

CS 380 - GPU and GPGPU Programming

Lecture 18: GPU Texturing, Pt. 4

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



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

Magnification (Bi-linear Filtering Example)



Original image



Nearest neighbor



Bi-linear filtering



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 (α_1, α_2) :

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

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




- MIP Mapping (“Multum In Parvo”)
 - Texture size is reduced by factors of 2 (*downsampling* = “many things in a small place”)
 - Simple (4 pixel average) and memory efficient
 - Last image is only ONE texel



- MIP Mapping (“Multum In Parvo”)
 - Texture size is reduced by factors of 2
(*downsampling* = “many things in a small place”)
 - Simple (4 pixel average) and memory efficient
 - Last image is only ONE texel



geometric series:

$$a + ar + ar^2 + ar^3 + \dots + ar^{n-1} =$$
$$= \sum_{k=0}^{n-1} ar^k = a \left(\frac{1 - r^n}{1 - r} \right)$$

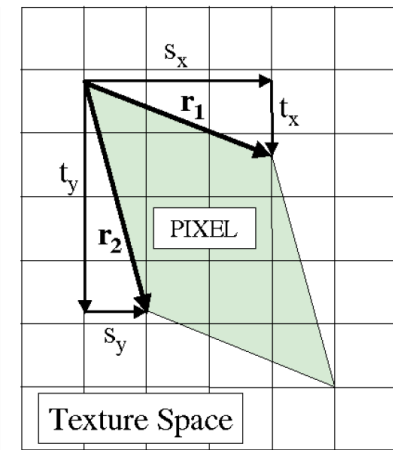
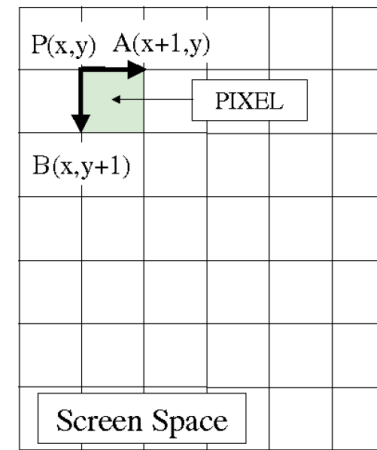



Texture Anti-Aliasing: MIP Mapping

- MIP Mapping Algorithm
- $D := \text{ld}(\max(d_1, d_2))$ "Mip Map level"
- $T_0 := \text{value from texture } D_0 = \text{trunc}(D)$
 - Use *bilinear interpolation*



MIP-Map Level Computation



- Use the partial derivatives of texture coordinates with respect to screen space coordinates
- This is the Jacobian matrix
- Area of parallelogram is the absolute value of the Jacobian determinant (the *Jacobian*)

$$\begin{pmatrix} \partial u / \partial x & \partial u / \partial y \\ \partial v / \partial x & \partial v / \partial y \end{pmatrix} = \begin{pmatrix} s_x & s_y \\ t_x & t_y \end{pmatrix}$$

MIP-Map Level Computation (OpenGL)



- OpenGL 4.6 core specification, pp. 251-264

(3D tex coords!)

$$\lambda_{base}(x, y) = \log_2[\rho(x, y)]$$

$$\rho = \max \left\{ \sqrt{\left(\frac{\partial u}{\partial x}\right)^2 + \left(\frac{\partial v}{\partial x}\right)^2 + \left(\frac{\partial w}{\partial x}\right)^2}, \sqrt{\left(\frac{\partial u}{\partial y}\right)^2 + \left(\frac{\partial v}{\partial y}\right)^2 + \left(\frac{\partial w}{\partial y}\right)^2} \right\}$$

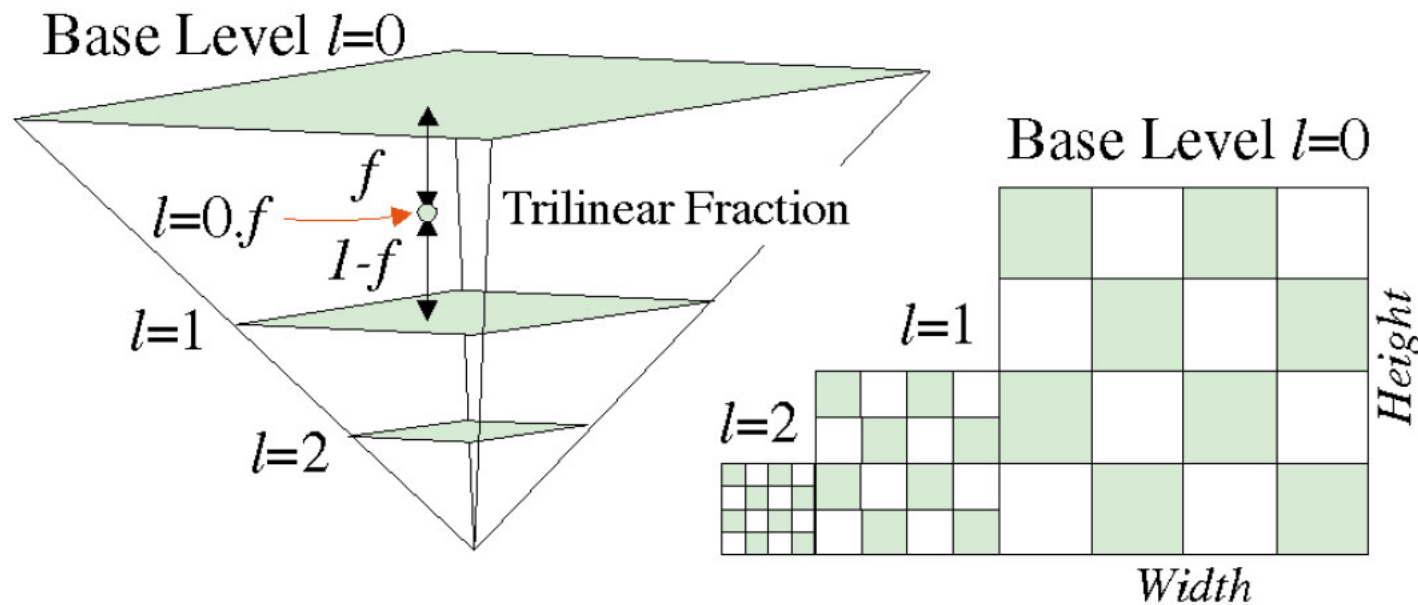
Does not use area of parallelogram but greater hypotenuse [Heckbert, 1983]

