

Visible Surface Determination


CS4600 Computer Graphics
From Rich Riesenfeld
Fall 2015

Class of Algorithms

- Object (Model) Space Algorithms
 - Work in the model data space
- Image Space Algorithms
 - Work in the projected space
 - Most common VSD domain


Utah School of Computing

Back Face Culling: Object Space



3

Back Face Culling: Object Space



4


Back Face Culling Test

- For Object Space look at sign of $v \cdot n$
- For Image Space look at sign of n_z

Utah School of Computing

5


Back Face Culling: Image Space



Utah School of Computing

6

Back Face Culling: Image Space





Utah School of Computing

7

Back Face Culling



- Completes the job for convex polyhedral objects
- Nonconvex objects need additional processing beyond back face culling



Utah School of Computing

8



Back Face Culling: Examples



Utah School of Computing

9



Back Face Culling: Examples



Utah School of Computing

10

Back Face Culling: Examples




Utah School of Computing

11

Back-Face Culling

- On the surface of a closed manifold, polygons whose normals point away from the camera are always occluded:



Note: backface culling alone doesn't solve the hidden-surface problem!

Back-Face Culling


- Not rendering backfacing polygons improves performance
 - *By how much?*
 - *Reduces by about half the number of polygons to be considered for each pixel*

Silhouettes

- For Object Space $v \cdot n = 0$
- For Image Space $n_z = 0$

Occlusion

- For most interesting scenes, some polygons will overlap:




- To render the correct image, we need to determine which polygons *occlude* which

Painter's Algorithm

- How do painter's solve this?

Painter's Algorithm

- Simple approach: render the polygons from back to front, "painting over" previous polygons:




- Draw blue, then green, then orange
- *Will this work in the general case?*

Painter's Algorithm

- How do painter's solve this?
- Sort the polygons in depth order
- Draw the polygons back-to-front
- QED

Painter's Algorithm: Problems

- *Intersecting polygons* present a problem
- Even non-intersecting polygons can form a cycle with no valid visibility order:



Analytic Visibility Algorithms

- Early visibility algorithms computed the set of visible polygon *fragments* directly, then rendered the fragments to a display:




– Now known as *analytic visibility algorithms*

Analytic Visibility Algorithms

- *What is the minimum worst-case cost of computing the fragments for a scene composed of n polygons?*

- Answer:
 $O(n^2)$



Analytic Visibility Algorithms

- So, for about a decade (late 60s to late 70s) there was intense interest in finding efficient algorithms for *hidden surface removal*
- We'll talk about two:
 - *Warnock's Algorithm*
 - *Binary Space-Partition (BSP) Trees*


Warnock's Algorithm (1969)

PIXAR uses a similar scheme

- Elegant scheme based on a powerful general approach common in graphics: *if the situation is too complex, **subdivide***
 - Start with a *root viewport* and a list of all primitives
 - Then recursively:
 - Clip objects to viewport
 - If number of objects incident to viewport is zero or one, visibility is trivial
 - Otherwise, subdivide into smaller viewports, distribute primitives among them, and recurse


Warnock's Algorithm

- What is the terminating condition?
- How to determine the correct visible surface in this case?




Warnock's Algorithm

- What is the terminating condition?
 - **One polygon per cell**
- How to determine the correct visible surface in this case?



Warnock's Algorithm

- What is the terminating condition?
 - **One polygon per cell**
- How to determine the correct visible surface in this case?
 - **Cell = single pixel**



Warnock's Algorithm

- Pros:
 - Very elegant scheme
 - Extends to any primitive type
- Cons:
 - Hard to embed hierarchical schemes in hardware
 - Complex scenes usually have small polygons and high *depth complexity*
 - Thus most screen regions come down to the single-pixel case


Binary Space Partition Trees (1979)

- BSP tree: organize all of space (hence *partition*) into a binary tree
 - *Preprocess*: overlay a binary tree on objects in the scene
 - *Runtime*: correctly traversing this tree enumerates objects from back to front
 - Idea: divide space recursively into half-spaces by choosing *splitting planes*
 - Splitting planes can be arbitrarily oriented
 - Notice: nodes are always convex

