How to rasterize a triangle!
Edge functions and interpolation

Michael Doggett

Department of Computer Science

Lund University

Thursday, 5 November 2009

Misc stuff

- Assignment 1 due Wednesday next week
  - Sign up for next Wednesday's Pluto lab
  - Link on web page
  - Meet during break to find partners
- Assignment 2 available on the web
  - Talk about it more next tuesday
- Do you check out the Online Discussion forum? You should...
  - Assignment questions and answers
- The following lectures will be about graphics hardware. There are notes (~100 pages) available for free on the course website.
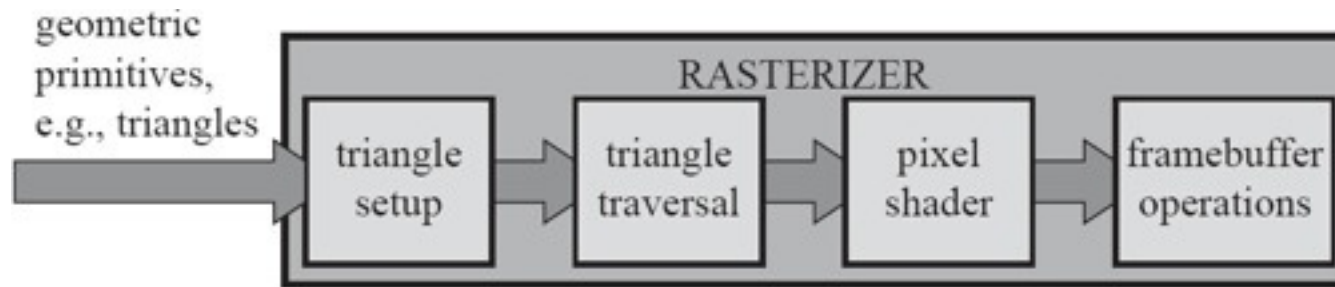
Thursday, 5 November 2009

Overview of today's lecture

- How graphics hardware can draw a triangle:
  - Edge functions
  - Triangle traversal
  - Interpolation

- Why do I need to know about this?
  - Before you can learn how to walk, you need to know...
  - If you want to design your own hardware
  - If you design new hardware algorithms
  - If you work at company XXX and need to buy graphics hardware for that company's mobile phones, for example
  - If you need to write your own rasterizer ...

Thursday, 5 November 2009

The pipeline is long, why rasterization?

- For any computer system, memory accesses are the operation that uses most energy!
- For mobile devices: consume as little energy as possible!
- Fact: pipeline is roughly two parts:
    - GEOMETRY (per-vertex)
    - RASTERIZER (per-pixel)
- Fact: RASTERIZER accesses memory much more than GEOMETRY
- **Conclusion**: focus on making RASTERIZER as efficient as possible

4

Thursday, 5 November 2009

Overview of rasterizer

geometric primitives, e.g., triangles

RASTERIZER

triangle setup → triangle traversal → pixel shader → framebuffer operations

- Triangle setup
  - Put everything we can factor out from the rest of the pipeline, i.e., all "per-triangle" computations

- Triangle traversal
  - Find pixels (or samples) that are inside the triangle

- Pixel shader: apply per-pixel computations, texture accesses, etc

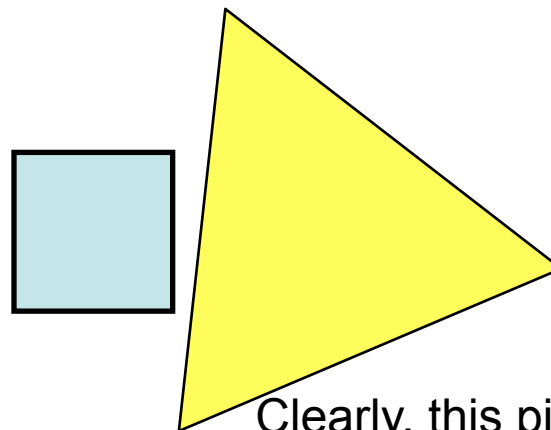- Frame buffer operations: depth buffer testing, stencil, write to buffers, e.g.

