

# Texturing


Peter Sikachev,  
Stefan Jeschke,  
Eduard Gröller

Institute of Computer Graphics and Algorithms




**Vienna University of Technology**




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




- Also possible: model very complex objects just by using simple textured geometry



# Texturing: General Approach



Texels




Texture space  $(u, v)$

Object space  $(\bar{x}, \bar{y}, \bar{z})$


Image Space  $(x_I, y_I)$

Parametrization

Rendering  
(Projection etc.)



- Sampled textures drawbacks
  - ◆ Take a lot of memory
  - ◆ Texture fetches are slow
  - ◆ Parameterization is always needed
  - ◆ Aliasing
- Procedural textures
  - ◆ Texture is generated on-the-fly per fragment



- A texture is a **function**:
  - ◆ *Procedural*: fast to evaluate, but limited variety
  - ◆ *Sampled*: most common method
    - Raster images that are taken with a digital camera, scanned or synthesized
- Textures can be defined:
  - ◆ In 3D object space: “*3D texturing*“, “*cube mapping*”
  - ◆ On the 2D object surface: “*texture mapping*”
  - ◆ As a 1D function: “*lookup table*”



- Encoding of environment is needed
- 3D texture is an overkill: why?



- Encoding of environment is needed
- 3D texture is an overkill:
  - ◆ Memory





- Encoding of environment is needed
- 3D texture is an overkill:
  - ◆ Memory
  - ◆ Takes time to generate



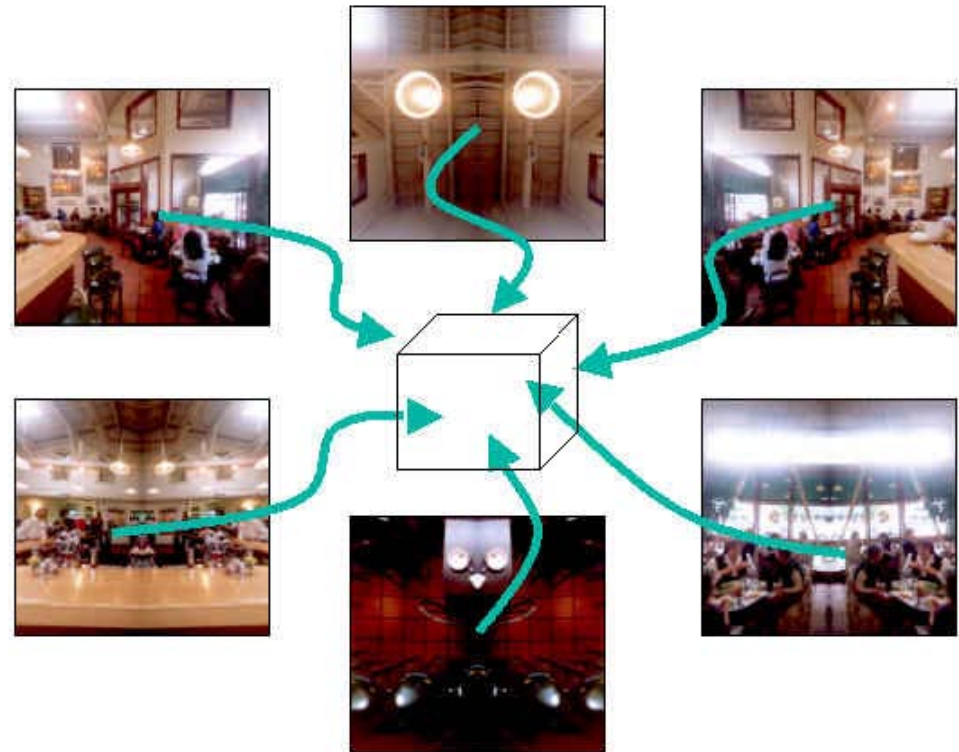
- Encoding of environment is needed
- 3D texture is an overkill:
  - ◆ Memory
  - ◆ Takes time to generate
  - ◆ Takes time to sample



- Encoding of environment is needed
- 3D texture is an overkill:
  - ◆ Memory
  - ◆ Takes time to generate
  - ◆ Takes time to sample
- Solutions?



