


Full-time wrapup

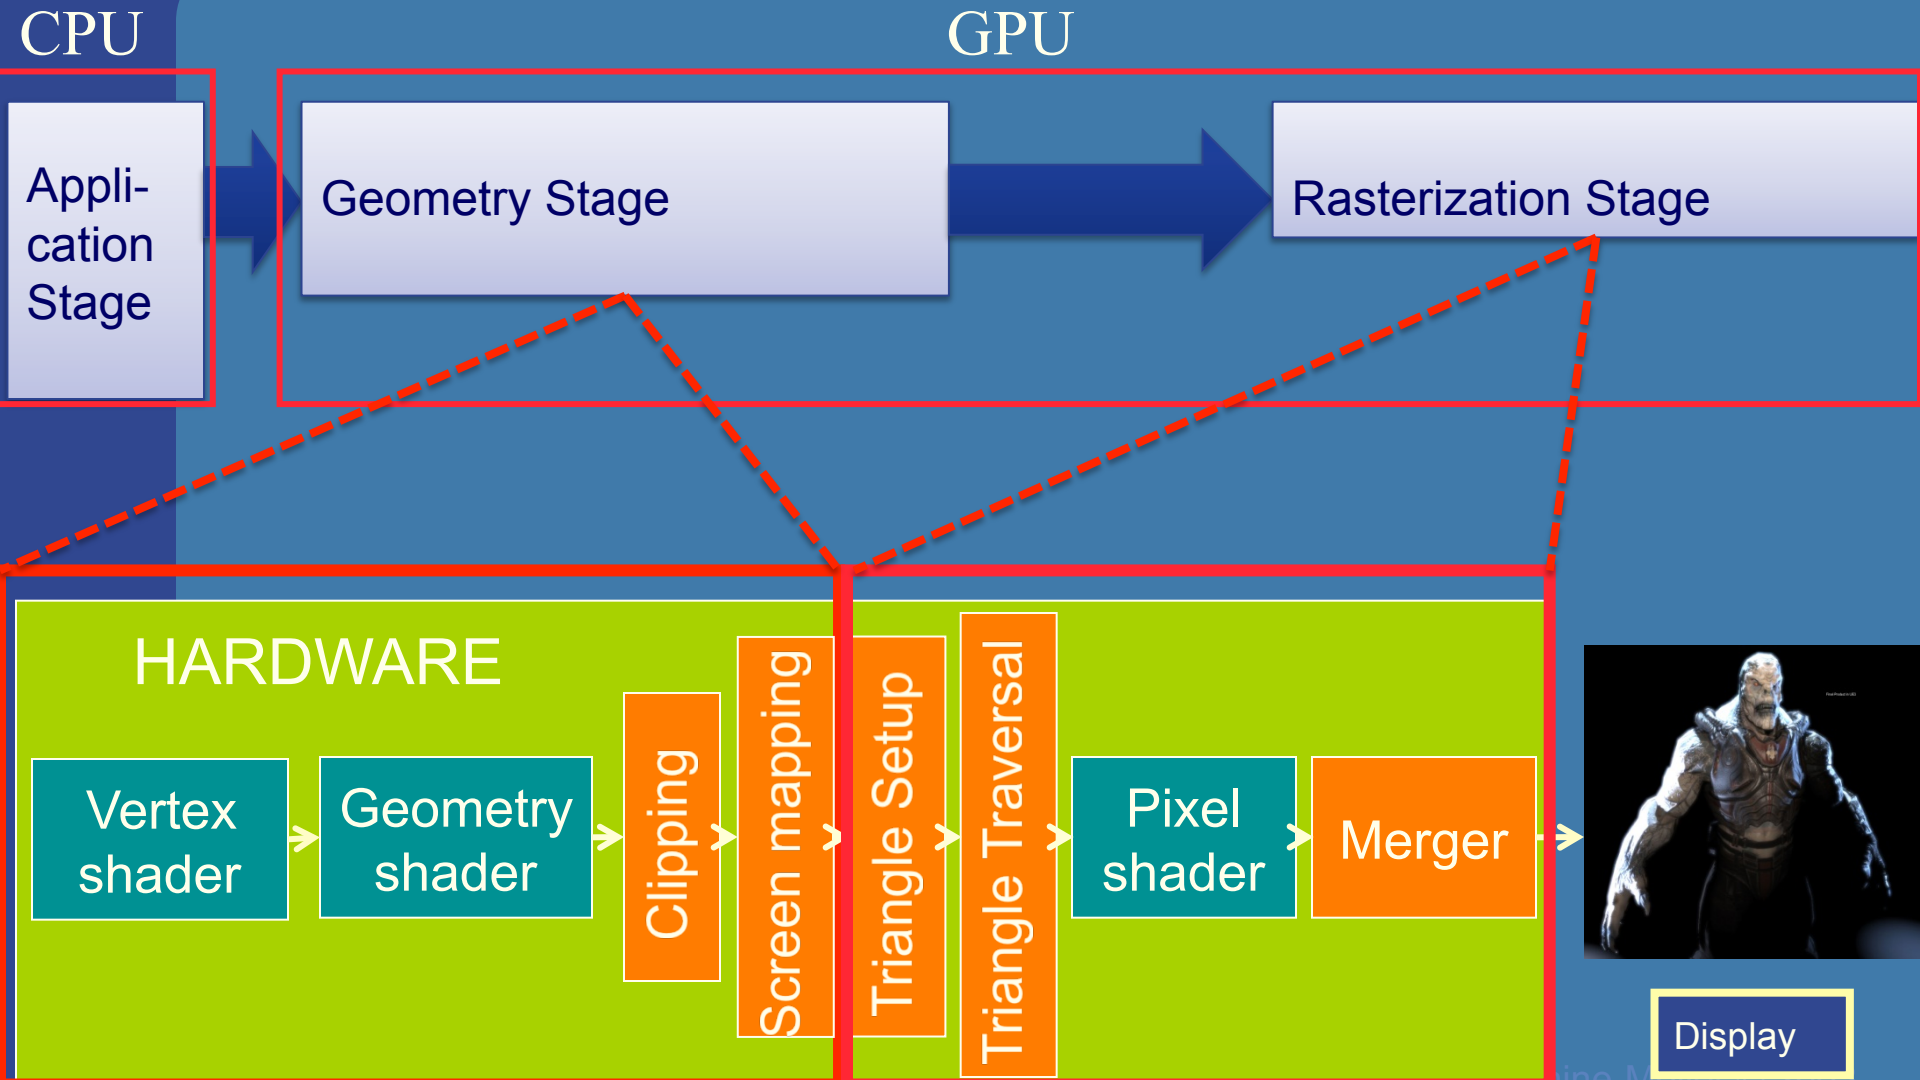
Lecture 1: Real-time Rendering

The Graphics Rendering Pipeline

- Three conceptual stages of the pipeline:
 - Application (executed on the CPU)
 - logic, speed-up techniques, animation, etc...
 - Geometry
 - Executing vertex and geometry shader
 - Vertex shader:
 - lighting computations per triangle vertex
 - Project onto screen (3D to 2D)
 - Rasterizer
 - Executing fragment shader
 - Interpolation of parameters (colors, texcoords etc) over triangle
 - Z-buffering, fragment merge (i.e., blending), stencil tests...



Rendering Pipeline and Hardware





Display

Hardware design


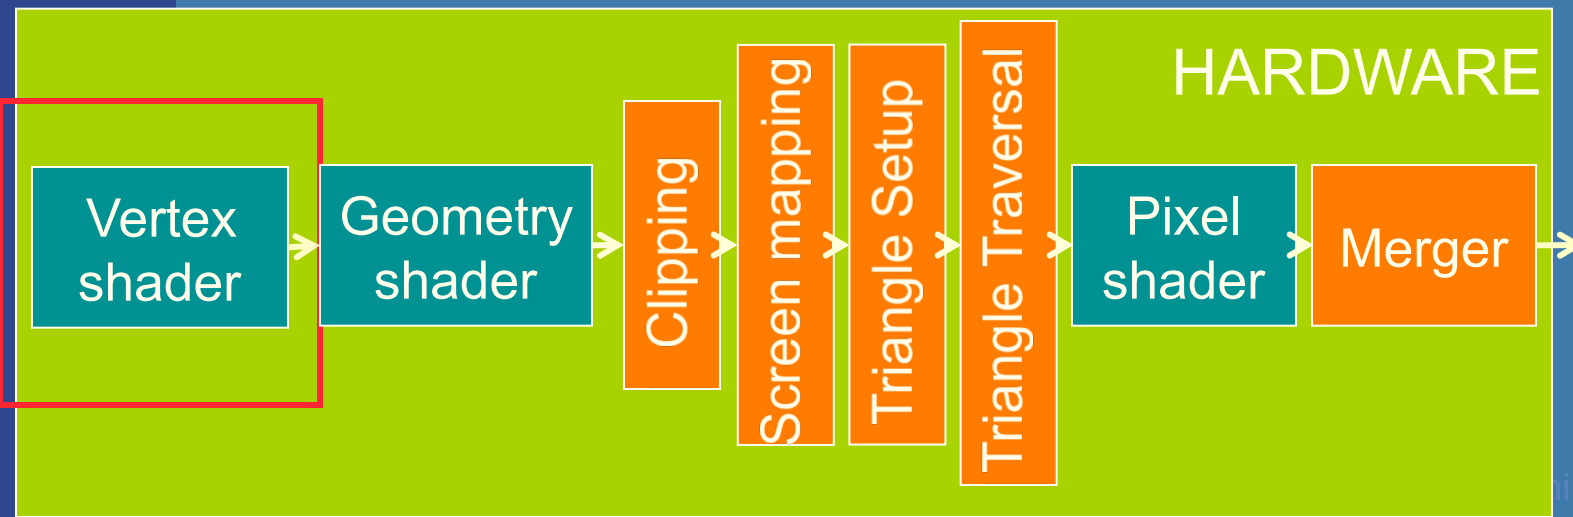

Geometry Stage

Vertex shader:

