bmatrix} x_n w_n \\ y_n w_n \\ z_n w_n \\ w_n \end{bmatrix} = \begin{bmatrix} x_c \\ y_c \\ z_c \\ w_c \end{bmatrix} = \begin{bmatrix} s_x & 0 & -c_x & 0 \\ 0 & s_y & -c_y & 0 \\ 0 & 0 & \frac{f+n}{f-n} & -\frac{2fn}{f-n} \\ 0 & 0 & -1 & 0 \end{bmatrix} \begin{bmatrix} x_e \\ y_e \\ z_e \\ 1 \end{bmatrix}$$

Min H. Kim (KAIST) Foundations of 3D Computer Graphics, S. Gortler, MIT Press, 2012

Chapter 12

RASTERIZATION

Min H. Kim (KAIST) Foundations of 3D Computer Graphics, S. Gortler, MIT Press, 2012

Path from vertex to pixel

- Three vertices have passed through the vertex shader
- Follow it journey to be a bunch of pixels

Min H. Kim (KAIST) Foundations of 3D Computer Graphics, S. Gortler, MIT Press, 2012

Varying variables

- Varying variables provides an interface between the vertex and fragment shader.
- When the primitives are assembled and fragments computed, for each fragment there is a set of variables that are interpolated automatically and provided to the fragment shader.
- An example is the color of the fragment. The color that arrives at the fragment shader is the result of the interpolation of the colors of the vertices that make up the primitive.
 - varying float intensity;

Min H. Kim (KAIST) Foundations of 3D Computer Graphics, S. Gortler, MIT Press, 2012

Clipping

- What if there is a vertex behind you?

Min H. Kim (KAIST) Foundations of 3D Computer Graphics, S. Gortler, MIT Press, 2012 9

Clipping

- Not only front vertex but also back vertex will be projected
- We will be drawing in completely the wrong area.
- Beware interpolation nature of OpenGL: Projection of vertices followed by interpolation

Min H. Kim (KAIST) Foundations of 3D Computer Graphics, S. Gortler, MIT Press, 2012 10

Clipping

- Processing for triangles that are fully or partially out of view
- We don't want to see behind us
- We want to minimize processing
- The tricky part will be to deal with eye-spanning triangles.

Min H. Kim (KAIST) Foundations of 3D Computer Graphics, S. Gortler, MIT Press, 2012 11

Eye spanners

- Back vertex projects higher up in the image
- Filling in the in-between pixels will fill in the wrong region.
- Solution: slice up the geometry by the six faces of the view frustum

Min H. Kim (KAIST) Foundations of 3D Computer Graphics, S. Gortler, MIT Press, 2012 12

Clipping coordinates

- Eye coordinates (projected) → clip coordinates → normalized device coordinates (NDCs)
- (reminder) Dividing clip coordinates (x_c, y_c, z_c, w_c) by the w_c ($w_c = w_n$) component (the fourth component in the homogeneous coordinates) yields normalized device coordinates (NDCs).

$$\begin{bmatrix} x_n w_n \\ y_n w_n \\ z_n w_n \\ w_n \end{bmatrix} = \begin{bmatrix} x_c \\ y_c \\ z_c \\ w_c \end{bmatrix} = \begin{bmatrix} s_x & 0 & -c_x & 0 \\ 0 & s_y & -c_y & 0 \\ 0 & 0 & \frac{f+n}{f-n} & -\frac{2fn}{f-n} \\ 0 & 0 & -1 & 0 \end{bmatrix} \begin{bmatrix} x_e \\ y_e \\ z_e \\ 1 \end{bmatrix}$$

Min H. Kim (KAIST) Foundations of 3D Computer Graphics, S. Gortler, MIT Press, 2012 13

Clipping coordinates

- Suppose that x_c / w_c is the same as $-x_c / -w_c$. What if w_c is zero?
- If you wait for normalized device coordinates (NDCs) where the vertex has flipped, and it's too late to do the clipping.
- We could do in the eye space, but then would need to use the camera parameters
- The solution is to use clip coordinates: post-matrix-multiply but pre-divide.
- No divide = no flipping

Min H. Kim (KAIST) Foundations of 3D Computer Graphics, S. Gortler, MIT Press, 2012 14

Clipping coordinates

- Recall that we want points in the range:

$$\begin{aligned} -1 < x_n < 1 \\ -1 < y_n < 1 \\ -1 < z_n < 1 \end{aligned}$$
- In clip coordinates this is:

$$\begin{bmatrix} x_n w_n \\ y_n w_n \\ z_n w_n \\ w_n \end{bmatrix} = \begin{bmatrix} x_c \\ y_c \\ z_c \\ w_c \end{bmatrix} \quad \begin{aligned} -w_c < x_c < w_c \\ -w_c < y_c < w_c \\ -w_c < z_c < w_c \end{aligned}$$

Min H. Kim (KAIST) Foundations of 3D Computer Graphics, S. Gortler, MIT Press, 2012 15

Clipping coordinates

- Primitives totally inside the clipping volume are not altered. Primitives outside the viewing volume are discarded. Primitives whose edges intersect the boundaries of the clipping volume are clipped (learn later).