- Encoding of environment is needed
- 3D texture is an overkill:
  - ◆ Memory
  - ◆ Takes time to generate
  - ◆ Takes time to sample
- Solution:
  - ◆ Take 6 photos from scene center
  - ◆ Texture is addressed via 3D vector, which has semantics of direction, not position




- Application: specular object in the environment




Cube map


Rendered scene



- Not directly applied for polygonal models
- Representation of some 3D value field
- Applications:
  - Volume rendering
  - Media (gas/liquid) rendering
  - Hybrid rendering
  - etc






- **1D** texture: parameter can have arbitrary domain (along one axis, incident angle, etc.)




- **2D** texture

- ◆ Projection of 2D data
- ◆ Many possibilities for definition








- *Other 2D texture: cylindrical parametrization*
  - ◆ Depending on cylindrical coordinates of each point




- *Other 2D texture: spherical parametrization*
  - ◆ Depending on spherical coordinates of each point




- Difficulty: how to minimize texture distortions?




## Other problems?



- Non-convex objects: mapping is not bijective




- Parametrization per vertex
  - ◆ Apply  $(u, v)$  texture coordinates to every vertex
- At runtime, ***(bi)linearly interpolate*** between these coordinates during rasterization



- Problem: perspective correctness
  - ◆ Affine texture mapping does not take into account the depth information


$$u_{\alpha} = (1 - \alpha)u_0 + \alpha u_1; 0 \leq \alpha \leq 1$$




## ■ Solution

- ◆ Interpolate coordinates, divided by depth and depth reciprocals


$$u_{\alpha} = \frac{(1-\alpha)\frac{u_0}{z_0} + \alpha\frac{u_1}{z_1}}{(1-\alpha)\frac{1}{z_0} + \alpha\frac{1}{z_1}}$$



Flat



Affine






Correct






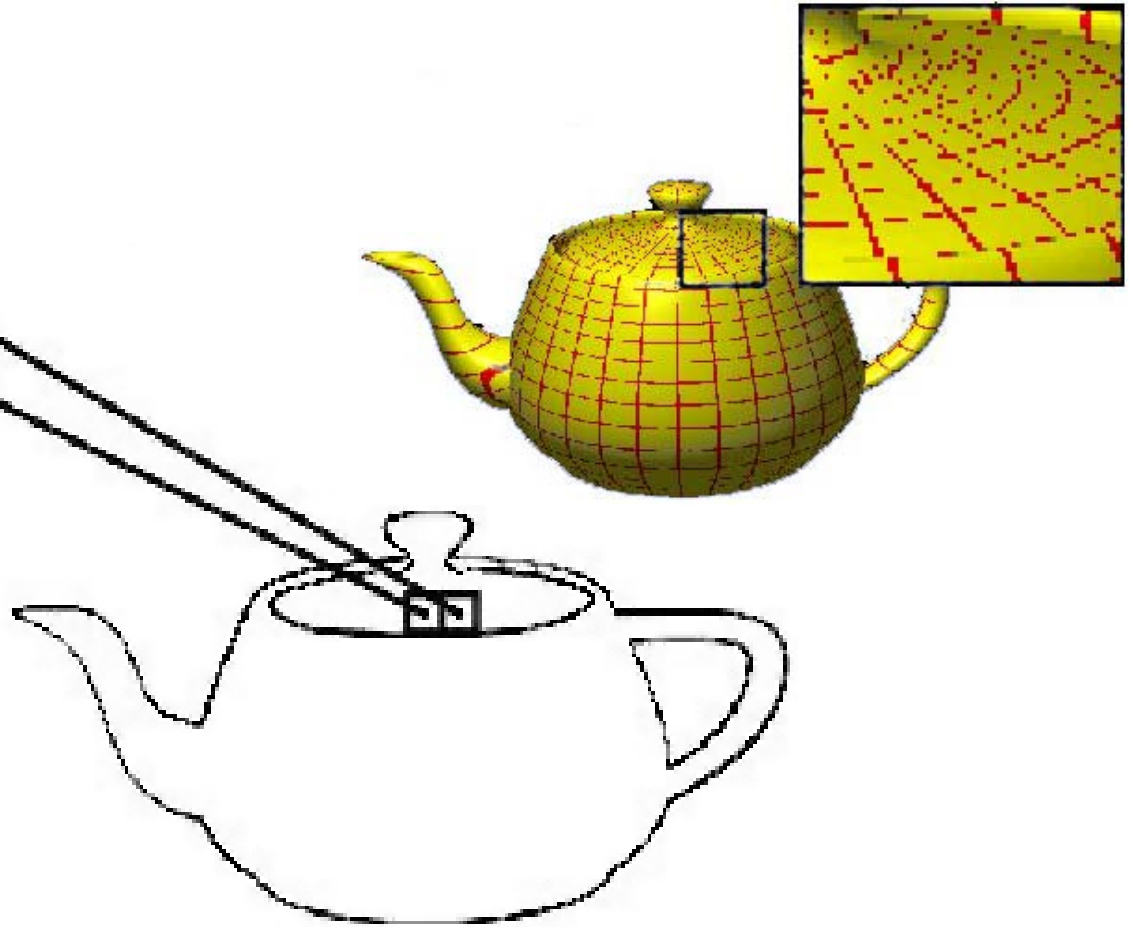
- Problem: aliasing
  - ◆ One pixel in image space covers many texels



