

Lecture 12: Interactive Ray Tracing and Acceleration Structures

Fall 2004
Kavita Bala
Computer Science
Cornell University

HW 1

- Add whatever you need to
 - Get color in the materials
 - Diffuse, etc.
- Only direct lighting
- Only hard shadows
- So why spheres? So that radiosity/radiance conversions etc. work out.

© Kavita Bala, Computer Science, Cornell University

Interactive Software Rendering

- Interactive
 - User-driven, not pre-scripted animation
 - At least a few frames per second (fps)
- Software
 - Major shading done in software
 - Can use hardware to help
- Rendering
 - Online, not pre-computed or captured
 - Eg, lightfields are pre-computed

© Kavita Bala, Computer Science, Cornell University


Why Software Rendering?

- Global Illumination: Non-local information
- Extremely high complexity
- Arbitrary shading models
- Portability
 - No tweaking: just works
 - No scene dependent optimizations

© Kavita Bala, Computer Science, Cornell University

Performance Results II


- Comparison to Rasterization-Hardware
 - Ray tracing scales well for large environments



© Kavita Bala, Computer Science, Cornell University

Rendering as Sampling

- Ray tracers compute radiance at each pixel




- Rendering = Sampling radiance

© Kavita Bala, Computer Science, Cornell University

Coherence

- Within one frame: spatial coherence




- Across many frames: temporal coherence

© Kavita Bala, Computer Science, Cornell University

Strategy


- Insight: radiance is mostly smooth -- use sparse sampling and reconstruction



- Radiance samples are very expensive
- Goal: reconstruct most pixels by interpolation
- Issues: discontinuities, non-linear variations

© Kavita Bala, Computer Science, Cornell University

Modified Visual Feedback Loop



© Kavita Bala, Computer Science, Cornell University

Display Process

- Automatically exploit spatial and temporal coherence
- Layered on top of an existing (slow) global illumination renderer
- Provide interactive performance

© Kavita Bala, Computer Science, Cornell University

Aside: Frameless Rendering

- Update pixels as they are computed
 - Don't wait for full frame to finish



© Kavita Bala, Computer Science, Cornell University

Frameless Rendering




Figure 1. The bottom row shows 7 frames of a 15 Hz frameless rendering sequence with 33% of the pixels updated in each frame. The middle row shows 3 frames of a double-buffering sequence updated at 5 Hz.

© Kavita Bala, Computer Science, Cornell University

Sparse Sampling Approaches

- 4D:
 - Radiance Interpolants
 - Holodeck
- 2D: Image based
 - Post-rendering Warp
 - Render Cache
 - Edge and Point Rendering
 - Corrective Texturing

© Kavita Bala, Computer Science, Cornell University


Post-Rendering 3D Warp

- Render subset of frames
 - E.g, every 6th frame is rendered
- Use standard image warping techniques to compute the other frames

© Kavita Bala, Computer Science, Cornell University

Aside: Pixel Reprojection


- Goal: Want image at new viewpoint
- Reproject points from input images



© Kavita Bala, Computer Science, Cornell University

Aside: Pixel Reprojection

- Assume have depth/disparity per pixel
- If pixel (x,y) sees point P,
 - $P = C + t D$
 - C is camera position,
 - D is direction from C through (x,y)
 - t is distance along D




© Kavita Bala, Computer Science, Cornell University

Aside: Pixel Reprojection

- Direction D

$$D = C + x i + y j + d k$$

(x, y) = pixel



- C = camera center
- d = distance of image plane from C
- C, d are known


© Kavita Bala, Computer Science, Cornell University

Aside: Pixel Reprojection

$$P = C_0 + t_0 D_0(x,y)$$

$$C_0 + t_0 D_0 = C_1 + t_1 D_1$$

$$t_1 D_1 = (C_0 - C_1) + t_0 D_0$$




C_0, C_1, D_0, t_0 are known

$t_1 D_1$ defines the reprojected pixel

© Kavita Bala, Computer Science, Cornell University

Aside: Pixel Reprojection


- $D_1 = C_1 + x_1 i + y_1 j + d_1 k$
- Solve for x_1 , y_1 and t_1



© Kavita Bala, Computer Science, Cornell University

Post-Rendering 3D Warp

- Problem:
 - Pixels do not project to pixel centers
 - Multiple pixels project to same pixel in new view
 - Holes and missing data




Reference frame Warped frame
The camera is moving to the left in this example.

