

**Advanced Computer Graphics**  
**CS 563: “Real-Time Volumetric**  
**Shadows using 1D Min-Max Mipmaps”**

---

Rob Martin

*Computer Science Dept.*  
*Worcester Polytechnic Institute (WPI)*







# Introduction

- We want to model the effects of light scattering in a *participating medium*
- Dust, moisture, etc.
- Difficult to do in real time since scattering occurs within a volume; i.e., all points along a camera ray
- This algorithm is an attempt at recreating this effect through the use of approximation techniques, epipolar shadow maps, min-map mipmaps, and parallelism.

# Introduction







# Assumptions

- Use single-scattering model [Blinn 1982]
  - Approximates more complex physics
  - Simple, realistic results
- Assume the scattering medium is homogeneously distributed.



# Naïve approach: Ray marching

- Render shadow map from light source
- Shoot a ray from the camera for each pixel
- Walk down the ray at equal intervals from the camera to the first intersection
- Take sample at each interval
- Compute final radiosity from samples



# Algorithm Overview

- Render depth map from the camera and a shadow map from the light source
- Perform ‘epipolar rectification’ on the shadow map
- Compute approximate scattering integral
- Build min-max mipmaps from shadow map
- Traverse mipmaps for each camera ray to determine lit segments
- (If epipole is on or near screen, use ray marching)

# Single Scattering



$$L(\mathbf{v}) = \int_0^d e^{-\sigma_t s} V(s\mathbf{v}) \sigma_s \rho(\theta) L_{\text{in}}(s\mathbf{v}) ds, \quad (1)$$

$$L_{\text{in}}(s\mathbf{v}) = \frac{I e^{-\sigma_t d(s\mathbf{v})}}{d(s\mathbf{v})^2}, \quad d(s\mathbf{v}) = \|\mathbf{x} - s\mathbf{v}\|, \quad (2)$$

$\mathbf{v}$  : normalized direction vector

$d$  : distance to first occluder

$\sigma_s$  : scattering coef.

$\sigma_t$  :  $\sigma_s - \sigma_a$  extinction coef.

$V(s\mathbf{v})$  : 1 if in light, 0 otherwise

$L_{\text{in}}(s\mathbf{v})$  : radiance assuming no occlusion

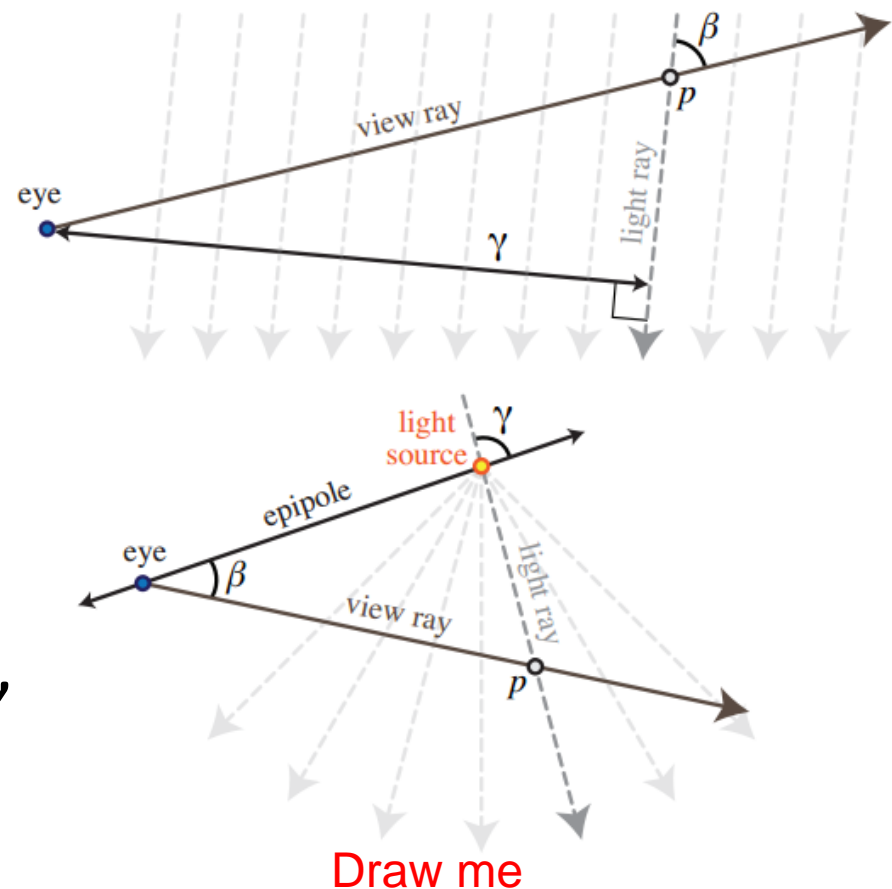
$\rho(\theta)$  : scattering phase function