Let's study triangle traversal

- Critical operation in rasterizer
  - without it, we have nothing
  - Once we have it, we will study algorithms to dramatically reduce memory accesses
- When does a pixel belong to a triangle?

Clearly, this pixel belongs to the triangle

Clearly, this pixel does **NOT** belong to the triangle

Thursday, 5 November 2009

When does a pixel belong to a triangle?

How about this case? And this? And this? And this?

- It all depends on where you *sample*!
- For (low-quality) normal rasterizaion, you sample in the center of the triangle:

If sample point is inside triangle, then pixel is inside (belongs to) the triangle.

In a later lecture, we will see how quality can be improved by using more than one sample per pixel

Thursday, 5 November 2009

Which coordinate system? Where are we?

$$\mathbf{M}\mathbf{v} = \mathbf{h} = \begin{pmatrix} h_x \\ h_y \\ h_z \\ h_w \end{pmatrix} \implies \begin{pmatrix} h_x/h_w \\ h_y/h_w \\ h_z/h_w \\ h_w/h_w \end{pmatrix} = \begin{pmatrix} h_x/h_w \\ h_y/h_w \\ h_z/h_w \\ 1 \end{pmatrix} = \mathbf{p}$$

$\mathbf{M}$ is a projection matrix

- After projection matrix, and homogenization, we have $\mathbf{p} = (p_x, p_y, p_z, 1)^T$

- (p_x, p_y) are screen-space coordinates.

- Those are used for triangle traversal

- We assume that triangle have been clipped to screen-space, that is, (p_x, p_y) are in [0,w] x [0,h] (and w x h is screen resolution)

8

Thursday, 5 November 2009

Screen-space coordinates

- This is according to the specification of OpenGL

9

Thursday, 5 November 2009

What happens if you round off floating point vertices to nearest pixel center?

Triangle edge using floating point coords

Frame 1

Frame 2

Frame 3

With sub-pixel coordinates this will get solved

Edge with "snapped" vertex coordinates

Big jump here... looks really bad.

10

After projection: sub-pixel coords

- Projected points are floating-point

- Due to limited range of the coordinates, we can use fixed point math (integer)

- However, as we've seen, we cannot round off to nearest pixel center

- Instead use *sub-pixel* coordinates

- With 2 subpixel fractional bits per x, and y, we get:

Sub-pixel Sample points

Floating point coordinates snaps to the closest sub-pixel sample

Remember: integer coords at pixel corners!

Thursday, 5 November 2009

Edge functions

- Hardware uses these to find pixels inside a triangle
- For each edge of the triangle, create a line equation (implicit form): $ax + by + c = 0$
- Edge function for two points $\mathbf{p}^0$ and $\mathbf{p}^1$:

$$e(x,y) = -(p_y^1 - p_y^0)(x - p_x^0) + (p_x^1 - p_x^0)(y - p_y^0)$$

$$e(x,y) = ax + by + c = \mathbf{n} \cdot (x,y) + c.$$

- Can be interpreted as the "normal" of the line

- By definition, $e(x,y)=0$, if (x,y) is exactly **on** the (infinite) line through $\mathbf{p}^1$ and $\mathbf{p}^0$
- How about other points?

Thursday, 5 November 2009

Intuition about edge functions

- $e(\mathbf{u})$ must be >0 or <0? Why?
- $e(\mathbf{v})$ must be >0 or <0? Why?

Recall this one:

$$e(x, y) = ax + by + c = \mathbf{n} \cdot (x, y) + c.$$

13

Three edge funcs per triangle

$$e_0(x, y) = -(p_y^2 - p_y^1)(x - p_x^1) + (p_x^2 - p_x^1)(y - p_y^1) = a_0 x + b_0 y + c_0$$
$$e_1(x, y) = -(p_y^0 - p_y^2)(x - p_x^2) + (p_x^0 - p_x^2)(y - p_y^2) = a_1 x + b_1 y + c_1$$
$$e_2(x, y) = -(p_y^1 - p_y^0)(x - p_x^0) + (p_x^1 - p_x^0)(y - p_y^0) = a_2 x + b_2 y + c_2$$

- A sample point (x,y) is inside the triangle if $e_i(x,y) >= 0$ for $i = 0,1,2$

Thursday, 5 November 2009

Not really that simple, though :-(

- ## What happens to pixels exactly on an edge?

