

Real-Time Rendering (Echtzeitgraphik)



Dr. Michael Wimmer
wimmer@cg.tuwien.ac.at



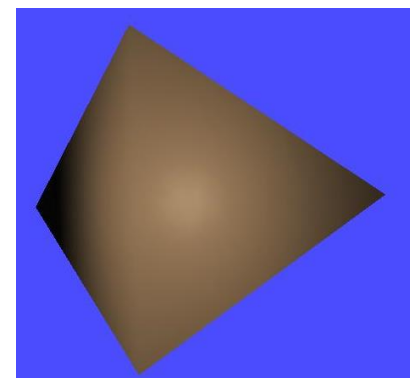

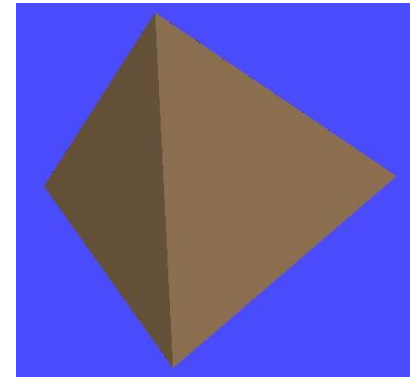
Texturing




- OpenGL lighting refresher
- Texture Spaces
- Texture Aliasing and Filtering
- Multitexturing
 - Lightmapping
- Texture Coordinate Generation
- Projective Texturing
- Multipass Rendering



- Flat shading
 - compute light interaction per polygon
 - the whole polygon has the same color
- Gouraud shading
 - compute light interaction per vertex
 - interpolate the colors
- Phong shading
 - interpolate normals per pixel
- Remember: difference between
 - Phong Light Model
 - Phong Shading




- Phong light model at each vertex (glLight, ...)
- Local model only (no shadows, radiosity, ...)
- ambient + diffuse + specular (glMaterial!)



- Fixed function: Gouraud shading
 - Note: need to interpolate specular separately!
- Phong shading: calculate Phong model in fragment shader



- Idea: enhance visual appearance of plain surfaces by applying fine structured details




- Basis for most real-time rendering effects
- Look and feel of a surface
- Definition:
 - A *regularly sampled function* that is *mapped* onto every *fragment* of a surface
 - Traditionally an image, but...
- Can hold arbitrary information
 - Textures become general data structures
 - Will be interpreted by fragment programs
 - Can be rendered into → important!




- Spatial Layout
 - 1D, 2D, 3D
 - Cube Maps
- Formats (too many), e.g. OpenGL
 - LUMINANCE16_ALPHA16: 32bit = 2 x 16 bit bump map
 - RGBA4: 16bit = 4 x 4 colors
 - RGBA_FLOAT32: 128 bit = 4 x 32 bit float
 - compressed formats, high dynamic range formats, ...



Texturing: General Approach




Texture space (u, v)



Object space (x_0, y_0, z_0)



Image Space (x_I, y_I)



Modeling

Object space
(x, y, z, w)

Parameter Space
(s, t, r, q)

Texture Space
(u, v)

Rendering

Texture
projection


Texture
function



Where do texture coordinates come from?

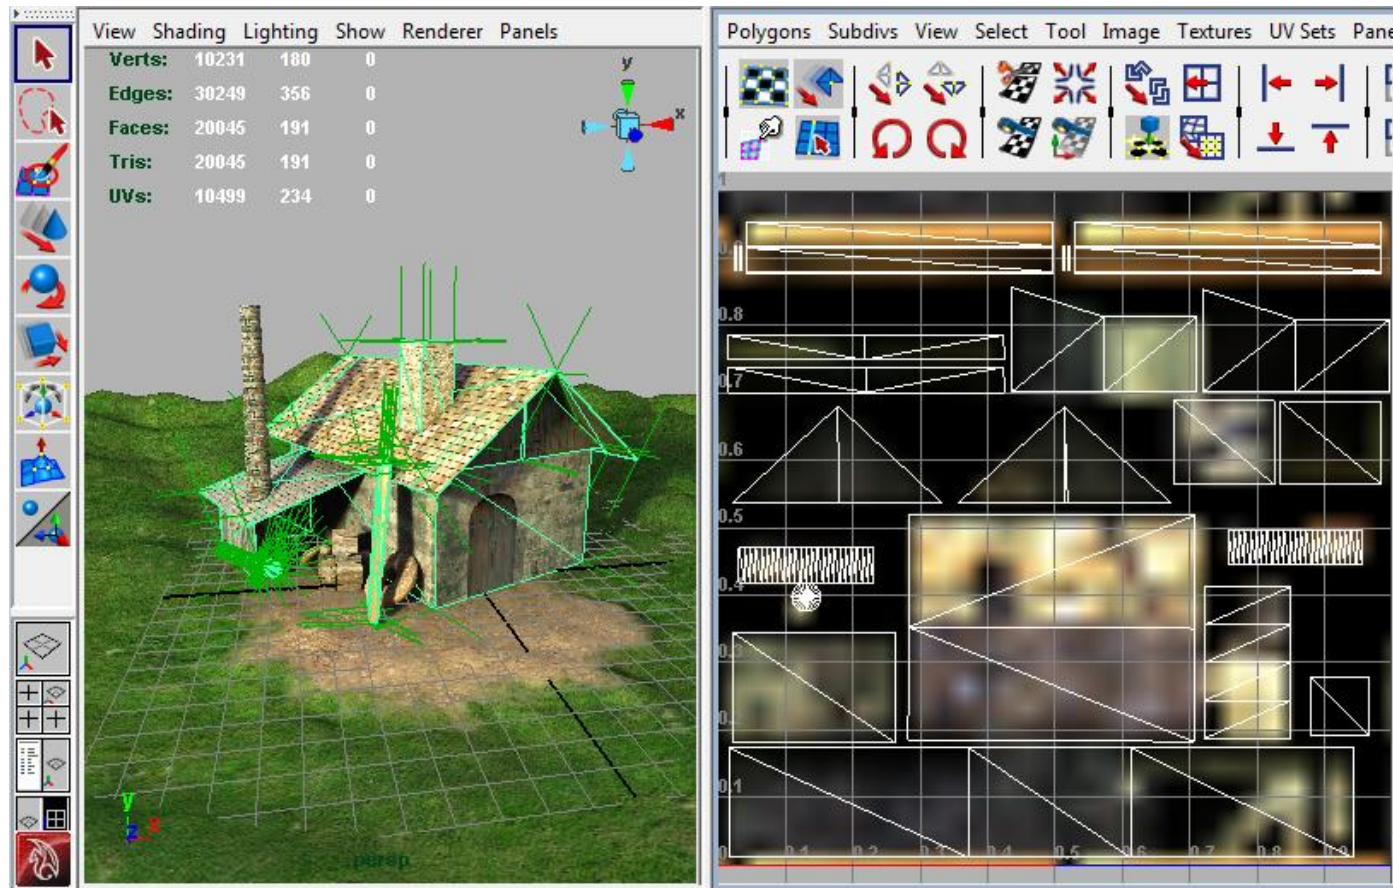
- **Online**: texture matrix/texcoord generation
- **Offline**: manually (or by modeling prog)