- Caused by *undersampling*: texture information is lost




**Texture space**



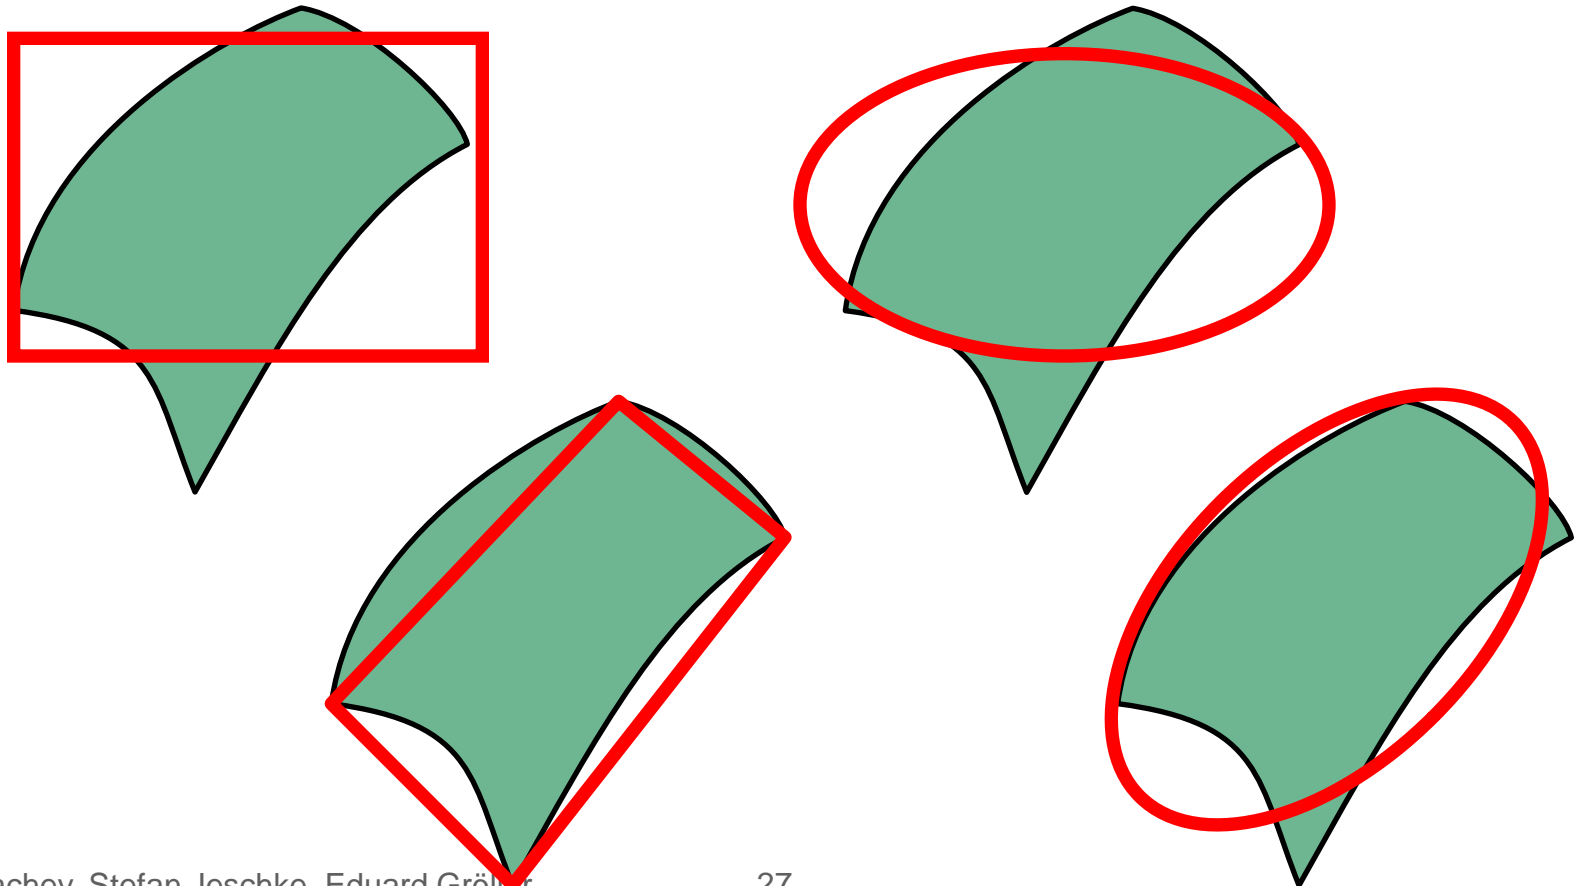
**Image space**



- A good pixel value is the weighted mean of the pixel area projected into texture space
  - ◆ Calculation at *runtime* or in a *preprocess*