Does the pixel belong to A or B, or both or neither of them?

- ## Correct answer: one and only one of A or B
- ## Why?
  - Neither → cracks between triangles
  - Both → unnecessary work is done, and transparency & shadow volumes (e.g.) give incorrect results

Thursday, 5 November 2009

A solution (by McCool et al)

bool INSIDE(e, x, y)

1   **if** $e(x,y) > 0$ **return** true;
2   **if** $e(x,y) < 0$ **return** false;
3   **if** $a > 0$ **return** true;
4   **if** $a < 0$ **return** false;
5   **if** $b > 0$ **return** true;
6   **return** false;

- Another way to think about it:
  - We exclude shadowed edges

light

How about when a vertex coincides with the sampling point?

- You get the same kind of problems!

- One solution: offset the subpixel grid so that sampling points never coincide with sub-pixel grid

Thursday, 5 November 2009

Another solution

- Don't change the sub-pixel grid, allow vertices to coincide with sample point
- So, now, two edge functions will be zero for this pixel: $e_1(x,y)=0$ and $e_2(x,y)=0$
- How would you do it?
- Choose one direction, say southwards:

- The sampling point should only belong to the triangle that has the arrow in it
- Can be determined from looking at the "normals" of the edge functions

- Edges sharing sample points is the most common problem, so solve that first...

© 2009 Tomas Akenine-Möller and Michael **[Idea by John Owens, UC Davis]**

Thursday, 5 November 2009
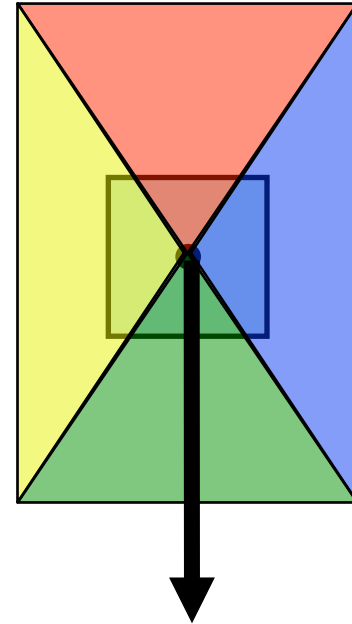
Incremental updates

- Assume we have evaluated an edge function at (x,y), so we already know $e(x,y)$

- Reduce operations when moving from one pixel to the next

$$e(x+1, y) = a(x+1) + by + c = e(x,y) + a$$

- Similar for the other neighbors, and in general:

$$e(\mathbf{s} + \mathbf{t}) = e(\mathbf{s}) + \mathbf{n} \cdot \mathbf{t}$$

- Finally, we have the tool (edge functions) needed for triangle traversal!

Thursday, 5 November 2009

Triangle traversal strategies

- Simple (and stupid): execute `Inside()` for every pixel on screen, and for every edge
- Little better: compute bounding box first
  - Called "bounding box traversal"

Visits all light and dark gray pixels

But only dark gray pixels are inside, and sent down the pipeline for further processing

20

Thursday, 5 November 2009

Backtrack traversal

- Have been used on mobile devices
  - by a Korean research group (KAIST)
- Advantage: only traverse from left to right
  - Could make for more efficient memory accesses
  - Could backtrack at a faster pace (because no mem acc)

Thursday, 5 November 2009

Zigzag traversal

- Simple technique that avoids backtracking
  - Otherwise, very similar
  - Can still visit a bunch of pixels outside
    - see next to most bottom scanline
  - Can be solved, but requires more `Inside()`-testing