spherical cylindrical planar natural




Where do texture coordinates come from?

- **Offline:** manual UV coordinates by DCC program
- **Note:** a modeling Problem!




- How to extend texture beyond the border?
- Border and repeat/clamp modes
- Arbitrary $(s,t,\dots) \rightarrow [0,1] \rightarrow [0,255] \times [0,255]$


repeat




mirror/repeat






clamp




border



- Problem: One pixel in image space covers many texels



- Caused by *undersampling*: texture information is lost



Texture space

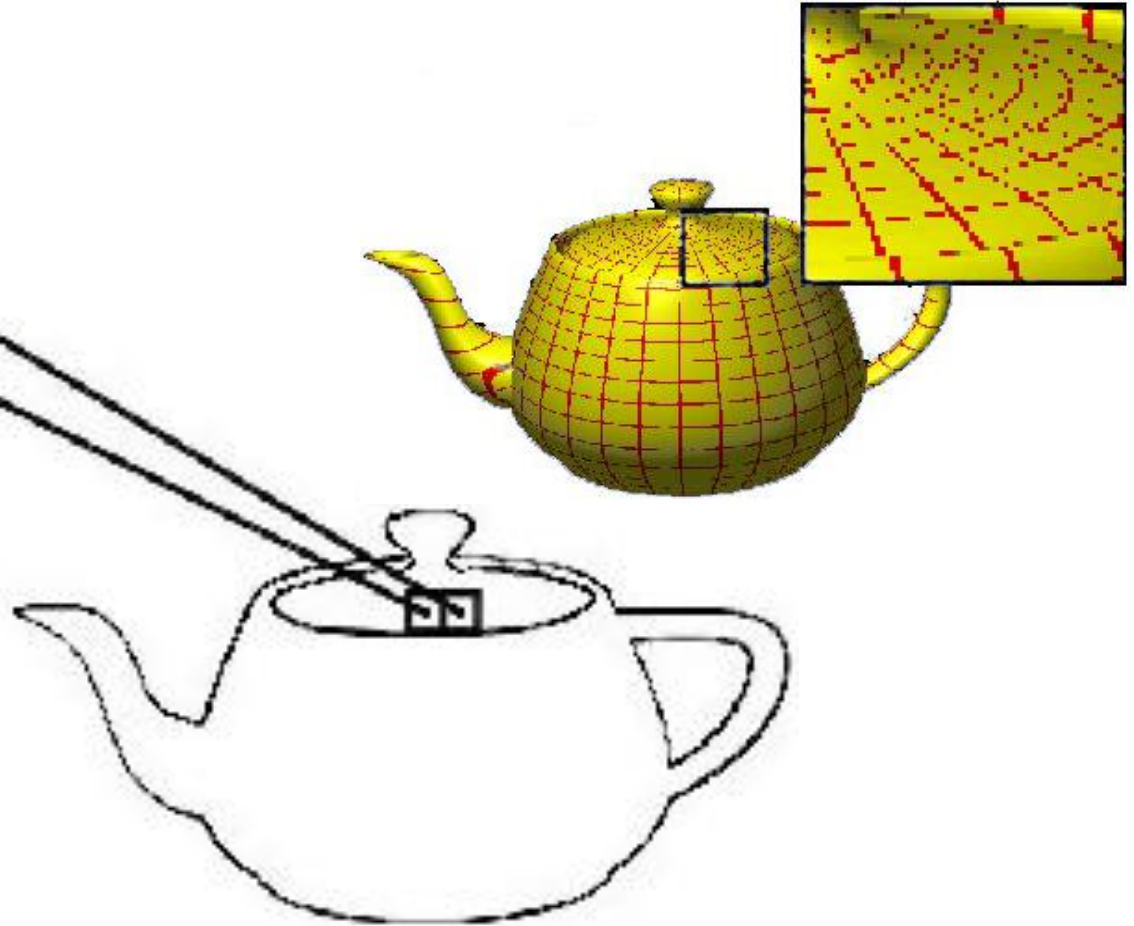
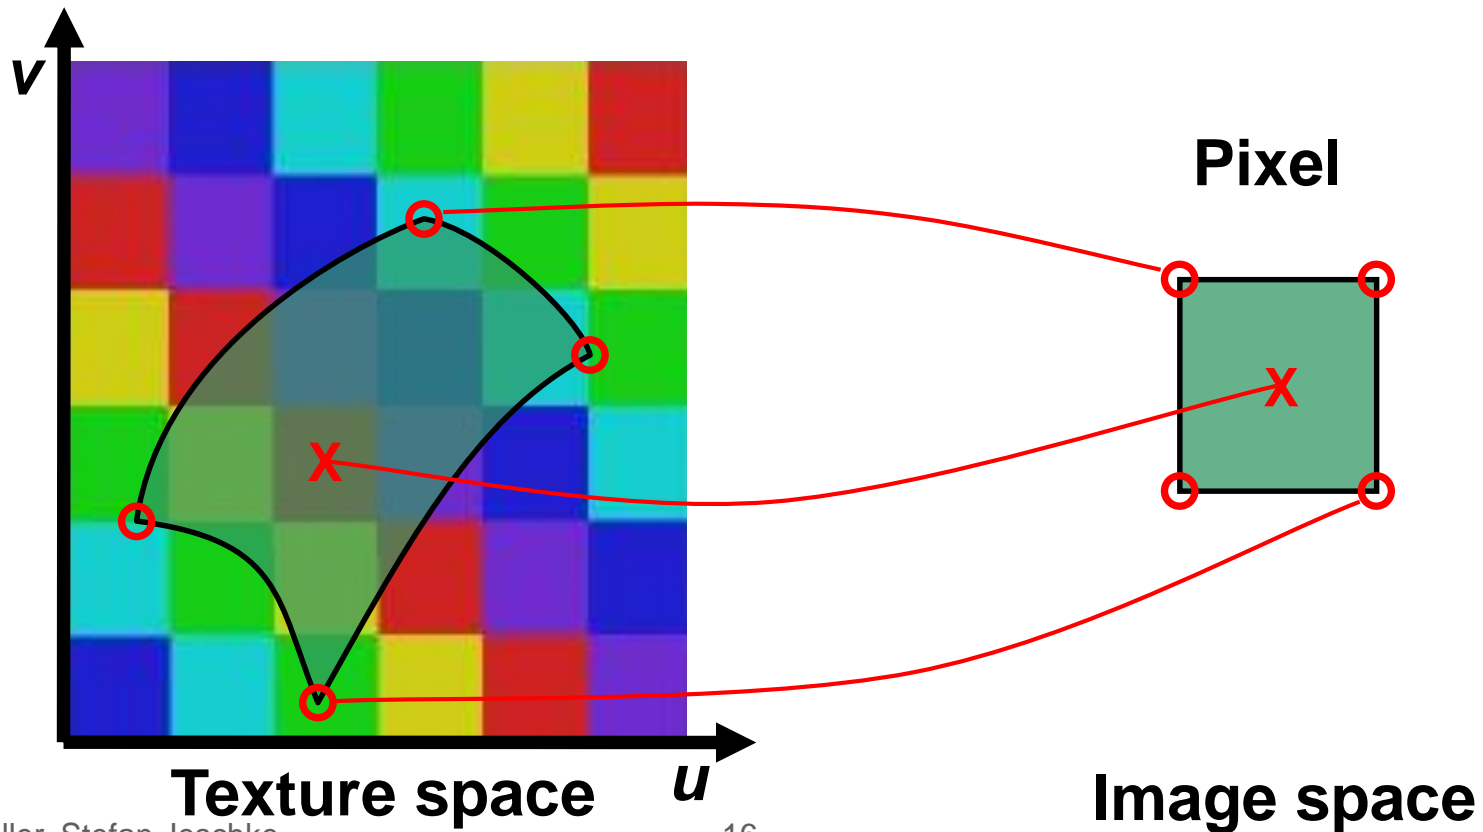