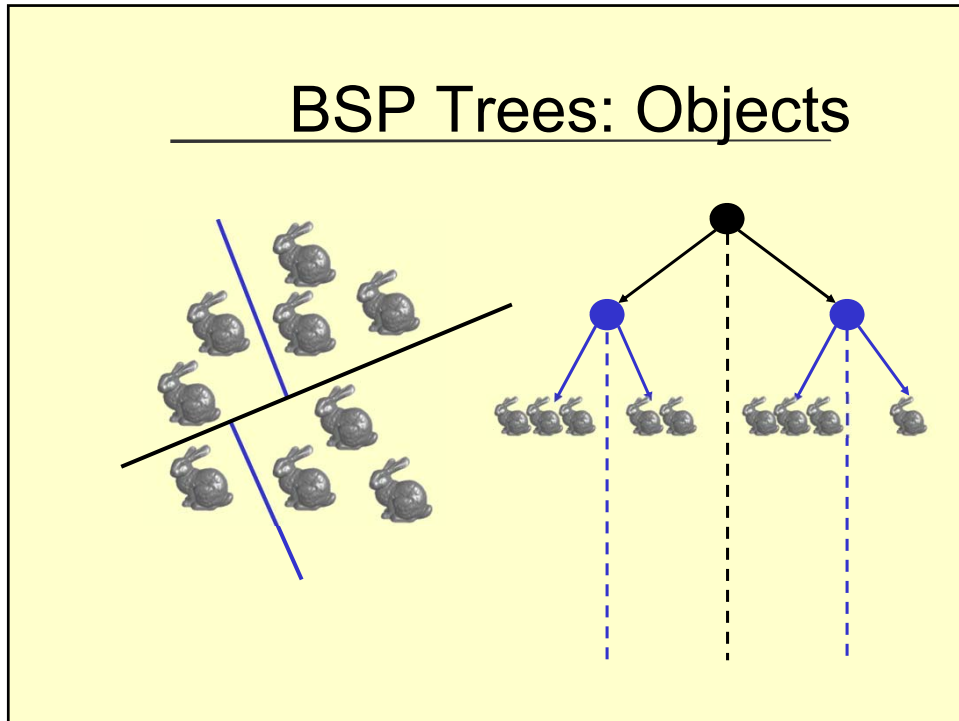
BSP Trees: Objects



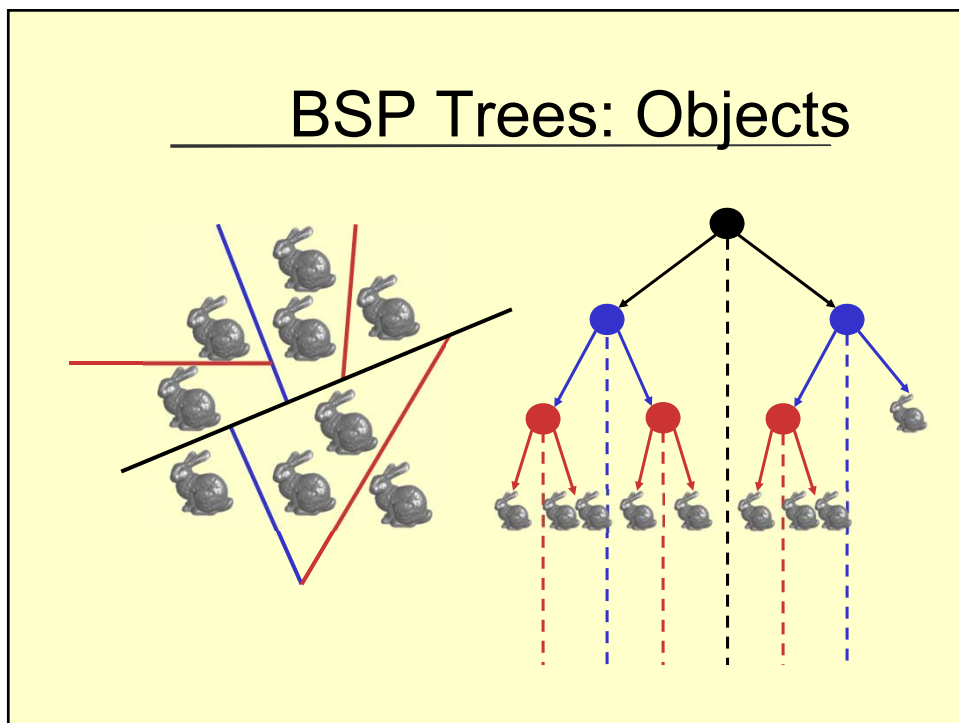
BSP Trees: Objects




BSP Trees: Objects



BSP Trees: Objects



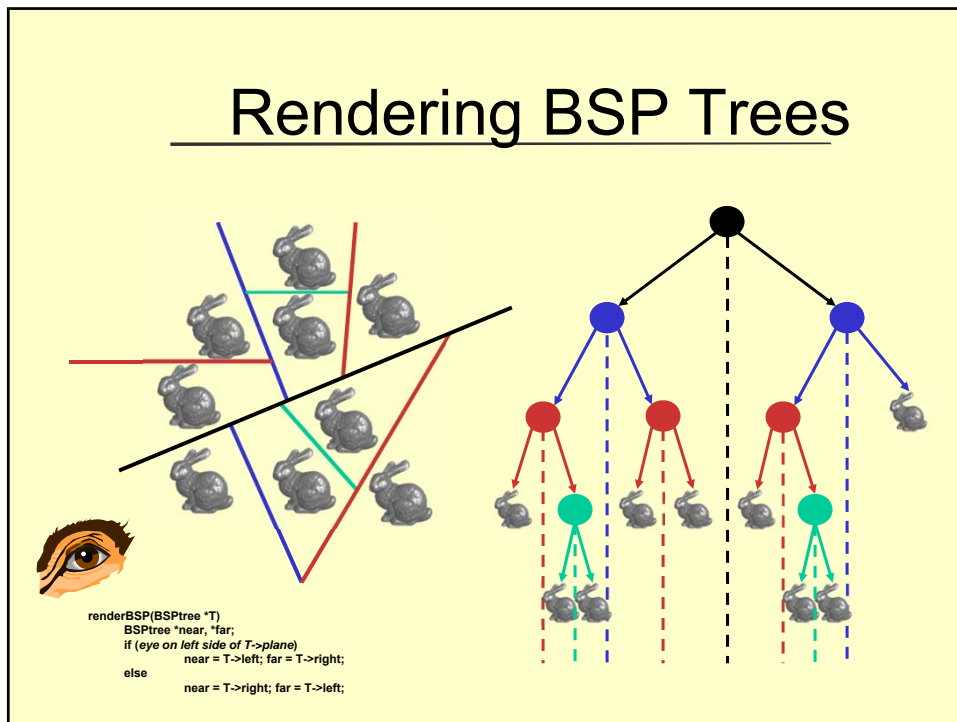
BSP Trees: Objects



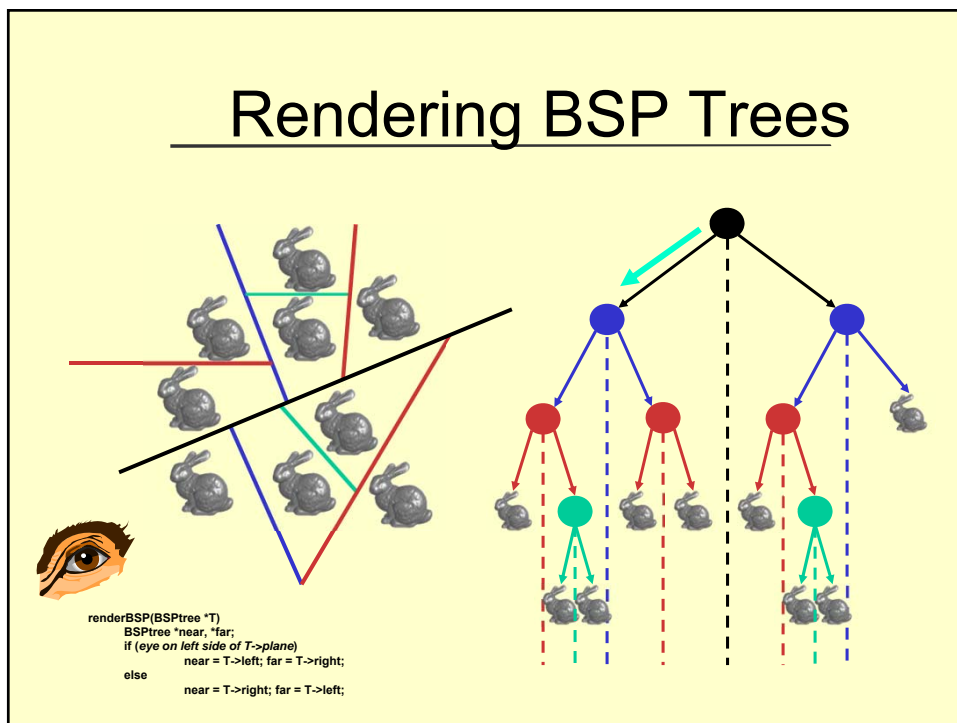
Rendering BSP Trees

```
renderBSP(BSPtree *T)
    BSPtree *near, *far;
    if (T is a leaf node)
        renderObject(T)
    else {
        if (eye on left side of T->plane)
            near = T->left; far = T->right;
        else
            near = T->right; far = T->left;
        renderBSP(far);
        renderBSP(near);
    }
```