Thursday, 5 November 2009

Side by side comparison Backtrack vs zigzag

Backtrack never visits unnecessary pixels to the left

Zigzag never visits unnecessary pixels to the left on even scanlines and to the right on odd scanlines
(and avoids backtracking)

We only cover simple schemes here. There are more sophisticated schemes that visits fewer pixels, but those are more expensive

Thursday, 5 November 2009

Tiled traversal

- General idea: divide screen space into non-overlapping tiles (a tile is wxh pixels)
  - Traverse one tile at a time, and finish visiting pixels in tile before moving to next tile

4x4 tile

8x8 tile size is common in desktop graphics cards

- Better because (all topics will be treated in later lextures):
  - Gives better texture cache performance
  - Enables simple culling (Zmin & Zmax)
  - Real-time buffer compression (color and depth)

Is tiled traversal that different?

- No, not really. We need:
  - $\mathbf{I}$: Traverse to tiles overlapping triangle
  - $\mathbf{II}$: Test if tile overlaps with triangle
  - $\mathbf{III}$: Traverse pixels inside tile
- Previous algorithms can handle $\mathbf{I}$ and $\mathbf{III}$
- $\mathbf{II}$ needs to be handled
  - Easily solved using....edge functions!
  - See next slides...

Thursday, 5 November 2009

Tile/triangle overlap test (1)

- Reuse knowledge from 3D intersection testing—Haines and Wallace's trick:

Only need to test the corners whose diagonal aligns best with the normal of the plane

26

Thursday, 5 November 2009

Tile/triangle overlap test (2)

- A tile is fully **outside** a triangle if:
  - Either the tile is outside the bounding box of the triangle (trivial)
  - Or if tile is fully outside at least one edge func (concentrate on this case)

$\mathbf{n}=(a,b)$

$e(x,y)>0$

$e(x,y)<0$

y

x

We only need to test the black corner of the tile!!!

How can we know?

27

Tile/triangle overlap test (3)

- Assume we have evaluated the edge function for the lower left corner, $\mathbf{s} = (s_x, s_y)$, of the tile, i.e., $e(\mathbf{s})$

- We can evaluate the tile corners with:

$$e(\mathbf{s} + \mathbf{t}) = e(\mathbf{s}) + \mathbf{n} \cdot \mathbf{t}$$

- Where $\mathbf{t} = (0,0)$, or $\mathbf{t} = (w,0)$, or $\mathbf{t} = (0,h)$, or $\mathbf{t} = (w,h)$
  - Since $\mathbf{n} = (a,b)$ is known from the edge function, and w & h in all realistic scenarios are powers of two (e.g., 2^g)...
  - ...this reduces to adding shifted versions of a and b

- How compute $\mathbf{t}$?

Thursday, 5 November 2009

Tile/triangle overlap test (4)

- How compute **t**?
- Try example to the right

$\mathbf{n}=(a,b)$ $e(x,y)>0$

$e(x,y)<0$

- Look at normal direction $\mathbf{n}=(a,b)=(n_x,n_y)$
- In general:

$$t_x = \begin{cases} w, & n_x \geq 0 \\ 0, & n_x < 0 \end{cases} \qquad t_y = \begin{cases} h, & n_y \geq 0 \\ 0, & n_y < 0 \end{cases}$$

- **t** is computed as part of the "triangle setup", since the info is constant all over the triangle

Thursday, 5 November 2009

Interpolation

- Now, we can find pixels inside triangles

- Next, interpolate parameters across triangles

- Why?

- Example 1; interpolation of vertex colors:

