

Ray Tracing

- What was the rendering equation? Motivate & list the terms.
- Relate the rendering equation to forward ray tracing. Why is forward ray tracing not good for image formation?
- What is the difference between ray-"casting" and ray-"tracing". When would you only need raycasting?
- Given incoming ray (s,d) hit surface point p (with normal $'n'$) describe secondary rays you would emit to compute:
 - direct lighting from light at point $'l'$.
 - simple (mirror-like) reflection.
 - glossy reflection.
 - refraction.
 - glossy refraction.
 - full indirect illumination.
- How would you implement scattering media in a raytracer?
- Review: Describe how to perform ray intersection tests against:
 - Plane.
 - Sphere.
 - Triangle.
 - Mesh.
 - CSG object (assume you have lists of the child intersections already).
- Given a camera at $(0,0,0)$ facing $(1,0,0)$ with up-vector $(0,0,1)$, horizontal field-of-view 60 degrees, square pixels, and a 640×480 image, what ray passes through pixel $(10,40)$? (NOTE: assume the top-left pixel is $(0,0)$.)
- Write pseudo-code for a backward ray-tracer that uses a recursive function.
- Write the contribution of light intensity at a point from various sources, including direct lighting (ambient, diffuse, and specular), a single reflection ray, and a single transmission ray.
- Draw a simple scene and sketch all rays that will be traced from the eye through one pixel in the following cases:
 - Ray casting
 - Ray tracing with one bounce
 - Ray tracing with two bounces

- Consider a partially transparent object that is illuminated with two lights, one visible from each side of the object. Start with a ray from the eye that hits one side of this object at approximately a 45 degree angle from the surface normal and sketch all of the rays that will be traced to determine the color of that intersection point.
- Give Snell's Law and use it to derive an expression for the refraction ray in terms of the surface normal, the incoming direction, and the indices of refraction of the two media.
- Under certain circumstances, there will not be a solution to the equation used to compute the refraction ray. What are these circumstances? What happens to the light ray? This phenomenon is called total internal reflection.
- In ray tracing we choose to trace very few rays through the scene in order to keep computation times reasonable. List some of the paths that light can travel that are not captured by ray tracing.
- Consider the following effects: mirror reflection, refraction, caustics, and color bleeding. Which of these can be captured through ray tracing? For which does ray tracing not perform well?
- We can define a view frustum with the following inputs:
 - \mathbf{e} = eye point
 - \mathbf{v} = up vector
 - $-\mathbf{w}$ = look vector
 - l = distance of left side of image from the look vector
 - r = distance of right side of image from the look vector
 - t = distance of top of image from look vector
 - b = distance of bottom of image from the look vector
 - n = distance of near clipping plane along look vector
 - f = distance of far clipping plane along look vector
 - n_x = number of pixels in a row of the image
 - n_y = number of pixels in a column of the image

Here we assume that the image plane is the near clipping plane. Sketch a diagram showing all of these input variables. Give an expression for the ray cast from the eye point through pixel (i, j) based on these input parameters.

- Show how to compute ray triangle intersections, including identification of barycentric coordinates for use in interpolation of normals and texture coordinates.

- Distributed ray tracing can be used to capture a variety of effects and create less sharp and more realistic images. Explain how to use distributed ray tracing to capture each of the following effects:
 - antialiasing
 - gloss
 - translucency
 - soft shadows
 - depth of field
 - motion blur
- If we use distributed ray tracing for all of these effects at once, how can we keep the number of rays that are traced in a scene to some reasonable number and still get good even sampling of each of these effects?