© Kavita Bala, Computer Science, Cornell University

How to fill holes?


- Warp from both past and future reference frames
 - Heuristics for combining pixel results



© Kavita Bala, Computer Science, Cornell University

Problem: Post-Rendering Warp

- Must predict the locations of future frames
 - Longer predictions become rapidly less accurate



© Kavita Bala, Computer Science, Cornell University

Sparse Sampling Approaches

- 4D:
 - Radiance Interpolants
 - Holodeck
- 2D: Image based
 - Post-rendering Warp
 - Render Cache
 - Edge and Point Rendering
 - Corrective Texturing

© Kavita Bala, Computer Science, Cornell University


Render Cache (Walter et al.)

- Interactivity is important
 - Maintain relatively constant framerate
 - e.g., > 5 fps
 - Degrade gracefully as rendering becomes more expensive
- Cache shaded pixels as 3D colored points
- Render new image
 - Project points onto current image plane
 - Filter to reduce artifacts
- Prioritize future rendering
 - Identify problem pixels
 - Sparse sampling for limited render budget

© Kavita Bala, Computer Science, Cornell University

Approach

- Data: Cloud of unordered points with:
 - 3D position, color, age, object id
- Approach: reproject points into image plane
 - Occlusion errors, holes in data




Initial view After reprojection

© Kavita Bala, Computer Science, Cornell University

Image Estimation

- Depth cull heuristic
 - Problem: occluded points may be visible
 - Z-buffering only works within a pixel
 - Clean up using nearby depth information
- Interpolate 3x3




Raw projection depth cull interpolation

© Kavita Bala, Computer Science, Cornell University

Sampling

- Choose pixels for rendering: sparse sampling
- Requested pixels sent to renderer(s)
 - Results returned at some later frame




Displayed image Priority image Requested pixels

© Kavita Bala, Computer Science, Cornell University

Render Cache Adv and Limitations

- Improved interactivity
- Independent display process
- Drawback: pixel artifacts



Ray tracing

Path tracing

© Kavita Bala, Computer Science, Cornell University


Sparse Sampling Approaches

- 4D:
 - Radiance Interpolants
 - Holodeck
- 2D: Image based
 - Post-rendering Warp
 - Render Cache
 - Edge and Point Rendering
 - Corrective Texturing

© Kavita Bala, Computer Science, Cornell University

Edge-and-Point [Bala'03]

- Goal: Interactive high-quality rendering
 - Expensive shading: e.g., global illumination
 - But, mostly smooth (coherent)





points

naïve reconstruction

© Kavita Bala, Computer Science, Cornell University

Edge-and-Point Motivation: Performance


- Discontinuities are perceptually important
 - Artifacts are disturbing




- Finding discontinuities by sampling is expensive

© Kavita Bala, Computer Science, Cornell University

Edge-and-Point Rendering



Edges: important discontinuities
 – Silhouettes and shadows
 Points: sparse shading samples





edge-and-point reconstruction

© Kavita Bala, Computer Science, Cornell University

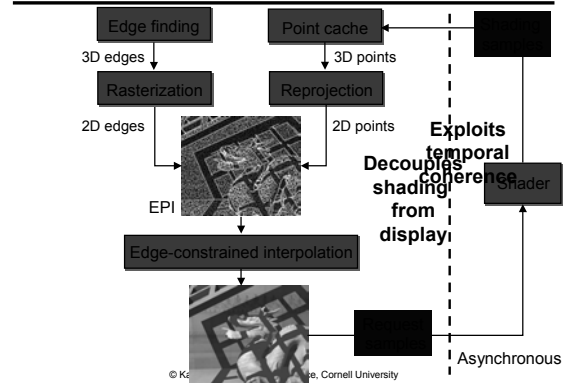
Edge-and-Point Image

- Alternative display representation
- Edge-constrained interpolation preserves sharp features
- Fast anti-aliasing