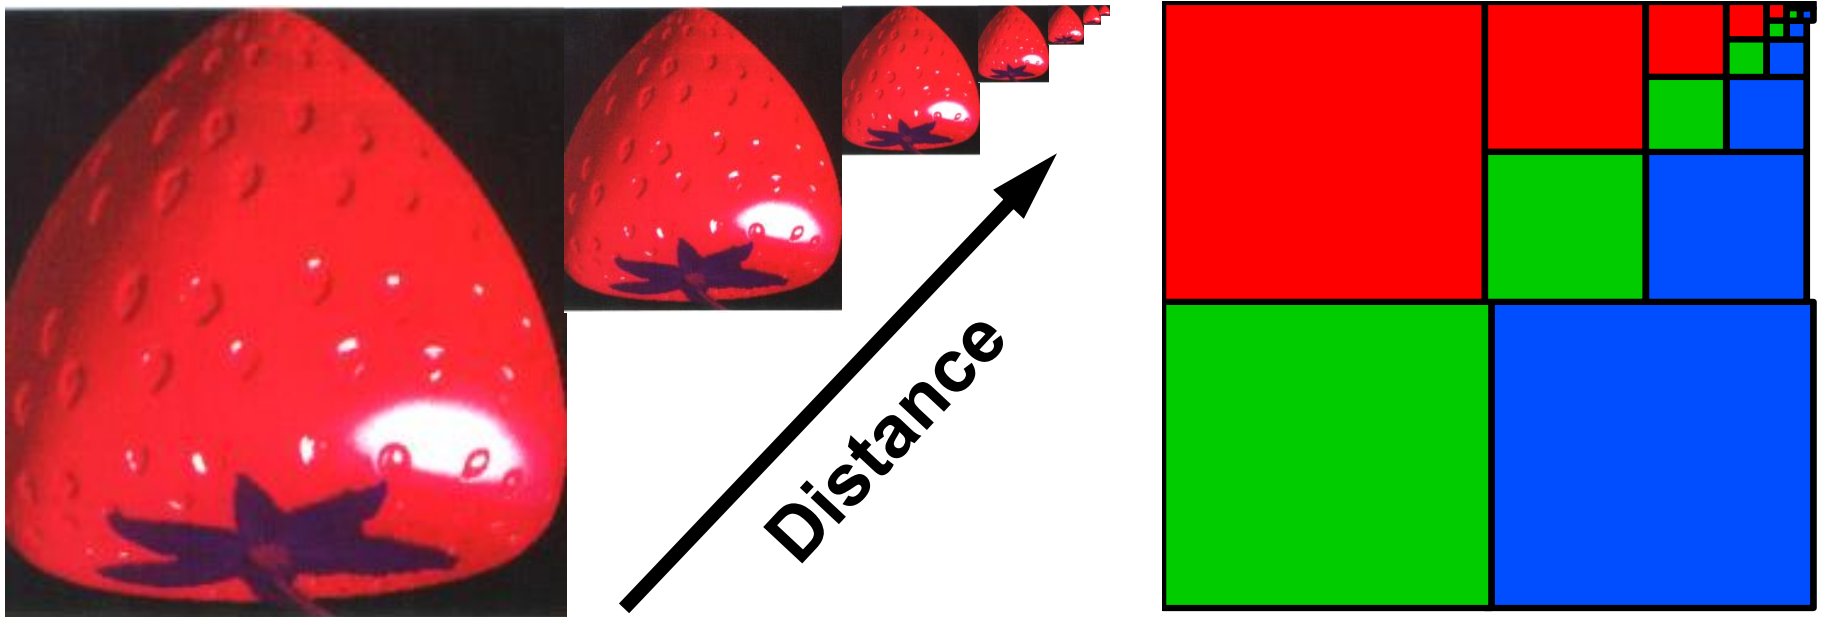
Image space



- 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* = "much info on a small area")
 - Simple (4 pixel average) and memory efficient
 - Last image is only ONE texel





