

# **CS 380 - GPU and GPGPU Programming**

## **Lecture 17: GPU Texturing, Pt. 3**

Markus Hadwiger, KAUST

# Reading Assignment #10 (until Nov 6)



## Read (required):

- MIP-Map Level Selection for Texture Mapping

<https://ieeexplore.ieee.org/stamp/stamp.jsp?arnumber=765326>

## Read (optional):

- Vulkan Tutorial

<https://vulkan-tutorial.com>

# Quiz #2: Nov 9



## Organization


- First 30 min of lecture
- No material (book, notes, ...) allowed

## Content of questions


- Lectures (both actual lectures and slides)
- Reading assignments
- Programming assignments (algorithms, methods)
- Solve short practical examples

# GPU Texturing

# Texturing: General Approach



Texels




Texture space  $(u, v)$

Object space  $(x_O, y_O, z_O)$

Image Space  $(x_I, y_I)$

Parametrization

Rendering  
(Projection etc.)



# Interpolation #1



Interpolation Type + Purpose #1:  
**Interpolation of Texture Coordinates**  
*(Linear / Rational-Linear Interpolation)*

# Homogeneous Coordinates (1)



## Projective geometry

- (Real) projective spaces  $\mathbb{RP}^n$ :  
Real projective line  $\mathbb{RP}^1$ , real projective plane  $\mathbb{RP}^2$ , ...
- A point in  $\mathbb{RP}^n$  is a line through the origin (i.e., all the scalar multiples of the same vector) in an  $(n+1)$ -dimensional (real) vector space



## Homogeneous coordinates of 2D projective point in $\mathbb{RP}^2$

- Coordinates differing only by a non-zero factor  $\lambda$  map to the same point  
 $(\lambda x, \lambda y, \lambda)$       dividing out the  $\lambda$  gives  $(x, y, 1)$ , corresponding to  $(x,y)$  in  $\mathbb{R}^2$
- Coordinates with last component = 0 map to “points at infinity”  
 $(\lambda x, \lambda y, 0)$       division by last component not allowed; but again this is the same point if it only differs by a scalar factor, e.g., this is the same point as  $(x, y, 0)$



# Texture Mapping

---

2D (3D) Texture Space

| Texture Transformation

2D Object Parameters

| Parameterization

3D Object Space

| Model Transformation


3D World Space

| Viewing Transformation

3D Camera Space

| Projection

2D Image Space



# Perspective-correct linear interpolation

---

Only projected values interpolate correctly, so project  $A$

- Linearly interpolate  $A_1/w_1$  and  $A_2/w_2$

Also interpolate  $1/w_1$  and  $1/w_2$

- These also interpolate linearly in screen space

Divide interpolants at each sample point to recover  $A$


- $(A/w) / (1/w) = A$
- Division is expensive (more than add or multiply), so
  - Recover  $w$  for the sample point (reciprocate), and
  - Multiply each projected attribute by  $w$

Barycentric triangle parameterization:

$$A = \frac{aA_1/w_1 + bA_2/w_2 + cA_3/w_3}{a/w_1 + b/w_2 + c/w_3} \quad a + b + c = 1$$

# Perspective Texture Mapping

- Solution: interpolate  $(s/w, t/w, 1/w)$
- $(s/w) / (1/w) = s$  etc. at every fragment



Heckbert and Moreton




# Perspective-Correct Interpolation Recipe





$$r_i(x, y) = \frac{r_i(x, y) / w(x, y)}{1 / w(x, y)}$$

- (1) Associate a record containing the  $n$  parameters of interest  $(r_1, r_2, \dots, r_n)$  with each vertex of the polygon.
- (2) For each vertex, transform object space coordinates to homogeneous screen space using  $4 \times 4$  object to screen matrix, yielding the values  $(xw, yw, zw, w)$ .
- (3) Clip the polygon against plane equations for each of the six sides of the viewing frustum, linearly interpolating all the parameters when new vertices are created.
- (4) At each vertex, divide the homogeneous screen coordinates, the parameters  $r_i$ , and the number 1 by  $w$  to construct the variable list  $(x, y, z, s_1, s_2, \dots, s_{n+1})$ , where  $s_i = r_i/w$  for  $i \leq n$ ,  $s_{n+1} = 1/w$ .
- (5) Scan convert in screen space by linear interpolation of all parameters, at each pixel computing  $r_i = s_i/s_{n+1}$  for each of the  $n$  parameters; use these values for shading.