- Lighting (colors)
- Screen space positions



Geometry




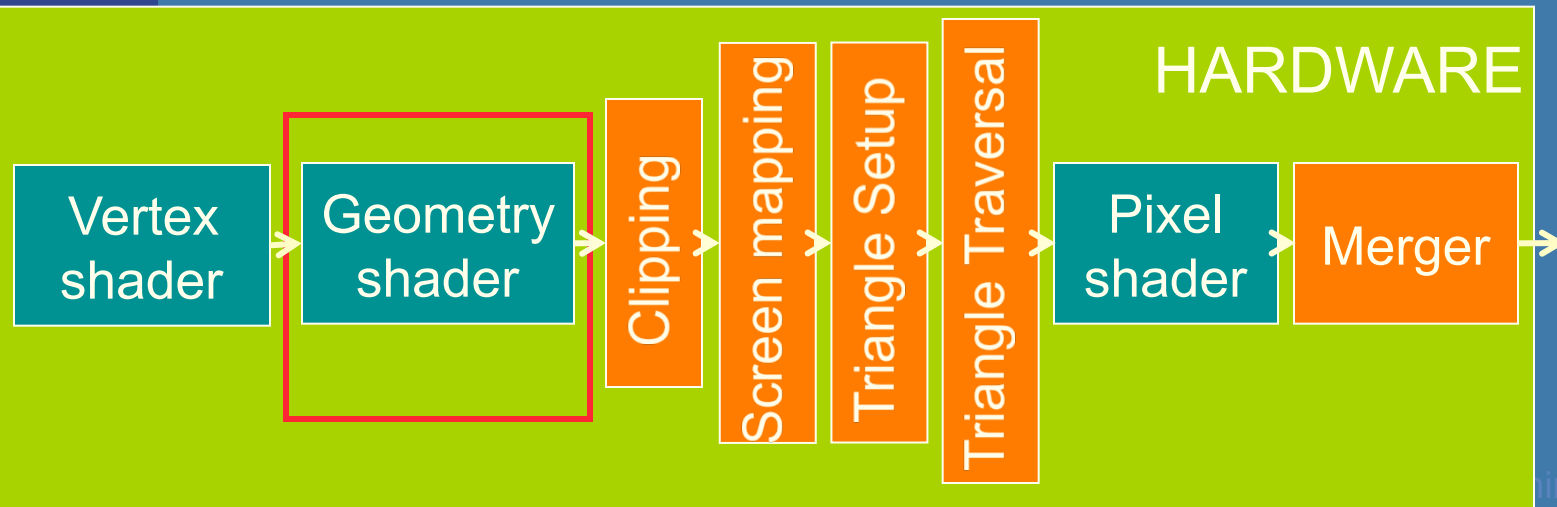

Display

Hardware design

Geometry Stage

Geometry shader:

- One input primitive
- Many output primitives

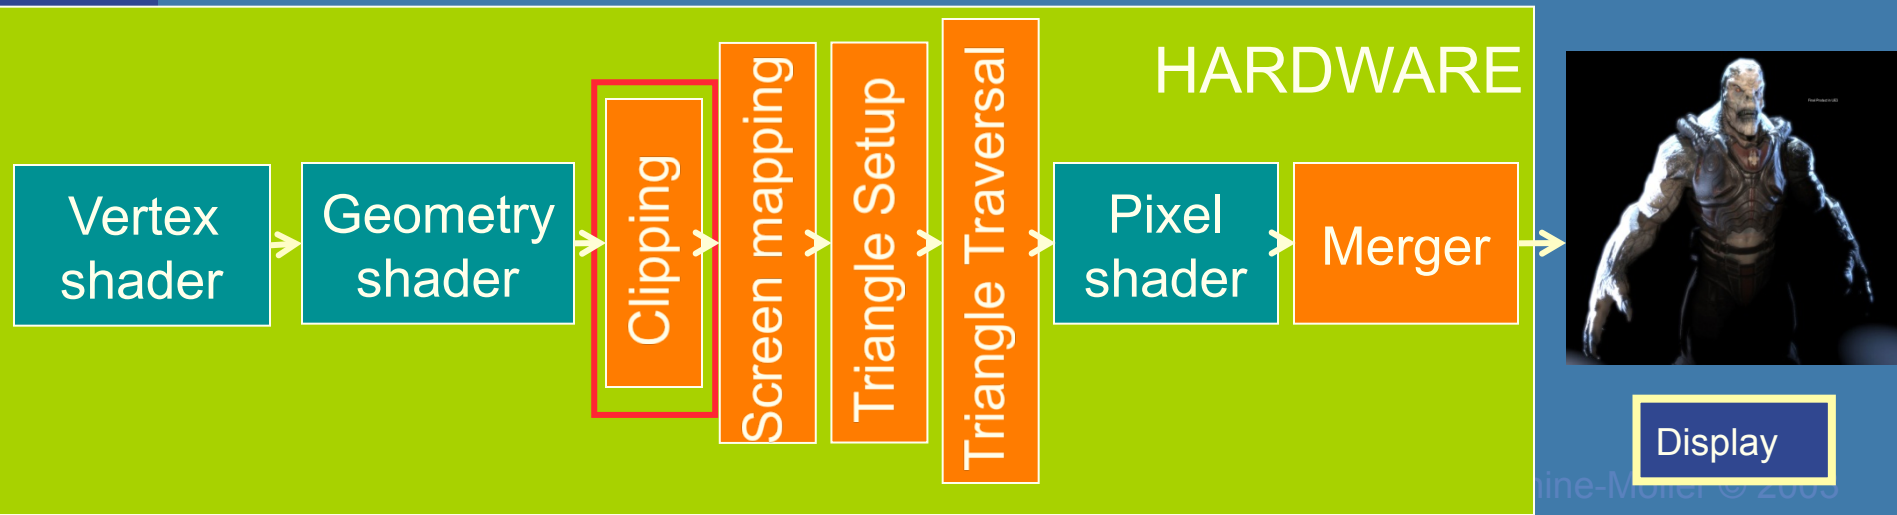



Display

Hardware design

Geometry Stage

Clips triangles against the unit cube (i.e., "screen borders")

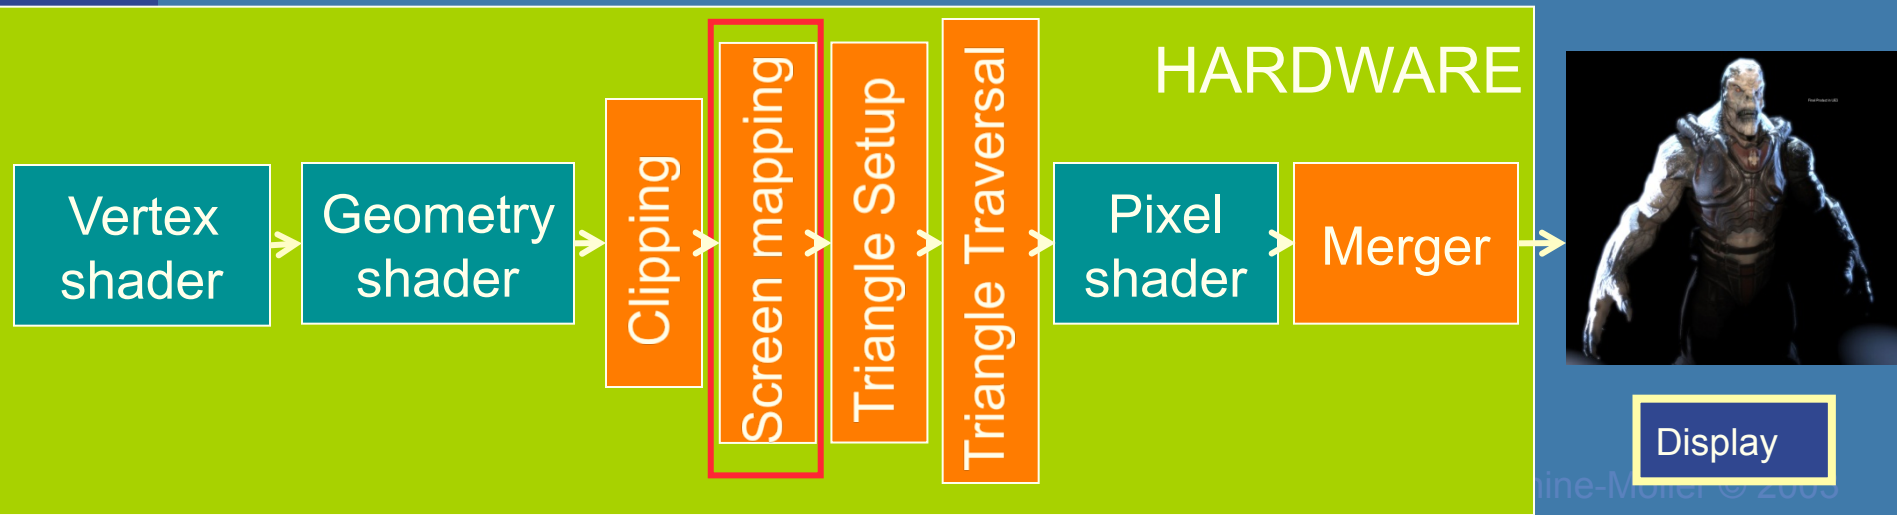



Hardware design

Rasterizer Stage

Maps window size to unit cube


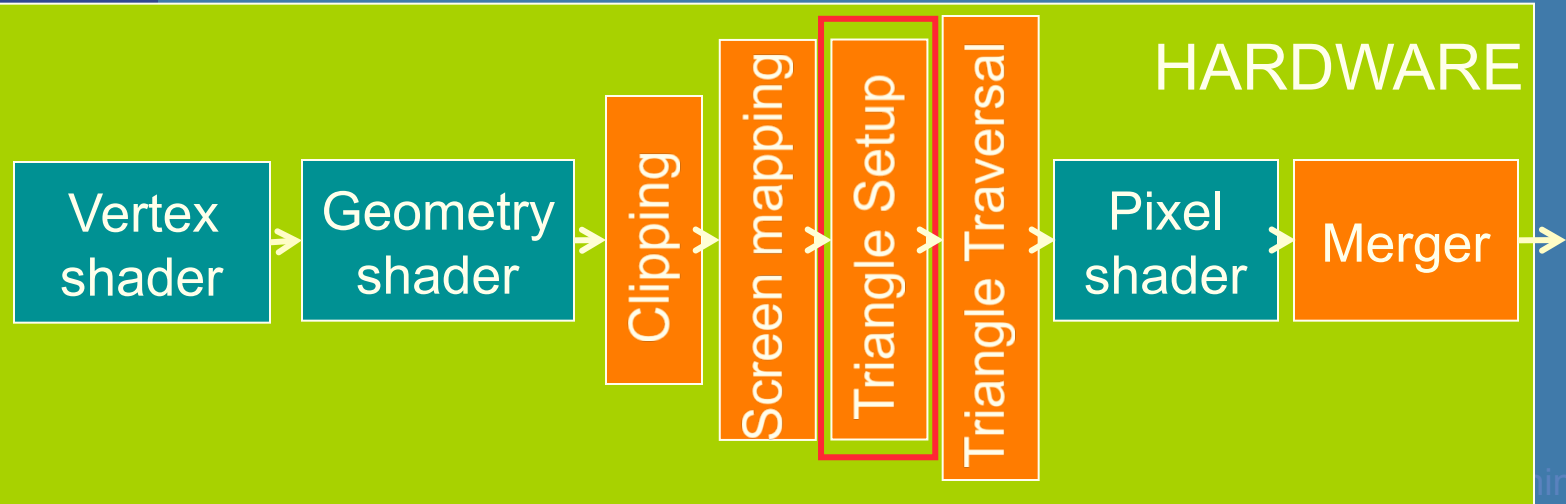

Geometry stage always operates inside a unit cube $[-1,-1,-1]-[1,1,1]$
Next, the rasterization is made against a draw area corresponding to window dimensions.



