



**TAKE
CONTROL**

www.gdconf.com

MARCH 5-9
2007
SAN FRANCISCO

MOSCONE
CENTER





Real-time Atmospheric Effects in Games Revisited

Carsten Wenzel



WWW.GDCONF.COM



The deal

- ⊕ Follow up to a talk I gave at SIGGRAPH 2006
- ⊕ Covers material presented at the time plus recent additions and improvements



Overview

- ⊕ Introduction
- ⊕ Scene depth based rendering
- ⊕ Atmospheric effects breakdown
 - Sky light rendering
 - Fog approaches
 - Soft particles
 - Cloud rendering (updated/new)
 - Volumetric lightning approximation
 - River and Ocean rendering (updated/new)
- ⊕ Scene depth based rendering and MSAA (new)
- ⊕ Conclusions



Introduction

- ⊕ Atmospheric effects are important cues of realism (especially outdoors)
- ⊕ Why...
 - Create sense of depth
 - Increase level of immersion



Motivation

- ④ Atmospheric effects are mathematically complex (so far usually coarsely approximated if any)
- ④ Programmability and power of today's GPUs allow implementation of sophisticated models
- ④ How to can these be mapped efficiently?



Related Work

- ④ Deferred Shading (Hargreaves 2004)
- ④ Atmospheric Scattering (Nishita et al 1993)
- ④ Cloud Rendering (Wang 2003)
- ④ Real-time Atmospheric Effects in Games (Wenzel 2006)



Scene Depth Based Rendering: Motivation

- ④ Many atmospheric effects require accessing scene depth
- ④ Similar to Deferred Shading [Hargreaves04]
- ④ Mixes well with traditional style rendering
 - Deferred shading is **not** a must!
 - Think of it as writing a pixel shader with scene depth available
- ④ Requires laying out scene depth first and making it available to following rendering passes



Scene Depth Based Rendering: Benefits

- ④ Decouple rendering of opaque scene geometry and application of other effects
 - Atmospheric effects
 - Post-processing
 - More
- ④ Apply complex models while keeping the shading cost moderate
 - Features are implemented in separate shaders
 - Helps avoiding hardware shader limits (can support older HW)



Scene Depth Based Rendering: Concerns

- ⊕ Alpha-transparent objects
 - Only one color / depth value stored
 - However, per-pixel overdraw due to alpha transparent objects potentially unbound
 - Workaround for specific effects needed (will be mentioned later)



Scene Depth Based Rendering: API and Hardware Concerns

- ⊕ Usually cannot directly bind Z-Buffer and reverse map
- ⊕ Write linear eye-space depth to texture instead
- ⊕ Float format vs. RGBA8
- ⊕ Supporting Multi-Sample Anti-Aliasing is tricky (more on that later)



Recovering World Space Position from Depth

- ④ Many deferred shading implementations transform a pixel's homogenous clip space coordinate back into world space
 - 3 dp4 or mul/mad instructions
- ④ There's often a simpler / cheaper way
 - For full screen effects have the distance from the camera's position to its four corner points at the far clipping plane interpolated
 - Scale the pixel's normalized linear eye space depth by the interpolated distance and add the camera position (one mad instruction)



Sky Light Rendering

- ⊕ Mixed CPU / GPU implementation of [Nishita93]
- ⊕ Goal: Best quality possible at reasonable runtime cost
 - Trading in flexibility of camera movement
- ⊕ Assumptions and constraints:
 - Camera is always on the ground
 - Sky infinitely far away around camera
 - Win: Sky update is view-independent, update only over time



Sky Light Rendering: CPU

- ④ Solve Mie / Rayleigh in-scattering integral

For 128x64 sample points on the sky hemisphere solve...

$$I_v(\lambda) = I_s(\lambda)K(\lambda)F(\theta, g) \int_{P_a}^{P_b} \left(e^{\frac{-h}{H_0}} e^{(-t(P P_c, \lambda) - t(P P_a, \lambda))} \right) \quad (1)$$

Using the current time of day, sunlight direction, Mie / Rayleigh scattering coefficients

Store the result in a floating point texture

- ④ Distribute computation over several frames

Each update takes several seconds to compute



Sky Light Rendering: GPU

- ③ Map float texture onto sky dome
- ③ Problem: low-res texture produces blocky results even when filtered
 - Solution: Move application of phase function to GPU ($F(\theta, g)$ in Eq.1)
 - High frequency details (sun spot) now computed per-pixel
- ③ SM3.0/4.0 could solve Eq.1 via pixel shader and render to texture
 - Integral is a loop of ~200 asm instructions iterating 32 times
 - Final execution ~6400 instructions to compute in-scattering for each sample point on the sky hemisphere



Global Volumetric Fog

- ④ Nishita's model still too expensive to model fog/aerial perspective
- ④ Want to provide an atmosphere model
 - To apply its effects on arbitrary objects in the scene
- ④ Developed a simpler method to compute height/distance based fog with exponential fall-off