- 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*




Bilinear interpolation




Trilinear interpolation

- Trilinear interpolation:


- $T_1 :=$ value from texture $D_1 = D_0 + 1$ (bilin.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)



- Other example for bilinear vs. trilinear filtering



- Bilinear reconstruction for texture magnification ($D < 0$) ("upsampling")
 - Weight adjacent texels by distance to pixel position



$$\begin{aligned} T(u+du, v+dv) &= du \cdot dv \cdot T(u+1, v+1) \\ &+ du \cdot (1-dv) \cdot T(u+1, v) \\ &+ (1-du) \cdot dv \cdot T(u, v+1) \\ &+ (1-du) \cdot (1-dv) \cdot T(u, v) \end{aligned}$$





Original image



Nearest neighbor





Bilinear filtering



■ Anisotropic Filtering


- View dependent filter kernel
- Implementation: *summed area table*, "*RIP Mapping*", "*footprint assembly*", "*sampling*"



Texture space



- Example



- Everything is done in hardware, nothing much to do!
- `gluBuild2DMipmaps ()` generates MIPmaps
- Set parameters in `glTexParameter ()`
 - `GL_LINEAR_MIPMAP_NEAREST`
 - `GL_TEXTURE_MAG_FILTER`
- Anisotropic filtering is an extension:
 - `GL_EXT_texture_filter_anisotropic`
 - Number of samples can be varied (4x,8x,16x)
 - Vendor specific support and extensions



- Fourier Transform of signal \rightarrow frequency space („spectrum“)
- Multiplication (mul) in primary space = Convolution (conv) in frequency space
- Typical signals and their spectra:
 - Box \leftrightarrow $\sin(x)/x$ („sinc“)
 - Gaussian \leftrightarrow Gaussian
 - Impulse train \leftrightarrow Impulse train
 - Width inverse proportional!



- Initial Sampling
- Resampling
- Display



- Input: continuous signal
 - Nature or computer generated
- Bandlimiting: remove high frequencies
 - conv sinc \leftrightarrow mul box
 - Happens in camera optics, lens of eye, or antialiasing (direct convolution, supersampling)
- Sampling:
 - mul impulse train \leftrightarrow conv impulse train
 - Leads to replica of spectra!
- Result: image or texture



- Input: Samples = discrete signal (usually texture)
- Reconstruction:
 - conv sinc \leftrightarrow mul box
 - „Removes“ replica of spectrum in sampled repr.
- Bandlimiting:
 - Only required if new sampling frequency is lower!
 - Typically through mipmapping
- Sampling
- Result: another texture or final image (=frame buffer)



- Input: Samples (from frame buffer)
- Reconstruction
 - Using display technology (e.g. CRT: Gaussian!)
- Result: continuous signal (going to eye)





- Practice: substitute sinc by Gaussian
 - sinc has negative values
 - Gaussian can be cut off gracefully
- „Reconstruction“ is really an interpolation!
 - Reconstruction \neq Antialiasing!
- Aliasing: overlap of signal replica in sampling
 - Bandlimiting = Antialiasing
- Magnification \rightarrow reconstruction only
- Minification \rightarrow bandlimiting + reconstruction



- Supersampling
- Multisampling (MSAA): combines
 - Supersampling (for edges)
 - Texture filtering (for textures)
 - Only one shader evaluation per final pixel
- Morphological Antialiasing (FXAA, SMAA, ...):
 - Postprocess
 - Analyzes image, recovers edges, antialiases them



- 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



- Simple(!) texture environment example:

```
glActiveTexture(GL_TEXTURE1);  
glTexEnvf(GL_TEXTURE_ENV, ...  
...   GL_TEXTURE_ENV_MODE, GL_COMBINE);  
...   GL_COMBINE_RGB, GL_MODULATE);  
...   { GL_SOURCE1_RGB, GL_TEXTURE);  
...     GL_OPERAND1_RGB, GL_SRC_COLOR);  
...     GL_SOURCE2_RGB, GL_PREVIOUS);  
...     GL_OPERAND2_RGB, GL_SRC_COLOR); }
```



$$C = CT_1 \cdot CT_0$$

- Programmable shading makes this easier!



- 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





Original scene




Light-mapped




- Precomputation based on non-realtime methods
 - Radiosity
 - Raytracing
 - Monte Carlo Integration
 - Pathtracing
 - Photonmapping





Lightmap



mapped






Original scene




Light-mapped

- Special case of light mapping
- Cos-weighted visibility to environment modulates intensity:



- Darker where more occluded
- Option: “per object” lightmap
 - Allows to move object






Model/Texture: Rendermonkey




- Map generation:
 - Use single map for group of coplanar polys
 - Lightmap UV coordinates need to be in $(0..1) \times (0..1)$
- Map application:
 - Premultiply textures by light maps
 - Why is this not appealing?
 - Multipass with framebuffer blend
 - Problems with specular
 - Multitexture
 - Fast, flexible




- Why premultiplication is bad...



Full Size Texture
(with Lightmap)



Tiled Surface Texture
plus Lightmap



→ use tileable surface textures and low resolution lightmaps



- DCC programs (*Blender*, *Maya*...)
- Game Engines (*Irrlicht*)
- Light Map Maker (free)

- Ambient Occlusion:
 - xNormal



- Specified manually (*glMultiTexCoord()*)
- Using classical OpenGL texture coordinate generation
 - Linear: from object or eye space vertex coords
 - Special texturing modes (env-maps)
 - Can be further modified with texture matrix
 - E.g., to add texture animation
 - Can use 3rd or 4th texture coordinate for projective texturing!
- Shader allows complex texture lookups!



- Specify a “plane” (i.e., a 4D-vector) for each coordinate (s,t,r,q)
- Example: $s = p_1 x + p_2 y + p_3 z + p_4 w$

```
GLfloat Splane[4] = { p1, p2, p3, p4 };  
glTexGenfv(GL_S, GL_EYE_PLANE, Splane);  
glEnable(GL_TEXTURE_GEN_S);
```

- Think of this as a matrix T with plane parameters as row vectors



- Object-linear:

$$\begin{bmatrix} s \\ t \\ r \\ q \end{bmatrix} = T \begin{bmatrix} x \\ y \\ z \\ w \end{bmatrix} \quad \text{object}$$

- Eye-linear:

$$T_e = T \cdot M^{-1}$$

(M...Modelview matrix at time of specification!)