- Want to simulate a beamer
  - ... or a flashlight, or a slide projector
- Precursor to shadows
- Interesting mathematics:  
2 perspective  
projections involved!
- Easy to program!



# Projective Texture Mapping



# Projective Shadows in Doom 3




- What about **homogeneous** texture coords?
- Need to do perspective divide also for projector!
  - $(s, t, q) \rightarrow (s/q, t/q)$  for every fragment
- How does OpenGL do that?
  - Needs to be perspective correct as well!
  - Trick: interpolate  $(s/w, t/w, r/w, q/w)$
  - $(s/w) / (q/w) = s/q$  etc. at every fragment
- Remember:  $s, t, r, q$  are equivalent to  $x, y, z, w$  in projector space!  $\rightarrow r/q = \text{projector depth!}$






- Apply multiple textures in one pass
- *Integral* part of programmable shading
  - e.g. diffuse texture map + gloss map
  - e.g. diffuse texture map + light map
- Performance issues
  - How many textures are free?
  - How many are available



- Used in virtually every commercial game
- Precalculate diffuse lighting on static objects
  - Only low resolution necessary
  - Diffuse lighting is view independent!
- Advantages:
  - No runtime lighting necessary
    - VERY fast!
  - Can take global effects (shadows, color bleeds) into account






Original LM texels





Bilinear Filtering



- Why premultiplication is bad...



Full Size Texture  
(with Lightmap)



Tiled Surface Texture  
plus Lightmap

→ use tileable surface textures and low resolution lightmaps





Original scene




Light-mapped




- Precomputation based on non-realtime methods
  - Radiosity
  - Ray tracing
    - Monte Carlo Integration
    - Path tracing
    - Photon mapping






Lightmap



mapped





Original scene



Light-mapped





# Interpolation #2



Interpolation Type + Purpose #2:  
**Interpolation of Samples in Texture Space**  
*(Multi-Linear Interpolation)*

- Spatial layout
  - Cartesian grids: 1D, 2D, 3D, 2D\_ARRAY, ...
  - Cube maps, ...
- Formats (too many), e.g. OpenGL
  - GL\_LUMINANCE16\_ALPHA16
  - GL\_RGB8, GL\_RGBA8, ...: integer texture formats
  - GL\_RGB16F, GL\_RGBA32F, ...: float texture formats
  - compressed formats, high dynamic range formats, ...
- External (CPU) format vs. internal (GPU) format
  - OpenGL driver converts from external to internal



# Magnification (Bi-linear Filtering Example)



Original image




Nearest neighbor




Bi-linear filtering



# Nearest-Neighbor vs. Bi-Linear Interpolation




nearest-neighbor



bi-linear

wikipedia





Bilinear patch (courtesy J. Han)

# Bi-Linear Interpolation



Consider area between 2x2 adjacent samples (e.g., pixel centers)

Example #2: 1 at top-left and bottom-right, 0 at bottom-left, 0.5 at top-right



# Bi-Linear Interpolation



Consider area between 2x2 adjacent samples (e.g., pixel centers):

Given any (fractional) position


$$\alpha_1 := x_1 - \lfloor x_1 \rfloor \quad \alpha_1 \in [0.0, 1.0)$$

$$\alpha_2 := x_2 - \lfloor x_2 \rfloor \quad \alpha_2 \in [0.0, 1.0)$$

and 2x2 sample values

$$\begin{bmatrix} v_{01} & v_{11} \\ v_{00} & v_{10} \end{bmatrix}$$

Compute:  $f(\alpha_1, \alpha_2)$



# Bi-Linear Interpolation



Consider area between 2x2 adjacent samples (e.g., pixel centers):

