



Shadow Mapping Algorithms

Gary King

Technical Developer Relations

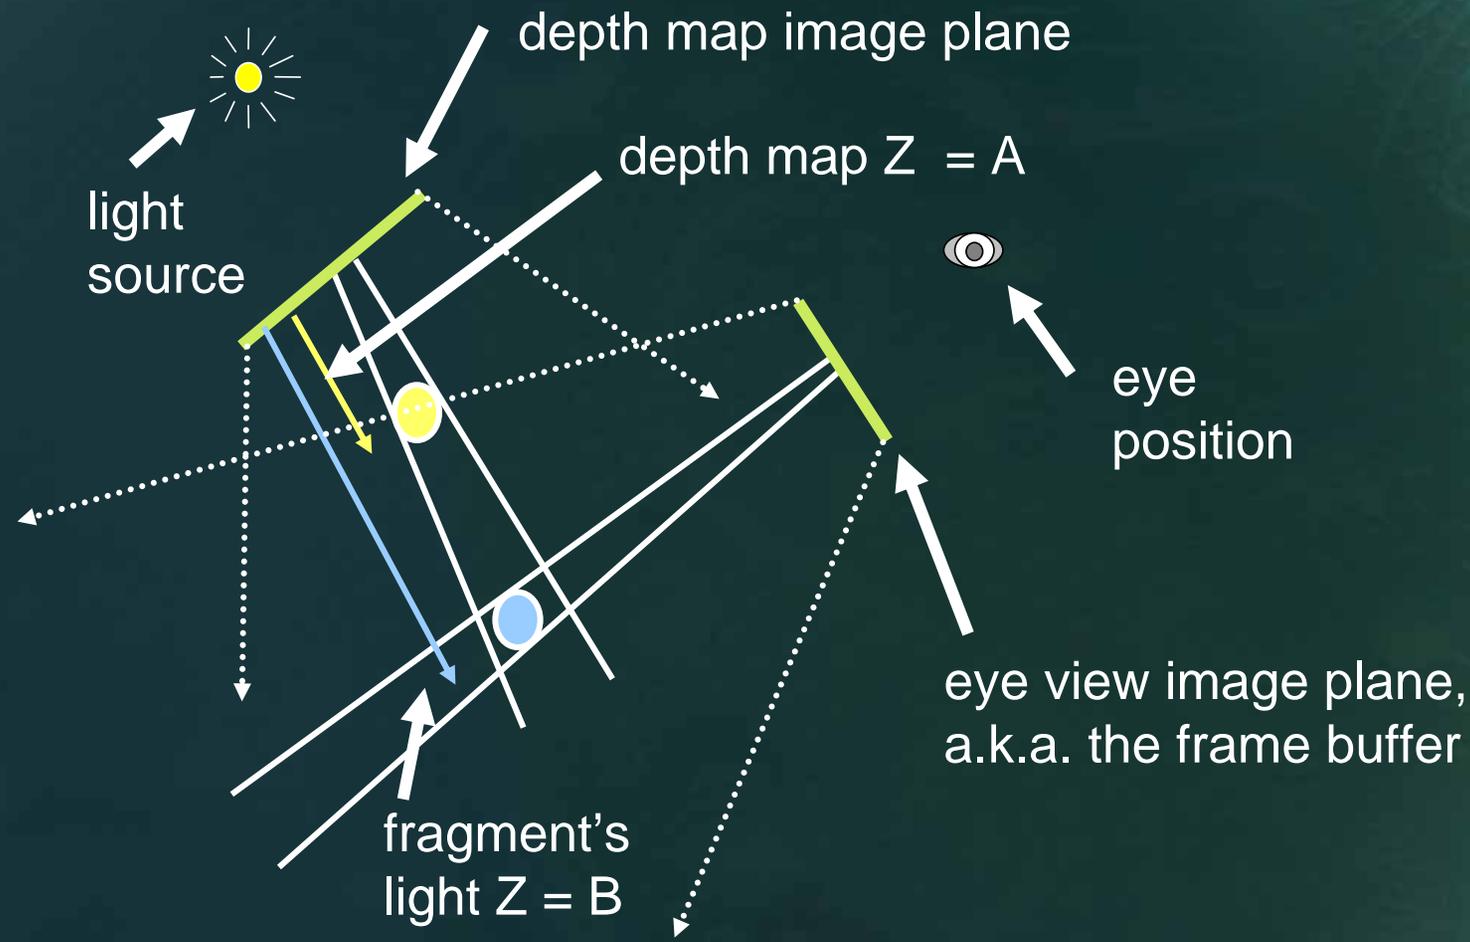


Shadow Mapping Introduction

- Casting Curved Shadows on Curved Surfaces, Lance Williams
- Render scene from light's point of view
 - Closest objects to light are occluders
- Render scene from eye's point of view
 - Project all fragments into shadow map
 - If $\text{depth} > \text{depth}(\text{shadow})$, object is in shadow



Shadow Mapping Diagram





Shadow Mapping on GPUs

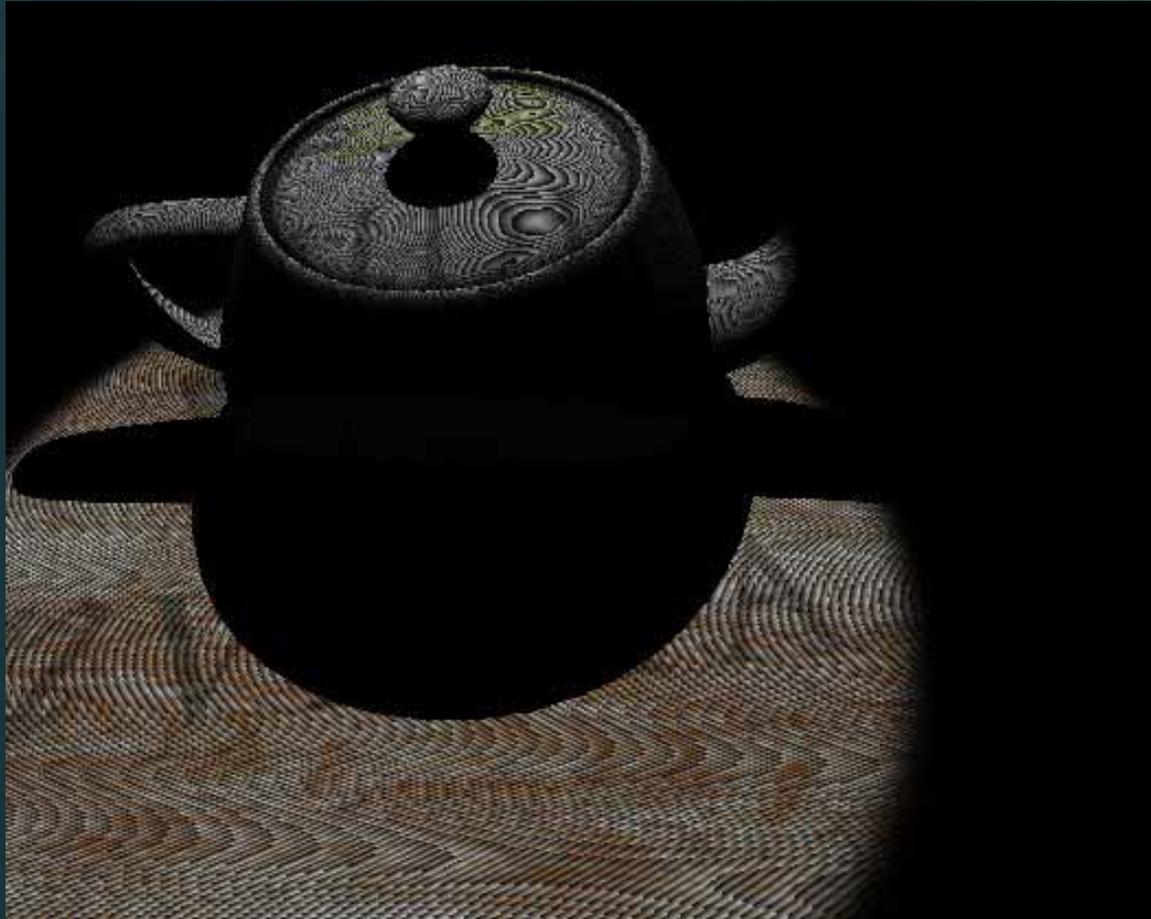
- On Radeon 9500+
 - Floating-point textures (R32F)
 - Pixel shader filtering and comparison
- On GeForce 3+
 - Native shadow map support (16 and 24-bit integer)
 - 2x2 bilinear percentage closer filtering for free
 - 2x speed rendering on GeForceFX and later GPUs
 - 32ppc on GeForce 6800 GT



Simple, Right?

- Shadow maps are a nice, elegant, easy-to-implement, technique for shadowing...
- Except for...