Hardware design

Rasterizer Stage

Collects three vertices into one triangle

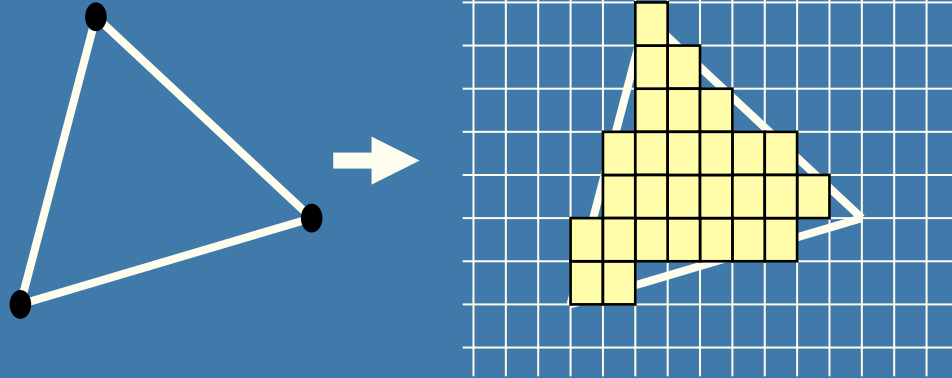


Display

Hardware design

Rasterizer Stage

Creates the fragments/
pixels for the triangle



Vertex
shader

Geometry
shader

Clipping

Screen mapping

Triangle Setup


Triangle Traversal

HARDWARE

Pixel
shader





Merger

Display



Hardware design

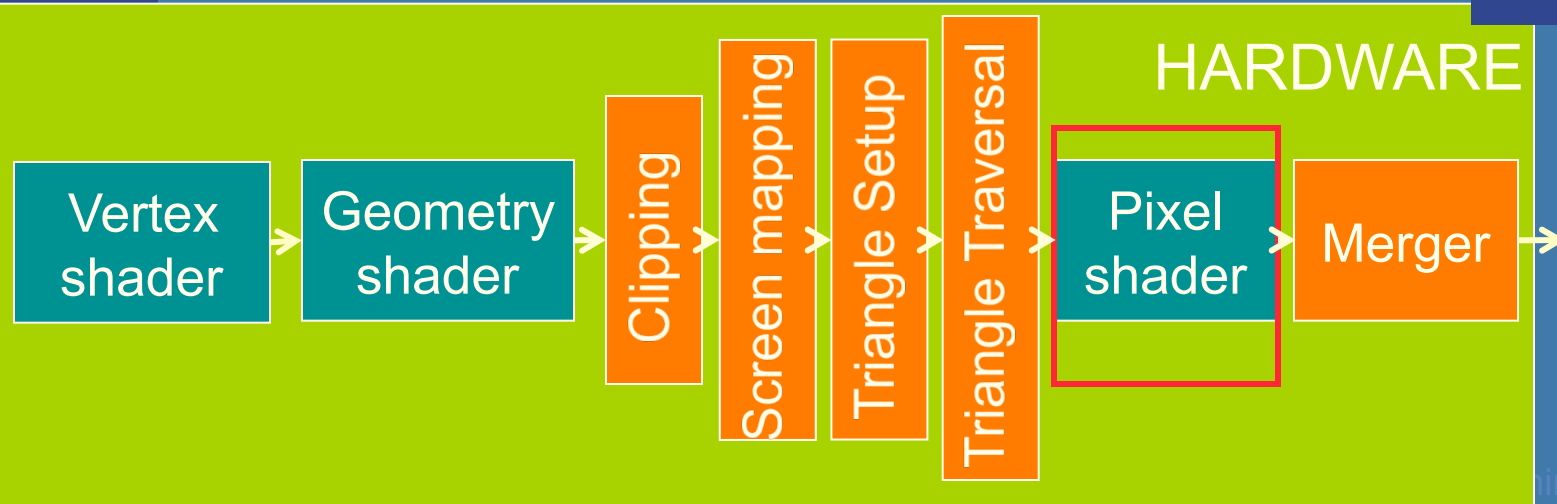
Rasterizer Stage



Rasterizer

Pixel Shader:
Compute color using:

- Textures
- Interpolated data (e.g. Colors + normals) from vertex shader



Display

Hardware design

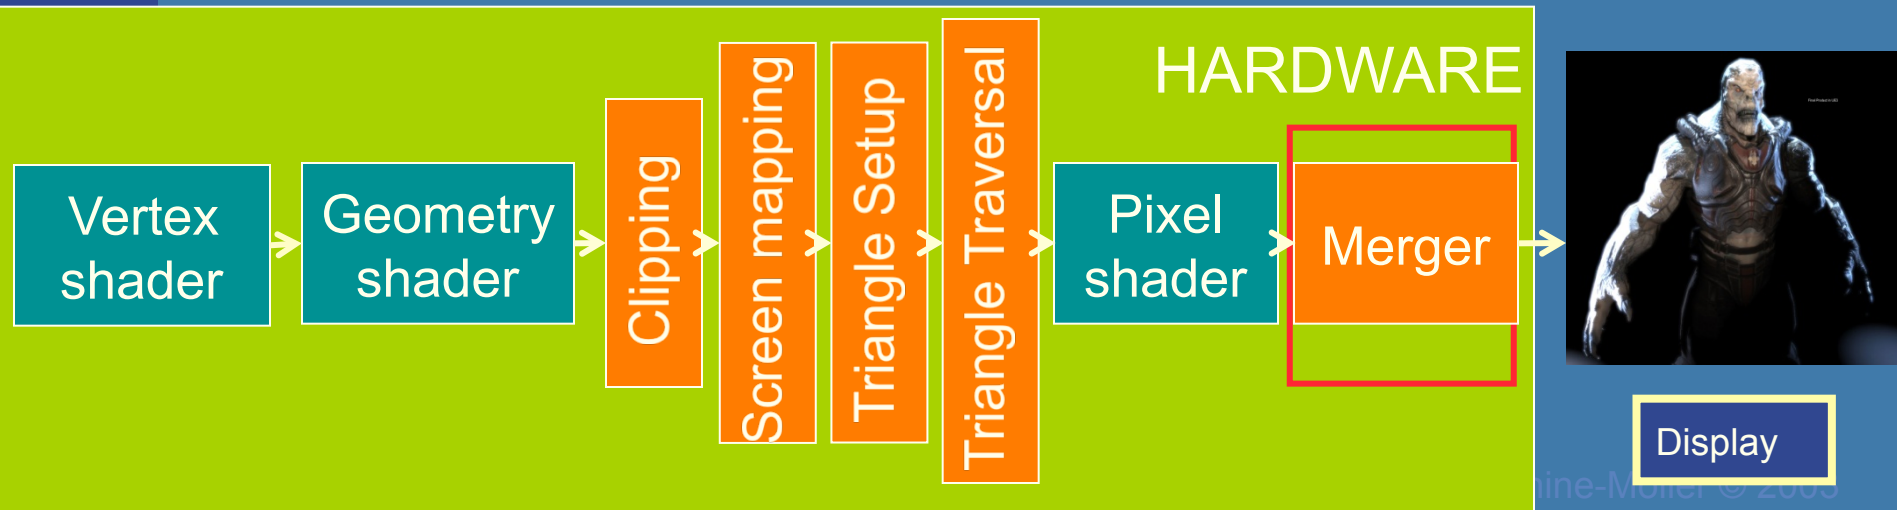
Rasterizer Stage

The merge units update the frame buffer with the pixel's color




Frame buffer:

- Color buffers
- Depth buffer
- Stencil buffer



z-Buffer Algorithm

- Use a buffer called the z or depth buffer to store the depth of the closest object at each pixel found so far
- As we render each polygon, compare the depth of each pixel to depth in z buffer
- If less, place shade of pixel in color buffer and update z buffer




Also know double buffering!

Screen Tearing



Painter's Algorithm

- Render polygons a back to front order so that polygons behind others are simply painted over



B behind A as seen by viewer

Fill B then A

- Requires ordering of polygons first
 - $O(n \log n)$ calculation for ordering
 - Not every polygon is either in front or behind all other polygons


I.e., : Sort all triangles and render them back-to-front.

Application

Geometry


Rasterizer


GEOMETRY – transformation summary



Lecture 2:


Transformation Pipeline





Model space


Model to World
Matrix



World space

World to
View
Matrix

$$\text{ModelViewMtx} = \text{Model to View Matrix}$$



View space

Lecture 2: Transforms

- Scaling, rotations, translations
- Cannot use same matrix to transform normals

Use: $\mathbf{N} = (\mathbf{M}^{-1})^T$ instead of \mathbf{M}


$(\mathbf{M}^{-1})^T = \mathbf{M}$
if rigid-body
transform

- Homogeneous notation
- Projections
- Quaternions $\hat{\mathbf{q}} = (\sin \phi \mathbf{u}_q, \cos \phi)$

– Know what they are good for. Not knowing the mathematical rules.

$$\hat{\mathbf{q}} \hat{\mathbf{p}} \hat{\mathbf{q}}^{-1}$$

- ...represents a rotation of 2ϕ radians around axis \mathbf{u}_q of point \mathbf{p}



Homogeneous notation

- A point: $\mathbf{p} = (p_x \ p_y \ p_z \ 1)^T$
- Translation becomes:

$$\underbrace{\begin{pmatrix} 1 & 0 & 0 & t_x \\ 0 & 1 & 0 & t_y \\ 0 & 0 & 1 & t_z \\ 0 & 0 & 0 & 1 \end{pmatrix}}_{\mathbf{T}(\mathbf{t})} \begin{pmatrix} p_x \\ p_y \\ p_z \\ 1 \end{pmatrix} = \begin{pmatrix} p_x + t_x \\ p_y + t_y \\ p_z + t_z \\ 1 \end{pmatrix}$$

Translation part

Rotation part

- A vector (direction): $\mathbf{d} = (d_x \ d_y \ d_z \ 0)^T$
- Translation of vector: $\mathbf{T}\mathbf{d} = \mathbf{d}$

Change of Frames

- How to get the $M_{\text{model-to-world}}$ matrix:

$$\mathbf{P} = (0, 5, 0, 1) \bullet$$

$$M_{\text{model-to-world}} = \begin{bmatrix} a_x & b_x & c_x & o_x \\ a_y & b_y & c_y & o_y \\ a_z & b_z & c_z & o_z \\ 0 & 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} 0 \\ 5 \\ 0 \\ 1 \end{bmatrix}$$

The basis vectors **a, b, c** are expressed in the world coordinate system

(Both coordinate systems are right-handed)


$$\mathbf{E.g.:} \quad \mathbf{p}_{\text{world}} = M_{\text{m} \rightarrow \text{w}} \mathbf{p}_{\text{model}} = M_{\text{m} \rightarrow \text{w}} (0, 5, 0, 1)^T = 5 \mathbf{b} (+ \mathbf{o})$$

Same example, just explained differently:

Change of Frames

$$\mathbf{p}_{\text{modelspace}} = (\mathbf{p}_x, \mathbf{p}_y, \mathbf{p}_z)$$

$$\mathbf{M}_{\text{model-to-world}} = \begin{bmatrix} a_x & b_x & c_x & o_x \\ a_y & b_y & c_y & o_y \\ a_z & b_z & c_z & o_z \\ 0 & 0 & 0 & 1 \end{bmatrix}$$



Let's initially disregard the translation \mathbf{o} . I.e., $\mathbf{o}=[0,0,0]$

X: One step along \mathbf{a} results in \mathbf{a}_x steps along world space axis x.