- Example 2; interpolation of "images" (textures)

image
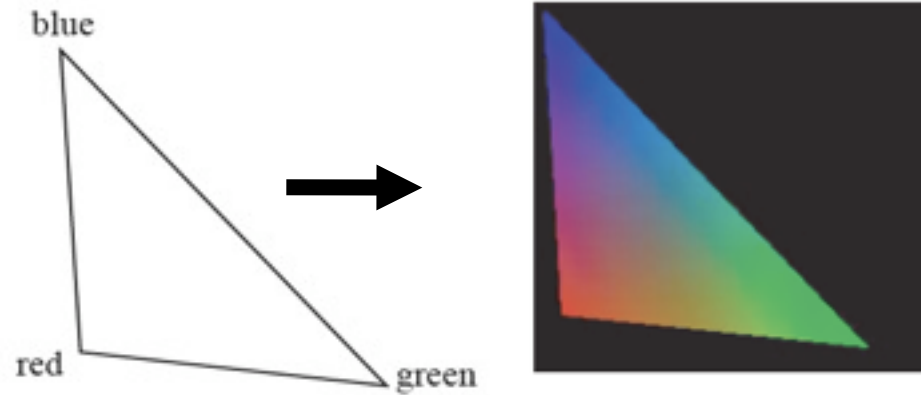
Interpolation

- What is s at **p**?

- $s(x,y)$ should vary smoothly over the triangle!

- Can be done using *barycentric coordinates*, $(\bar{u}, \bar{v}, \bar{w})$

$\mathbf{p}^2, s^2$

$\mathbf{p}^0, s^0$

A_1

p

A_0

A_2

$\mathbf{p}^1, s^1$

Thursday, 5 November 2009

Barycentric coordinates

- Are proportional to the signed areas of the subtriangles formed by **p** and the vertices

- Area computed using cross product, e.g.:

$$A_1 = \frac{1}{2}((p_x - p_x^0)(p_y^2 - p_y^0) - (p_y - p_y^0)(p_x^2 - p_x^0))$$

- In graphics, we always use barycentric coordinates normalized with respect to triangle area:

$$(\bar{u}, \bar{v}, \bar{w}) = \frac{(A_1, A_2, A_0)}{A_\Delta} \qquad A_\Delta = A_0 + A_1 + A_2$$

$$\bar{u} + \bar{v} + \bar{w} = 1 \qquad \bar{w} = 1 - \bar{u} - \bar{v}$$

Thursday, 5 November 2009

What are those barycentric coordinates?

- Some examples

$\bar{u}=0$
$\bar{u}=0.34$
A_1
$\bar{u}=1$

$\bar{v}=1$
$\bar{v}=0.37$
A_2
$\bar{v}=0$

- Why are they constant on lines parallel to an edge?

- Because height of subtriangle is constant!

Thursday, 5 November 2009

How are they used?
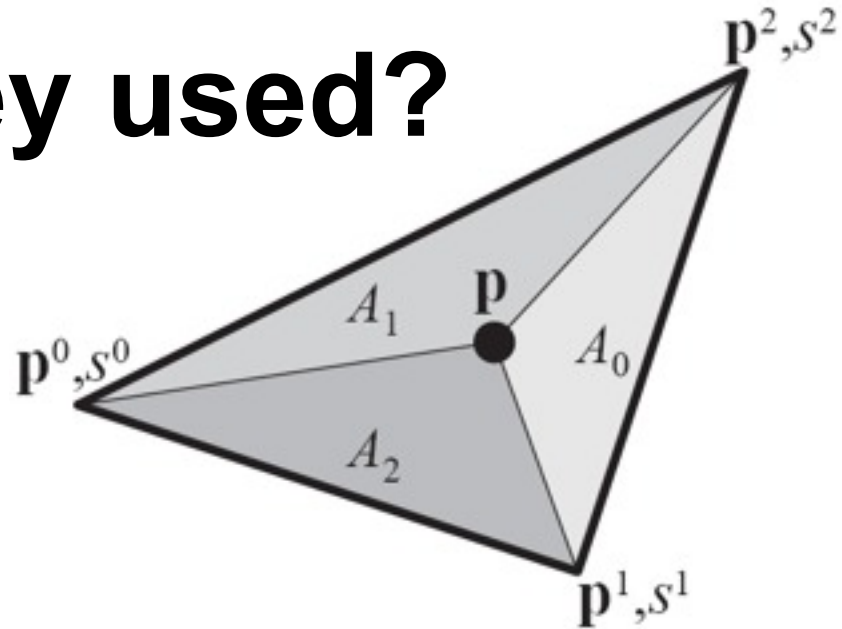
- 1: Compute barycentric coordinates for a pixel
- 2: Interpolate vertex parameters, s^0, s^1, s^2:

$$s = \bar{w}s_0 + \bar{u}s_1 + \bar{v}s_2 = (1 - \bar{u} - \bar{v})s_0 + \bar{u}s_1 + \bar{v}s_2$$
$$= s_0 + \bar{u}(s_1 - s_0) + \bar{v}(s_2 - s_0).$$