Shadow Acne



- Surface incorrectly self-shadows

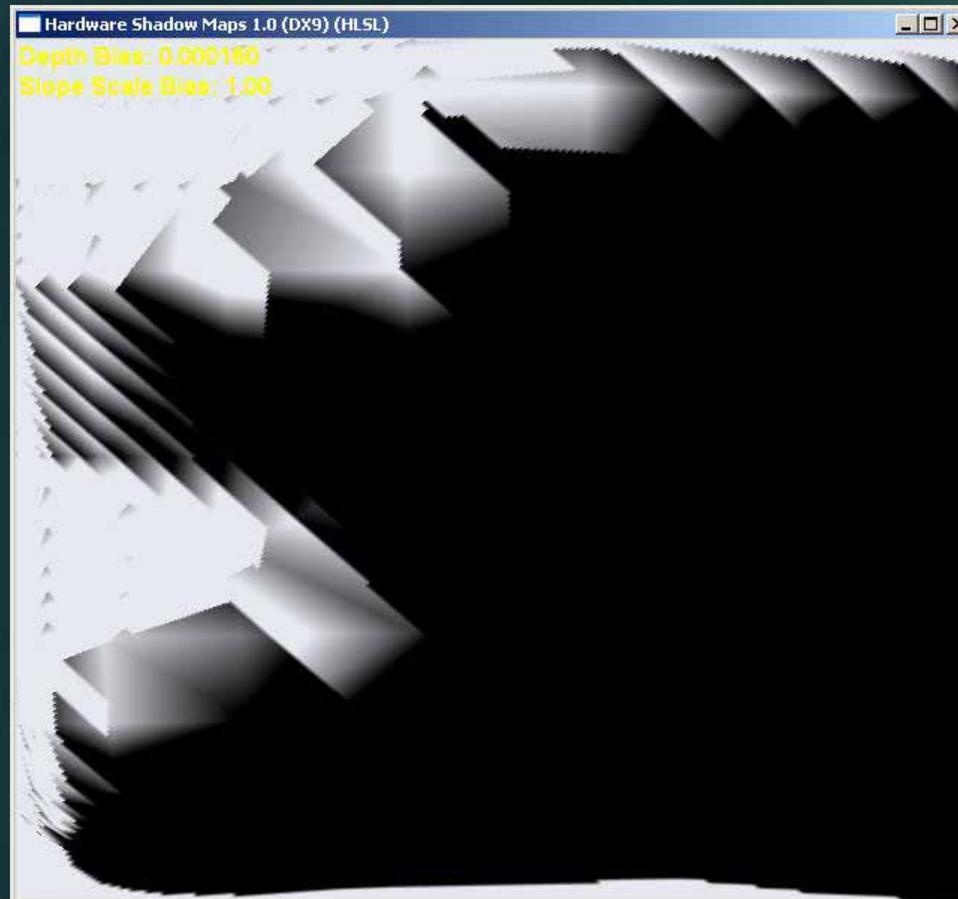


Perspective Aliasing



- Blocky shadows cast onto nearby receivers

Projective Aliasing



- Streaked shadows cast onto ~parallel receivers

Omnidirectional Lights



- Point lights cast shadows in all directions



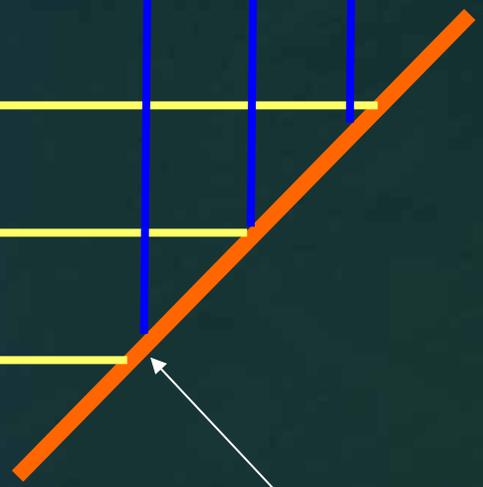
Source of Shadow Acne



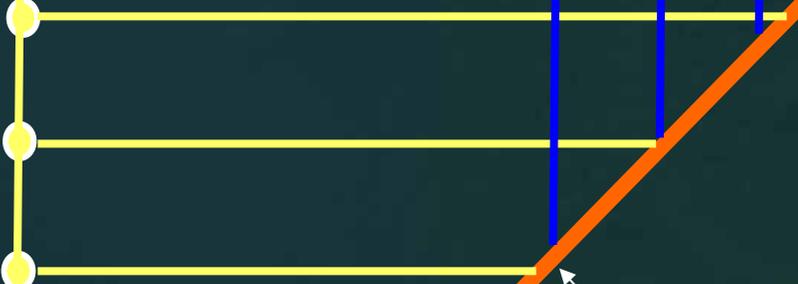
Light Plane



Polygon



Eye Plane



Eye's distance > shadow depth

Incorrect Self-Shadowing

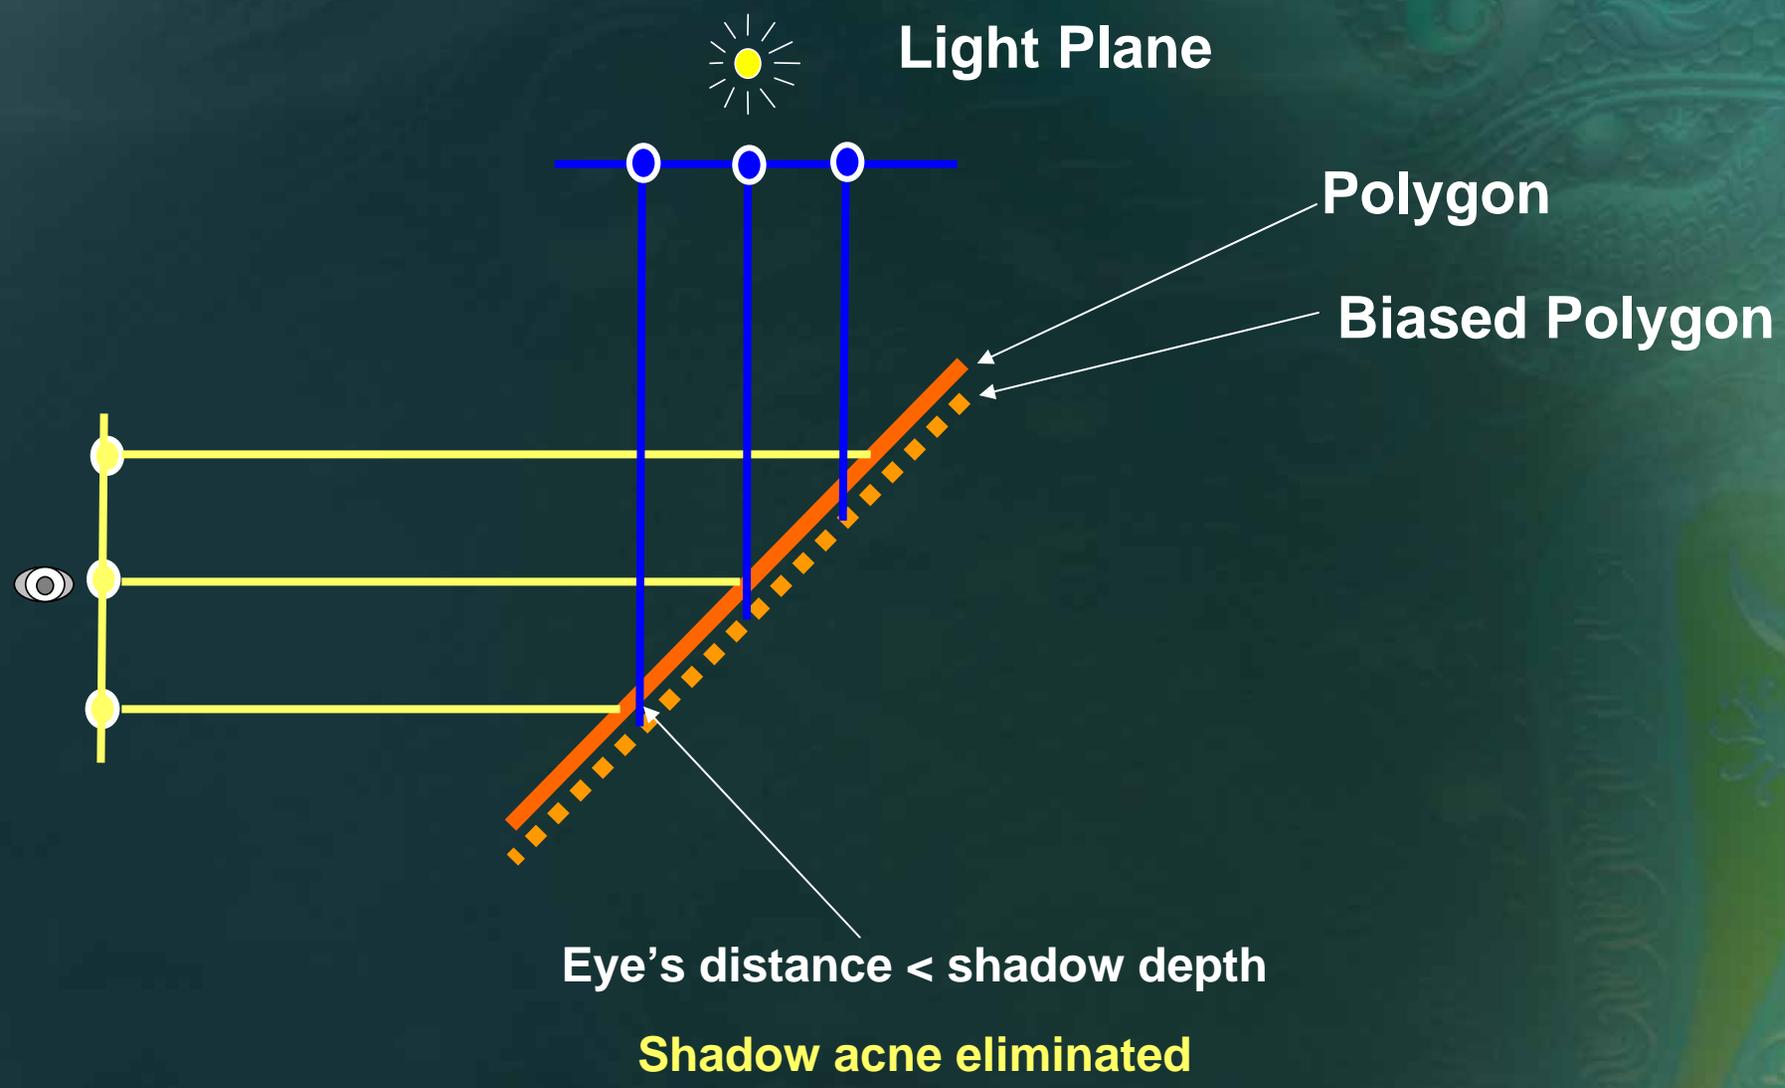


Depth Bias

- Add epsilon to shadow depths, such that
 $\text{eye depth} \leq \text{shadow depth}$
For all fragments from the same polygon
- How much epsilon is needed?
 - Art, not science
 - Movie houses spend hours/days tweaking shadow parameters frame-by-frame to avoid artifacts



Depth Bias Diagram





Using Depth Bias in Direct3D9

- Ignore DirectX SDK Depth Bias documentation
- Depth bias is a floating point number that is added to the interpolated fragment Z value
 - Can be positive or negative