# Epipolar Rectification

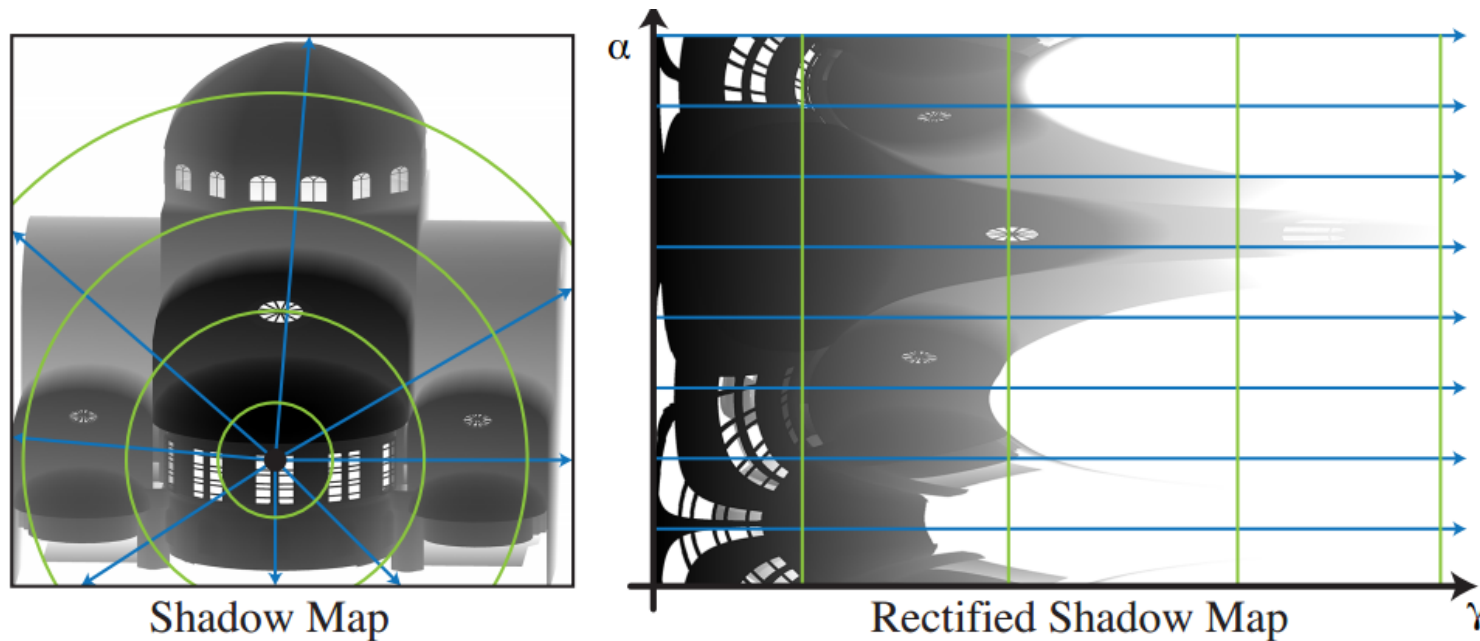
- Conversion from cartesian coordinate system to epipolar coordinates
- 3-dimensional space
- Every point has  $\alpha, \beta, \gamma$
- Camera rays are indexed by  $(\alpha, \beta)$
- Light rays by  $(\alpha, \gamma)$
- $\alpha$  determines the 'slice'