- Depth $d = p_z = h_z/h_w$ should be interpolated like this

- Note also:
  - And they are negative outside the triangle,
  - Or >1

$$\bar{u} \geq 0,$$
$$\bar{v} \geq 0, \text{ and}$$
$$\bar{w} \geq 0 \Leftrightarrow \bar{u} + \bar{v} \leq 1$$

34

Barycentric coordinates from edge functions (1)

- The a and b parameters of an edge function must be proportional to the normal
  - However, if defined as we have, then we can use the edge functions directly to compute barycentric coordinates as well!

- Focus on edge, e_2:

$$e_2(x,y) = e_2(\mathbf{p}) = \mathbf{n}_2 \cdot (\mathbf{p} - \mathbf{p}^0)$$

Thursday, 5 November 2009

Barycentric coordinates from edge functions (2)

- From definition of dot product:

$$e_2(x, y) = e_2(\mathbf{p}) = \mathbf{n}_2 \cdot (\mathbf{p} - \mathbf{p}^0) \quad \Longleftrightarrow$$

$$e_2(\mathbf{p}) = ||\mathbf{n}_2|| \, ||\mathbf{p} - \mathbf{p}^0|| \cos \alpha$$

- $||\mathbf{n}_2||$ must be exactly b (base of triangle)
- $||\mathbf{p}\text{-}\mathbf{p}^0|| \cos \alpha$ is the length of projection of $\mathbf{p}\text{-}\mathbf{p}^0$ onto $\mathbf{n}_2$ i.e., h (height of triangle)

Thursday, 5 November 2009

Barycentric coordinates from edge functions (3)

- This means:

$$\bar{u} = \frac{e_1(x,y)}{2A_\triangle}$$

$$\bar{v} = \frac{e_2(x,y)}{2A_\triangle}$$

- And $1/(2A_\triangle)$ can be computed in the triangle setup (once per triangle)

Thursday, 5 November 2009

Resulting interpolation

With barycentric coordinates,
i.e., without perspective correction

With perspective correction

- Looks even worse when animated...
- Clearly, perspective correction is needed!

Which is which?

Perspective-correct interpolation

- Why?
  - Things farther away appear smaller!

- And even inside objects, of course:

gett   39

Per-pixel division is required for perspective correct interpolation

- Surprisingly difficult to explain...
- Recall perspective projection: $\mathbf{Me}=\mathbf{h}$
  - Where $\mathbf{M}$ is projection matrix, $\mathbf{e}$ is coord in eye space, $\mathbf{h}$ is the result (but no homogenization done)

Thursday, 5 November 2009

Here is the matrix M for your convenience

$$\mathbf{Me} = \begin{pmatrix} \dfrac{2n}{r-l} & 0 & 0 & 0 \\ 0 & \dfrac{2n}{t-b} & 0 & 0 \\ 0 & 0 & -\dfrac{f+n}{f-n} & -\dfrac{2fn}{f-n} \\ 0 & 0 & -1 & 0 \end{pmatrix} \begin{pmatrix} e_x \\ e_y \\ e_z \\ 1 \end{pmatrix} = \begin{pmatrix} h_x \\ h_y \\ h_z \\ h_w \end{pmatrix}$$

- h_x and h_y are simply scaled (and translated, in the general case) eye coordinates in x and y
- h_w is simply the z-coordinate in eye-space
- Thus $(h_x,\ h_y,\ h_w)$ is **e** but in a slightly distorted space!