- To bias by 1 LSB of a 24-bit depth buffer, use

```
float fBias = 1.f / 16777215.f;
DWORD dwBias = *(DWORD*)&fBias;
SetRenderState(D3DRS_DEPTHBIAS, dwBias);
```



Resampling Error

- Projected pixels are snapped to nearest texel
 - This is a window of $[\pm 0.5, \pm 0.5]$ in the texture
 - Up to $|0.5 \frac{\partial z}{\partial s}| + |0.5 \frac{\partial z}{\partial t}|$
- Since polygons are planar, bilinearly filtering the depth values is *almost* correct
 - $\frac{\partial z}{\partial s}$ and $\frac{\partial z}{\partial t}$ can be inferred from nearby texels
 - Bilinear weights give ∂s and ∂t , **BUT**
 - Polygon edges break the planar assumption
 - Perspective Z is hyperbolic, not linear

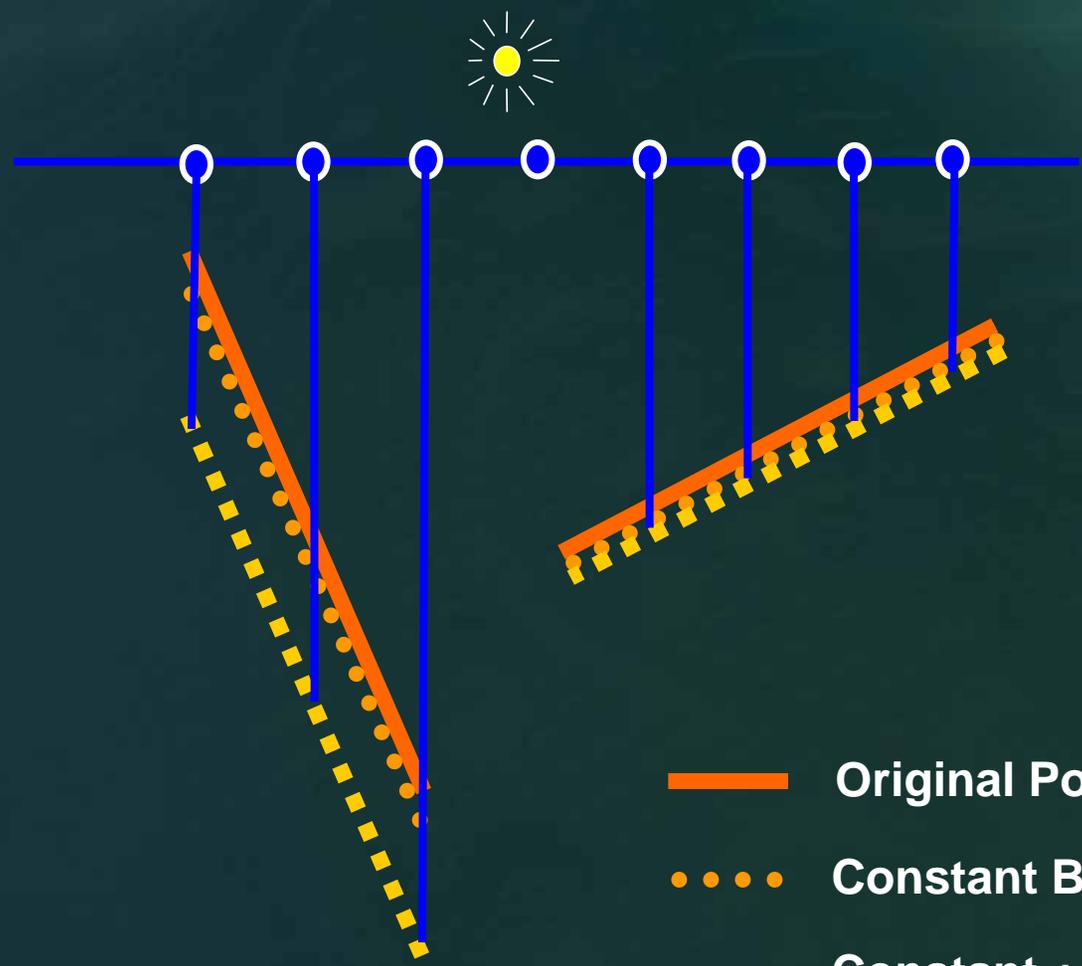


Slope-Scale Depth Bias

- Instead, bias the shadow depth based on the post-projective Z-slope
- Slope-scale depth bias lets GPU compute this
 - In D3D, use `D3DRS_SLOPESCALEDEPTHBIAS`
- Total bias is **$m * \text{SLOPESCALE} + \text{DEPTHBIAS}$**
 - Where $m = \max(| \partial z / \partial x | , | \partial z / \partial y |)$