- Effect: uses coordinate space at time of specification!

- Eye: M=identity
- World: M=view-matrix

$$\begin{bmatrix} s \\ t \\ r \\ q \end{bmatrix} = T_e \begin{bmatrix} x \\ y \\ z \\ w \end{bmatrix} \quad \text{eye}$$




- Classic OpenGL
 - Can specify an arbitrary 4x4 Matrix, each frame!
 - `glMatrixMode(GL_TEXTURE);`
 - There is also a texture matrix stack!
- Shaders allow arbitrary dynamic calculations with uv-coordinates
 - Many effects possible:
 - Flowing water, conveyor belts, distortions etc.





Projective Texturing




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

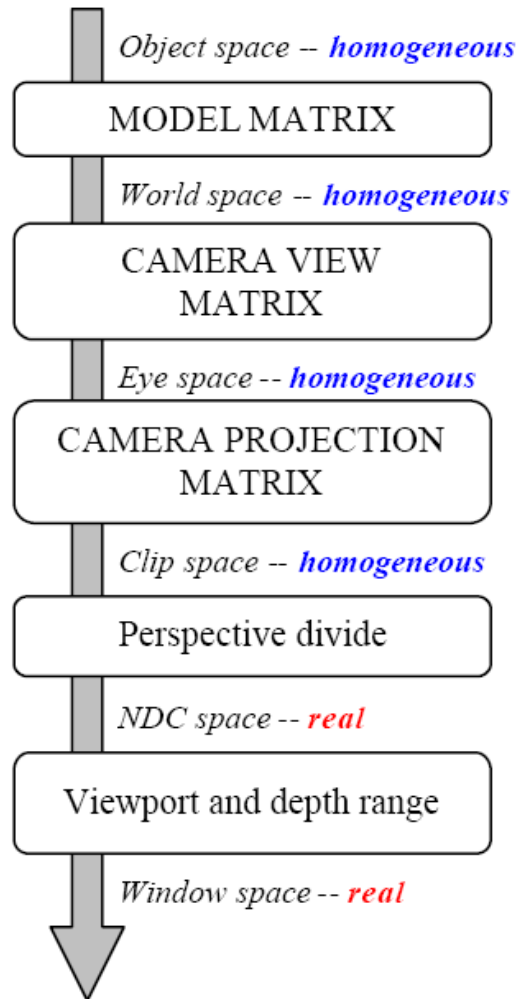




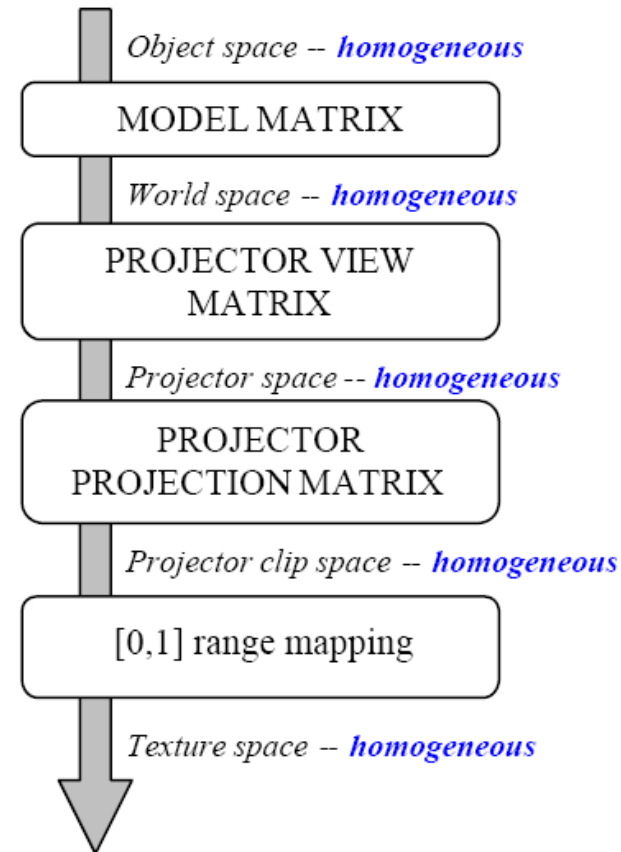
- Map vertices to light frustum
 - Option 1: from object space
 - Option 2: from eye space
- Projection
(perspective transform)




Camera



Projector




- OpenGL does not store Modeling Matrix
- No notion of world space!



- Version 1: transforming object space coordinates
 - Disadvantage: need to provide model matrix for each object in shader!
 - Classic OpenGL: even more difficult!

$$\begin{bmatrix} s \\ t \\ r \\ q \end{bmatrix} = \begin{bmatrix} 1/2 & 1/2 \\ & 1/2 & 1/2 \\ & & 1/2 & 1/2 \\ & & & 1 \end{bmatrix} \begin{bmatrix} \text{Light} \\ \text{(projection)} \\ \text{matrix} \end{bmatrix} \begin{bmatrix} \text{Light} \\ \text{view} \\ \text{(look at)} \\ \text{matrix} \end{bmatrix} \begin{bmatrix} \text{Modeling} \\ \text{matrix} \end{bmatrix} \begin{bmatrix} x_o \\ y_o \\ z_o \\ w_o \end{bmatrix}$$

Map [-1..1]
to [0..1]