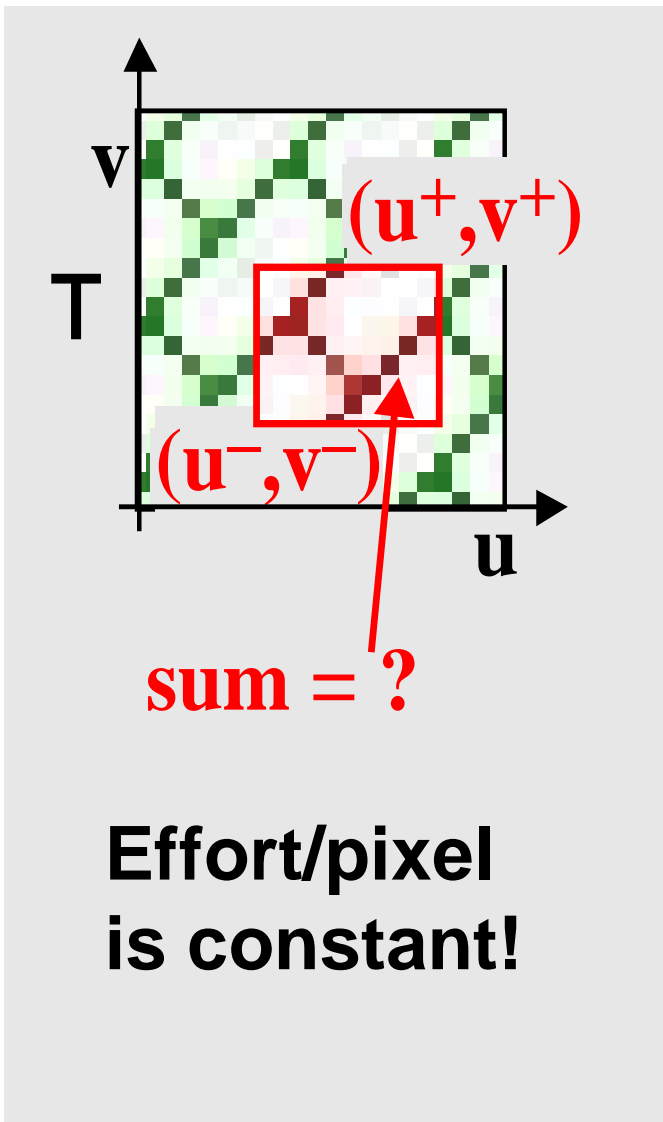
- Calculation of the weighted mean at runtime (called "*direct convolution*")
  - ◆ Often used approximations

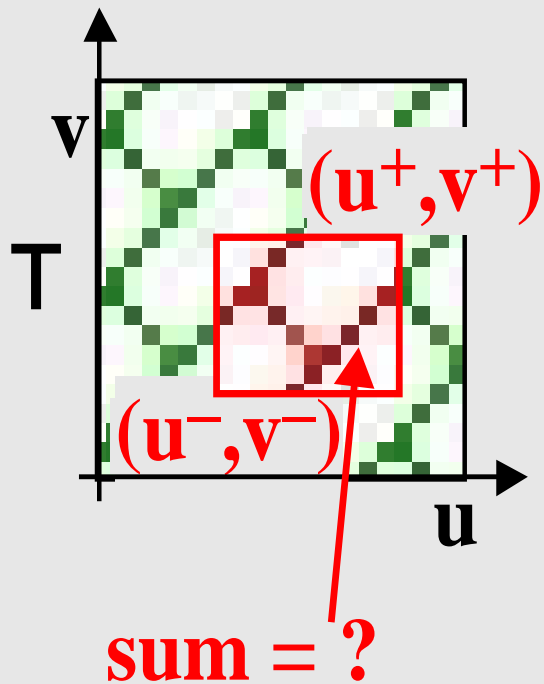


- Calculation of the weighted mean in a *preprocess* (called "*prefiltering*")
  - ◆ *Mip-mapping* (+ optional *anisotropic filtering*)
- Summed area table  $S$ 
  - ◆ Precalculation of rectangle sums for each pixel