Slope-Scale Depth Bias Diagram

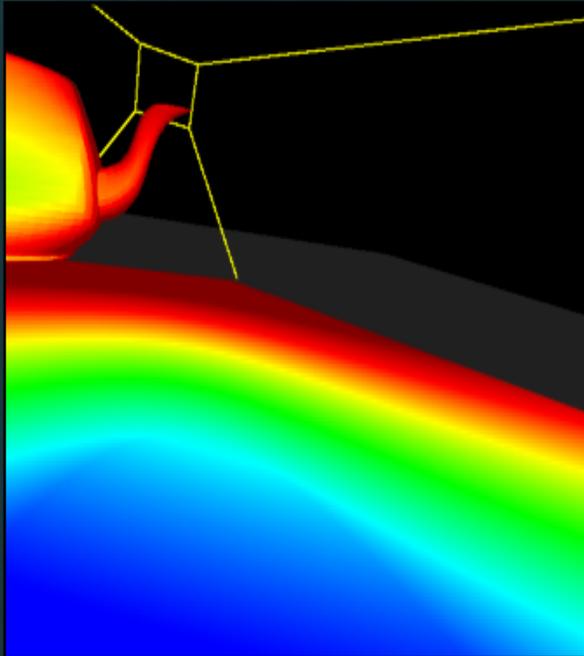


- Original Polygon
- Constant Bias
- ■ ■ Constant + Slope Bias

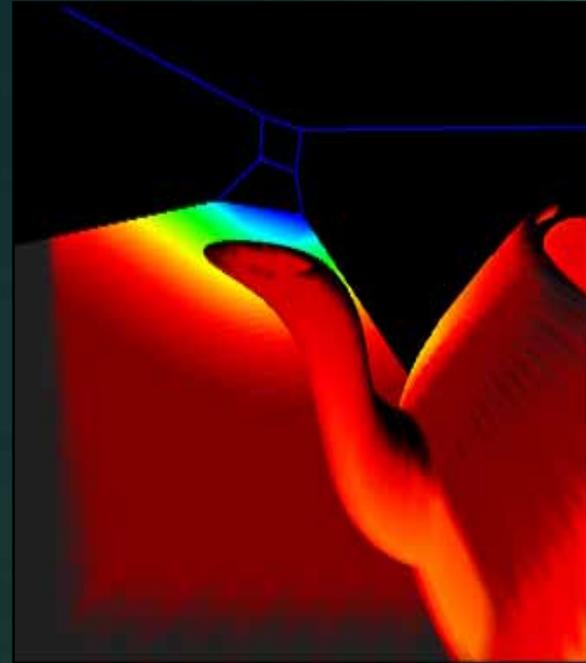


Source of Perspective Aliasing

Eye View



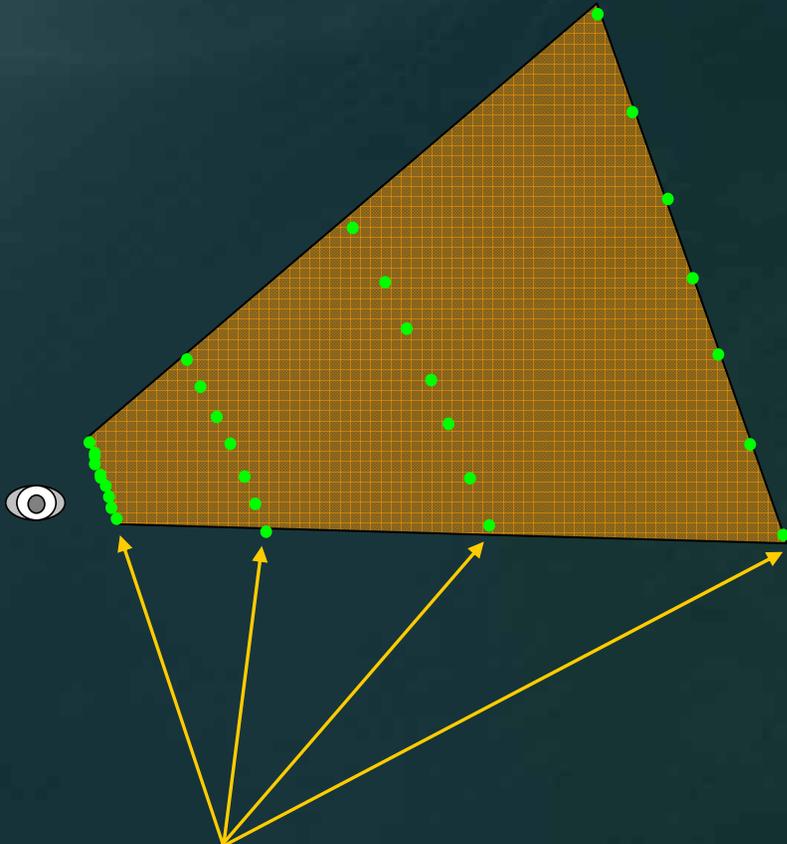
Shadow map



Blue and green region occupy significant portion of eye's view, but receive comparatively few shadow texels



