

## Arbitrary BRDFs [Kautz02]



$$r_p(\vec{v}) = \int L^{tran}(\vec{s}) f(\vec{v}, \vec{s}) H_N(\vec{s}) d\vec{s}$$

$$r_p(\vec{v}) = \int \left[ \sum l_i y_i(\vec{s}) \right] f(\vec{v}, \vec{s}) H_N(\vec{s}) d\vec{s}$$

$$r_p(\vec{v}) = \sum l_i \int y_i(\vec{s}) f(\vec{v}, \vec{s}) H_N(\vec{s}) d\vec{s}$$

$$r_{pv} = b_{pv}^T R_p M_p l$$



## Arbitrary BRDF Results



Anisotropic BRDFs

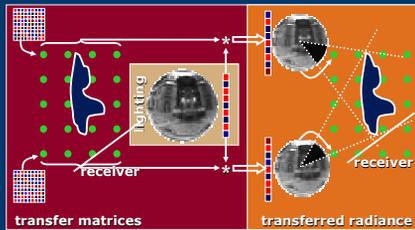
Other BRDFs

Spatially Varying

## Neighborhood Transfer



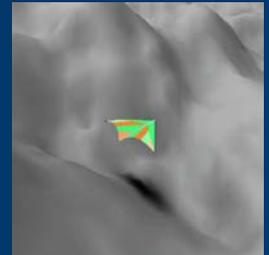
- Allows to cast shadows/caustics onto arbitrary receivers
- Store how object scatters/blocks light around itself (transfer matrices on grid)



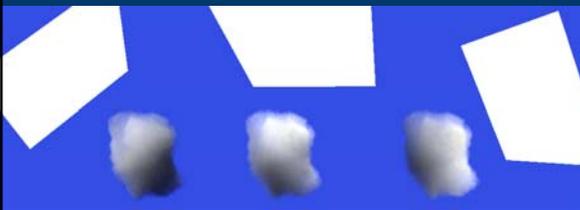
## Neighborhood Transfer Results



- 64x64x8 neighborhood
- diffuse receiver
- timings on 2.2Ghz P4, ATI Radeon 8500
- 4fps if light changes
- 120fps for constant light

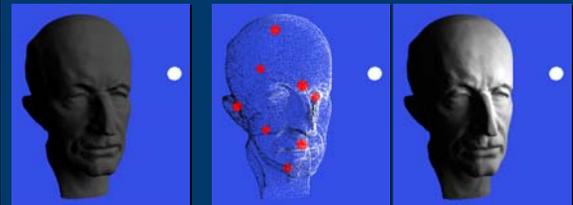


## Volumes



- Diffuse volume: 32x32x32 grid
- Runs at 40fps on 2.2Ghz P4, ATI 8500
- Here: dynamic lighting

## Local Lighting using Radiance Sampling



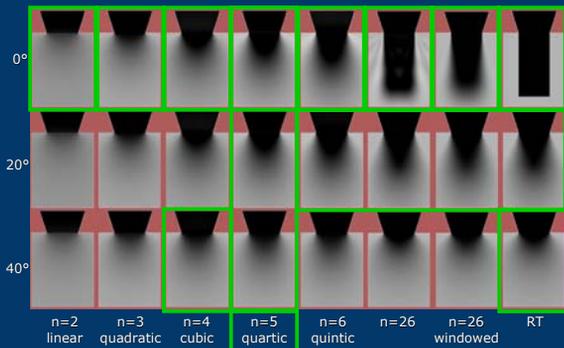
single sample  
(at center = light at  $\infty$ )

multi-sample  
locations

multi-sample  
result

- Sample incident radiance at **multiple** points
- Choose sample points over object using ICP from VQ
- Correct for shadows but not interreflections

## Light Size vs. SH Order



## Results



Live Demo

## Conclusions



### Pros:

- Fast, arbitrary dynamic lighting
  - on surfaces or in volumes
- Includes shadows and interreflections
- Works for diffuse and glossy BRDFs

### Cons:

- Works only for low-frequency lighting
- Rigid objects only, no deformation

## Future Work



- Practical glossy transfer
  - Eliminate frozen view/light constraints
  - Compress matrices/vectors
- Enhanced preprocessing
  - Subsurface scattering, dispersion
  - Simulator optimization
  - Adaptive sampling of transfer over surface
- Deformable objects

## Acknowledgements



- Thanks to:
  - Jason Mitchell & Michael Doggett (ATI)
  - Matthew Papakipos (NVidia)
  - Paul Debevec for light probes
  - Stanford Graphics Lab for Buddha model
  - Michael Cohen, Chas Boyd, Hans-Peter Seidel for early discussions and support

Questions?