© Kavita Bala, Computer Science, Cornell University

System



© Kavita Bala, Computer Science, Cornell University

Silhouettes




$$N_1 \cdot V > 0 \text{ (forward facing)}$$

$$N_2 \cdot V < 0 \text{ (backward facing)}$$


© Kavita Bala, Computer Science, Cornell University

Shadows: Hard and Soft



© Kavita Bala, Computer Science, Cornell University

Umbra and Penumbra Conditions




- Event plane tangential to light and blocker
 $L \cdot N_{\text{blocker}} = L \cdot N_{\text{light}} = 0$
 $N_{\text{light}} \cdot N_{\text{blocker}} = 1 \text{ (umbra)}, -1 \text{ (penumbra)}$

© Kavita Bala, Computer Science, Cornell University

Edge Finding


- Hierarchical trees: fast edge finding
 - Interval-based



© Kavita Bala, Computer Science, Cornell University

Soft Shadow Edges

Black: silhouettes,
Red: umbral edges, Blue: penumbral edges



© Kavita Bala, Computer Science, Cornell University

Pixel types

- Pixels can have arbitrary edge complexity
- Classify pixels into 3 groups
 - Empty: no edges
 - Simple: can be approximated by 1 edge
 - Complex: everything else




- Typical pixel classification statistics
 - empty (85-95%), simple (4-10%), complex (1-4%)

© Kavita Bala, Computer Science, Cornell University

Edge-and-Point Image (EPI)

- Goal: compact and fast
 - Store at most one edge and one point per pixel
 - Limited sub-pixel precision




- Combine edges and points in image space
 - View-driven, lazy evaluation

© Kavita Bala, Computer Science, Cornell University

Reachability


- Reachable samples
 - Pixel's 5x5 neighborhood
 - Connected without crossing any edges (or complex pixels)
- Propagated outward from each pixel



© Kavita Bala, Computer Science, Cornell University

Results: Quality

- Global illumination
- 3 lights
- 150k polygons





Without Edges With Edges

© Kavita Bala, Computer Science, Cornell University

Tea Stand

- Global illumination



© Kavita Bala, Computer Science, Cornell University

Sparse Sampling Approaches

- 4D:
 - Radiance Interpolants
 - Holodeck
- 2D: Image based
 - Post-rendering Warp
 - Render Cache
 - Edge and Point Rendering
 - Corrective Texturing


© Kavita Bala, Computer Science, Cornell University

Corrective Texturing

- Start with a standard hardware rendering of scene
 - Graphics hardware very good at interactive display
 - Start with a radiosity solution
- Compare to underlying renderer
 - Apply corrections where they differ
 - Corrections applied as projective textures

© Kavita Bala, Computer Science, Cornell University

Corrective Texturing



Radiosity
solution

Corrected
image

Corrective
texture

© Kavita Bala, Computer Science, Cornell University


Corrective Texturing

- Sparse rendered samples compared to hardware displayed results
 - Differences splatted into textures
 - More samples generated near points which had large differences
 - Samples which are likely to have changed are deprecated so that can be overwritten by future results

© Kavita Bala, Computer Science, Cornell University

Corrective Texturing

- Corrective textures are dynamically assigned to objects



© Kavita Bala, Computer Science, Cornell University

Comparison

	Target renderer speed	Sparseness ratio	Typical frame rates
Warp	< 1s	4 - 10	20 - 60 fps
Corrective Tex.	20 - 200s	250 - 1000	5 - 10 fps
RC	.5 - 10s	8 - 100	10 - 20 fps
EPI	.5 - 10s	8 - 100	10 - 20 fps

© Kavita Bala, Computer Science, Cornell University

Comparison

	Hardware accelerated	Independent of scene complexity	Moving objects	Quality
Warp	No	Yes	No	?
Corrective Tex.	Yes	No	No	Not really
RC	No	Yes	Yes	No
EPI	Yes	Yes	Yes	Yes