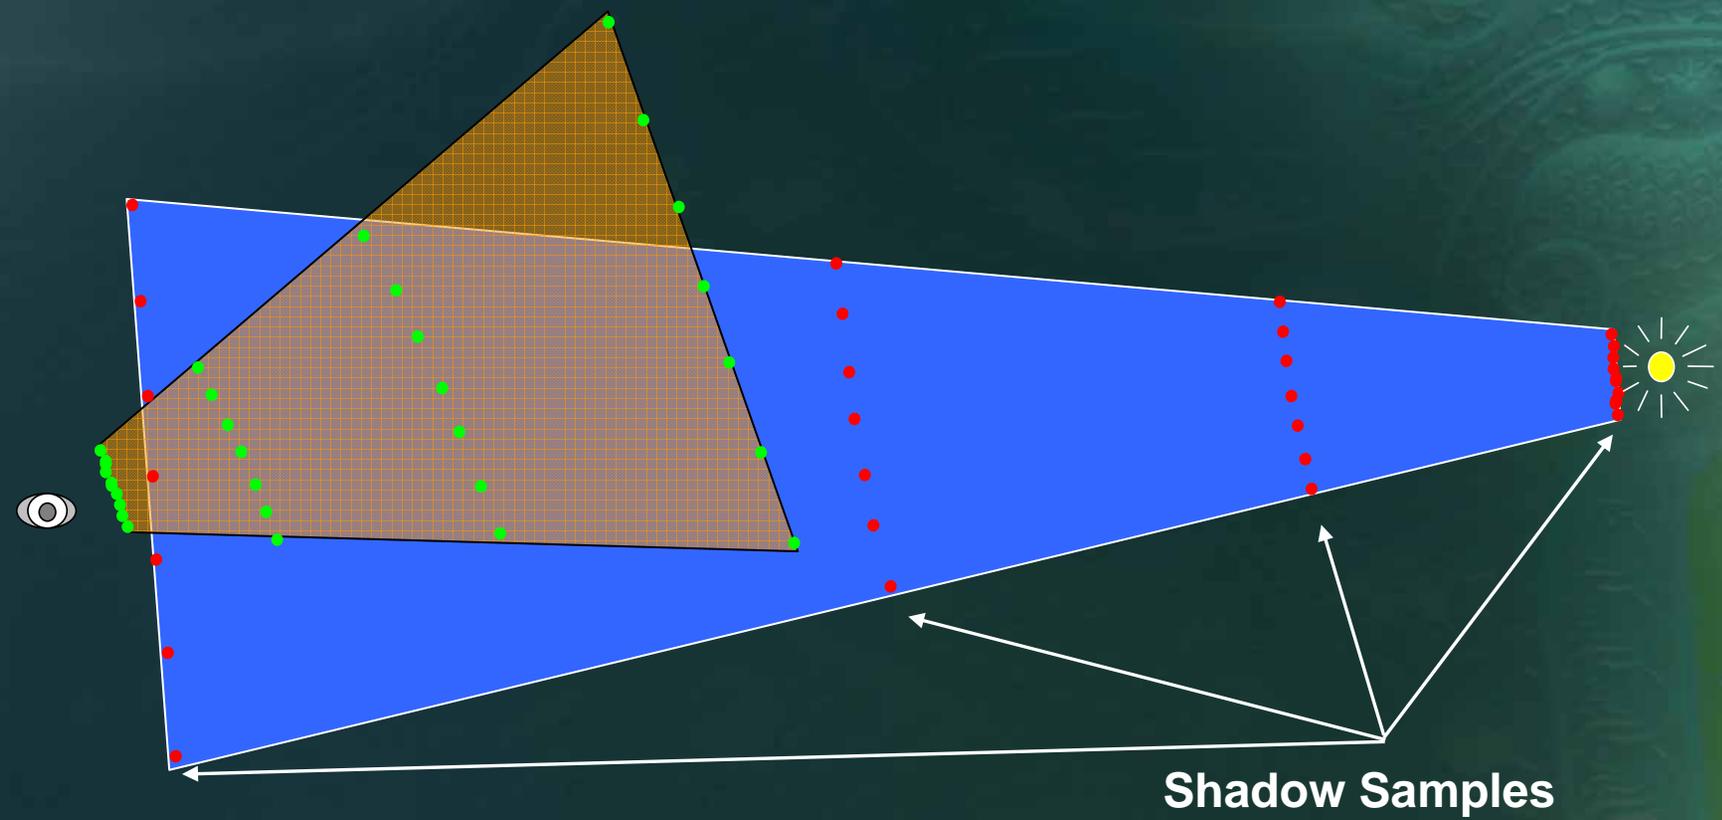
Perspective Aliasing in 2D



Eye Samples



Perspective Aliasing in 2D

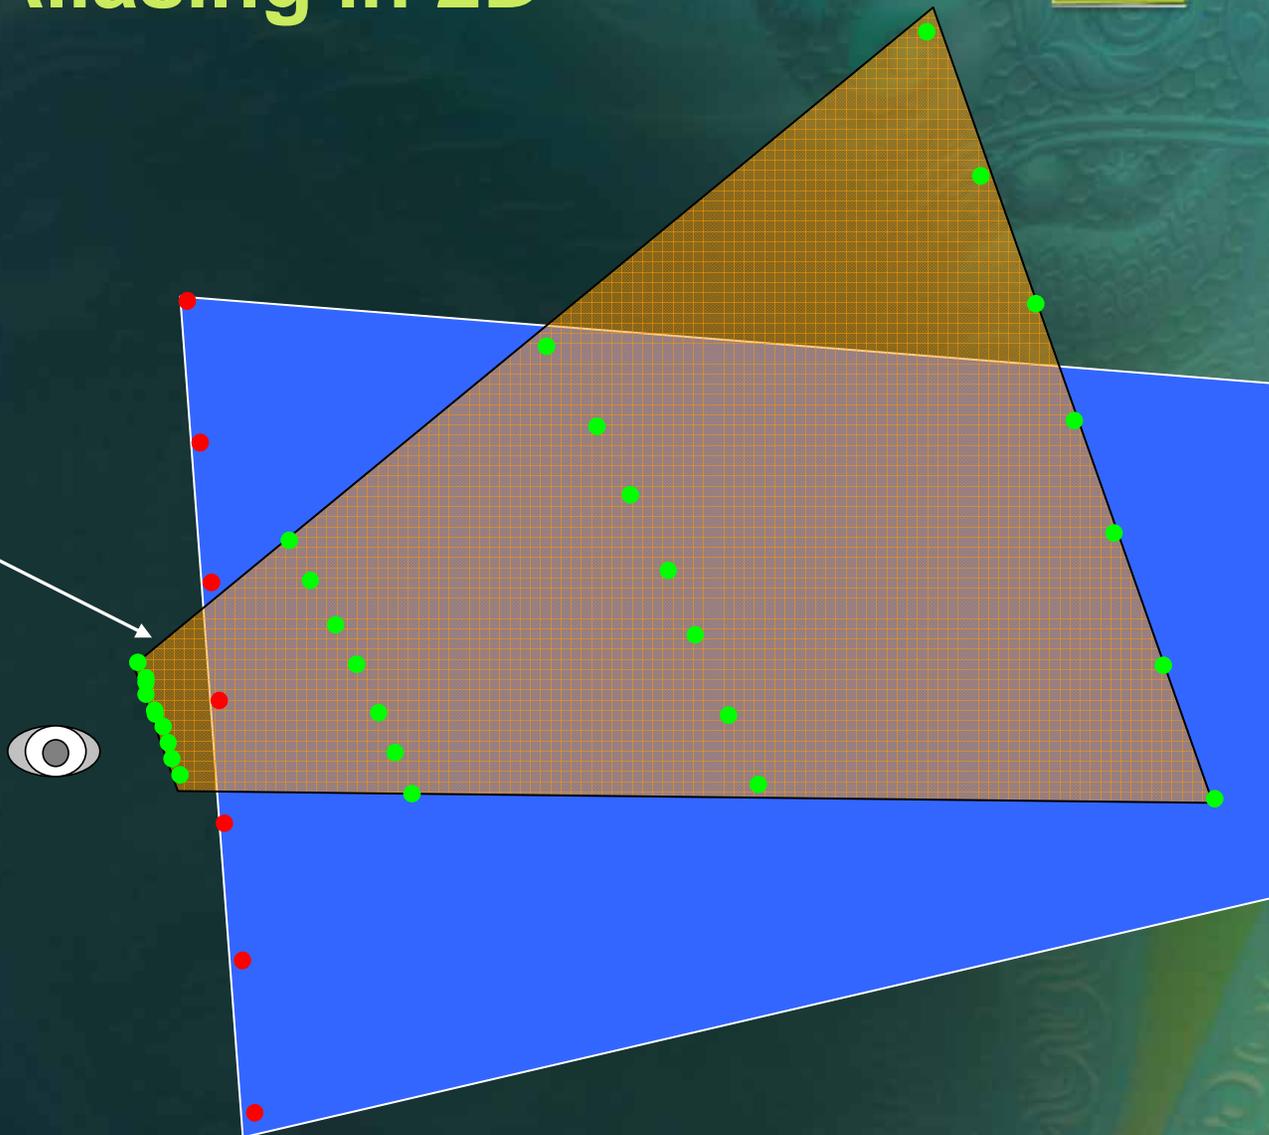


Duelling Frusta : Each frustum points toward the other, resulting in perfectly nonoptimal shadow maps

Perspective Aliasing in 2D



All pixels project to 3 shadow texels!





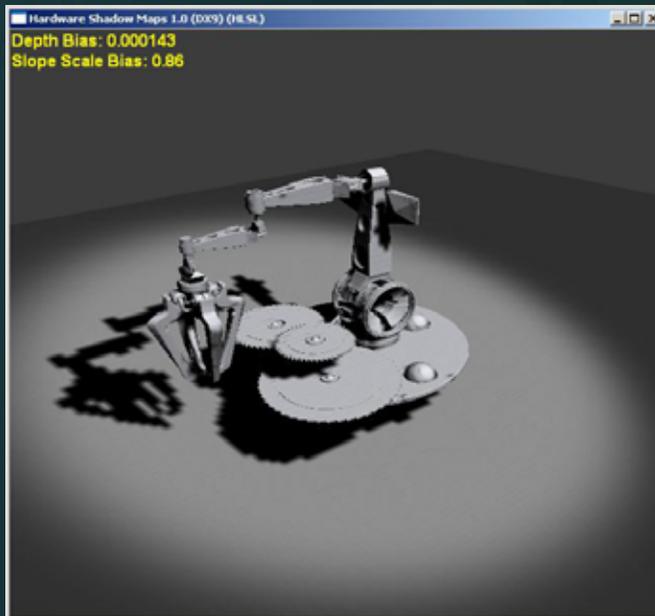
The Sea of Aliasing “Solutions”

- Many ways proposed to reduce aliasing
 - Increase shadow map resolution
 - Unit cube clipping
 - Cascaded Shadow Maps (????)
 - Perspective Shadow Maps (Stamminger, 2002)
 - Light-Space PSMs (Wimmer, 2004)
 - Trapezoidal Shadow Maps (Martin, 2004)
 - ... Not counting filtering & soft shadow algorithms!
- What works, and what are the problems?

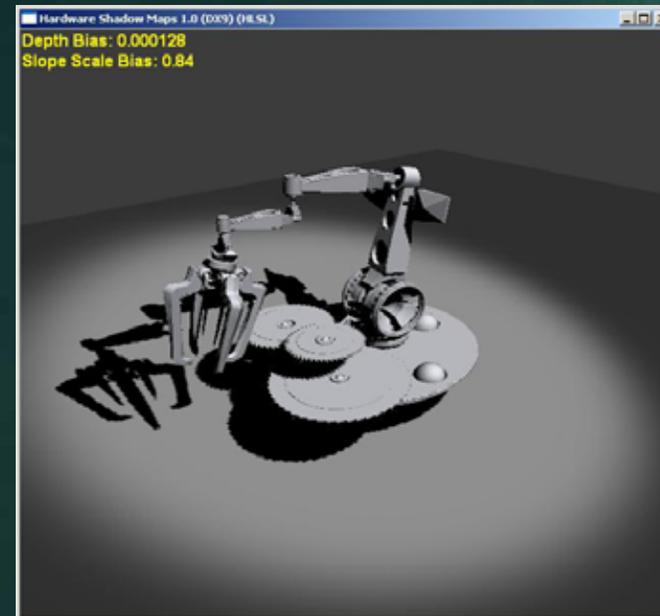


Increasing Resolution

- More samples = less aliasing



128x128



512x512

- But how many samples is enough?



Sample Density Requirements

- Ideally, we want at least one unique shadow sample for every pixel on the screen
 - Ray tracing and stencil shadow volumes give this

- So, we want to satisfy

$$\text{Area}(\text{eyeProj} * \text{eyeView} * P) \leq \text{Area}(\text{lightProj} * \text{lightView} * P)$$

for all P , and all view combinations

- This is 22 degrees of freedom



Increasing Resolution

- Not possible to satisfy this for arbitrary situations with one shadow map and one linear projection.
- But not all situations are arbitrary
 - Flashlights
 - $\text{lightView} * P \sim \text{eyeView} * P$
 - Static lights in enclosed environments
 - $\text{Area}(\text{lightProj} * \text{lightView} * P)$ is bounded at level-design time
- When DoF can be reduced, resolution tweaking may be sufficient

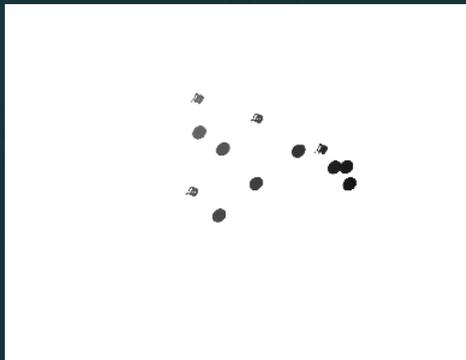
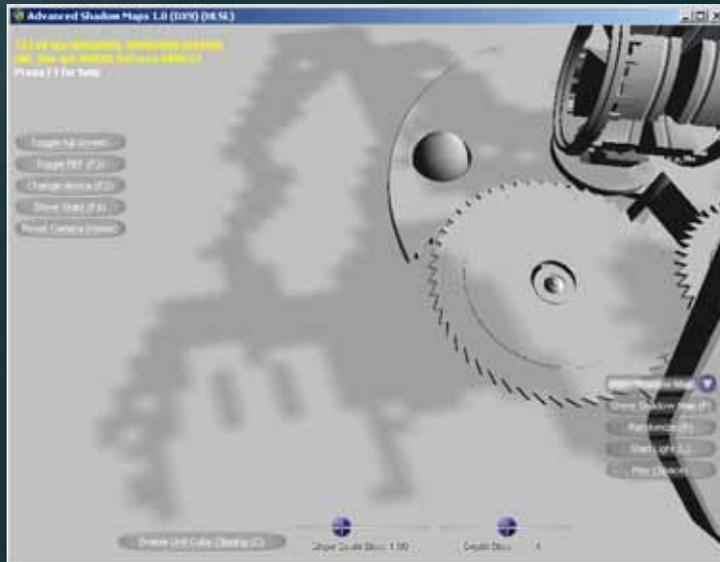