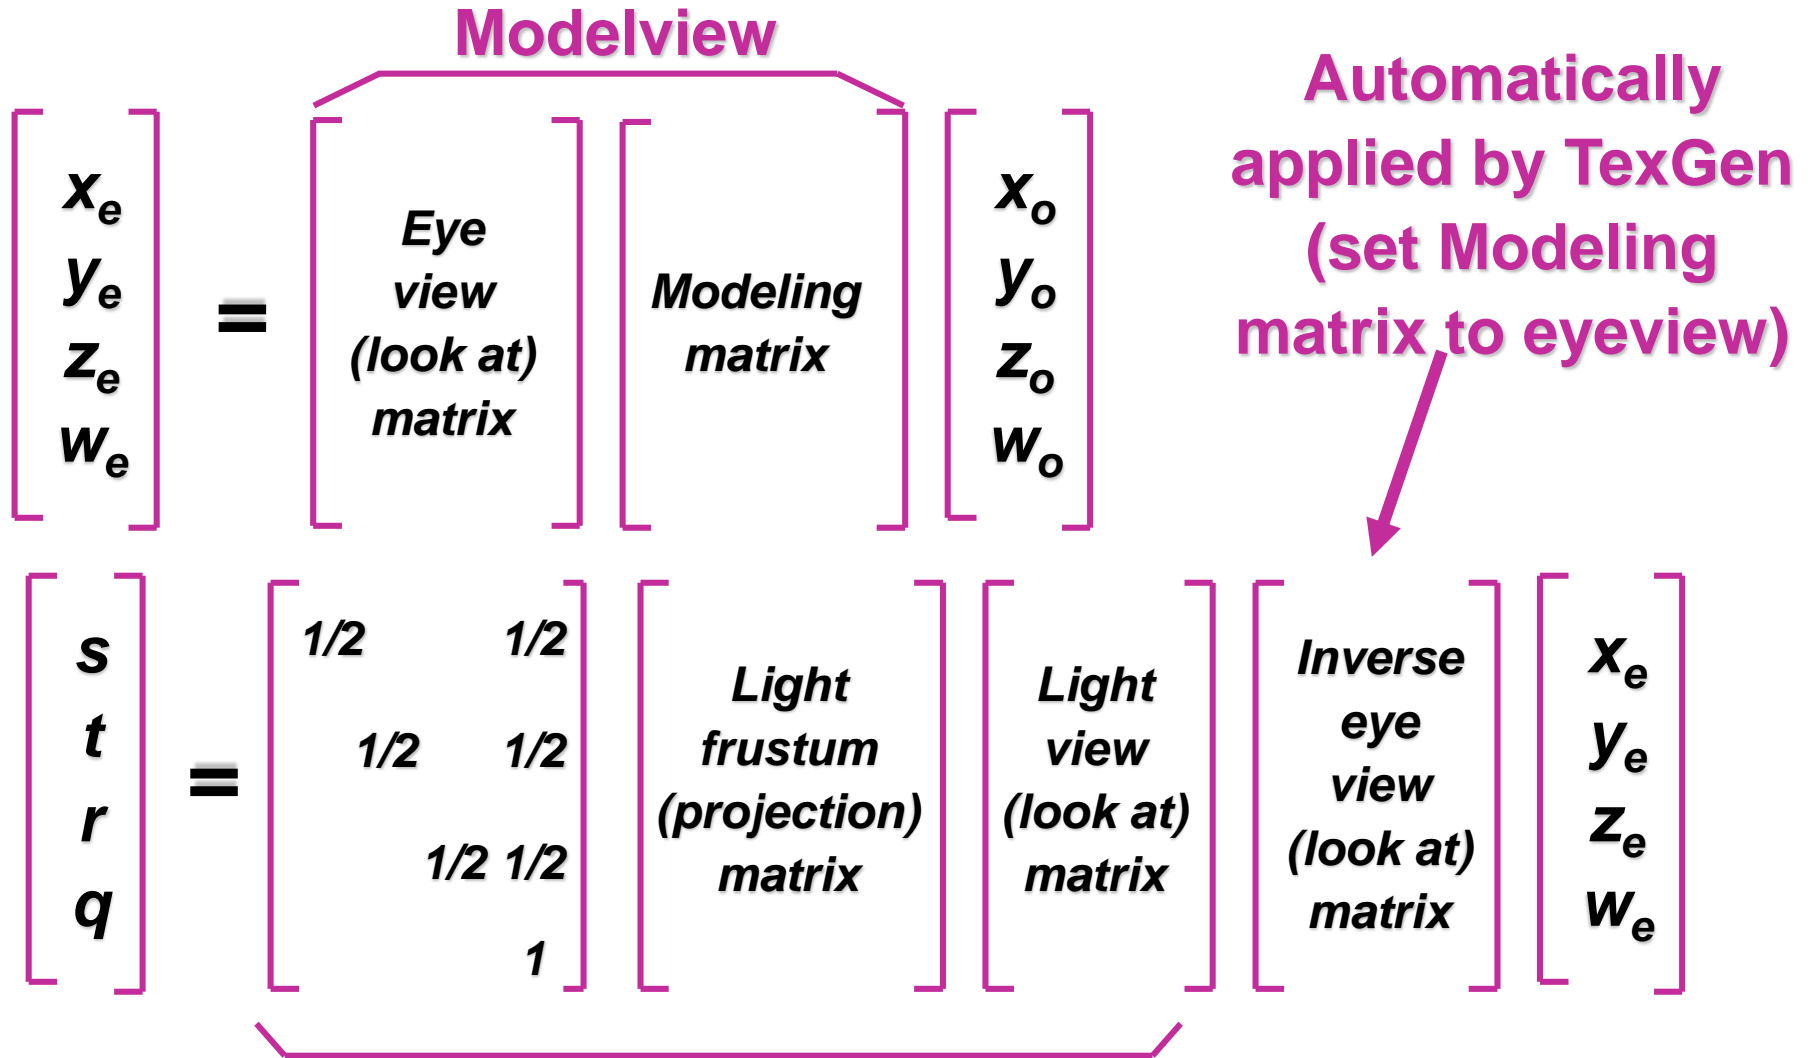
T



- Version 2: transforming eye space coordinates
 - Advantage: matrix works for all objects!

$$\begin{bmatrix} s \\ t \\ r \\ q \end{bmatrix} = \underbrace{\begin{bmatrix} 1/2 & & & & \\ & 1/2 & & & \\ & & 1/2 & & \\ & & & 1/2 & \\ & & & & 1 \end{bmatrix} \begin{bmatrix} \text{Light} \\ \text{(projection)} \\ \text{matrix} \end{bmatrix} \begin{bmatrix} \text{Light} \\ \text{view} \\ \text{(look at)} \\ \text{matrix} \end{bmatrix} \begin{bmatrix} \text{Inverse} \\ \text{eye} \\ \text{view} \\ \text{(look at)} \\ \text{matrix} \end{bmatrix} \begin{bmatrix} x_e \\ y_e \\ z_e \\ w_e \end{bmatrix} }_T$$