© Kavita Bala, Computer Science, Cornell University

Prediction

- Hardware
 - Speed
 - Programmability
- Software
 - High-complexity data sets
 - Complex GI
- Hybrid techniques

© Kavita Bala, Computer Science, Cornell University

Dealing with High Complexity

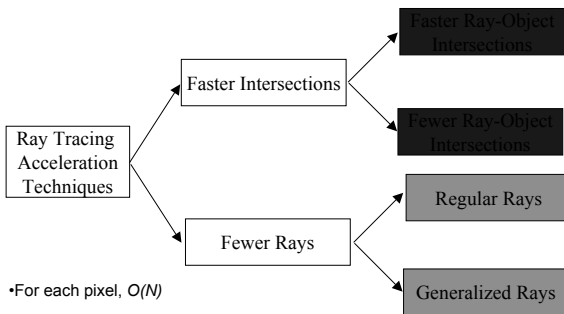
- Many Lights
- Display systems
 - Point-based approaches
- Visibility pre-processing systems
- Image-Based Rendering

© Kavita Bala, Computer Science, Cornell University

Acceleration Data Structures

CS 665

Making RT faster



•For each pixel, $O(N)$


•For each light, k shadow rays

•For GI and antialiasing: many rays per pixel


© Kavita Bala, Computer Science, Cornell University

Fewer Rays: Regular Rays

- Regular rays
 - Adaptive tree-depth control
 - Adaptive antialiasing




$K_t \cdot K_s \cdot K_r$ Shading



© Kavita Bala, Computer Science, Cornell University

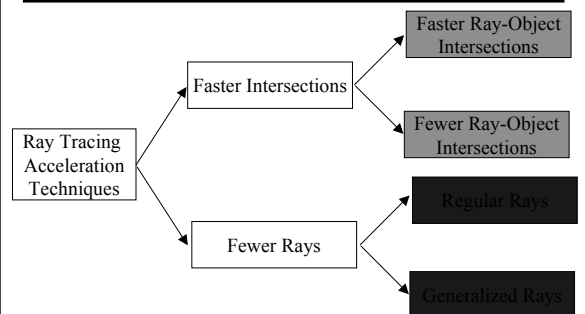
Generalized Rays

- Generalized rays represent a set of rays
 - Cone
 - Beam
- Pros
 - Good for anti-aliasing
 - Decreases number of rays
- Cons
 - More complex intersection tests
 - Reflections and refractions get hairy



© Kavita Bala, Computer Science, Cornell University



Making RT faster



© Kavita Bala, Computer Science, Cornell University

Faster Ray-Object Intersections

- Object bounding volumes



- Avoid intersection tests for expensive objects: e.g., polygon sets, spline surfaces
 - Ray/sphere or ray/cuboid test is fast

© Kavita Bala, Computer Science, Cornell University

Intersection: sphere

Assume sphere $x^2 + y^2 + z^2 = 1$

Point of intersection $p = O + t_{\text{intersection}} D$


p lies on sphere

Solve $A t_{\text{intersection}}^2 + B t_{\text{intersection}} + C = 0$

$A = 1, B = 2(O \cdot D), C = (O \cdot O - 1)$

$t_{\text{intersection}} = (-B \pm S)/(2A),$

$S = \text{sqrt}(B^2 - 4AC)$



© Kavita Bala, Computer Science, Cornell University

Intersection: cube

$t_{Near} = -inf$, $t_{Far} = +inf$

For each pair of planes for the x,y,z axes {

Solve for $O[i] + D[i] t1 = Min[i]$

Solve for $O[i] + D[i] t2 = Max[i]$

What if $t1 > t2$? swap


$t_{Near} = \max(t1, t_{Near})$

$t_{Far} = \min(t2, t_{Far})$

}


if ($t_{Near} > t_{Far}$) missed box

else hit box



© Kavita Bala, Computer Science, Cornell University

Tight Fit to Bounding Volume



© Kavita Bala, Computer Science, Cornell University