Unit Cube Clipping

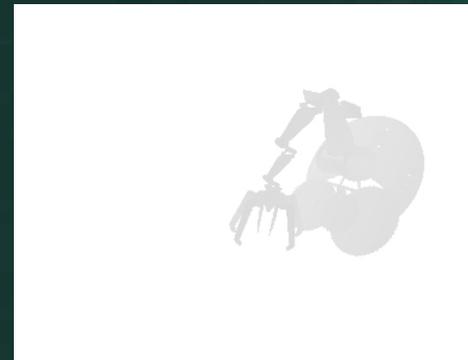
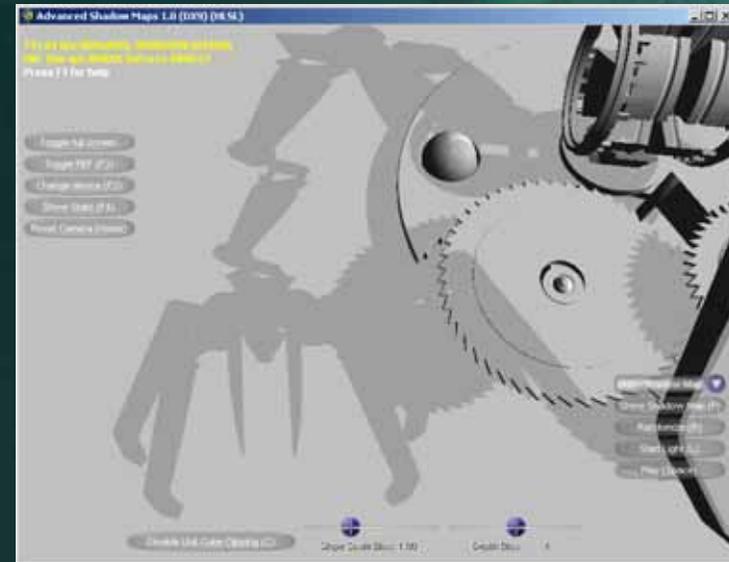
- Only focus shadow map where it is used
 - On the visible receivers
- For large light sources / scenes, this may be a small percentage of the total scene
- Shadow aliasing will depend on the viewer, but quality improvement can be substantial
 - Quality is always at least as good as *not* clipping



Unit Cube Clipping



Without Clipping

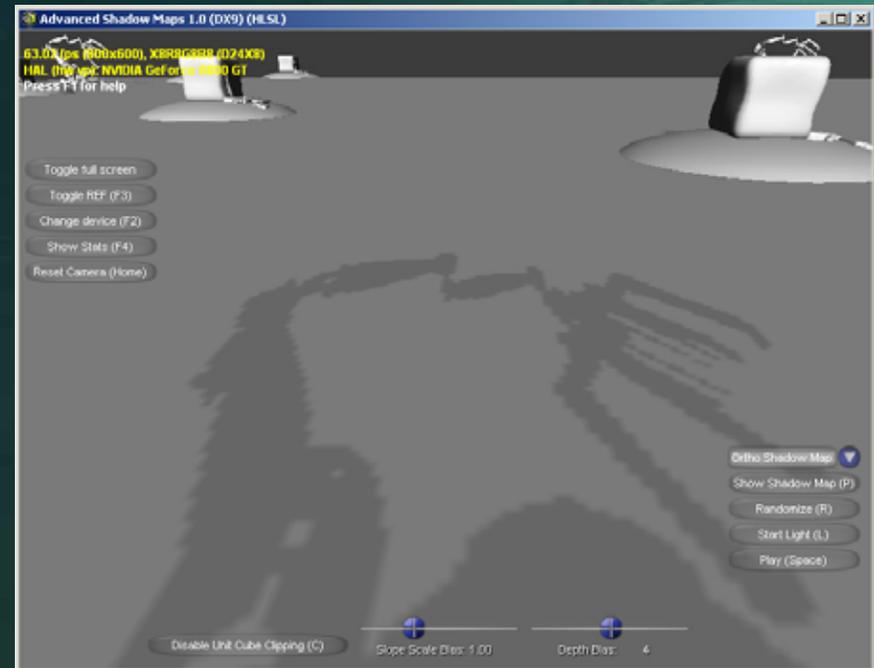


With Clipping

Unit Cube Clipping View Dependence



Viewer focused on small area



Viewer looking at distance



Implementation

- For each light, determine which visible objects will receive shadows from it
- Build light projection matrix to maximize the shadow map area used by these objects
 - Transform receiver AABBs by unmodified `lightProj`
 - Merge AABBs into full bounding rectangle
 - Compute ortho transform to bias & scale resulting AABB to cover the entire shadow map
 - `lightProj' = ortho * lightProj;`



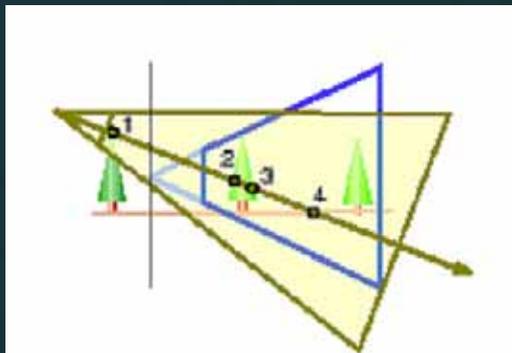
Cascaded Shadow Maps

- Satisfy area constraint using multiple maps
 - Extremely popular technique
- For each receiver in the scene, compute optimal $lightProj * lightView$, given $eyeView * P$.
- Many variants/trade-offs possible, to reduce number of maps required (at expense of quality)
- Handles duelling frusta and omnidirectional lights

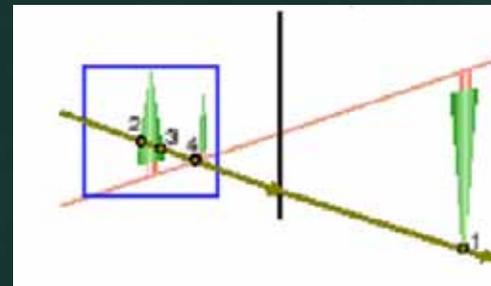


PSMs & Variants

- Why not warp shadow map by view projection?
- Perspective transform makes objects close to viewer larger than objects farther away
 - Increasing number of shadow texels focused on them



Eye Space



Post-Projective Space



Perspective Variants

- Perspective Shadow Maps
 - Stamminger & Drettakis, 2002
- Light-Space Perspective Shadow Maps
 - Wimmer, 2004
- Trapezoidal Shadow Maps & variants
 - Martin, 2004



Implementing PSMs

- Transform scene into post-projective view space
- Construct a projection transform for the light
 - From light's position in post-projective space
- Compose light projection matrix with scene
- Simple in concept, hard in practice
 - Post-projective space is filled with singularities
 - Spends too many shadow texels on nearby objects



Implementing LSPSMs

- Construct “Light Space” transformation based on view & light directions
- Compute a 3DH projection point based on frustum Z_{near} and Z_{far}
- Compose perspective transform with light space
- Perspective xform bounds based on CSG ops
 - Easy for directional lights, clipper required for others

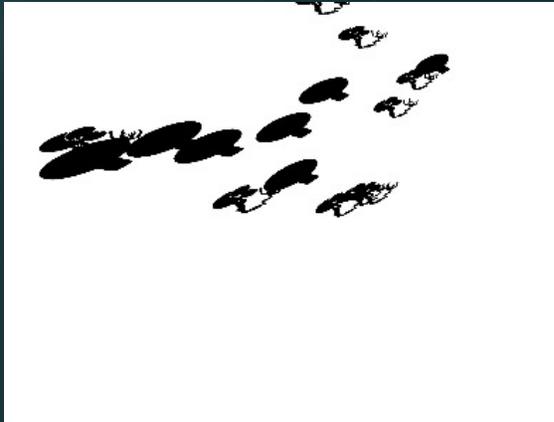


Implementing TSMs

- Project view frustum extremities by `lightProj`
 - `lightProj` is conventional shadow mapping matrix
 - Clip view frustum to light frustum for local lights
- Find 2D convex hull of projected frustum
- Construct a trapezoid that tightly bounds hull
- Use 2DH transform to convert trapezoid to square