One step along \mathbf{b} results in \mathbf{b}_x steps along world space axis x.

One step along \mathbf{c} results in \mathbf{c}_x steps along world space axis x.

The x-coord for \mathbf{p} in *world space* (instead of modelspace) is thus $[\mathbf{a}_x \ \mathbf{b}_x \ \mathbf{c}_x]\mathbf{p}$.





The y-coord for \mathbf{p} in world space is thus $[\mathbf{a}_y \ \mathbf{b}_y \ \mathbf{c}_y]\mathbf{p}$.


The z-coord for \mathbf{p} in world space is thus $[\mathbf{a}_z \ \mathbf{b}_z \ \mathbf{c}_z]\mathbf{p}$.

With the translation \mathbf{o} we get $\mathbf{p}_{\text{worldspace}} = \mathbf{M}_{\text{model-to-world}} \mathbf{p}_{\text{modelspace}}$

Projections

- Orthogonal (parallel) and Perspective







Orthogonal projection

- Simple, just skip one coordinate
 - Say, we're looking along the z-axis
 - Then drop z, and render

$$\mathbf{M}_{ortho} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix} \Rightarrow \mathbf{M}_{ortho} \begin{pmatrix} p_x \\ p_y \\ p_z \\ 1 \end{pmatrix} = \begin{pmatrix} p_x \\ p_y \\ 0 \\ 1 \end{pmatrix}$$



DDA Algorithm



- Digital Differential Analyzer

- DDA was a mechanical device for numerical solution of differential equations

- Line $y=kx+m$ satisfies differential equation


$$dy/dx = k = \Delta y / \Delta x = y_2 - y_1 / x_2 - x_1$$

- Along scan line $\Delta x = 1$


```
y=y1;  
For (x=x1; x<=x2, ix++) {  
    write_pixel(x, round(y), line_color)  
    y+=k;  
}
```

Using Symmetry

- Use for $1 \geq k \geq 0$
- For $k > 1$, swap role of x and y
 - For each y , plot closest x



Otherwise we get
problem for steep
slopes



02. Rasterization, Depth Sorting and Culling:

Very Important!



- The problem with DDA is that it uses floats which was slow in the old days
- Bresenham's algorithm only uses integers

You do not need to know Bresenham's algorithm by heart. It is enough that you **understand** it if you see it.

Lecture 3: Shading

Lighting

$$\mathbf{i} = \mathbf{i}_{amb} + \mathbf{i}_{diff} + \mathbf{i}_{spec} + \mathbf{i}_{emission}$$



+



+



=




Know how to compute components.
Also, Blinns and Phong's highlight model

Lighting

Light:


- Ambient (r,g,b,a)
- Diffuse (r,g,b,a)
- Specular (r,g,b,a)



Material:

- Ambient (r,g,b,a)
- Diffuse (r,g,b,a)
- Specular (r,g,b,a)
- Emission (r,g,b,a) = "självlysande färg"

DIFFUSE	Base color
SPECULAR	Highlight Color
AMBIENT	Low-light Color
EMISSION	Glow Color
SHININESS	Surface Smoothness



03. Shading:

Lighting

$$\mathbf{i} = \mathbf{i}_{amb} + \mathbf{i}_{diff} + \mathbf{i}_{spec} + \mathbf{i}_{emission}$$

I.e.:

$$\mathbf{i} = \mathbf{i}_{amb} + \mathbf{i}_{diff} + \mathbf{i}_{spec} + \mathbf{i}_{emission}$$

$$\mathbf{i}_{amb} = \mathbf{m}_{amb} \otimes \mathbf{s}_{amb}$$

$$\mathbf{i}_{diff} = (\mathbf{n} \cdot \mathbf{l}) \mathbf{m}_{diff} \otimes \mathbf{s}_{diff}$$

$$((\mathbf{n} \cdot \mathbf{l}) < 0) \Rightarrow \mathbf{i}_{spec/diff} = 0$$

Phong's reflection model:

$$\mathbf{i}_{spec} = \max(0, (\mathbf{r} \cdot \mathbf{v}))^{m_{shi}} \mathbf{m}_{spec} \otimes \mathbf{s}_{spec}$$

Blinn's reflection model:

$$\mathbf{i}_{spec} = \max(0, (\mathbf{h} \cdot \mathbf{n}))^{m_{shi}} \mathbf{m}_{spec} \otimes \mathbf{s}_{spec}$$

$$\mathbf{i}_{emission} = \mathbf{m}_{emission}$$

03. Shading:

Diffuse component : i_{diff}



- $\mathbf{i} = \mathbf{i}_{amb} + \mathbf{i}_{diff} + \mathbf{i}_{spec} + \mathbf{i}_{emission}$



- Diffuse is Lambert's law: $i_{diff} = \mathbf{n} \cdot \mathbf{l} = \cos \phi$

- Photons are scattered equally in all directions

$$\mathbf{i}_{diff} = (\mathbf{n} \cdot \mathbf{l}) \mathbf{m}_{diff} \otimes \mathbf{s}_{diff}$$

(Note that \mathbf{n} and \mathbf{l} need to be normalized)

○ light source




Lighting

Specular component : i_{spec}



- Diffuse is dull (left)
- Specular: simulates a highlight

○ light source






Specular component: Phong

(**n** needs to be normalized)
 \mathbf{n}

- Phong specular highlight model
- Reflect \mathbf{l} around \mathbf{n} :

$$\mathbf{r} = -\mathbf{l} + 2(\mathbf{n} \cdot \mathbf{l})\mathbf{n}$$

$$i_{spec} = (\mathbf{r} \cdot \mathbf{v})^{m_{shi}} = (\cos \rho)^{m_{shi}}$$



$$\mathbf{i}_{spec} = \max(0, (\mathbf{r} \cdot \mathbf{v}))^{m_{shi}} \mathbf{m}_{spec} \otimes \mathbf{s}_{spec}$$

- Next: Blinns highlight formula: $(\mathbf{n} \cdot \mathbf{h})^m$



Halfway Vector

Blinn proposed replacing $\mathbf{v} \cdot \mathbf{r}$ by $\mathbf{n} \cdot \mathbf{h}$ where

$$\mathbf{h} = (\mathbf{l} + \mathbf{v}) / \|\mathbf{l} + \mathbf{v}\|$$

$(\mathbf{l} + \mathbf{v})/2$ is halfway between \mathbf{l} and \mathbf{v}


If \mathbf{n} , \mathbf{l} , and \mathbf{v} are coplanar:

$$\psi = \phi/2$$

Must then adjust exponent

so that $(\mathbf{n} \cdot \mathbf{h})^{e'} \approx (\mathbf{r} \cdot \mathbf{v})^e$

$$(e' \approx 4e)$$




$$\mathbf{i}_{spec} = \max(0, (\mathbf{h} \cdot \mathbf{n})^{m_{shi}}) \mathbf{m}_{spec} \otimes \mathbf{s}_{spec}$$

03. Shading:


Shading

- Three common types of shading:
 - Flat, Gouraud, and Phong
- In standard Gouraud shading the lighting is computed per triangle vertex and for each pixel, the color is interpolated from the colors at the vertices.
- In Phong Shading the lighting is not computed per vertex. Instead the normal is interpolated per pixel from the normals defined at the vertices and full lighting is computed per pixel using this normal. This is of course more expensive but looks better.

Flat





Gouraud



Phong





- Color of fog: \mathbf{c}_f color of surface: \mathbf{c}_s

$$\mathbf{c}_p = f\mathbf{c}_s + (1 - f)\mathbf{c}_f \quad f \in [0,1]$$

- How to compute f ?
- E.g., linearly:

$$f = \frac{Z_{end} - Z_p}{Z_{end} - Z_{start}}$$

Transparency and alpha

- Transparency
 - Very simple in real-time contexts
- The tool: alpha blending (mix two colors)
- Alpha (α) is another component in the frame buffer, or on triangle
 - Represents the opacity
 - 1.0 is totally opaque
 - 0.0 is totally transparent

- The over operator:

(Blending)

$$\mathbf{c}_o = \alpha \mathbf{c}_s + (1 - \alpha) \mathbf{c}_d$$


Rendered object

Transparency


- Need to sort the transparent objects
 - **First, render all non-transparent triangles as usual.**
 - **Then, sort all transparent triangles and render back-to-front with blending enabled. (and using standard depth test)**
 - **The reason is to avoid problems with the depth test and because the blending operation (i.e., over operator) is order dependent.**

Gamma correction

$$C = C_i^{(1/\gamma)}$$



- (Standard is 2.2)
 1. Screen has non-linear color intensity
 1. We often want linear output (e.g. antialiasing)
 2. Also happens to give more efficient color space when compressing intensity from 32-bit floats to 8-bits. Thus, often desired when storing textures.



better distribution for humans

Gamma of 2.2

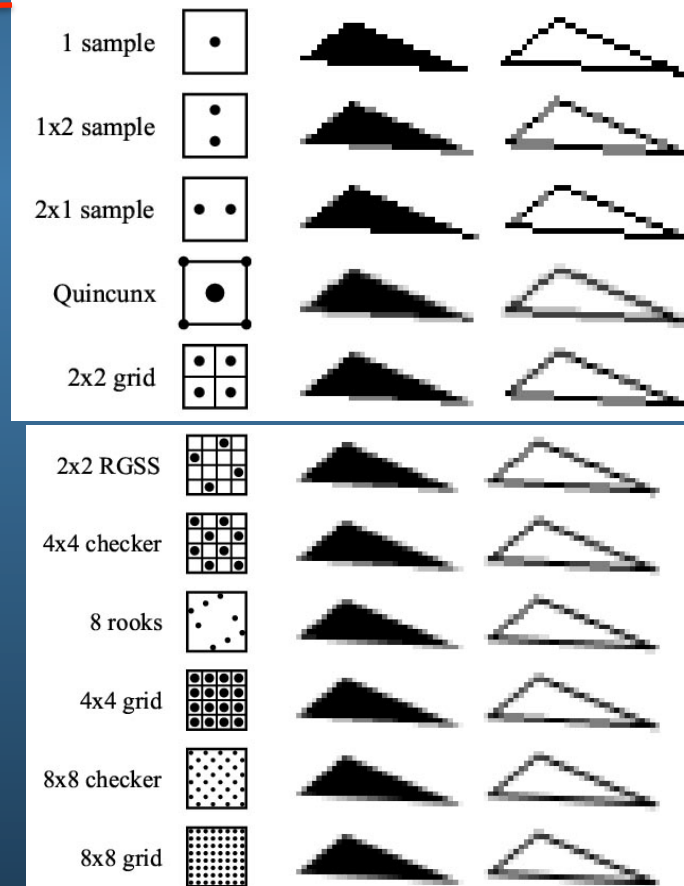

Truth

On most displays (those with gamma of about 2.2), one can observe that the linear intensity output (bottom) has a large jump in perceived brightness between the intensity values 0.0 and 0.1, while the steps at the higher end of the scale are hardly perceptible. A linear input that has a nonlinearly-increasing intensity (upper), will show much more even steps in perceived brightness.