Tutorial on homogenization

$$\mathbf{Mv} = \mathbf{h} = \begin{pmatrix} h_x \\ h_y \\ h_z \\ h_w \end{pmatrix} \implies \begin{pmatrix} h_x/h_w \\ h_y/h_w \\ h_z/h_w \\ h_w/h_w \end{pmatrix} = \begin{pmatrix} h_x/h_w \\ h_y/h_w \\ h_z/h_w \\ 1 \end{pmatrix} = \mathbf{p}$$

M is a projection matrix

(p_x, p_y) **are screen space coords**

öller and Michael Doggett   41

Our goal

- Interpolate with perspective
- Let's try to use linear interpolation as much as possible

Thursday, 5 November 2009

Perspective correct interpolation (2)

- Now, we have parameters, s^i, per vertex that we want to interpolate.

- In distorted eye space, this is straightforward, because we can write the following for each vertex:

$$s^i = kh_x^i + lh_y^i + mh_w^i \qquad i=0,1,2$$

- Three equations, three unknowns (k,l,m) → solvable:

$$s(h_x, h_y, h_w) = kh_x + lh_y + mh_w$$

- We need to express this in screen space: (p_x, p_y)
  - But $(p_x, p_y) = (h_x/h_w,\ h_y/h_w)$

Thursday, 5 November 2009

Perspective correct interpolation (3)

$$s(h_x, h_y, h_w) = kh_x + lh_y + mh_w$$

- Divide equation by h_w! $\rightarrow$

$$s(h_x, h_y, h_w) = kh_x + lh_y + mh_w \Leftrightarrow$$

$$k\frac{h_x}{h_w} + l\frac{h_y}{h_w} + m\frac{h_w}{h_w} = kp_x + lp_y + m = \hat{s}(p_x, p_y)$$

- That is, *linear* interpolation in **screen space**!

- But unfortunately, we interpolate s/h_w ☹

- Simple to correct for though:
  - Linearly interpolation of: $1/h_w$ Call func: $\hat{o}(p_x, p_y)$
  - And then…

Thursday, 5 November 2009

Perspective correct interpolation (4)

- Correction is done as: $s(p_x, p_y) = \dfrac{\hat{s}(p_x, p_y)}{\hat{o}(p_x, p_y)}$

- Using sloppy notation (simpler to understand, perhaps): $\dfrac{s/w}{1/w} = \dfrac{sw}{w} = s$

- So, to sum up:
  – Linearly interpolate s/w in screen space
  – Linearly interpolate 1/w in screen space
  – Once s/w and 1/w have been computed for a pixel (or sample point), recover s as: (s/w)/(1/w)
  – Linear interpolation can be done with barycentric coordinates: $(\bar{u}, \bar{v})$
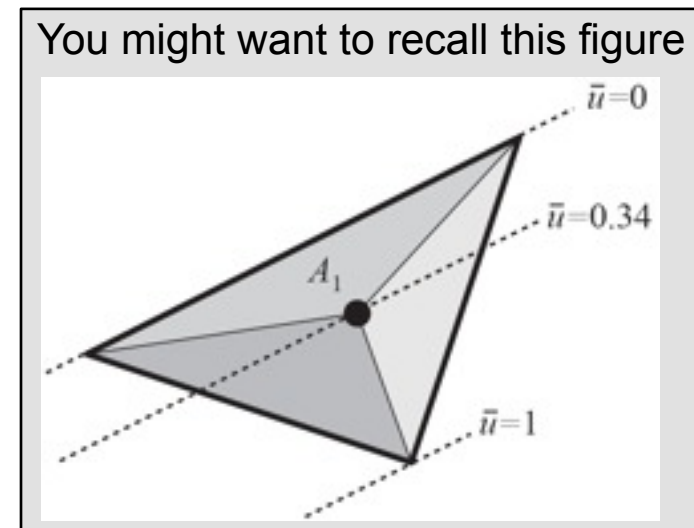
Thursday, 5 November 2009

"Perspectively-correct interpolation coordinates" (1)

- Background: many parameters to interpolate (texture coordinates (1-8), fog, color, etc) and expensive to compute parameter/h_w