Spatial Data Structures

- Why do we want to build a hierarchical data structure such as a bounding box hierarchy?
- Describe the main difference between the following two approaches: (1) construct a bounding volume hierarchy (e.g. bounding boxes or bounding spheres), (2) construct a hierarchy of splitting planes (e.g. KD trees or BSP trees).
- Describe a technique for constructing a bounding box hierarchy. What data structure will you use to store this hierarchy?
- Is your technique top-down or bottom up? Give an algorithm for the alternative approach (i.e., if your first approach was top-down, give a bottom-up approach to constructing a bounding box hierarchy).
- How do we do intersection testing between a ray and a bounding box? Note that we do not need to find the point of intersection, but we only need to determine whether or not the ray intersects the bounding box.
- How do we do intersection testing between a ray and a bounding sphere? Again, we only need to know whether the ray intersects the sphere. We do not need to find the specific intersection point.
- Explain in detail how to use a bounding box hierarchy to identify the intersection between a ray and the closest object in the environment. Your algorithm should, of course, be more efficient (in the general case) than the brute force process of checking the ray for intersection with all objects.
- How would you update the bounding box hierarchy if there are moving objects in the environment? Is there a way to do this efficiently, i.e., without rebuilding the entire tree?
- An alternative to the bounding box or bounding sphere hierarchy is to use splitting planes to divide space. Octrees, KD trees, and BSP trees are all splitting plane algorithms. Describe the differences between these approaches.
- Describe how to construct a KD tree. Describe how to perform ray-object intersection using this data structure. Your ray-object intersection algorithm should check regions from front to back order, so that the search for an intersection may be halted when an intersection point is found.
- In the KD tree scenario, what problem is introduced when the splitting planes intersect some of the objects? Assume that you do not want to use the plane to split the objects (e.g., splitting a sphere into two parts may not be desirable for ray tracing). How can you solve or work around this problem?
- Describe how to construct a BSP tree and how to do ray-object intersection testing in front to back order with this data structure.

Radiosity

- Explain why diffuse reflection produces a color on a surface that is independent of viewing direction.
- What is the main idea behind radiosity that takes advantage of this property?
- How do we store the results of radiosity rendering in such a way that they can be used for many camera viewpoints (e.g., as in an architectural walkthrough)?
- Why is radiosity often used in indoor architectural environments?
- The radiosity energy balance equation is shown below. Describe the terms in this equation.

$$B_i = E_i + \rho_i \sum_{j=1}^n F_{i,j} B_j$$

- What types of light sources are easiest to model for radiosity?
- Radiosity is typically solved using either a shooting or a gathering approach. Explain the key ideas behind these two approaches.
- The energy balance equation can be written in matrix form as shown below. What are the elements of the R matrix? Explain how this equation represents a gathering approach.

$$B^{k+1} = E + RB^k$$

- Why is a shooting approach typically more practical than the gathering approach? Consider as one aspect of your answer the number of form factors that must be computed for the gathering approach and for each step of the shooting approach.
- How do we avoid visible seams at edges of single-color patches used for radiosity?
- How can we get sharp shadows?
- Give some of the pros and cons of ray tracing vs. radiosity.

Photon Mapping

- The contributions of light to an image can be considered as light collected over a huge collection of paths through an environment. When thinking about the possible paths, we need to consider all types of bounces, both specular (predominant in ray tracing) and diffuse (predominant in radiosity). If we use a regular expression for the paths considered in ray tracing, we would have:

$$L(D|S)E \quad (\text{direct lighting})$$

$$L(D|S)S^*E \quad (\text{with specular bounces})$$

where	L	indicates that the path starts at the light
	D	refers to a diffuse bounce
	S	refers to a specular bounce
	E	indicates that the path ends at the eye

Explain these expressions and show how they result from the ray tracing equation:

$$I = k_a L_a + k_d (l \cdot n) L_d + k_s (r \cdot v)^\alpha L_s + k_r I_r + k_t I_t$$

- What is the equivalent regular expression for paths represented in radiosity?
- What light paths are missing (not captured by either technique)? Think of an effect captured by such paths that neither ray tracing nor radiosity will handle well.
- What is the key idea behind photon mapping? In particular, describe the steps of photon shooting and photon collection. The following reference may be useful: <http://graphics.ucsd.edu/~henrik/>
- In photon mapping, we collect two separate maps: (1) a global map, with rays expressed as $[L(S|D)^* D - LS^+ D]$ and (2) a caustic map, with rays expressed as $[LS^+ D]$. Explain the caustic paths and explain why this map is separated from the global map.
- How are shadows captured in photon mapping?
- Give some pros and cons of photon mapping compared to ray tracing and radiosity.