Lecture 3.2: Sampling, filtering, and Antialiasing



- When does it occur?
 - In 1) pixels, 2) time, 3) texturing
- Supersampling schemes
- Jittered sampling
 - Why is it good?



04. Texturing



Most important:

- Texturing
 - Mipmapping: bi/tri-linear filtering, anisotropic filtering
- Environment mapping
 - Cube mapping
- Bump mapping
- 3D-textures,
- Particle systems
- Sprites and billboards

Filtering


FILTERING:

- For magnification: Nearest or Linear (box vs Tent filter)






- For minification: nearest, linear and...
 - Bilinear – using mipmapping
 - Trilinear – using mipmapping
 - Anisotropic – up to 16 mipmap lookups along line of anisotropy


Interpolation



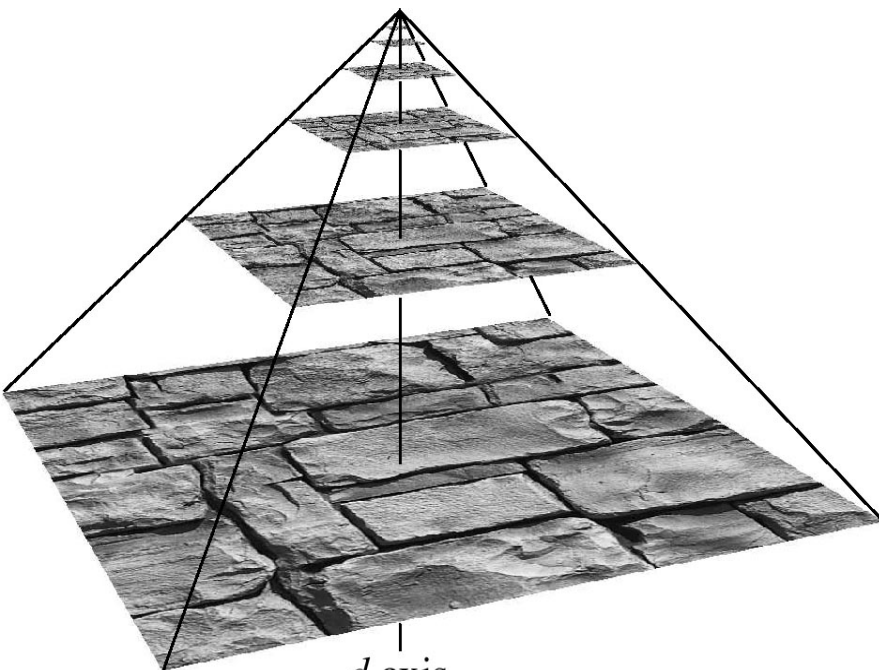

Magnification



Minification




Bilinear filtering using Mipmapping




Mipmapping

- Image pyramid
- Half width and height when going upwards
- Average over 4 "parent texels" to form "child texel"
- Depending on amount of minification, determine which image to fetch from
- Compute d first, gives two images
 - Bilinear interpolation in each



Mipmapping



- Interpolate between those bilinear values
 - Gives trilinear interpolation



- Constant time filtering: 8 texel accesses


Mipmapping: Memory requirements

- Not twice the number of bytes...!




- Rather 33% more – not that much


Anisotropic texture filtering



pixel space




texture space





See page
168-169

Environment mapping



- Assumes the environment is infinitely far away
- Sphere mapping, or Cube mapping
- Cube mapping is the norm nowadays



Cube mapping




- Simple math: compute reflection vector, \mathbf{r}
- Largest abs-value of component, determines which cube face.
 - Example: $\mathbf{r}=(5,-1,2)$ gives POS_X face
- Divide \mathbf{r} by $\text{abs}(5)$ gives $(u,v)=(-1/5,2/5)$
- Also remap from $[-1,1]$ to $[0,1]$ by $(u,v) = ((u,v)+\text{vec2}(1,1))*0.5$;
- Your hardware does all the work for you. You just have to compute the reflection vector.

Bump mapping

- by Blinn in 1978






- Inexpensive way of simulating wrinkles and bumps on geometry
 - Expensive to model these geometrically
- Instead let a texture modify the normal at each pixel, and then use this normal to compute lighting per pixel



3D Textures

- 3D textures:
 - Texture filtering is no longer trilinear
 - Rather quadlinear (linear interpolation 4 times)
 - Enables new possibilities
 - Can store light in a room, for example



05. Texturing:

Just know what “sprites” is and that they are very similar to a billboard

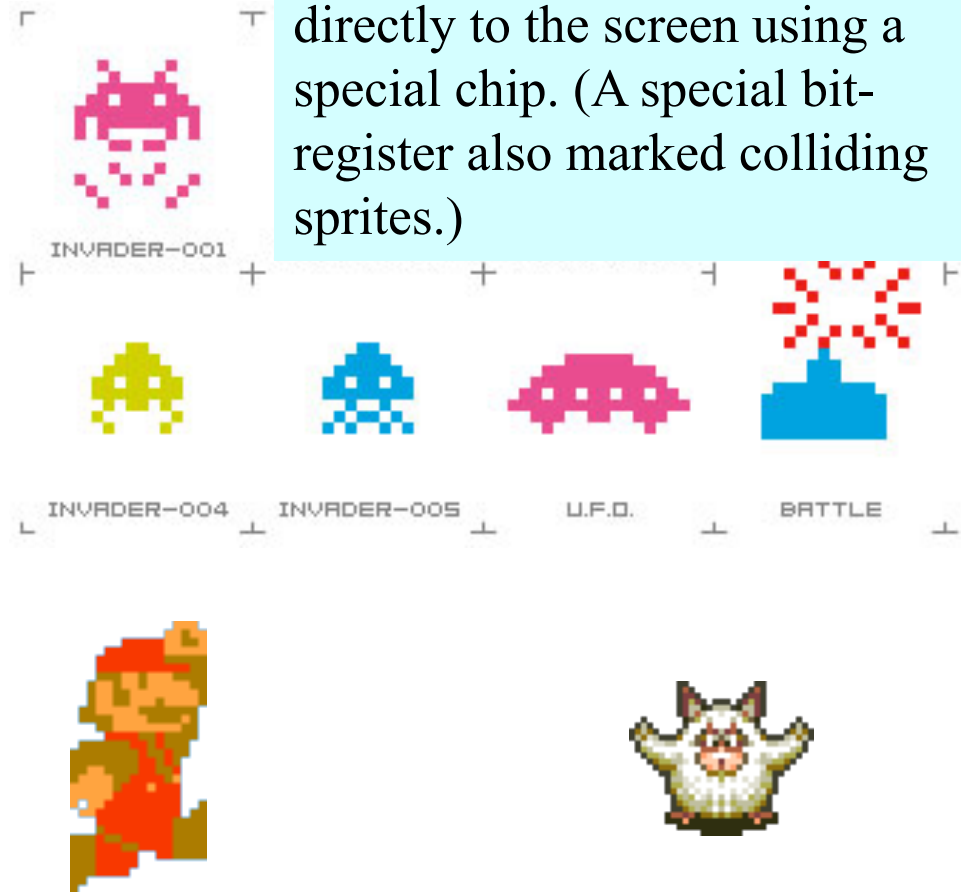
Sprites

Sprites (=älvor) was a technique on older home computers, e.g. VIC64. As opposed to billboards sprites does not use the frame buffer. They are rasterized directly to the screen using a special chip. (A special bit-register also marked colliding sprites.)

```
GLbyte M[64]=
{ 127,0,0,127, 127,0,0,127,
  127,0,0,127, 127,0,0,127,
  0,127,0,0, 0,127,0,127,
  0,127,0,127, 0,127,0,0,
  0,0,127,0, 0,0,127,127,
  0,0,127,127, 0,0,127,0,
  127,127,0,0, 127,127,0,127,
  127,127,0,127, 127,127,0,0};
```

```
void display(void) {
  glClearColor(0.0,1.0,1.0,1.0);
  glClear(GL_COLOR_BUFFER_BIT);
  glEnable (GL_BLEND);
  glBlendFunc (GL_SRC_ALPHA,
               GL_ONE_MINUS_SRC_ALPHA);
  glRasterPos2d(xpos1,ypos1);
  glPixelZoom(8.0,8.0);
  glDrawPixels(width,height,
               GL_RGBA, GL_BYTE, M);

  glPixelZoom(1.0,1.0);
  glutSwapBuffers();
}
```





Billboards

- 2D images used in 3D environments
 - Common for trees, explosions, clouds, lens flares



Billboards



- Rotate them towards viewer
 - Either by rotation matrix or
 - by orthographic projection

Billboards

- Fix correct transparency by blending AND using alpha-test

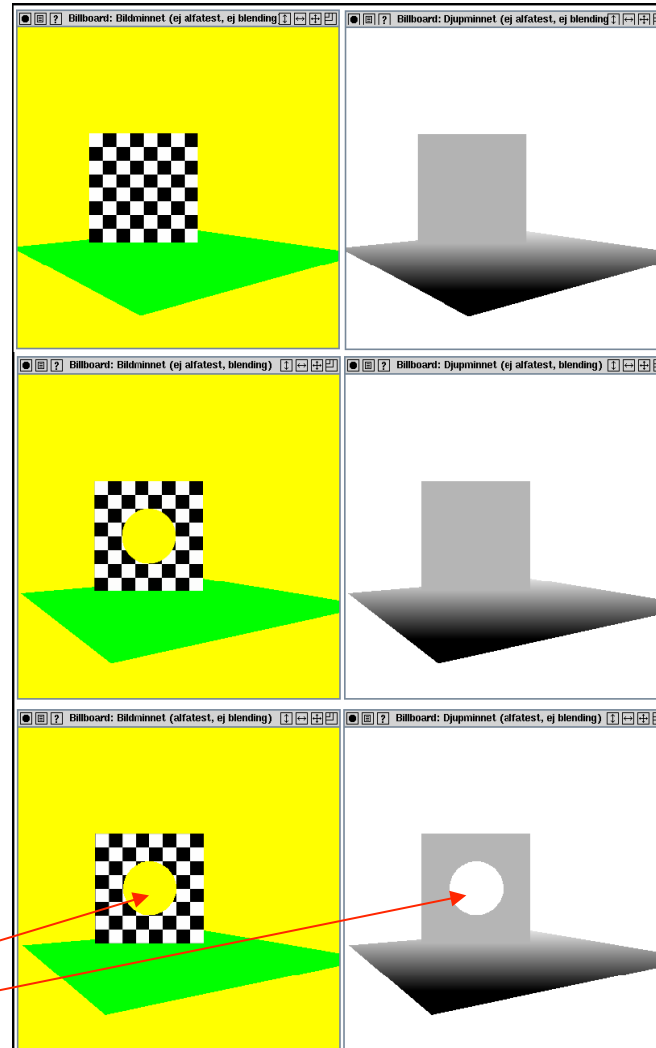
- In fragment shader:
if (color.a < 0.1) discard;

If alpha value in texture is lower than this threshold value, the pixel is not rendered to. I.e., neither frame buffer nor z-buffer is updated, which is what we want to achieve.