## Performance



Model	# Verts	GF4 4600 FPS	R300 FPS	Pre- compute
Max	50,060	215	304	1.1h
Buddha	49,990	191	269	2.5h
Tweety	48,668	240	326	1.2h
Tyra	100,000	118	179	2.4h
Teapot	152,413	93	154	4.4h

## Matrix Formulation



$$M_i = \int L(\vec{s}) B_i(\vec{s}) V(\vec{s}) d\vec{s}$$

$$M_i = \int \left( \sum_j l_j B_j(\vec{s}) \right) B_i(\vec{s}) V(\vec{s}) d\vec{s}$$

$$M_i = \sum_j l_j \int B_j(\vec{s}) B_i(\vec{s}) V(\vec{s}) d\vec{s}$$

$$M_{ij} = \int B_i(\vec{s}) B_j(\vec{s}) V(\vec{s}) d\vec{s}$$

## Results – Preprocessing

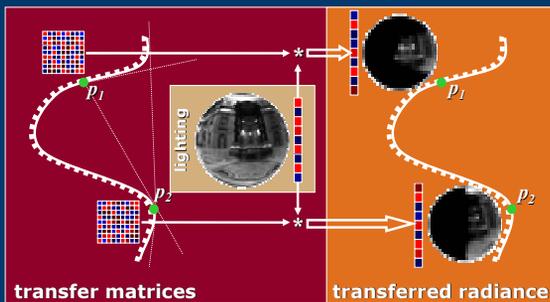


Model	Type	Sampling	Preproc.	FPS
head	diffuse	50K vert.	1.1h	129
ring	diffuse	256x256 t.	8m	94
buddha	diffuse	50K vert.	2.5h	125
buddha	glossy	50K vert.	2.5h	..125
tyra	diffuse	100K vert.	2.4h	83
tyra	glossy	100K vert.	2.4h	..83
teapot	glossy	150K vert.	4.4h	..49
cloud	diffuse	32x32x32	15m	40
glider	neighb.	64x64x8	3h	..120

## Transfer Matrix



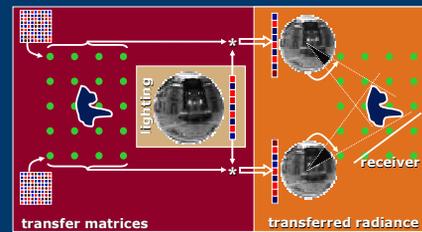
Precompute how global lighting  $\Rightarrow$  transferred lighting



## Neighborhood Transfer



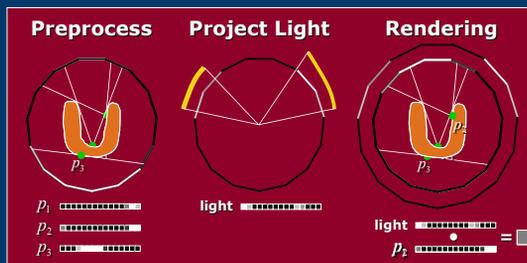
- Store how object transfers radiance around itself
- Transfer matrices on grid



## Diffuse Self-Transfer



2D example, piecewise constant basis, shadows only



## Previous Work – Precomputed Transport

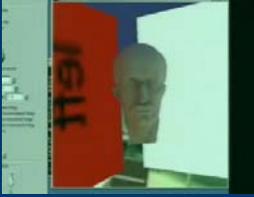


- [Greger96] irradiance volumes  
move diffuse object through precomputed lighting
- [Miller98, Wood00, Chen02] surface lightfields  
frozen lighting environments
- [Ashikmin02] steerable illumination textures  
steers small light source over diffuse object
- [Matusik02] image-based 3D photography  
surface lightfield + reflectance field – not interactive

## Dynamic Lighting



- Sample incident lighting  $L_p$  on-the-fly



- Results
  - low-resolution cube maps sufficient: 6x16x16
  - average error: 0.2%, worst-case: 0.5%
  - takes 1.16 ms on P3-933Mhz, ATI 8500