Direct / Indirect Separation

- Explain the technique presented in class for separating direct and global illumination effects from a scene. How is direct lighting recovered? How is global illumination recovered?
- When will this approach fail? Give one or more specific examples of possible failures.
- When is this information useful for computer graphics – give some applications.
- Were there direct vs. global lighting effects that you found surprising in the examples presented in class? If so, which effects and why.

Animation and Motion Capture

- List and describe some of the principles of traditional animation.
- There are three techniques typically used for animation: keyframing, procedural animation, and motion capture. Explain each of these techniques.
- List the pros and cons of each of these three techniques for the purpose of animating human characters.
- List the pros and cons of these techniques for animating passive objects such as clothing and water.
- What is your opinion on the value of using physical simulation for animating human characters? When does it work well? When does it work less well?

Differential Equations, Particle Systems, and Fluids

- Draw and label a block diagram for a simple physically based simulation.
- What is Euler integration? What is the midpoint method? What is RK4?
- Explain the problems that may arise with Euler integration. Use diagrams and/or examples to illustrate your point.
- Why may RK4 help to address these problems?
- What are the basic steps to be accomplished in any rigid body simulation?
- How do we keep track of orientations? Give pros and cons of two techniques for maintaining and updating orientations.
- (Bonus) What are the basic steps to be accomplished in any fluid simulation?

Image Processing and High Dynamic Range Images

- Give an example of a 3x3 filter that could be used to blur an image.
- Give an example of a 3x3 filter to detect edges. How does this filter work (describe it in words)?
- Be prepared to compute the results of (or write pseudocode for) convolving a given filter with a given image.
- To make an image smaller, we could simply downsample (e.g., save every other pixel). Why is this a bad idea? What should we do to obtain better results?
- What is the purpose of environment mapping? What are the advantages compared to standard ways of specifying light sources?
- Describe how to modify your ray tracer to handle environment mapping.
- Using ray tracing, one can only apply environment mapping to specular surfaces. Why? In Paul Debevec's Rendering with Natural Light movie, the diffuse surfaces are also lit with an environment map. How did he do that?
- What approximations are made when an object is lit with an environment map instead of the true lighting?
- Ignoring sampling issues, will an object lit with a spherical env. map look the same as if lit with a cubic env. map? Why, or why not?
- When placing a synthetic object into a real scene, why is it important to capture a high dynamic range environment map?
- One way to capture an environment map is by photographing a mirrored ball. Does this capture light from all directions? Almost all directions? Only a hemisphere of directions? Explain by drawing a picture.
- If mirrored balls are so great for capturing environment maps, why aren't people using the same idea to capture panoramic photographs?
- What is the goal of tone mapping?
- Why is tone mapping even possible? That is, why can we ever hope that a sunset and a picture of a sunset could be made to look the same to a human observer?

Fall 2012

15-462 Computer Graphics -- Srinivas Narasimhan

15-463 Computational Photography -- Alexei Efros

15-466 Computer Game Programming -- Maxim Likhachev

15-869 Special Topics in Graphics: Human Motion Modeling and Analysis --
Yaser Sheikh, Leon Sigal

Spring 2013

15-462 Computer Graphics -- Nancy Pollard

15-464 Technical Animation -- Nancy Pollard

15-465/60-414 Animation Art and Technology -- Jessica Hodgins

15-869 The Animation of Natural Phenomena -- Adrien Treuille