E.g. here: so that objects behind is visible through the hole

Color Buffer

Depth Buffer



With
blending

With
alpha test



(Also called *Impostors*)



axial billboarding

The rotation axis is fixed and
disregarding the view position

Lecture 5: OpenGL

- How to use OpenGL (or DirectX)
 - Will not ask about syntax. Know how to use.
 - I.e. functionality
 - E.g. how to achieve
 - Blending and transparency
 - Fog – how would you implement in a fragment shader?
 - pseudo code is enough
 - Specify a material, a triangle, how to translate or rotate an object.
 - Triangle – vertex order and facing

Reflections with environment mapping

- Understand at pseudo code level!

VERTEX SHADER


```
in vec3      vertex;
in  vec3     normalIn;    // The normal
out vec3     normal;
out vec3     eyeVector;
uniform mat4 normalMatrix;
uniform mat4 modelViewMatrix;
uniform mat4 modelViewProjectionMatrix;

void main()
{
    gl_Position = modelViewProjectionMatrix * vec4(vertex,1);
    normal = (normalMatrix * vec4(normalIn,0.0)).xyz;
    eyeVector = (modelViewMatrix * vec4(vertex, 1)).xyz;
}
```

FRAGMENT SHADER

```
in vec3 normal;
in vec3 eyeVector;
uniform samplerCube tex1;
out vec4 fragmentColor;

void main()
{
    vec3 reflectionVector = normalize(reflect(normalize(eyeVector),
                                                normalize(normal)));
    fragmentColor = texture(tex1, reflectionVector);
}
```



Buffers

- Frame buffer
 - Back/front/left/right – **glDrawBuffers()**
 - Offscreen buffers (e.g., framebuffer objects, auxiliary buffers)

Frame buffers can consist of:

- Color buffer - rgb(a)
- Depth buffer (z-buffer)
 - For correct depth sorting
 - Instead of BSP-algorithm or painters algorithm...
- Stencil buffer
 - E.g., for shadow volumes or only render to frame buffer where stencil = certain value (e.g., for masking).

Lecture 6: Intersection Tests


- 4 techniques to compute intersections:
 - Analytically
 - Geometrically – e.g. ray vs box (3 slabs)
 - SAT (Separating Axis Theorem) for convex polyhedra
Test:
 1. axes orthogonal to face of A,
 2. axes orthogonal to face of B
 3. All different axes formed by crossprod of one edge of A and one of B
 - Dynamic tests – know what it means.
- E.g., describe an algorithm for intersection between a **ray** and a
 - Polygon, triangle, sphere and plane.
- Know equations for ray, sphere, cylinder, plane, triangle

Analytical: Ray/plane intersection

- Ray: $r(t) = o + td$
- Plane formula: $n \cdot p + d = 0$



- Replace p by $r(t)$ and solve for t :
 $n \cdot (o + td) + d = 0$
 $n \cdot o + tn \cdot d + d = 0$
 $t = (-d - n \cdot o) / (n \cdot d)$

Here, one scalar equation and one unknown \rightarrow just solve for t .



Analytical: Ray/sphere test

- Sphere center: \mathbf{c} , and radius r
- Ray: $\mathbf{r}(t) = \mathbf{o} + t\mathbf{d}$
- Sphere formula: $\|\mathbf{p} - \mathbf{c}\| = r$
- Replace \mathbf{p} by $\mathbf{r}(t)$: $\|\mathbf{r}(t) - \mathbf{c}\| = r$



$$(\mathbf{r}(t) - \mathbf{c}) \cdot (\mathbf{r}(t) - \mathbf{c}) - r^2 = 0$$

$$(\mathbf{o} + t\mathbf{d} - \mathbf{c}) \cdot (\mathbf{o} + t\mathbf{d} - \mathbf{c}) - r^2 = 0$$


$$(\mathbf{d} \cdot \mathbf{d})t^2 + 2((\mathbf{o} - \mathbf{c}) \cdot \mathbf{d})t + (\mathbf{o} - \mathbf{c}) \cdot (\mathbf{o} - \mathbf{c}) - r^2 = 0$$

$$t^2 + 2((\mathbf{o} - \mathbf{c}) \cdot \mathbf{d})t + (\mathbf{o} - \mathbf{c}) \cdot (\mathbf{o} - \mathbf{c}) - r^2 = 0 \quad \|\mathbf{d}\| = 1$$

This is a standard quadratic equation. Solve for t .

Geometrical: Ray/Box Intersection (2)

- Intersect the 2 planes of each slab with the ray



- Keep max of t^{\min} and min of t^{\max}
- If $t^{\min} < t^{\max}$ then we got an intersection
- Special case when ray parallel to slab

$$\text{Plane : } \pi : \mathbf{n} \cdot \mathbf{p} + d = 0$$

Point/Plane

- Insert a point \mathbf{x} into plane equation:

$$f(\mathbf{x}) = \mathbf{n} \cdot \mathbf{x} + d = ?$$

$$f(\mathbf{x}) = \mathbf{n} \cdot \mathbf{x} + d = 0 \quad \text{for } \mathbf{x}'\text{s on the plane}$$

$$f(\mathbf{x}) = \mathbf{n} \cdot \mathbf{x} + d < 0 \quad \text{for } \mathbf{x}'\text{s on one side of the plane}$$

$$f(\mathbf{x}) = \mathbf{n} \cdot \mathbf{x} + d > 0 \quad \text{for } \mathbf{x}'\text{s on the other side}$$

Negative
half space

Positive
half space




Sphere/Plane AABB/Plane

$$\text{Plane : } \pi : \mathbf{n} \cdot \mathbf{p} + d = 0$$

$$\text{Sphere : } \mathbf{c} \quad r$$

$$\text{Box : } \mathbf{b}^{\min} \quad \mathbf{b}^{\max}$$


- Sphere: compute $f(\mathbf{c}) = \mathbf{n} \cdot \mathbf{c} + d$
- $f(\mathbf{c})$ is the signed distance (\mathbf{n} normalized)
- $\text{abs}(f(\mathbf{c})) > r$ no collision
- $\text{abs}(f(\mathbf{c})) = r$ sphere touches the plane
- $\text{abs}(f(\mathbf{c})) < r$ sphere intersects plane
- Box: insert all 8 corners
- If all f 's have the same sign, then all points are on the same side, and no collision



AABB/plane

$$\text{Plane : } \pi : \mathbf{n} \cdot \mathbf{p} + d = 0$$
$$\text{Sphere : } \mathbf{c} \quad r$$
$$\text{Box : } \mathbf{b}^{\min} \quad \mathbf{b}^{\max}$$

- The smart way (shown in 2D)
- Find the two vertices that have the most positive and most negative value when tested against the plane



$$\mathbf{v}_{pos_x} = (\mathbf{n}_x > 0) ? \mathbf{b}_{max_x} : \mathbf{b}_{min_x}$$

$$\mathbf{v}_{pos_y} = (\mathbf{n}_y > 0) ? \mathbf{b}_{max_y} : \mathbf{b}_{min_y}$$

$$\mathbf{v}_{pos_z} = (\mathbf{n}_z > 0) ? \mathbf{b}_{max_z} : \mathbf{b}_{min_z}$$

$$\mathbf{v}_{neg_x} = (\mathbf{n}_x < 0) ? \mathbf{b}_{max_x} : \mathbf{b}_{min_x}$$


$$\mathbf{v}_{neg_y} = (\mathbf{n}_y < 0) ? \mathbf{b}_{max_y} : \mathbf{b}_{min_y}$$

$$\mathbf{v}_{neg_z} = (\mathbf{n}_z < 0) ? \mathbf{b}_{max_z} : \mathbf{b}_{min_z}$$

OBB almost as easy. Just first project \mathbf{n} on OBB's axes – see p: 757

Another analytical example: Ray/ Triangle in detail



- Ray: $\mathbf{r}(t) = \mathbf{o} + t\mathbf{d}$
- Triangle vertices: $\mathbf{v}_0, \mathbf{v}_1, \mathbf{v}_2$
- A point in the triangle:
- $\mathbf{t}(u, v) = \mathbf{v}_0 + u(\mathbf{v}_1 - \mathbf{v}_0) + v(\mathbf{v}_2 - \mathbf{v}_0) =$
 $= (1 - u - v)\mathbf{v}_0 + u\mathbf{v}_1 + v\mathbf{v}_2 \quad [u, v \geq 0, u + v \leq 1]$
- Set $\mathbf{t}(u, v) = \mathbf{r}(t)$, and solve!




$$\begin{pmatrix} | & | & | \\ -\mathbf{d} & \mathbf{v}_1 - \mathbf{v}_0 & \mathbf{v}_2 - \mathbf{v}_0 \\ | & | & | \end{pmatrix} \begin{pmatrix} t \\ u \\ v \end{pmatrix} = \begin{pmatrix} | \\ \mathbf{o} - \mathbf{v}_0 \\ | \end{pmatrix}$$

Ray/Polygon: very briefly

- Intersect ray with polygon plane
- Project from 3D to 2D
- How?
- Find $\max(|n_x|, |n_y|, |n_z|)$
- Skip that coordinate!
- Then, count crossing in 2D



View frustum testing example



- Algorithm:

- if sphere is outside any of the 6 frustum planes -> report "outside".
- Else report intersect.

- Not exact test, but not incorrect, i.e.,

- A sphere that is reported to be inside, can be outside
- Not vice versa, so test is conservative


Lecture 7.1: Spatial Data Structures and Speed-Up Techniques

- Speed-up techniques

- Culling


- Backface
- View frustum (hierarchical)
- Portal
- Occlusion Culling
- Detail

- Levels-of-detail:





- How to construct and use the spatial data structures

- BVH, BSP-trees (polygon aligned + axis aligned)



Axis Aligned Bounding Box Hierarchy - an example

- Assume we click on screen, and want to find which object we clicked on



•
click!