- Approximation without square-roots

$$m_u = \max \left\{ \left| \frac{\partial u}{\partial x} \right|, \left| \frac{\partial u}{\partial y} \right| \right\} \quad m_v = \max \left\{ \left| \frac{\partial v}{\partial x} \right|, \left| \frac{\partial v}{\partial y} \right| \right\} \quad m_w = \max \left\{ \left| \frac{\partial w}{\partial x} \right|, \left| \frac{\partial w}{\partial y} \right| \right\}$$


$$\max\{m_u, m_v, m_w\} \leq f(x, y) \leq m_u + m_v + m_w$$

MIP-Map Level Interpolation




- Level of detail value is fractional!
- Use fractional part to blend (lin.) between two adjacent mipmap levels

- Trilinear interpolation:
 - T_1 := value from texture $D_1 = D_0 + 1$ (bilinear interpolation)
 - Pixel value := $(D_1 - D) \cdot T_0 + (D - D_0) \cdot T_1$
 - Linear interpolation between successive MIP Maps
 - Avoids "Mip banding" (but doubles texture lookups)





Texture Anti-Aliasing: MIP Mapping

- Other example for bilinear vs. trilinear filtering



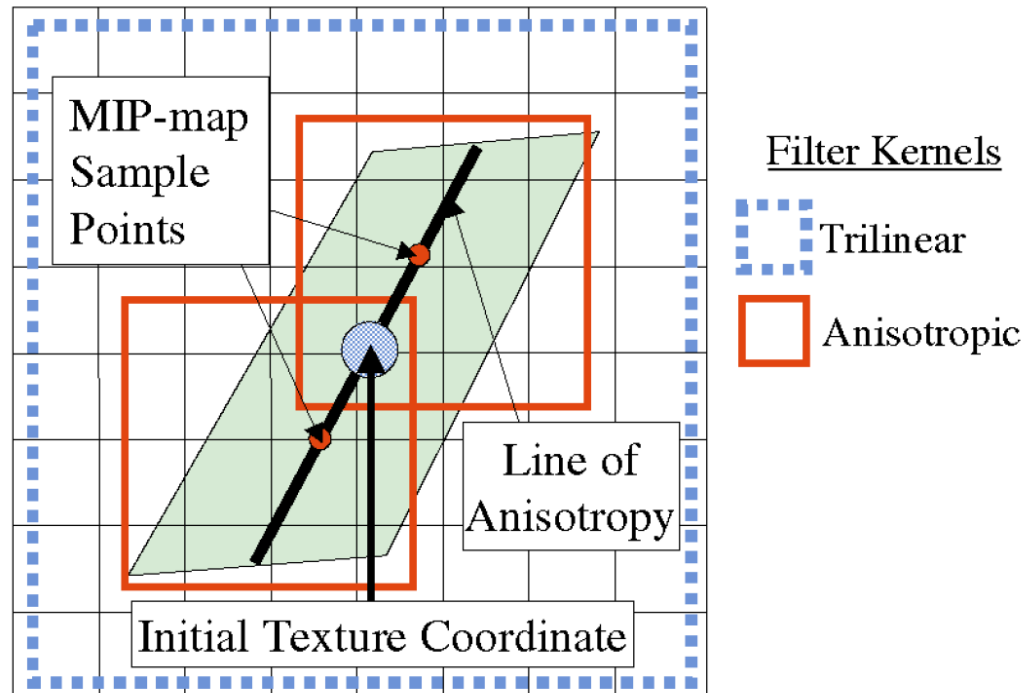
- Anisotropic filtering
 - View-dependent filter kernel
 - Implementation: *summed area table*, "*RIP Mapping*", *footprint assembly*, *elliptical weighted average (EWA)*



Texture space




Anisotropic Filtering: Footprint Assembly



Anti-Aliasing: Anisotropic Filtering

- Example



- Basically, everything done in hardware
- `gluBuild2DMipmaps()` generates MIPmaps
- Set parameters in `glTexParameter()`
 - `GL_TEXTURE_MAG_FILTER: GL_NEAREST, GL_LINEAR, ...`
 - `GL_TEXTURE_MIN_FILTER: GL_LINEAR_MIPMAP_NEAREST`
- Anisotropic filtering is an extension:
 - `GL_EXT_texture_filter_anisotropic`
 - Number of samples can be varied (4x,8x,16x)
 - Vendor specific support and extensions



Thank you.