## Epipolar Rectification, cont.

- To rectify the shadow map, rows are indexed by  $\alpha$ , columns by  $\gamma$ ; each pixel contains the  $\beta$
- Do not need to rectify the camera



# Revamped Scattering Integral



$$L(\alpha, \beta) = \int_0^{D(\alpha, \beta)} e^{-\sigma_t s(\beta, \gamma)} V(\alpha, \beta, \gamma) \sigma_s \rho(\beta, \gamma) L_{\text{in}}(\beta, \gamma) \frac{ds}{d\gamma} d\gamma, \quad (3)$$

$$L(\alpha, \beta) = \int_0^{D(\alpha, \beta)} V(\alpha, \beta, \gamma) \mathcal{I}(\beta, \gamma) d\gamma, \quad (4)$$

$D(\alpha, \beta)$  : the  $\gamma$  coordinate where the camera ray hits the first occluder



# Integral approximations

$$L(\alpha, \beta) \approx \sum_i^N \left( B_i(\beta) \int_0^{D(\alpha, \beta)} V(\alpha, \beta, \gamma) \Gamma_i(\gamma) d\gamma \right). \quad (5)$$

$$L(\alpha, \beta) \approx \sum_i^N \left( B_i(\beta) \sum_{\substack{\gamma < D(\alpha, \beta) \\ S[\alpha, \gamma] > \beta}} \Gamma_i(\gamma) \Delta\gamma \right), \quad (6)$$

- The I term is too complex to compute, so approximate by sampling in a 64x64 grid
- Approximate integral in equation 5 by using a Riemann sum; integrate in the visibility term (V)

# Textured Lights



$$L(\alpha, \gamma) \approx \sum_i^N \left( B_i(\beta) \sum_{\substack{\gamma < D(\alpha, \beta) \\ S[\alpha, \gamma] > \beta}} \Gamma_i(\gamma) T(\alpha, \gamma) \Delta\gamma \right). \quad (7)$$

- To enable textured lights, must also rectify the texture map
- New texture coordinates given by  $T(\alpha, \gamma)$
- Must compute prefix sums of  $\Gamma_i(\gamma) T(\alpha, \gamma)$
- Must do this for each  $\alpha$




# 1D Min-Max Mipmaps

- Binary tree structure
- Each node holds the minimum and maximum value of all child nodes
- Each row of the rectified shadow map corresponds to an  $\alpha$  slice in epipolar coordinates
- Create a min-max mipmap for each row/slice



# Min-Max Mipmaps, cont.



- Use the mipmaps to quickly find lit/unlit areas of the camera ray
- For a ray with angle  $\beta$ , if  $\beta$  is less than minimum, then the area is lit. If greater than the maximum, the area is in shadow.





# Parallelism

- With the rectified shadow maps, min-max mipmaps, and precomputed scattering integrals, each camera ray can be computed in parallel on a fragment/pixel shader
- Key for algorithm performance



# Results





## Results, cont.

- Incremental integration is a very similar algorithm
- Main difference is the data structure used to detect lit/unlit areas
- Uses a dynamic data structure called a partial sum tree
- Is not static, nor is it parallelizable



# Limitations

- Homogeneous scattering medium
- Anisotropic mediums mess up the integral approximations
- Shadow maps are susceptible to aliasing
- Must brute force near the epipole



# Conclusions

- Effective real-time algorithm for producing the effect of shadow volumes and participating medium
- Mathematical techniques simplify computations by pre-computing static values used in the integration
- Simplified data structure allows for parallelism
- Algorithm can run on most hardware (DX9)



## References

- Blinn, J. F. 1982. Light reflection functions for simulation of clouds and dusty surfaces. 21–29
- Jiawen Chen, Ilya Baran, Fredo Durand, Wojciech Jarosz. *Real-Time Volumetric Shadows using 1D Min-Max Mipmaps*. Proceedings of the 2011 Symposium on Interactive 3D Graphics and Games, 2011.