- 1) Test the root first
- 2) Descend recursively as needed
- 3) Terminate traversal when possible

In general: get $O(\log n)$ instead of $O(n)$


How to create a BVH? Example: using AABBs

AABB = Axis Aligned
Bounding Box
BVH = Bounding Volume
Hierarchy



- Find minimal box, then split along longest axis





x is longest




Find minimal
boxes



Split along
longest axis



Find minimal
boxes


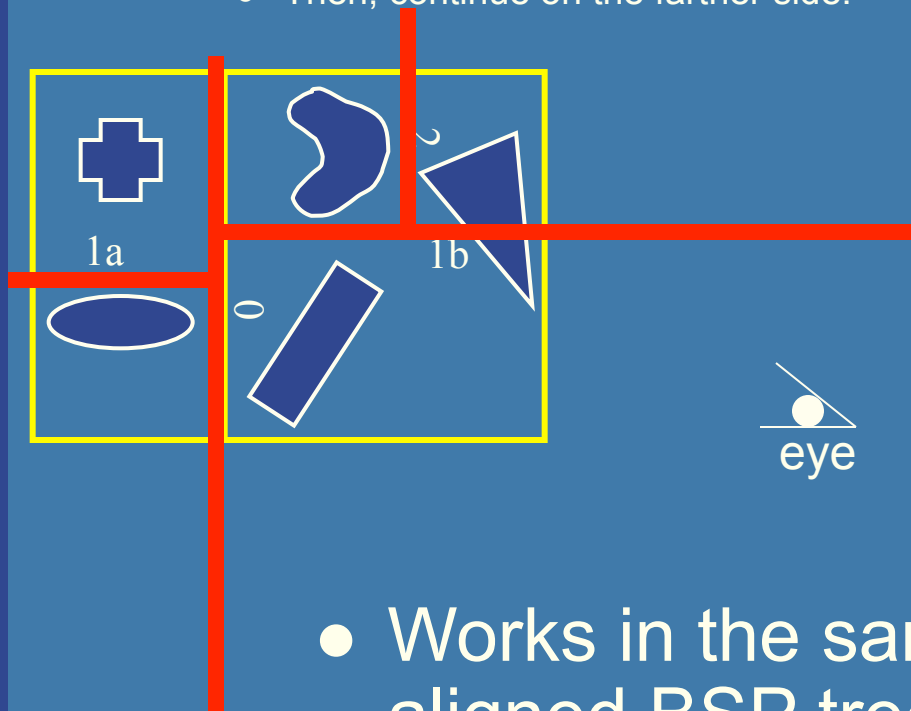


Called TOP-DOWN method
Similar for other BVs

Axis-aligned BSP tree

Rough sorting

- Test the planes, recursively from root, against the point of view. For each traversed node:
 - If node is leaf, draw the node's geometry
 - else
 - Continue traversal on the "hither" side with respect to the eye to sort front to back
 - Then, continue on the farther side.




- Works in the same way for polygon-aligned BSP trees --- but that gives exact sorting

Polygon-aligned BSP tree

- Allows exact sorting
- Very similar to axis-aligned BSP tree
 - But the splitting plane are now located in the planes of the triangles

```
Drawing Back-to-Front {  
  recurse on farther side of P;  
  Draw P;  
  Recurse on hither side of P;  
} // farther/hither is with respect to eye pos.
```





Know how to build it and how to traverse back-to-front or front-to-back with respect to the eye position (here: v)

A Scene Graph is a hierarchical scene description – more typically a **logical** hierarchy (than e.g. **spatial**)


Scene graphs

- BVH is the data structure that is used most often
 - Simple to understand
 - Simple code
- However, BVH stores just geometry
 - Rendering is more than geometry
- The scene graph is an extended BVH with:
 - Lights
 - Materials
 - Transforms
 - several connections to a node
 - And more



Lecture 7.2: Collision Detection

- 3 types of algorithms:
 - With rays
 - Fast but not exact
 - With BVH
 - Slower but exact
 - You should be able to write pseudo code for BVH/BVH test for coll det between two objects.
 - For many many objects.
 - why? Course pruning of "obviously" non-colliding objects
 - Sweep-and-prune



Lecture 8: Ray tracing

- Adaptive Super Sampling

- Jittering 

- How to stop ray tracing recursion?

- Speedup techniques




- Spatial data structures

- Optimizations for BVHs: skippointer tree
- Ray BVH-traversal
- (You do not need to learn the **ray traversal** algorithms for Grids nor AA-BSP trees)

- Shadow cache

- Material (Fresnel: metall, dielectrics)


- Constructive Solid Geometry – how to implement




Adaptive Supersampling

Pseudo code:

```
Color AdaptiveSuperSampling() {  
  - Make sure all 5 samples exist  
    • (Shoot new rays along diagonal if necessary)  
  - Color col = black;  
  - For each quad i  
    • If the colors of the 2 samples are fairly similar  
      - col += (1/4)*(average of the two colors)  
    • Else  
      - col +=(1/4)* adaptiveSuperSampling(quad[i])  
  - return col;  
}
```





Summary of the Ray tracing- algorithm:




- **main()-calls trace()** for each pixel
- **trace():** should return color of closest hit point along ray.
 1. calls findClosestIntersection()
 2. If any object intersected → call shade().
- **Shade():** should compute color at hit point
 1. For each light source, shoot shadow ray to determine if light source is visible
If not in shadow, compute diffuse + specular contribution.
 2. Compute ambient contribution
 3. Call trace() recursively for the reflection- and refraction ray.

Data structures

- Octree





- Kd tree




Kd-tree = Axis-Aligned BSP tree with fixed recursive split plane order (e.g. x,y,z,x,y,z...)

- Grids

Including mail boxing




Hierarchical grid



Recursive grid

- Bounding box hierarchies




Lecture 10 – Global Illumination


- Global illumination:

$$L_o = L_e + \int_{\Omega} f_r(\mathbf{x}, \omega, \omega') L_i(\mathbf{x}, \omega') (\omega' \cdot \mathbf{n}) d\omega'$$


- Why is not standard ray tracing enough?
- rendering eq., BRDFs
- Monte Carlo Ray Tracing
- Path tracing
- Photon mapping



Ray tracing




Global
Illumination



Lecture 10: What you need to know

- The rendering equation
 - Be able to explain all its components
- Path tracing
 - Why it is good, compared to naive monte-carlo sampling
 - The overall algorithm (on a high level as in these slides).
- Photon Mapping
 - The overall algorithm. See the summary slide on:
 - Creating Photon Maps...
 - Ray trace from eye...
 - Growing spheres...
- Final Gather
 - Why it is good. How it works:
 - At the first diffuse hit, instead of using global map directly, sample indirect slow varying light around \mathbf{p} by sampling the hemisphere with ~ 1000 rays and use the two photon maps where those rays hit a diffuse surface.
- Bidirectional Path Tracing, Metropolis Light Transport
 - Just their names. Don't need to know the algorithms.


Monte Carlo Ray Tracing (naively)



- Sample indirect illumination by shooting sample rays over the hemisphere, at each hit.
 - At some recursion depth, stop and compute standard local lighting (i.e., without indirect illumination)

Monte Carlo Ray Tracing (naively)


- This gives a ray tree with most rays at the bottom level. This is bad since these rays have the lowest influence on the pixel color.



PathTracing


– a smarter Monte-Carlo ray-tracing method

- Path Tracing instead only traces one of the possible ray paths at a time. This is done by randomly selecting only one sample direction at a bounce. Hundreds of paths per pixel are traced.



Equally number of rays are traced at each level

Path Tracing – indirect + direct illumination



- Shoot many paths per pixel (the image just shows one light path).
 - At each intersection,
 - Shoot one shadow ray per light source
 - at random position on light, for area/volumetric light sources
 - and randomly select one new ray direction.

Photon Mapping

- **Creating Photon Maps:**

- Trace photons (~100K-1M) from light source. Store them in kd-tree when they hit diffuse surface. Then, use russian roulette to decide if the photon should be absorbed or specularly or diffusively reflected. Create both global map and caustics map. For the Caustics map, we send more of the photons towards reflective/refractive objects.


- **Ray trace from eye:**

- As usual: I.e., shooting primary rays and recursively shooting reflection/refraction rays, and at each intersection point \mathbf{p} , compute direct illumination (shadow rays + shading).
- Also grow sphere around each \mathbf{p} in caustics map to get caustics contribution and in global map to get slow-varying indirect illumination.
- If final gather is used: At the first diffuse hit, instead of using global map directly, sample indirect slow varying light around \mathbf{p} by sampling the hemisphere with ~100 – 1000 rays and use the two photon maps where those rays hit a surface.

- **Growing sphere:**

- Uses the kd-tree to expand a sphere around \mathbf{p} until a fixed amount (e.g. 50) photons are inside the sphere. The radius is a measure of the intensity of indirect light at \mathbf{p} . The BRDF at \mathbf{p} could also be used to get a more accurate color and intensity value.

A modification for indirect illumination – Final Gather



- Too noisy to use the global map for direct visualization
- Remember: eye rays are recursively traced (via reflections/refractions) until a diffuse hit, p . There, we want to estimate slow-varying indirect illumination.
 - Instead of growing sphere in global map at p , Final Gather shoots 100-1000 indirect rays from p and grows sphere in the global map and also caustics map, where each of those rays end at a diffuse surface.

Lecture 11: Shadows + Reflection

- Point light / Area light
- Three ways of thinking about shadows
 - The basis for different algorithms.
- Shadow mapping
 - Be able to describe the algorithm
- Shadow volumes
 - Be able to describe the algorithm
 - Stencil buffer, 3-pass algorithm, Z-pass, Z-fail,
 - Creating quads from the silhouette edges as seen from the light source, etc
- Pros and cons of shadow volumes vs shadow maps
- Planar reflections – how to do. Why not using environment mapping?



Ways of thinking about shadows

- As separate objects (like Peter Pan's shadow) **This corresponds to planar shadows**
- As volumes of space that are dark
 - **This corresponds to shadow volumes**
- As places not seen from a light source looking at the scene. **This corresponds to shadow maps**
- Note that we already "have shadows" for objects facing away from light

Shadow Maps - Summary

Shadow Map Algorithm:

- Render a z-buffer from the light source
 - Represents geometry in light
- Render from camera
 - For every fragment:
 - Transform(warp) its 3D-pos (x,y,z) into shadow map (i.e. light space) and compare depth with the stored depth value in the shadow map
 - If depth greater \rightarrow point in shadow
 - Else \rightarrow point in light
 - Use a bias at the comparison




Shadow Map (=depth buffer)

Understand z-fighting and light leaks



Bias

- Need a tolerance threshold (depth bias) when comparing depths to avoid surface self shadowing





Bias

- Need a tolerance threshold (depth bias) when comparing depths to avoid surface self shadowing




Bias

- Need a tolerance threshold (depth bias) when comparing depths to avoid surface self shadowing



Shadow volumes

Create shadow quads for all silhouette edges (as seen from the light source).
(The normals are pointing outwards from the shadow volume.)




Edges between one triangle front facing the light source and one triangle back facing the light source are considered silhouette edges.

Then...