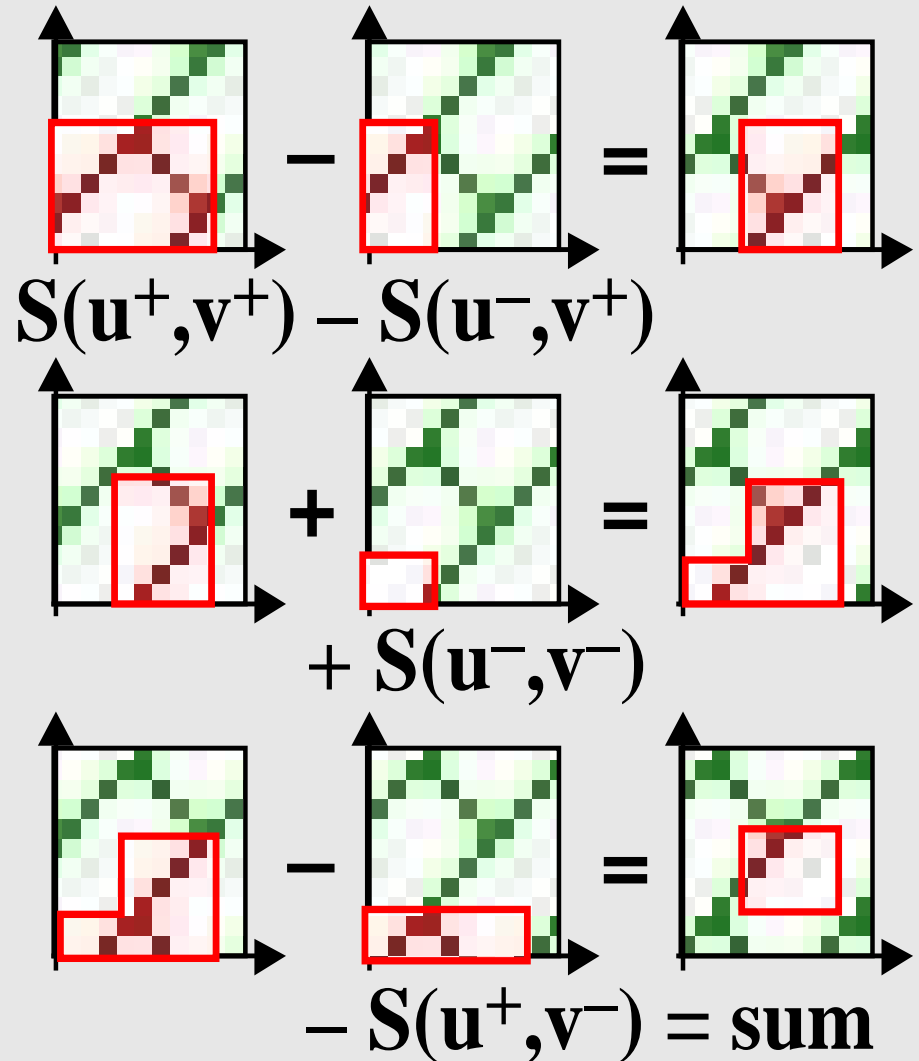
$$S(u_0, v_0) = \sum_{\substack{u \leq u_0 \\ v \leq v_0}} T(u, v)$$








**Effort/pixel  
is constant!**




# Anti-Aliasing: Example (Summed Area T.)







- 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



# Anti-Aliasing: Example (Mip Mapping)





- MIP Mapping Algorithm

- $D := ld(max(d_1, d_2))$  "Mip Map level"

- $T_0 :=$  value from texture  $D_0 = trunc(D)$

  - ◆ Use *bilinear interpolation* (why aliasing?)





- MIP Mapping Algorithm


- $D := ld(max(d_1, d_2))$  "Mip Map level"

- $T_0 :=$  value from texture  $D_0 = trunc(D)$