Supply this combined transform to *glTexGen*

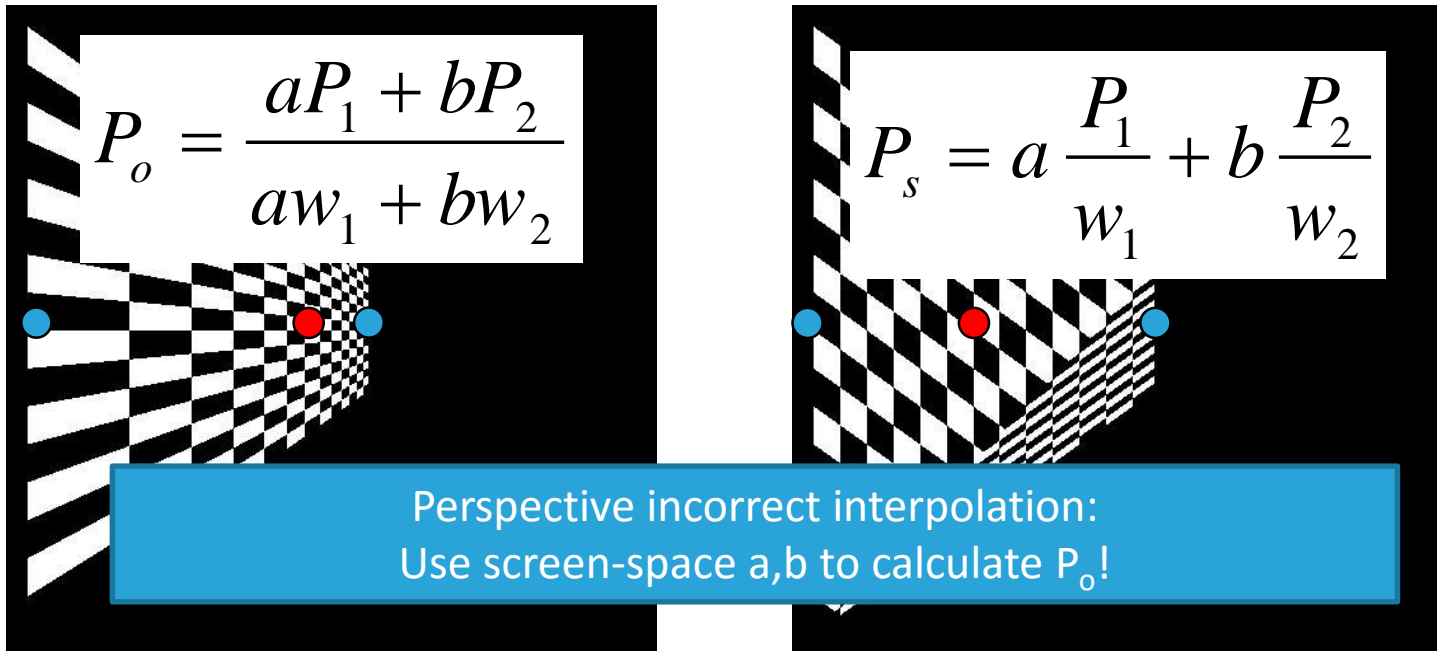


- Problem: texture coordinate interpolation
 - Texture coordinates are homogeneous!
- Look at perspective correct texturing first!



- Problem: linear interpolation in rasterization?


objectspace interpolation	$\frac{ax_1 + bx_2}{aw_1 + bw_2} \neq a \frac{x_1}{w_1} + b \frac{x_2}{w_2}$	screenspace interpolation
------------------------------	--	------------------------------



$$a = b = 0,5; P = (x,y,z,w,1,u,v,...)$$

Perspective Texture Mapping

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



- 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! $\rightarrow r/q = \text{projector depth!}$



- $[x,y,z,1,r,g,b,a]$
- texcoord generation $\rightarrow [x,y,z,1, r,g,b,a, s,t,r,q]$
- Modelviewprojection $\rightarrow [x',y',z',w,1, r,g,b,a, s,t,r,q]$
- Project (/w) \rightarrow
 $[x'/w, y'/w, z'/w, 1/w, r,g,b,a, s/w, t/w, r/w, q/w]^{\text{vert}}$
- Rasterize and interpolate \rightarrow
 $[x'/w, y'/w, z'/w, 1/w, r,g,b,a, s/w, t/w, r/w, q/w]^{\text{frag}}$
- Homogeneous: \rightarrow texture project (/ q/w) \rightarrow
 $[x'/w,y'/w,z'/w,1/w, r,g,b,a, s/q,t/q,r/q,1]$
- Or non-homogeneous: \rightarrow standard project (/ 1/w) \rightarrow
 $[x'/w, y'/w, z'/w, 1/w, r,g,b,a, s,t,r,q]$ (for normals)




■ Problem

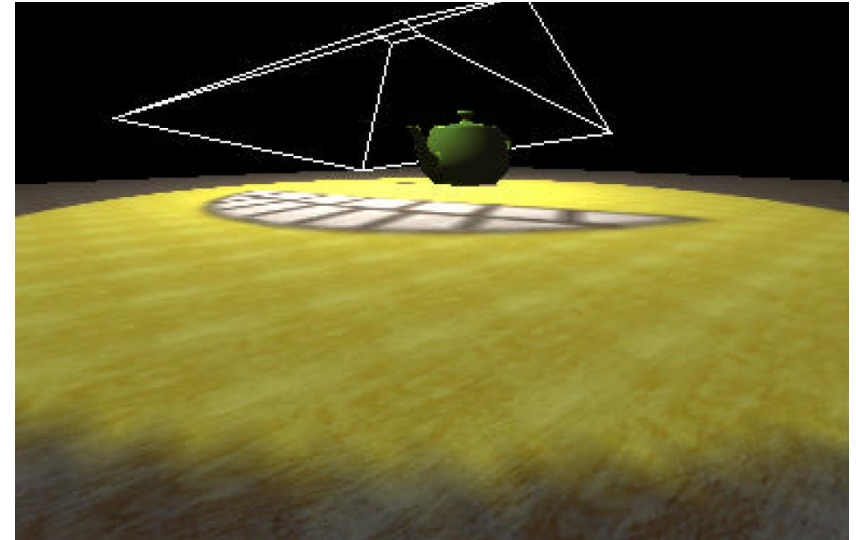
- reverse projection

■ Solutions

- Cull objects behind projector
- Use clip planes to eliminate objects behind projector
- Fold the back-projection factor into a 3D attenuation texture
- Use to fragment program to check $q < 0$



- Problems
 - Resolution problems
 - Projection behind shadow casters
- Shadow Mapping!