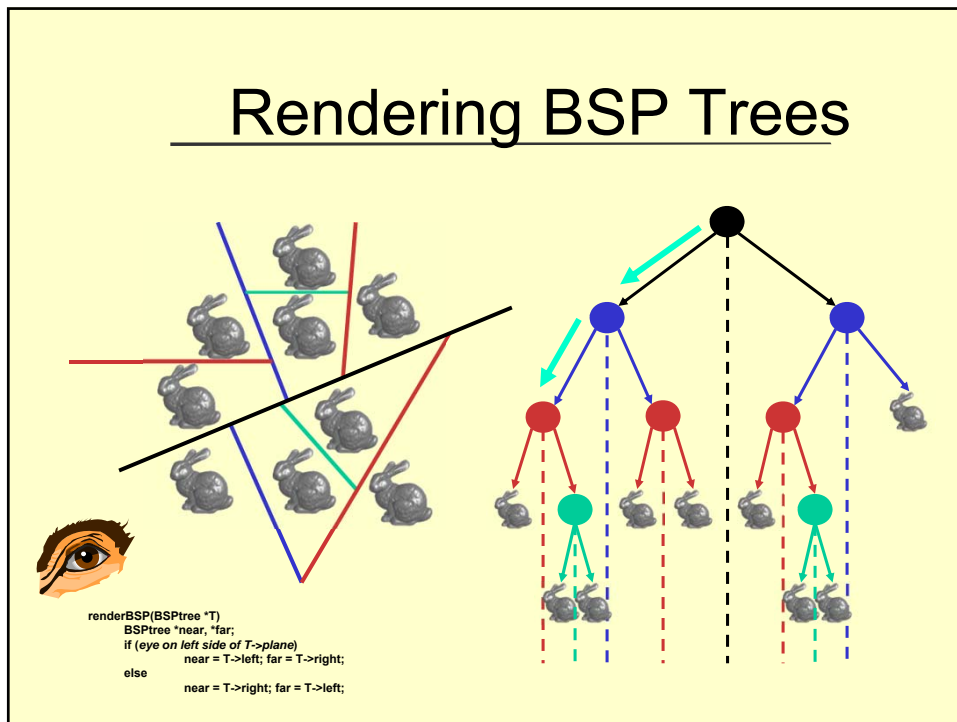
Rendering BSP Trees



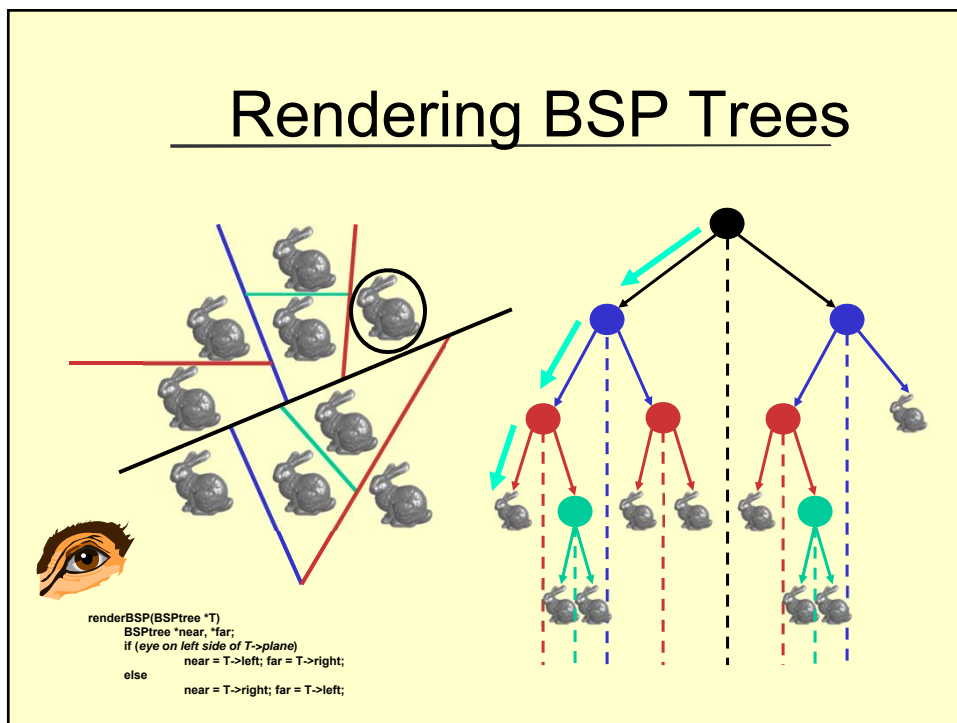
Rendering BSP Trees



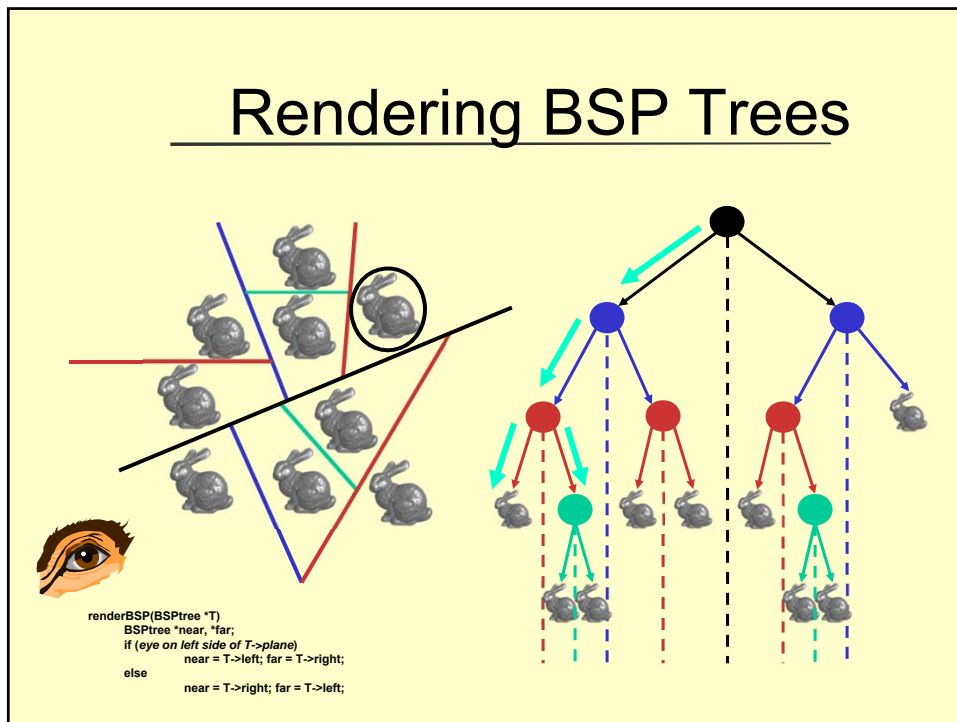
Rendering BSP Trees



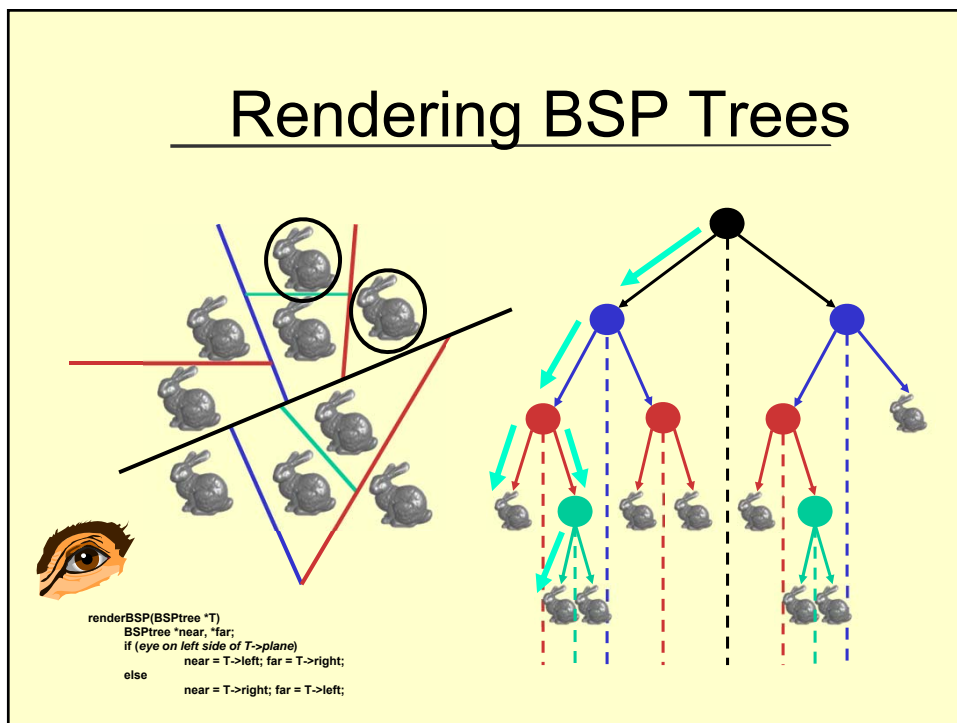
Rendering BSP Trees



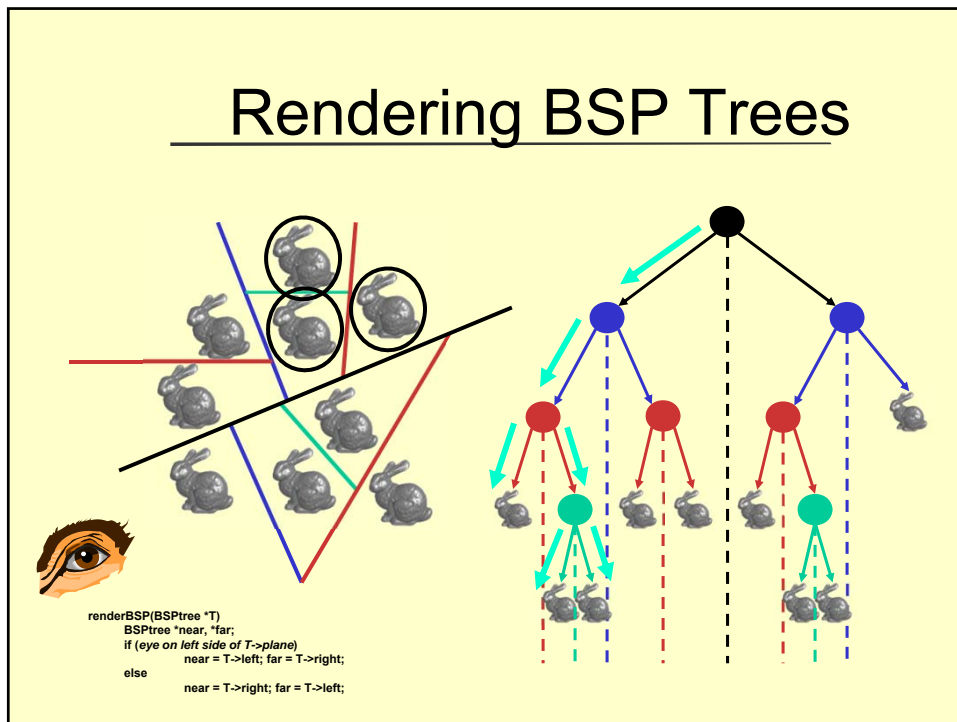
Rendering BSP Trees



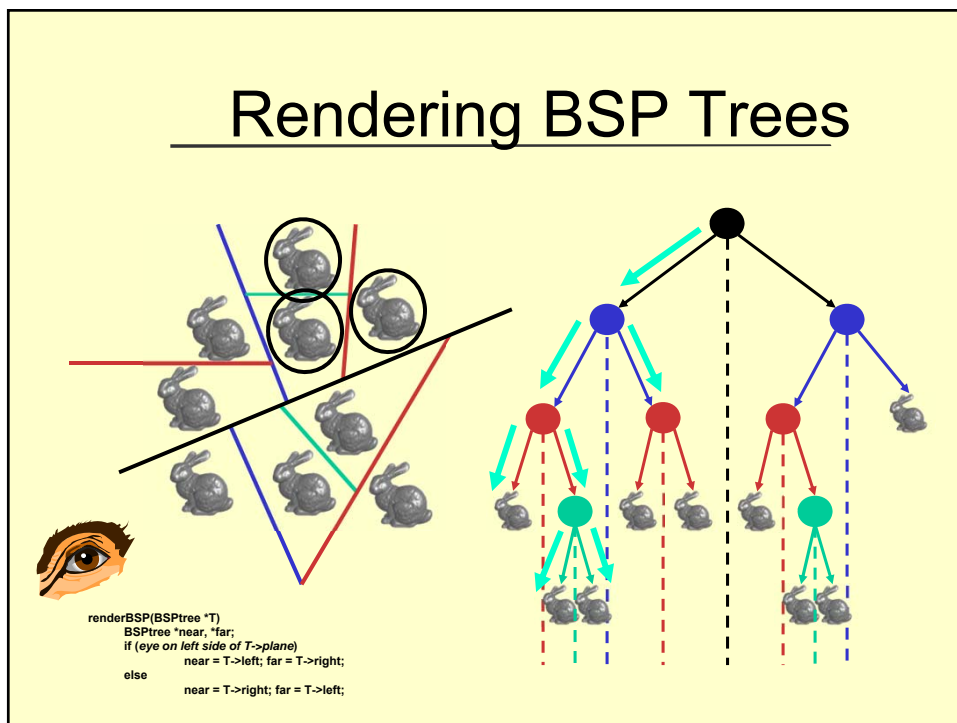
Rendering BSP Trees



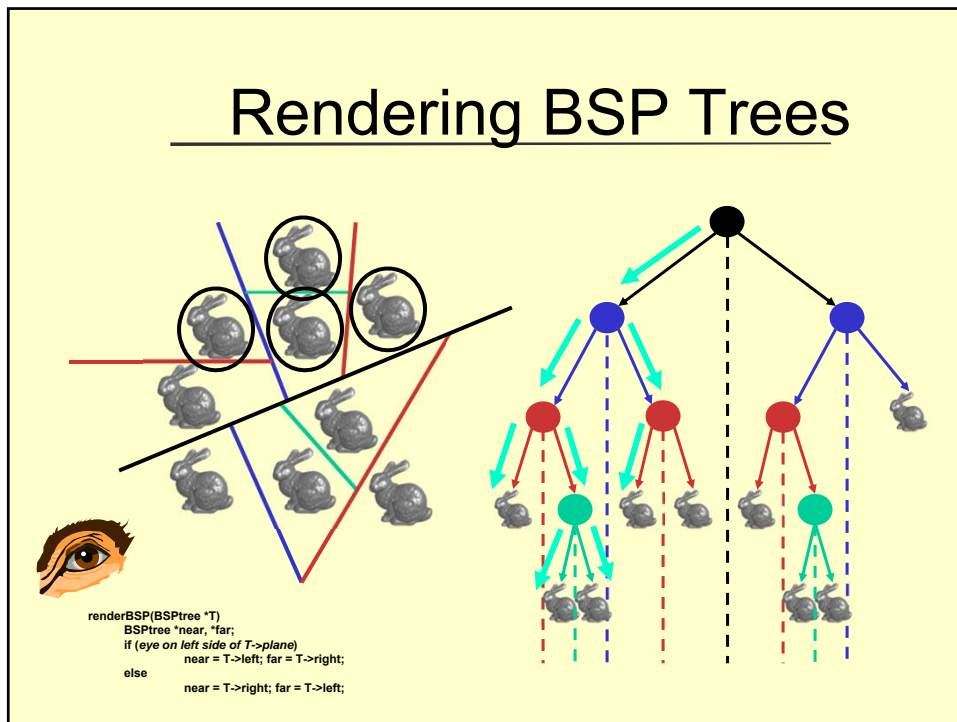
Rendering BSP Trees



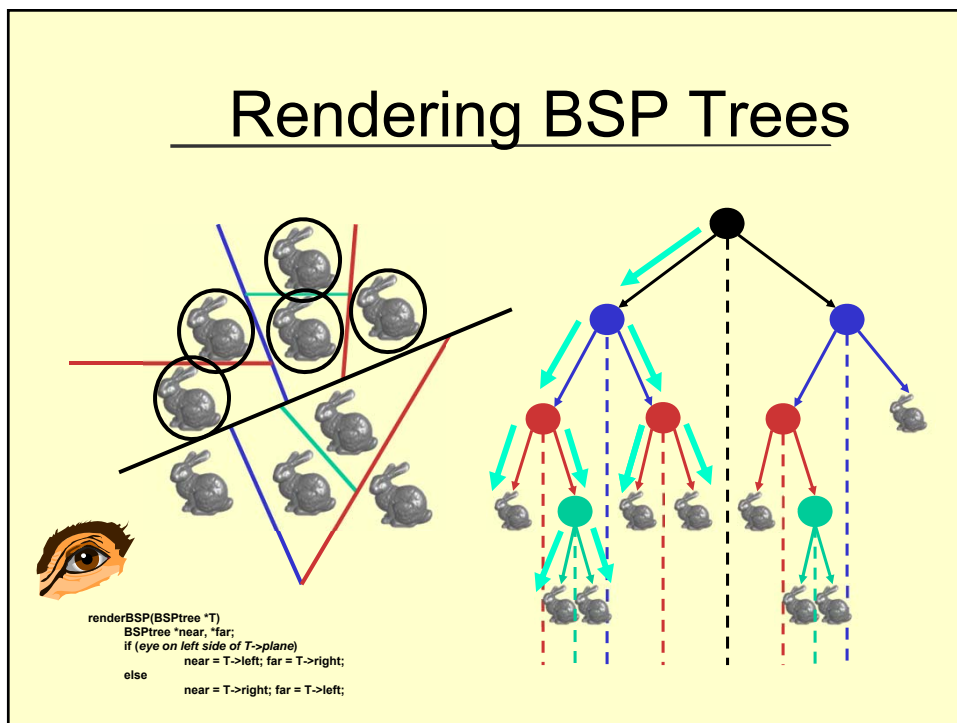
Rendering BSP Trees




Rendering BSP Trees




Rendering BSP Trees




Rendering BSP Trees




Rendering BSP Trees




Rendering BSP Trees




Rendering BSP Trees




Rendering BSP Trees




Rendering BSP Trees



Rendering BSP Trees



Rendering BSP Trees



3D Polygons: BSP Tree Construction

- Split along the plane containing any polygon
- Classify all polygons into positive or negative half-space of the plane
 - If a polygon intersects plane, split it into two