Global Volumetric Fog

$$f((x, y, z)^T) = be^{-cz}$$

$$\vec{v}(t) = \vec{o} + t\vec{d}$$

$$\oint f(\vec{v}(t))dt = \int_0^1 f((o_x + td_x, o_y + td_y, o_z + td_z)^T) \|\vec{d}\| dt$$

$$= be^{-co_z} \sqrt{d_x^2 + d_y^2 + d_z^2} \left[\frac{1 - e^{-cd_z}}{cd_z} \right]$$

$$F(\vec{v}(t)) = e^{-\oint f(\vec{v}(t))dt} \tag{2}$$

f – fog density distribution

b – global density

c – height fall-off

F – fog density along v

v – view ray from camera (o) to target pos (o+d), t=1



Global Volumetric Fog: Shader Implementation

Eq.2 translated into HLSL...

```
float ComputeVolumetricFog( in float3 cameraToWorldPos )
{
    // NOTE: cVolFogHeightDensityAtViewer = exp( -cHeightFalloff *
    cViewPos.z );

    float fogInt = length( cameraToWorldPos ) *
    cVolFogHeightDensityAtViewer;

    const float cSlopeThreshold = 0.01;
    if( abs( cameraToWorldPos.z ) > cSlopeThreshold )
    {
        float t = cHeightFalloff * cameraToWorldPos.z;
        fogInt *= ( 1.0 - exp( -t ) ) / t;
    }

    return exp( -cGlobalDensity * fogInt );
}
```



Combining Sky Light and Fog

- ⊕ Sky is rendered along with scene geometry
- ⊕ To apply fog...
 - Draw a full screen quad
 - Reconstruct each pixel's world space position
 - Pass position to volumetric fog formula to retrieve fog density along view ray
 - What about fog color?



Combining Sky Light and Fog

- ③ Fog color

 - Average in-scattering samples along the horizon while building texture

 - Combine with per-pixel result of phase function to yield approximate fog color

- ③ Use fog color and density to blend against back buffer



Combining Sky Light and Fog: Results



— *



Fog Volumes

- ⊕ Fog volumes via ray-tracing in the shader
- ⊕ Currently two primitives supported: Box, Ellipsoid
- ⊕ Generalized form of Global Volumetric Fog
 - Exhibits same properties (additionally, direction of height no longer restricted to world space up vector, gradient can be shifted along height dir)
- ⊕ Ray-trace in object space: Unit box, unit sphere
- ⊕ Transform results back to solve fog integral
- ⊕ Render bounding hull geometry
 - Front faces if outside, otherwise back faces
- ⊕ For each pixel...
 - Determine start and end point of view ray to plug into Eq.2




Fog Volumes

- ③ Start point
 - Either camera pos (if viewer is inside) or ray's entry point into fog volume (if viewer is outside)
- ③ End point
 - Either ray's exit point out of the fog volume or world space position of pixel depending which one of the two is closer to the camera
- ③ Render fog volumes back to front
- ③ Solve fog integral and blend with back buffer



Fog Volumes



Rendering of fog volumes: Box (top left/right), Ellipsoid (bottom left/right)



Fog and Alpha-Transparent Objects

- ④ Shading of actual object and application of atmospheric effect can no longer be decoupled
 - Need to solve both and combine results in same pass
- ④ Global Volumetric Fog
 - Approximate per vertex
 - Computation is purely math op based (no lookup textures required)
 - Maps well to older HW...
 - ④ Shader Models 2.x
 - ④ Shader Model 3.0 for performance reasons / due to lack of vertex texture fetch (IHV specific)



Fog and Alpha-Transparent Objects

⊕ Fog Volumes

Approximate per object, computed on CPU
Sounds awful but it's possible when designers know limitation and how to work around it

- ⊕ Alpha-Transparent objects shouldn't become too big, fog gradient should be rather soft

Compute weighted contribution by processing all affecting of fog volumes back to front w.r.t camera



Soft Particles

④ Simple idea

Instead of rendering a particle as a regular billboard, treat it as a camera aligned volume

Use per-pixel depth to compute view ray's travel distance through volume and use the result to fade out the particle

Hides jaggies at intersections with other geometry

Some recent publications use a similar idea and treat particles as spherical volumes

- ④ We found a volume box to be sufficient (saves shader instructions; important as particles are fill-rate hungry)

GS can setup interpolators so point sprites are finally feasible



Soft Particles: Results



Comparisons shots of particle rendering with soft particles disabled (left) and enabled (right) *



Clouds Rendering Using Per-Pixel Depth

- ④ Follow approach similar to [Wang03], Gradient-based lighting
- ④ Use scene depth for soft clipping (e.g. rain clouds around mountains) – similar to Soft Particles
- ④ Added rim lighting based on cloud density





Cloud Shadows



- ③ Cloud shadows are cast in a single full screen pass
- ③ Use depth to transform pixel position into shadow map space



Distance Clouds

- ⊗ Dynamic sky and pre-baked sky box clouds don't mix well
- ⊗ Real 3D cloud imposters can be expensive and are often not needed
- ⊗ Limited to 2D planes above the camera clouds can be rendered with volumetric properties
- ⊗ Sample a 2D texture (cloud density) along the view dir
 - For each sample point sample along the direction to sun
- ⊗ Adjust number of sample along both directions to fit into shader limits, save fillrate, etc.