Given any (fractional) position


$$\alpha_1 := x_1 - \lfloor x_1 \rfloor \quad \alpha_1 \in [0.0, 1.0)$$

$$\alpha_2 := x_2 - \lfloor x_2 \rfloor \quad \alpha_2 \in [0.0, 1.0)$$

and 2x2 sample values

$$\begin{bmatrix} v_{01} & v_{11} \\ v_{00} & v_{10} \end{bmatrix}$$

Compute:  $f(\alpha_1, \alpha_2)$





# Bi-Linear Interpolation



Weights in 2x2 format:

$$\begin{bmatrix} \alpha_2 \\ (1 - \alpha_2) \end{bmatrix} \begin{bmatrix} (1 - \alpha_1) & \alpha_1 \end{bmatrix} = \begin{bmatrix} (1 - \alpha_1)\alpha_2 & \alpha_1\alpha_2 \\ (1 - \alpha_1)(1 - \alpha_2) & \alpha_1(1 - \alpha_2) \end{bmatrix}$$

Interpolate function at (fractional) position  $(\alpha_1, \alpha_2)$ :

$$f(\alpha_1, \alpha_2) = \begin{bmatrix} \alpha_2 & (1 - \alpha_2) \end{bmatrix} \begin{bmatrix} v_{01} & v_{11} \\ v_{00} & v_{10} \end{bmatrix} \begin{bmatrix} (1 - \alpha_1) \\ \alpha_1 \end{bmatrix}$$

# Bi-Linear Interpolation



Interpolate function at (fractional) position  $(\alpha_1, \alpha_2)$  :

$$\begin{aligned} f(\alpha_1, \alpha_2) &= \begin{bmatrix} \alpha_2 & (1 - \alpha_2) \end{bmatrix} \begin{bmatrix} v_{01} & v_{11} \\ v_{00} & v_{10} \end{bmatrix} \begin{bmatrix} (1 - \alpha_1) \\ \alpha_1 \end{bmatrix} \\ &= \begin{bmatrix} \alpha_2 & (1 - \alpha_2) \end{bmatrix} \begin{bmatrix} (1 - \alpha_1)v_{01} + \alpha_1 v_{11} \\ (1 - \alpha_1)v_{00} + \alpha_1 v_{10} \end{bmatrix} \\ &= \begin{bmatrix} \alpha_2 v_{01} + (1 - \alpha_2)v_{00} & \alpha_2 v_{11} + (1 - \alpha_2)v_{10} \end{bmatrix} \begin{bmatrix} (1 - \alpha_1) \\ \alpha_1 \end{bmatrix} \end{aligned}$$

# Bi-Linear Interpolation



Interpolate function at (fractional) position  $(\alpha_1, \alpha_2)$  :

$$f(\alpha_1, \alpha_2) = \begin{bmatrix} \alpha_2 & (1 - \alpha_2) \end{bmatrix} \begin{bmatrix} v_{01} & v_{11} \\ v_{00} & v_{10} \end{bmatrix} \begin{bmatrix} (1 - \alpha_1) \\ \alpha_1 \end{bmatrix}$$

$$= (1 - \alpha_1)(1 - \alpha_2)v_{00} + \alpha_1(1 - \alpha_2)v_{10} + (1 - \alpha_1)\alpha_2v_{01} + \alpha_1\alpha_2v_{11}$$

$$= v_{00} + \alpha_1(v_{10} - v_{00}) + \alpha_2(v_{01} - v_{00}) + \alpha_1\alpha_2(v_{00} + v_{11} - v_{10} - v_{01})$$






## REALLY IMPORTANT:

this is a different thing (for a different purpose)  
than the linear (or, in perspective, rational-linear)  
interpolation of texture coordinates!!

# Texture Minification


# Texture Aliasing: Minification

- Problem: One pixel in image space covers many texels




# Texture Aliasing: Minification

- Caused by *undersampling*: texture information is lost



**Texture space**




**Image space**



# Texture Anti-Aliasing: Minification

- A good pixel value is the weighted mean of the pixel area projected into texture space





Thank you.