Shadow Volumes - concept

- Perform counting with the stencil buffer
 - Render front facing shadow quads to the stencil buffer
 - Inc stencil value, since those represents entering shadow volume
 - Render back facing shadow quads to the stencil buffer
 - Dec stencil value, since those represents exiting shadow volume




- No updating of z-buffer
- Z-test is enabled as usual

Shadow Volumes with the Stencil Buffer

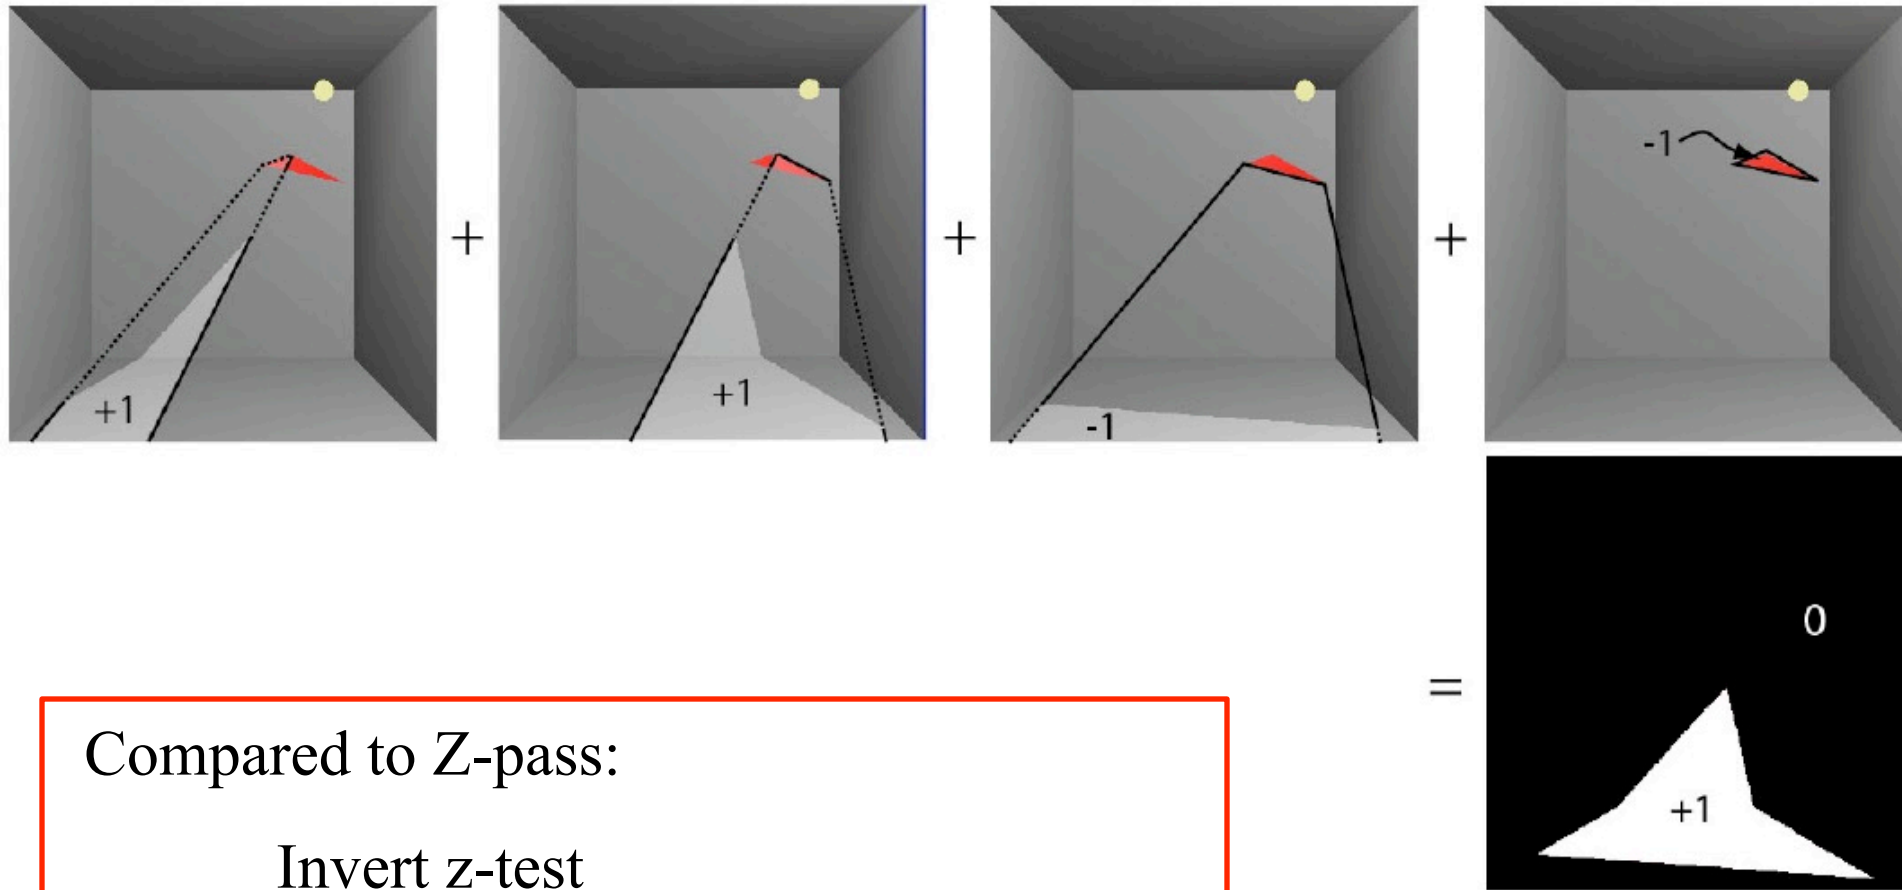
- A three pass process:
 - **1st pass:** Render *ambient* lighting
 - **2nd pass:**
 - Draw to stencil buffer only
 - Turn off updating of z-buffer and writing to color buffer but still use standard depth test
 - Set stencil operation to
 - » *incrementing* stencil buffer count for *frontfacing* shadow volume quads, and
 - » *decrementing* stencil buffer count for *backfacing* shadow volume quads
 - **3rd pass:** Render *diffuse and specular* where stencil buffer is 0.

The Z-fail Algorithm

- Z-pass must offset the stencil buffer with the number of shadow volumes that the eye is inside. Problematic.
- Count to infinity instead of to the eye
 - We can choose any reference location for the counting
 - A point in light avoids any offset
 - Infinity is always in light – if we cap the shadow volumes at infinity



Z-fail by example



Compared to Z-pass:


Invert z-test

Invert stencil inc/dec

I.e., count to infinity instead of from eye.


Merging Volumes

- Edge shared by two polygons facing the light creates front and backfacing quad.




This interior edge makes two quads, which cancel out

Silhouette Edges




From the light's view, caster interior edges do not contribute to the shadow volume.



Finding the silhouette edge gets rid of many useless shadow volume polygons.

Shadow Maps vs Shadow Volumes



Shadow Maps


- *Good:* Handles any rasterizable geometry, **constant cost** regardless of complexity, map can sometimes be reused. **Very fast.**
- *Bad:* Frustum limited. **Jagged shadows** if res too low, **biasing** headaches.
 - Solution:
 - 6 SM (cube map), high res., use filtering (huge topic)



Shadow Volumes

- *Good:* shadows are **sharp**. Handles omni-directional lights.
- *Bad:* **3 passes**, shadow polygons must be generated and rendered → lots of polygons & **fill**
 - Solution: culling & clamping

Planar reflections



Two methods:


1. Reflecting the object:

- If reflection plane is $z=0$ (else somewhat more complicated – see page 387)
 - Apply `glScalef(1,1,-1)` ;
- Backfacing becomes front facing!
 - i.e., use frontface culling instead of backface culling
- Lights should be reflected as well


2. Reflecting the camera in the reflection plane

Planar reflections

- Assume plane is $z=0$
- Then apply `glScalef(1,1,-1)` ;
- Effect:




Planar reflections



- How should you render?
- 1) the reflective ground plane polygons into the stencil buffer
- 2) the scaled $(1, 1, -1)$ model, but mask with stencil buffer
 - Reflect light pos as well
 - Use front face culling
- 3) the ground plane (semi-transparent)
- 4) the unscaled model

Continuity



- A) Non-continuous
- B) C^0 -continuous
- C) G^1 -continuous
- D) C^1 -continuous
- (C^2 -continuous)


See page 585-587 in
Real-time Rendering,
3rd ed.

Types of Curves

- The types of curves:


- Interpolating

- Blending polynomials (or cubic parametric polynomials) for interpolation of 4 control points (fit curve to 4 control points)




- Hermite

- fit curve to 2 control points + 2 derivatives (tangents)



- Bezier

- 2 interpolating control points + 2 intermediate points to define the tangents



- B-spline

- To get C² continuity

- NURBS

- Different weights of the control points and
- The control points can be at non-uniform intervalls

Goods and bads with these curves.

Splines and Basis

- If we examine the cubic B-spline from the perspective of each control (data) point, each interior point contributes (through the blending functions) to four segments
- We can rewrite $p(u)$ in terms of the data points as

$$p(u) = \sum B_i(u) p_i$$

defining the basis functions $\{B_i(u)\}$

B-Splines

These are our control points, p_0 - p_8 , to which we want to approximate a curve





Illustration of how the control points are evenly (uniformly) distributed along the parameterisation u of the curve $p(u)$.


In each point $p(u)$ of the curve, for a given u , the point is defined as a weighted sum of the closest 4 surrounding points. Below are shown the weights for each point along $u=0 \rightarrow 1$




SUMMARY

B-Splines

In each point $p(u)$ of the curve, for a given u , the point is defined as a weighted sum of the closest 4 surrounding points. Below are shown the weights for each point along $u=0 \rightarrow 1$



The weight function (blend function) $B_{p_i}(u)$ for a point p_i can thus be written as a translation of a basis function $B(t)$. $B_{p_i}(u) = B(u-i)$




Our complete B-spline curve $p(u)$ can thus be written as:

$$p(u) = \sum B_i(u) p_i$$

12. Curves and Surfaces:

NURBS



NURBS is similar to B-Splines except that:

1. The control points can have different weights, w_i , (higher weight makes the curve go closer to that control point)
2. The control points do not have to be at uniform distances ($u=0, 1, 2, 3, \dots$) along the parameterisation u . E.g.: $u=0, 0.5, 0.9, 4, 14, \dots$

NURBS = Non-Uniform Rational B-Splines

The NURBS-curve is thus defined as:


$$\mathbf{p}(u) = \frac{\sum_{i=0}^{n-1} B_i(u) w_i \mathbf{p}(i)}{\sum_{i=0}^{n-1} B_i(u) w_i}$$

Division with the sum of the weights, to make the combined weights sum up to 1, at each position along the curve. Otherwise, a translation of the curve is introduced (which is not desirable)

12. Curves and Surfaces:

NURBS

- Allowing control points at non-uniform distances means that the basis functions $B_{p_i}()$ are being stretched and non-uniformly located.
- E.g.:



Each curve $B_{p_i}()$ should of course look smooth and C^2 –continuous. But it is not so easy to draw smoothly by hand... (The sum of the weights are still =1 due to the division in previous slide)

Lecture 13:

- Perspective correct interpolation (e.g. for textures)

- Taxonomy:

- Sort first
- sort middle
- sort last fragment
- sort last image

- Bandwidth

- Why it is a problem
- How to "solve" it

- Be able to sketch the architecture of a modern graphics card

Linearly interpolate $(u_i/w_i, v_i/w_i, 1/w_i)$ in screenspace from each triangle vertex i .

Then at each pixel:

$$u_{ip} = (u/w)_{ip} / (1/w)_{ip}$$
$$v_{ip} = (v/w)_{ip} / (1/w)_{ip}$$

where ip = screen-space interpolated value between the triangle vertices.