Min H. Kim (KAIST) Foundations of 3D Computer Graphics, S. Gortler, MIT Press, 2012 16

Clipping coordinates → NDCs

- Clipping is done, we can now divide by w_c (typically z_c) to obtain normalized device coordinates.

$$\begin{bmatrix} x_n \\ y_n \\ z_n \\ 1 \end{bmatrix} = \begin{bmatrix} x_c / w_c \\ y_c / w_c \\ z_c / w_c \\ w_c / w_c \end{bmatrix}$$


Min H. Kim (KAIST)

Foundations of 3D Computer Graphics, S. Gortler, MIT Press, 2012

17

Backface culling (killing?)

- When drawing a closed solid object, we will only ever see one 'front' side of each triangle.
- For efficiency we can drop these from the processing.




Min H. Kim (KAIST)

Foundations of 3D Computer Graphics, S. Gortler, MIT Press, 2012

18

Without backface culling




Min H. Kim (KAIST)

Foundations of 3D Computer Graphics, S. Gortler, MIT Press, 2012

19

With backface culling



Min H. Kim (KAIST)

Foundations of 3D Computer Graphics, S. Gortler, MIT Press, 2012

20

Backface culling

- To do this, in OpenGL, we use the convention of ordering the three vertices in the draw call (IBO/VBO) so that they are counterclockwise (CCW) when looking at its front side.
- During setup, we call `glEnable(GL_CULL_FACE)`
- To implement culling, OpenGL does the following...


Min H. Kim (KAIST)

Foundations of 3D Computer Graphics, S. Gortler, MIT Press, 2012

21

Math of Backface Culling

- Let $\tilde{p}_1, \tilde{p}_2,$ and \tilde{p}_3 be the three vertices of the triangle projected down to the $(x_n, y_n, 0)$ plane.
- Define the vector $\vec{a} = \tilde{p}_3 - \tilde{p}_2$ and $\vec{b} = \tilde{p}_1 - \tilde{p}_2$
- Next compute the cross product $\vec{c} = \vec{a} \times \vec{b}$
- If the three vertices are counterclockwise in the plane, then \vec{c} will be in the $+z_n$ direction.




Min H. Kim (KAIST)

Foundations of 3D Computer Graphics, S. Gortler, MIT Press, 2012

22

Math of Backface Culling

- So the area of the triangle is: $Area = \frac{1}{2} \|\vec{a} \times \vec{b}\|$
 - Taking account of the direction, we could calculate the direction of the cross product of these two vectors.
- $$(x_n^3 - x_n^2)(y_n^1 - y_n^2) - (y_n^3 - y_n^2)(x_n^1 - x_n^2)$$
- If this is negative, the polygon looks backward.



NB There are typos at page 114 in our textbook. The negative and positive were flipped around in the textbook. These slides are typo-corrected.

Min H. Kim (KAIST)

Foundations of 3D Computer Graphics, S. Gortler, MIT Press, 2012

23

Viewport

- Now we want to position the vertices in the window. So it is time to move the NDCs to window coordinates.
 - in NDCs, lower left corner is $[-1, -1]^t$ and upper right corner is $[1, 1]^t$.
- Each pixel center has an integer coordinate.
 - This will make subsequent pixel computations more natural.
- We want the lower left pixel center to have 2D window coordinates of $[0, 0]^t$ and the upper right pixel center to have coordinates $[W - 1, H - 1]^t$

Min H. Kim (KAIST)

Foundations of 3D Computer Graphics, S. Gortler, MIT Press, 2012

24

Viewport

- We think of each pixel as owning the real estate which extends 0.5 pixel units in the positive and negative, horizontal and vertical directions from the pixel center.