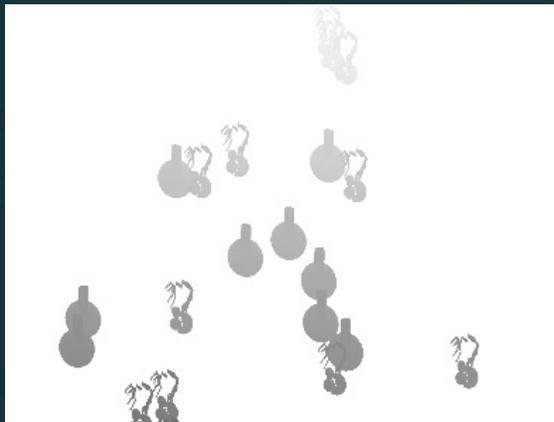
Shadow Map Comparison



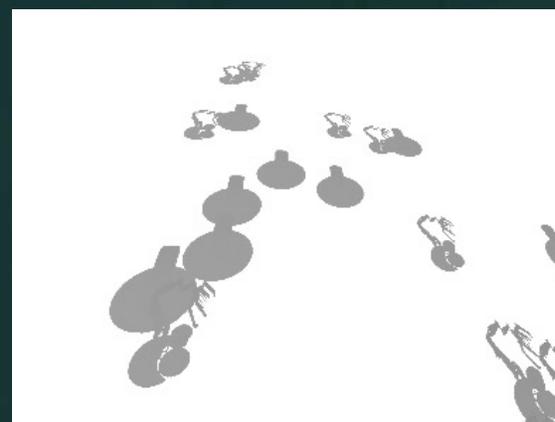
TSM



LSPSM



Uniform



PSM



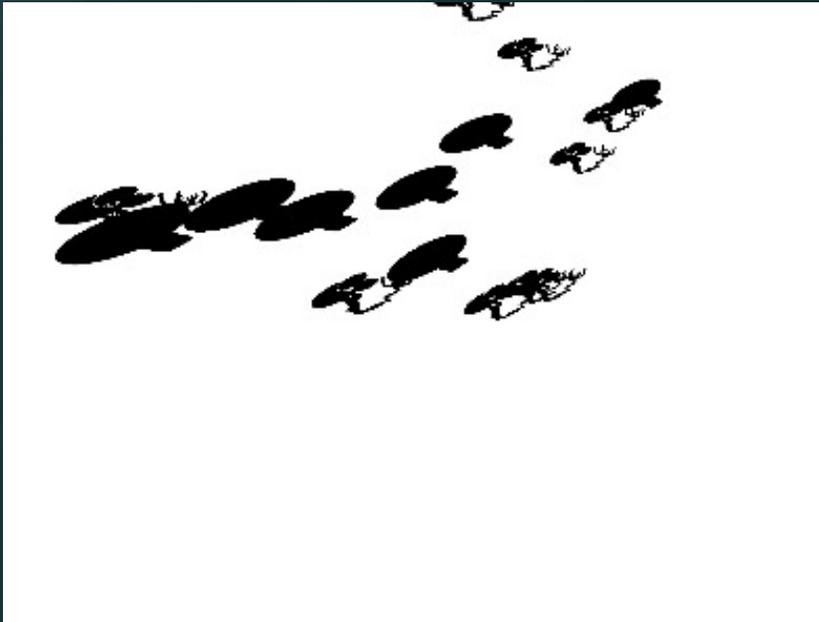
Implementation Issues

- No frame-to-frame caching
- Near/Far Shadow Map Quality
 - All variants have ways to balance this
- Depth Biasing
 - PSM variants much worse than uniform shadow maps
 - Slope/Scale frequently required, may not be enough
 - TSM is extreme case
 - So bad that authors recommend using depth replace to output non-warped Z instead

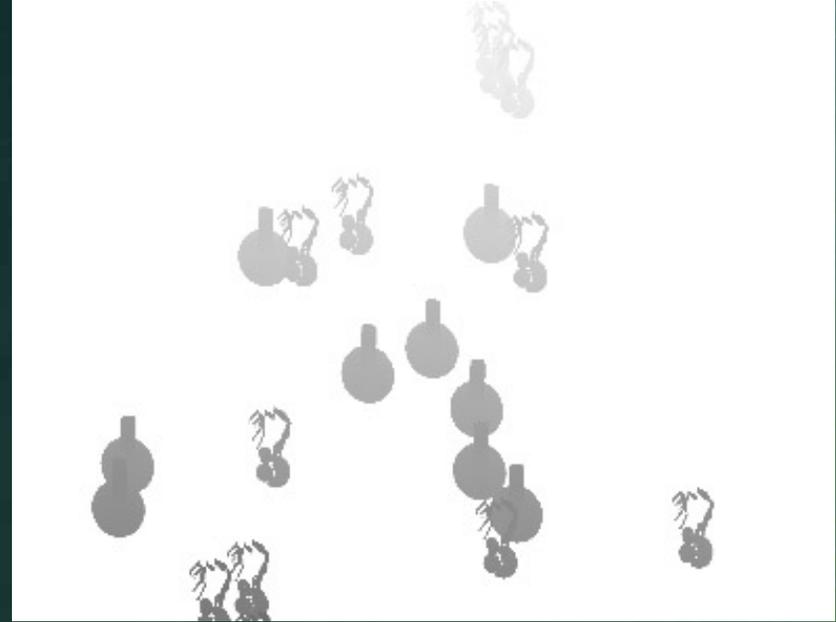


TSM Z vs Uniform Z

TSM



Uniform



8 bit depth values are all black for TSMs!



Optimization Parameters

- Each PSM variant has a way to trade-off near/far shadow quality
 - PSMs: Z_{infinity}
 - LSPSMs: N_{opt}
 - TSMs: 80% line
- All of these control the perspective distortion
- To get best quality, these terms must be adjusted dynamically (demo doesn't do this)

Infinite Light Demo



Advanced Shadow Maps 1.0 (DX9) (HLSL)

8.39 fps (800x600), X8R8G8B8 (D24X8)
HAL (hw vp): NVIDIA Quadro FX Go1000
Press F1 for help

Toggle full screen
Toggle REF (F3)
Change device (F2)
Show Stats (F4)
Reset Camera (Home)

TSM
Show Shadow Map (P)
Randomize (R)
Start Light (L)
Play (Space)

Disable Unit Cube Clipping (C)

Slope Scale Bias: 1.00
Depth Bias: 4
Focus Region Distance: 0.52

A 3D rendered scene showing a complex mechanical structure on the left, a unit cube on a circular platform in the center, and another unit cube on a circular platform on the right. The scene is lit from the top, creating long, soft shadows on the ground plane. The interface includes various control buttons and a status bar at the bottom.



Thoughts on PSM Variants

- Only practical for infinite lights
 - Local light implementations require clippers, 3D convex hulls, and many other expensive CPU ops
 - And none can handle duelling frusta
 - Not an issue for infinite lights
- LSPSMs and TSMs are much easier to implement than PSMs
 - Light and Trapezoid Space avoid major singularities
- TSM variants provide the best overall quality
 - Especially with unit-cube clipping

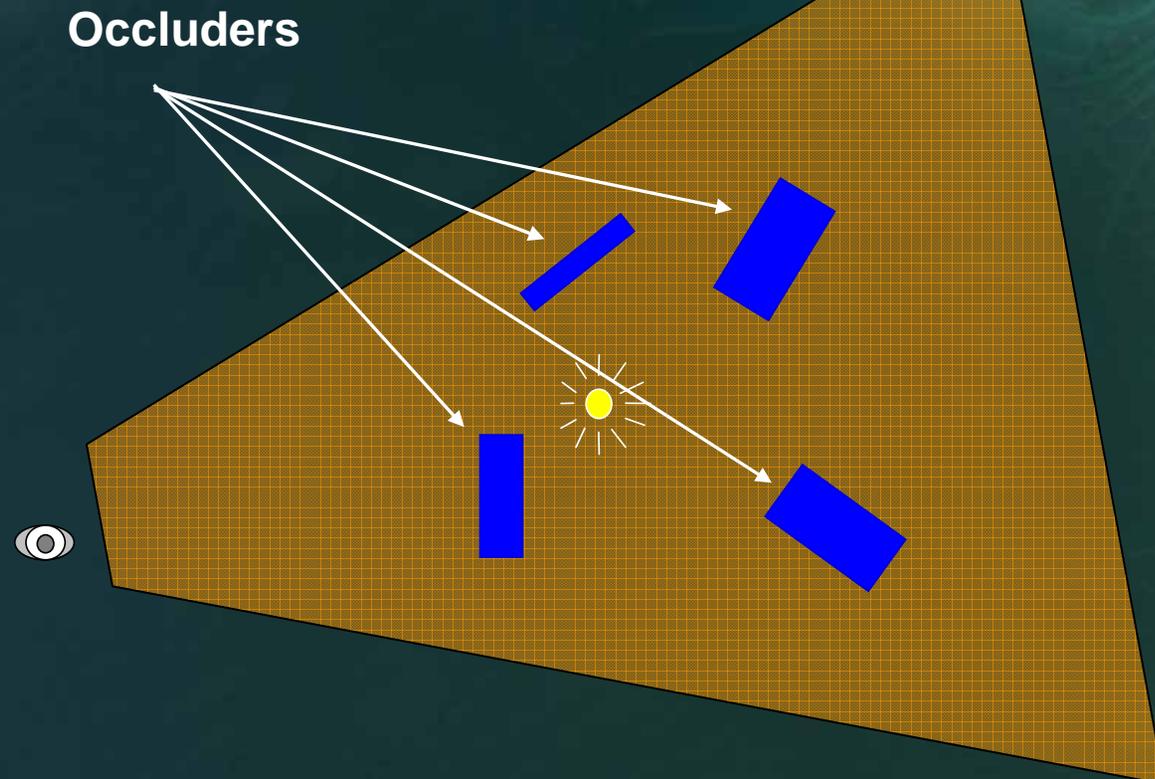


Cascaded PSM Variants

- For extreme outdoor scenes, a single PSM may be insufficient
- However, the texel redistribution is still beneficial
- Consider using cascaded PSMs/TSMs
 - Partition view frustum into chunks
 - Compute unique PSM/TSM transform for each chunk