- Recurse down the negative half-space
- Recurse down the positive half-space

Polygons: BSP Tree Traversal

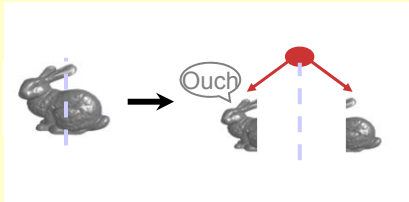
- Query: given a viewpoint, produce an ordered list of (possibly split) polygons from back to front:

```
BSPnode::Draw(Vec3 viewpt)
    Classify viewpt: in + or - half-space of node->plane?
    /* Call that the "near" half-space */
        farchild->draw(viewpt);
        render node->polygon; /* always on node->plane */
        nearchild->draw(viewpt);
```

- Intuitively: at each partition, draw the stuff on the farther side, then the polygon on the partition, then the stuff on the nearer side

Discussion: BSP Tree Cons

- No bunnies were harmed in my example
- But what if a splitting plane passes through an object?
 - Split the object; give half to each node:



- Worst case: can create up to $O(n^3)$ objects!

BSP Demo

- Nice demo:

<http://www.symbolcraft.com/graphics/bsp/>

Summary: BSP Trees


- Pros:
 - Simple, elegant scheme
 - Only writes to framebuffer (i.e., painters algorithm)
 - Once very popular for video games (but getting less so)
 - Widely used in ray-tracing
- Cons:
 - Computationally intense preprocess stage restricts algorithm to static scenes
 - Worst-case time to construct tree: $O(n^3)$
 - Splitting increases polygon count
 - Again, $O(n^3)$ worst case

The Z-Buffer Algorithm

- Both BSP trees and Warnock's algorithm were proposed when memory was expensive
 - Example: first 512x512 framebuffer > \$50,000!
- Ed Catmull (mid-70s) proposed a radical new approach called *z-buffering*.
- The big idea: resolve visibility *independently at each pixel*



The Z-Buffer Algorithm

- We know how to rasterize polygons into an image discretized into pixels:



The Z-Buffer Algorithm

- *What happens if multiple primitives occupy the same pixel on the screen? Which is allowed to paint the pixel?*



The Z-Buffer Algorithm

- Idea: retain depth (Z in eye coordinates) through projection transform
 - Use canonical viewing volumes (Clip / NDC frames)
 - Can transform clip perspective volume into NDC parallel volume with:

$$M = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & \frac{1}{1+z_{min}} & \frac{-z_{min}}{1+z_{min}} \\ 0 & 0 & -1 & 0 \end{bmatrix} = \begin{bmatrix} n & 0 & 0 & 0 \\ 0 & n & 0 & 0 \\ 0 & 0 & (n+f) & -nf \\ 0 & 0 & 1 & 0 \end{bmatrix}$$

z-Buffer (Depth Buffer)

Conceptually:

Sort $(\max z_{xy})$
 (x,y)

Sort $(\min z_{xy})$
 (x,y)

The Z-Buffer Algorithm

- Augment framebuffer with *Z-buffer* or *depth buffer* which stores Z value at each pixel
 - At frame beginning initialize all pixel depths to ∞
 - When rasterizing, interpolate depth (Z) across polygon and store in pixel of Z-buffer
 - Suppress writing to a pixel if its Z value is more distant than the Z value already stored there