- Example shown in CG Shading Language
 - CG is proprietary to NVIDIA
 - C-like syntax
 - HLSL (DirectX shading language) nearly the same syntax
- Shading languages have specialized calls for projective texturing:
 - CG/HLSL: `tex2Dproj`
 - GLSL: `texture2DProj`
 - They include perspective division



Input: float4 position,

float3 normal

Output: float4 oPosition,

float4 texCoordProj,

float4 diffuseLighting

Uniform: float Kd,

float4x4 modelViewProj,

float3 lightPosition,

float4x4 textureMatrix



```
oPosition =
    mul(modelViewProj, position);
texCoordProj =
    mul(textureMatrix, position);
float3 N = normalize(normal);
float3 L = normalize(lightPosition
    - position.xyz);
diffuseLighting =
    Kd * max(dot(N, L), 0);
```



Input: float4 texCoordProj,
float4 diffuseLighting

Output: float4 color

Uniform: sampler2D projectiveMap

```
float4 textureColor =  
    tex2Dproj (projectiveMap,  
              texCoordProj);
```

```
color = textureColor *  
        diffuseLighting;
```



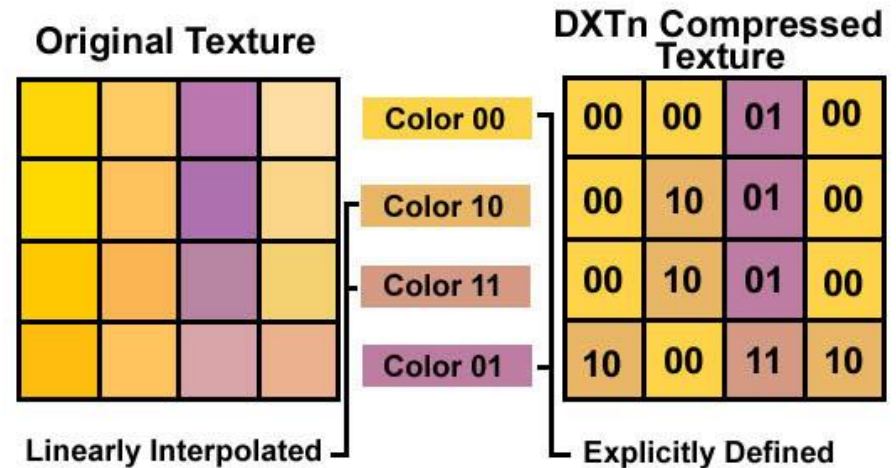
- Classic OpenGL:
 - Just supply correct matrix to glTexGen
- Projective texturing is easy to program and very effective method.
- Combinable with shadows



Projective Shadow in Doom 3



- S3TC texture compression (DXTn)
- Represent 4x4 texel block by two 16bit colors (5 red, 6 green, 5 blue)
- Store 2 bits per texel
- Uncompress
 - Create 2 additional Colors between c1 and c2
 - use 2 bits to index which color
- 4:1 or 6:1 compression



Multipass Rendering



- Recall 80 million triangle scene
- Games are NOT using $a = 0.5$
 - at least not yet
- Assume $a = 32$, $l = 1024 \times 768$, $d=4$
 - Typical for last generation games
 - $F = l * d = 3,1 \text{ MF/frame}$,
 - $T = F / a = 98304 \text{ T/frame}$
 - $60 \text{ Hz} \rightarrow \sim 189 \text{ MF/s}, \sim 5,6 \text{ MT/s}$



- Hardware underused with standard OpenGL lighting and texturing

What can we do with this power?

- Render scene more often:
multipass rendering
- Render more complex pixels:
multitexturing
 - 2 textures are usually for free
- Render more complex pixels and triangles:
programmable shading



- Conventional OpenGL allows for many effects using multipass
 - Still in use for mobile devices and last gen consoles
 - Modern form: render to texture
 - Much more flexible but same principle
- Programmable shading makes things easier
 - Specialized calls in shading languages



- OpenGL lighting model only
 - local
 - limited in complexity
- Many effects possible with multiple passes:
 - Dynamic environment maps
 - Dynamic shadow maps
 - Reflections/mirrors
 - Dynamic impostors
 - (Light maps)



- Render to auxiliary buffers, use result as texture
 - E.g.: environment maps, shadow maps
 - Requires pbuffer/fbo-support
- Redraw scene using fragment operations
 - E.g.: reflections, mirrors
 - Uses depth, stencil, alpha, ... tests
- “Multitexture emulation mode”: redraw
 - Uses framebuffer blending
 - (light mapping)



(assume redraw scene...)

■ First pass

- Establishes z-buffer (and maybe stencil)

```
glDepthFunc (GL_LEQUAL) ;
```

- Usually diffuse lighting

■ Second pass

- *Z-Testing* only

```
glDepthFunc (GL_LEQUAL) ;
```

- Render special effect using (examples):

- Blending

```
glStencilFunc (GL_EQUAL, 1, 1) ;
```



```
glEnable(GL_BLEND);  
glBlendEquation(GL_FUNC_ADD);
```

weighting factors

result color

$$C = C_s S + C_d D$$

incoming (source)
fragment color

framebuffer color

- Other equations: **SUBTRACT**, **MIN**, **MAX**




```
glBlendFunc (GL_SRC_ALPHA, GL_ONE_MINUS_SRC_ALPHA) ;
```


$$C = C_s \cdot \alpha + C_d \cdot (1 - \alpha)$$

- Example: transparency blending (window)
- Weights can be defined almost arbitrarily
- Alpha and color weights can be defined separately
- `GL_ONE`, `GL_ZERO`, `GL_DST_COLOR`,
`GL_SRC_COLOR`, `GL_ONE_MINUS_XXX`