Distance Clouds

- ③ Use the accumulated density to calc attenuation and factor in current sun / sky light



Distance Clouds at different times of day *



Volumetric Lightning Using Per-Pixel Depth

- ⊕ Similar to Global Volumetric Fog
 - Light is emitted from a point falling off radially
- ⊕ Need to carefully select attenuation function to be able to integrate it in a closed form
- ⊕ Can apply this lighting model just like global volumetric fog
 - Render a full screen pass



Volumetric Lightning Model

$$f((x, y, z)^T) = \frac{i}{1 + a \cdot \|l - (x, y, z)^T\|^2}$$

$$\vec{v}(t) = \vec{o} + t\vec{d}$$

$$\oint f(\vec{v}(t)) dt = \int_0^1 f((o_x + td_x, o_y + td_y, o_z + td_z)^T) \|\vec{d}\| dt$$

$$= 2i\sqrt{d_x^2 + d_y^2 + d_z^2} \left[\frac{\arctan\left(\frac{v + 2w}{\sqrt{4uw - v^2}}\right) - \arctan\left(\frac{v}{\sqrt{4uw - v^2}}\right)}{\sqrt{4uw - v^2}} \right]$$

$$= F(\vec{v}(t)) \tag{3}$$

- f – light attenuation function i – source light intensity
- l – lightning source pos a – global attenuation control value
- v – view ray from camera (o) to target pos (o+d), t=1
- F – amount of light gathered along v



Volumetric Lightning Using Per-Pixel Depth: Results





River shading

- ④ Rivers (and water areas in general)
- ④ Special fog volume type: Plane
- ④ Under water fog rendered as described earlier (using a simpler uniform density fog model though)
- ④ Shader for water surface enhanced to softly blend out at riverside (difference between pixel depth of water surface and previously stored scene depth)



River shading: Results



River shading –
Screens taken from a hidden section of the E3 2006 demo *



Ocean shading


- ④ Very similar to river shading, however...
- ④ Underwater part uses more complex model for light attenuation and in-scattering
- ④ Assume horizontal water plane, uniform density distribution and light always falling in top down
- ④ Can be described as follows...



Ocean shading



Ocean shading: Results



Underwater view: from ground up (1st row), from underneath the surface down (2nd row). Same lighting settings apply. Higher density on the right column. *



Scene depth based rendering and MSAA

- ⊕ Several problems
 - Cannot bind multi-sampled RT as texture
 - Shading per pixel and not per sample
- ⊕ Need to resolve depth RT which produces wrong values at silhouettes → potentially causes outlines in later shading steps
- ⊕ Two problems we ran into
 - Fog
 - River / Ocean



Scene depth based rendering and MSAA: Fog

- ⊕ Fog color doesn't changed drastically for neighboring pixel while density does
- ⊕ Have fog density computed while laying out depth (two channel RT)
- ⊕ During volumetric fog full screen pass only compute fog color and read density from resolved RT
- ⊕ Averaging density during resolve works reasonably well compared to depth



Scene depth based rendering and MSAA: River / Ocean

- ④ Shader assumes dest depth $>$ plane depth (otherwise pixel would have be rejected by z-test)
- ④ With resolved depth RT this cannot be guaranteed (depends on pixel coverage of object silhouettes)
- ④ Need to enforce original assumption by finding max depth of current pixel and all neighbors (direct neighbors suffice)



Scene depth based rendering and MSAA: Results



Fog full screen pass with MSAA disabled (left) / enabled (right)



River / Ocean shading artifact (left) and fix (right)



Conclusion

- ④ Depth Based Rendering offers lot's of opportunities
- ④ Demonstrated several ways of how it is used in *CryEngine2*
- ④ Integration issues (alpha-transparent geometry, MSAA)



Kualoa Ranch on Hawaii –

Real world photo (left), internal replica rendered with *CryEngine2* (right)

WWW.GDCONF.COM



References

- ④ [Hargreaves04] Shawn Hargreaves, “Deferred Shading,” Game Developers Conference, D3D Tutorial Day, March, 2004.
- ④ [Nishita93] Tomoyuki Nishita, et al., “Display of the Earth Taking into Account Atmospheric Scattering,” In Proceedings of SIGGRAPH 1993, pages 175-182.
- ④ [Wang03] Niniane Wang, “Realistic and Fast Cloud Rendering in Computer Games,” In Proceedings of SIGGRAPH 2003.
- ④ [Wenzel06] Carsten Wenzel, “Real-time Atmospheric Effects in Games,” SIGGRAPH 2006.



Questions

???

WWW.GDCONF.COM



Acknowledgements

WWW.GDCONF.COM



Acknowledgements

WWW.GDCONF.COM