Interpolating Z


- **Edge equations: Z is just another planar parameter:**

$$z = (-D - Ax - By) / C$$

If walking across scanline by (Δx)

$$z_{new} = z - (A/C)(\Delta x)$$


- Like drawing lines?
- Total cost:
 - 1 more parameter to increment in inner loop
 - 3x3 matrix multiply for setup







- **Edge walking: just interpolate Z along edges and across spans**



The Z-Buffer Algorithm



- *How much memory does the Z-buffer use?*
- *Does the image rendered depend on the drawing order?*
- *Does the time to render the image depend on the drawing order?*
- *How does Z-buffer load scale with visible polygons? With framebuffer resolution?*



























































← *z-Buffer: Pros* → z

- Simple algorithm
- Easy to implement in hardware
- Complexity is order N , for polygons
- No polygon processing order required
- Easily handles polygon interpenetration

Utah School of Computing

99

← *z-Buffer: Cons* → z

- Memory intensive
- Hard to do antialiasing
- Hard to simulate translucent polygons
- Precision issues
 - Scintillating = z-fighting
 - worse with perspective projection, why?

Utah School of Computing

100

← *z-Buffer Algorithm* z

- Initialize buffer
 - Set background *intensity, color* $\langle r, g, b \rangle$
 - Set *depth* to *max (min)* values

Utah School of Computing

101

← *z-Buffer Algorithm* z

- As a *polygon* P is scan converted
 - Calculate depth $z(x,y)$ at each pixel (x,y) being processed
 - Compare $z(x,y)$ with $z\text{-Buffer}(x,y)$
 - Replace $z\text{-Buffer}(x,y)$ with $z(x,y)$ if closer to eye

Utah School of Computing

102

← *z-Buffer* Algorithm → z

- *Convert all polygons*
- Correct image gets generated when *done*
- OpenGL: depth-buffer = z-Buffer

Utah School of Computing

103

← *a-Buffer* Algorithm → z

- Generates linked list for each pixel
- Memory of all contributions allows for proper handling of many advanced techniques
- Even *more* memory intensive
- Widely used for high quality rendering
- Widely used for deferred shading


Utah School of Computing

104

α -Buffer Algorithm: Example z



Utah School of Computing



105







Utah School of Computing



106














































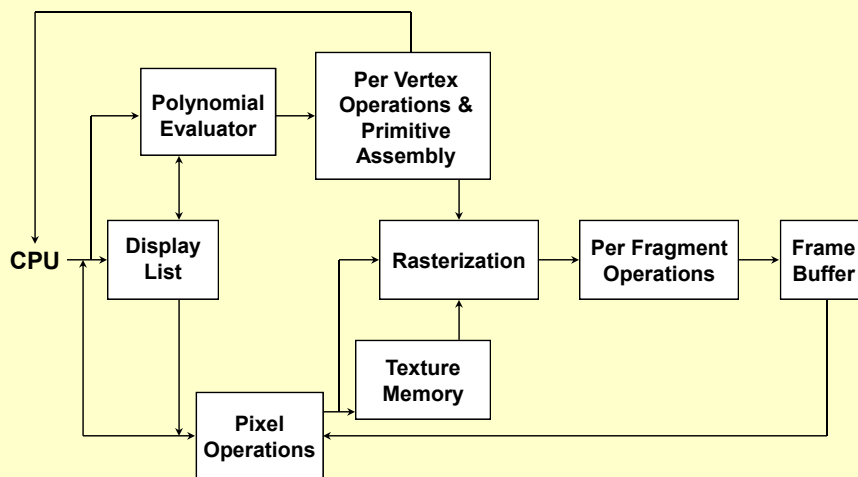









OpenGL Architecture



Depth Buffering and Hidden Surface Removal



Depth Buffering Using WebGL

- ① Enable depth buffering (typically in init)
`gl.enable(gl.DEPTH_TEST)`
- ② Clear color and depth buffers
`gl.clear(gl.COLOR_BUFFER_BIT | gl.DEPTH_BUFFER_BIT)`
- ③ Render scene
- ④ Swap color buffers (you know how)

Cubev.html (Chapter 4)

```
window.onload = function init()
{
    canvas = document.getElementById( "gl-canvas" );

    gl = WebGLUtils.setupWebGL( canvas );
    if ( !gl ) { alert( "WebGL isn't available" ); }

    gl.viewport( 0, 0, canvas.width, canvas.height );
    gl.clearColor( 1.0, 1.0, 1.0, 1.0 );

    gl.enable(gl.DEPTH_TEST);

    // shaders load/compile/link/use

    // set up vertex buffers and GUI

    render();
}
```

Cubev.html (Chapter 4)

```
function render()
{
    gl.clear( gl.COLOR_BUFFER_BIT | gl.DEPTH_BUFFER_BIT );

    theta[axis] += 2.0;
    gl.uniform3fv(thetaLoc, theta);

    gl.drawElements( gl.TRIANGLES, numVertices, gl.UNSIGNED_BYTE, 0 );

    requestAnimationFrame( render );
}
```

The End
*Visible Surface
Determination*

143