- Better: compute (u,v) that are similar to barycentric coordinates, $(\bar{u},\bar{v})$ but (u,v) have taken perspective into account

- Advantage: once (u,v) have been computed, we can interpolate all parameters with perspective as:

$$s(p_x, p_y) = (1 - u - v)s^0 + us^1 + vs^2 = s^0 + u(s^1 - s^0) + v(s^2 - s^0)$$

Thursday, 5 November 2009

"Perspectively-correct barycentric coordinates" (2)

- So, how compute (u,v)?
- Imagine, creating these per-vertex parameters, s^i :

$\mathbf{p}^2, s^2 = 0$

$\mathbf{p}^0, s^0 = 0$

$\mathbf{p}^1, s^1 = 1$

You might want to recall this figure

$\bar{u}=0$

$\bar{u}=0.34$

A_1

$\bar{u}=1$

- This means, that if we interpolate s with perspective in mind, we will get u, which is perspectively correct!
- Similar trick can be done for $v \rightarrow (u,v)$
- Can all be done with edge functions too!

Thursday, 5 November 2009

"Perspectively-correct barycentric coordinates" (3)

- Recall: $u(p_x, p_y) = \dfrac{\hat{s}(p_x, p_y)}{\hat{o}(p_x, p_y)}$

$$\hat{s}(p_x, p_y) = (1 - \bar{u} - \bar{v})\frac{0}{h_w^0} + \bar{u}\frac{1}{h_w^1} + \bar{v}\frac{0}{h_w^2}$$

$$\hat{o}(p_x, p_y) = (1 - \bar{u} - \bar{v})\frac{1}{h_w^0} + \bar{u}\frac{1}{h_w^1} + \bar{v}\frac{1}{h_w^2}$$

- Simplify:

$$u(p_x, p_y) = \frac{\dfrac{e_1}{h_w^1}}{\dfrac{e_0}{h_w^0} + \dfrac{e_1}{h_w^1} + \dfrac{e_2}{h_w^2}} \quad \rightarrow$$

Recall

$$\bar{u} = \frac{e_1(x, y)}{2A_\Delta}$$

$$\bar{v} = \frac{e_2(x, y)}{2A_\Delta}$$

Very important result

$$u = \frac{f_1}{f_0 + f_1 + f_2}$$

$$v = \frac{f_2}{f_0 + f_1 + f_2}$$

$$f_0 = \frac{e_0(x, y)}{h_w^0}, \qquad f_1 = \frac{e_1(x, y)}{h_w^1}, \qquad f_2 = \frac{e_2(x, y)}{h_w^2}$$

Triangle setup vs per-pixel computations

- ## Triangle setup

	Notation	Description
1	$a_i, b_i, c_i, i \in [0, 1, 2]$	Edge functions
2	$\dfrac{1}{2A_\Delta}$	Half reciprocal of triangle area
3	$\dfrac{1}{h_w^i}$	Reciprocal of w-coordinates

- ## Per pixel (only most basic stuff)

	Notation	Description
1	$e_i(x, y)$	Evaluate edge functions at (x, y)
2	$(\bar{u}, \bar{v})$	Barycentric coordinates (Equation 3.13)
3	$d(x, y)$	Per-pixel depth (Equation 3.14)
4	$f_i(x, y)$	Evaluation of per-pixel f-values (Equation 3.21)
5	(u, v)	Perspectively-correct interpolation coordinates (Equation 3.22)
6	$s(x, y)$	Interpolation of all desired parameters, s^i (Equation 3.15)

Thursday, 5 November 2009

Summary

- Edge functions are very powerful
  - Crackfree rasterization (and no unnecessary overdraw)
  - Can evaluate pixels in parallel if needed
  - Can interpolate using them as well
    - Perspective-correct and...
    - ...without perspective
  - Tiled traversal possible too

- Next week, the (even more) fun graphics hardware stuff starts!
  - Time to save memory accesses!!
  - And power!

- Before next lecture: read about caches
  - Chapter 5, section 5 in online notes