Min H. Kim (KAIST) 25

Viewport

- Thus the extent of 2D window rectangle covered by the union of all our pixels is the rectangle in window coordinates with lower left corner $[-0.5, -0.5]^t$ and upper right corner $[W - 0.5, H - 0.5]^t$

Min H. Kim (KAIST) 26

Viewport matrix

- We need a transform that maps the lower left corner to $[-0.5, -0.5]^t$ and upper right corner to $[W - 0.5, H - 0.5]^t$
- The appropriate scale and shift can be done using the viewport matrix:

$$\begin{bmatrix} x_w \\ y_w \\ z_w \\ 1 \end{bmatrix} = \begin{bmatrix} W/2 & 0 & 0 & (W-1)/2 \\ 0 & H/2 & 0 & (H-1)/2 \\ 0 & 0 & 1/2 & 1/2 \\ 0 & 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x_n \\ y_n \\ z_n \\ 1 \end{bmatrix}$$

Min H. Kim (KAIST) Foundations of 3D Computer Graphics, S. Gortler, MIT Press, 2012 27

Viewport matrix

- This does a scale and shift in both x and y
- You can verify that it maps the corners appropriately
- In OpenGL, we set up this viewport matrix with the call `glViewport(0,0,W,H)`
- The third row of this matrix is used to map the $[-1, -1]$ range of z_n values to the more convenient $[0...1]$ range.
- So now (in our conventions), $z_w = 0$ is far and $z_w = 1$ is near.

Min H. Kim (KAIST) Foundations of 3D Computer Graphics, S. Gortler, MIT Press, 2012 28

Viewport matrix

- So we must also tell OpenGL that when we clear the z-buffer, we should set it to 0; we do this with the call `glClearDepth(0.0)`

Min H. Kim (KAIST)

Foundations of 3D Computer Graphics, S. Gortler, MIT Press, 2012

29

Texture Viewport

- The abstract domain for textures is not the canonical square, but instead is the unit square.
- Its lower left corner is $[0,0]^t$ and upper right corner is $[1,1]^t$.
- In this case the coordinate transformation matrix is:

$$\begin{bmatrix} x_w \\ y_w \\ - \\ 1 \end{bmatrix} = \begin{bmatrix} W & 0 & 0 & -1/2 \\ 0 & H & 0 & -1/2 \\ - & - & - & - \\ 0 & 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x_t \\ y_t \\ - \\ 1 \end{bmatrix}$$


Min H. Kim (KAIST)

Foundations of 3D Computer Graphics, S. Gortler, MIT Press, 2012

30

Rasterization

- Starting from the window coordinates for the three vertices, the rasterizer needs to figure out which pixel centers are inside of the triangle.
- Each triangle on the screen can be defined as the intersection of three half-spaces.




Min H. Kim (KAIST)

Foundations of 3D Computer Graphics, S. Gortler, MIT Press, 2012

31

Rasterization

- Starting from the window coordinates for the three vertices, the rasterizer needs to figure out which pixel centers are inside of the triangle.
- Each triangle on the screen can be defined as the intersection of three half-spaces.



Min H. Kim (KAIST)

Foundations of 3D Computer Graphics, S. Gortler, MIT Press, 2012


32

Math of rasterization

- Each such halfspace is defined by a line that coincides with one of the edges of the triangle, and can be tested using an 'edge function' of the form:

$$edge = ax_w + by_w + c$$

where the (a,b,c) are constants that depend on the geometry of the edge.




Min H. Kim (KAIST)

Foundations of 3D Compute.

Math of rasterization

- A positive value of this function at a pixel with coordinates $[x_w, y_w]^t$ means that the pixel is inside the specified halfspace.
- If all three tests pass, then the pixel is inside the triangle.




Min H. Kim (KAIST)

Foundations of 3D Compute.

34

Speed up

- Only look at pixels in the bounding box of the triangle
- Test if a pixel block is entirely outside of triangle
- Use incremental calculations along a scanline.




Min H. Kim (KAIST)

Foundations of 3D Compute.

35

Interpolation

- As input to rasterization, each vertex also has some auxiliary data associated with it.
 - This data includes a z_w value,
 - As well as other data that is related, but not identical to the varying variables.
- It is also the job of the rasterizer to linearly interpolate this data over the triangle.



Min H. Kim (KAIST)

Foundations of 3D Computer Grap

Math of interpolation



- Each such value v to be linearly interpolated can be represented as an affine function over screen space with the form:

$$v = ax_w + by_w + c$$

- An affine function can be easily evaluated at each pixel by the rasterizer.
- Indeed, this is no different from evaluating the edge test functions just described.

Boundaries



- For pixel on edge or vertex it should be rendered exactly once.
- Need special care in the implementation to avoid duty edge representation.