- ◆ 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







No mip-mapping





simple  
mip-mapping



ray  
differentials





+ anisotropic  
filtering



## ■ Anisotropic Filtering


- ◆ View dependent filter kernel
- ◆ Implementation: *summed area table* (see above), "*RIP Mapping*", "*footprint assembly*"




**Texture space**




- *"RIP Mapping"*: like MIP Maps but half resolution in x and y direction *separately*
  - ◆ Fast, but needs more memory than MIP Mapping



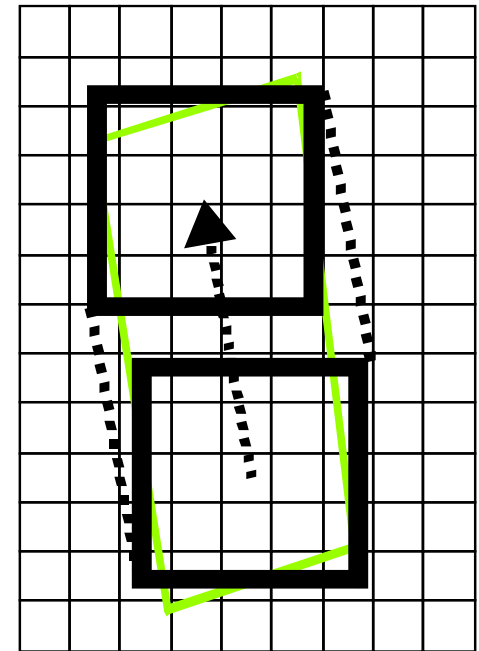
- *"Footprint assembly"*
  - ◆ Approximate pixel footprint by parallelogram
  - ◆ Sample along the parallelogram using MIP mapping