Omnidirectional Shadowing



No one linear transform can handle all of the occluders



Solutions

- Nonlinear projections
 - Dual-parabaloid
 - Hyperbolic
 - Spherical
- Multiple linear projections
 - Cube Maps



Nonlinear Projections

- Hardware interpolation of x/w y/w is wrong
- Requires high scene tessellation
 - Reducing number of interpolated fragments
- Doesn't work well with deferred renderers
 - Since exact projection per-pixel during lighting significantly differs from inexact projection per-vertex
- GPGPU “scatter” operations would help, here...



Cube Maps

- 6 simple linear projections
 - No interpolation problems
- Single-pass lighting
- Don't play well with depth textures, so you lose:
 - Free bilinear PCF
 - 32ppc rendering
 - Extra texture memory to support both 2D and omnidirectional shadows



Why No DSTs?

- DSTs support is lost since cubemap distance is typically radial, rather than planar.
- However, eye-space Z of a 90 degree projection is just an unsigned permutation of $[s \ t \ r]$
 - If Z_{near} and Z_{far} are equal for all faces, window- Z is:

$$Z_s = \frac{-1}{MA} \times \frac{Z_{far} Z_{near}}{Z_{far} - Z_{near}} + \frac{Z_{far}}{Z_{far} - Z_{near}}$$

- Where $MA = \max(|s|, |t|, |r|)$

Virtual Shadow Depth Cube Textures



- This allows us to use planar depth, which can be generated by rendering to a DST
- And if the cube faces are tiled in a DST...
 - Then we can render omnidirectional shadows in 1 pass, just like with cube maps!
- Need a way to efficiently map from (s,t,r) to (x,y)
 - Use a smaller, precomputed G16R16 cubemap to perform this indexing

VSDCT Demo

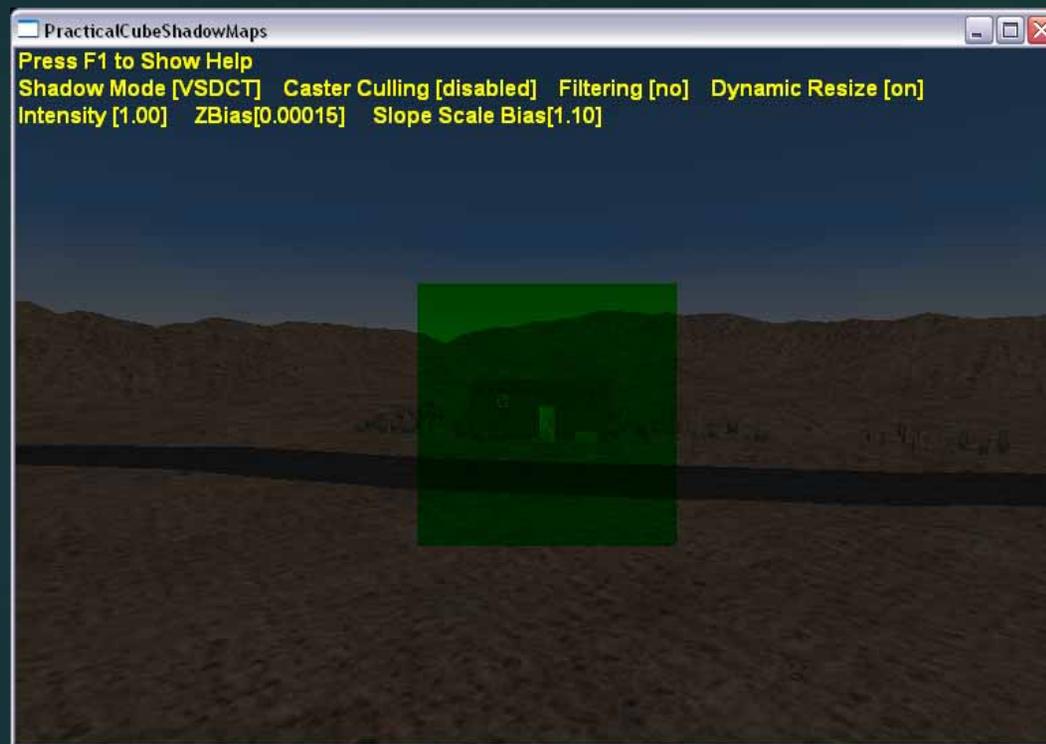


VSDCT filtering > 2x faster than cube filtering

Some General Performance Thoughts



- Use scissor to avoid shading pixels outside of light range; good when light is in view





General Performance Thoughts

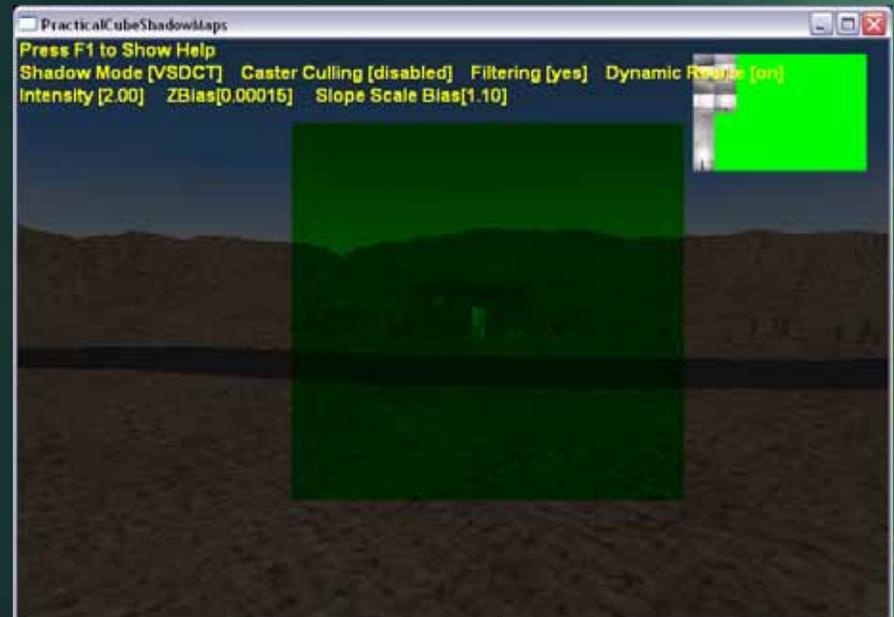
- Cull aggressively: if an omni light is outside the view frustum, at least one face can be culled





General Performance Thoughts

- When lights are small on-screen, they don't need high-res shadow maps.
 - Easily accomplished with VSDCTs, or mipmaps on cube maps





Summary

- Real-time shadow mapping isn't a solved problem
 - Lots of partial solutions for specific problems
- No one algorithm will work everywhere
 - And sometimes it is beneficial to combine them
- Always squeeze out optimizations
 - Shadows are expensive, don't spend perf when it doesn't help



Acknowledgements

- Sim, Ashu, Cass, Cem and other NVIDIAians whose work I've shamelessly stolen
- Will Newhall, for the key VSDCT insight
- Oles Shishkovstov, for convincing me to implement all of this
- Simon Kozlov, whose *GPU Gems* article made Perspective Shadow Maps clear to me
- And everyone else I'm forgetting, for everything I've taken for granted :)

Questions?



- gking@nvidia.com

The Source for GPU Programming

developer.nvidia.com

- Latest News
- Developer Events Calendar
- Technical Documentation
- Conference Presentations
- GPU Programming Guide
- Powerful Tools, SDKs and more ...

Join our FREE registered developer program for early access to NVIDIA drivers, cutting edge tools, online support forums, and more.



NVIDIA

developer.nvidia.com

©2004 NVIDIA Corporation. NVIDIA, and the NVIDIA logo are trademarks and/or registered trademarks of NVIDIA Corporation. Nalu is ©2004 NVIDIA Corporation. All rights reserved.