**Footprint in texture space**




**Approximation by parallelogram**




**Sampling using Mip Mapping**




- Example



- Preprocess: compute normal vectors from height field
- Runtime: use computed normals for illumination




# Bump Mapping Examples






- Improve bump mapping with (local) shadows
- Preprocess: compute  $n$  horizon values per texel
- Runtime:
  - ◆ Interpolate horizon values
  - ◆ Shadow accordingly



without




with



Shadows from bumps

No light on rear side of the object





- Parallax: apparent movement of close objects in front of further objects
  - ◆ Preprocess: calculate "*texture coordinate shifts*"
  - ◆ Runtime: "crd. shifts" alter the texture lookup position
    - Problem: occlusions cannot be modelled



**Bump mapping**



**Parallax + Bump mapping**



- Trace the ray locally until it crosses the surface




image courtesy of Natalya Tatarchuk



- Banding artifacts at acute angles
- Why do they occur?





image courtesy of Natalya Tatarchuk



- Banding artifacts at acute angles
- Solution: approximate intersection with point B





image courtesy of Natalya Tatarchuk

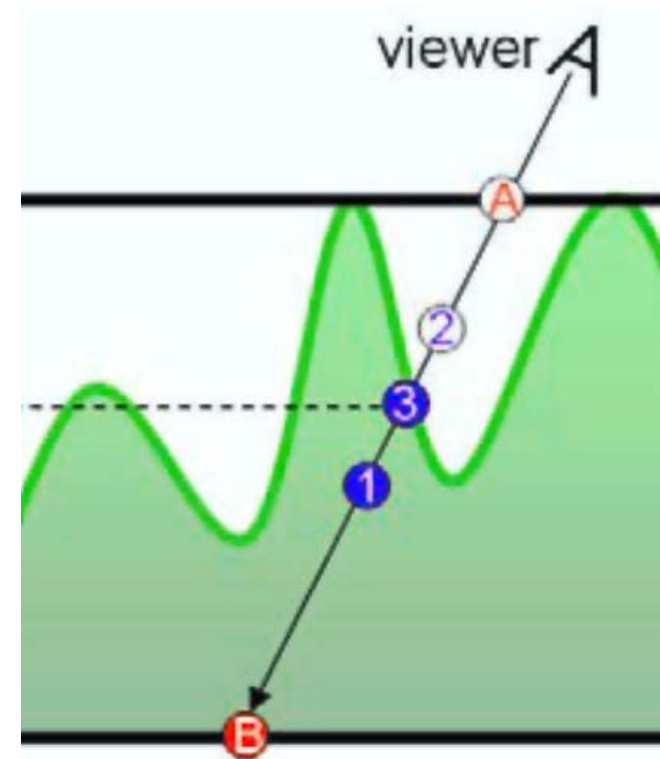



- At runtime: perform ray casting in the pixel shader
  - ◆ Calculate entry (**A**) and exit point (**B**)
  - ◆ March along ray until intersection with height field is found
  - ◆ How to accelerate?





- At runtime: perform ray casting in the pixel shader
  - ◆ Calculate entry (A) and exit point (B)
  - ◆ March along ray until intersection with height field is found
  - ◆ Binary search to refine the intersection position





**Bump mapping**



**Parallax mapping**



**Relief mapping**



# Relief Mapping Examples (2)

