

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 may 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.

NPR

- For what kinds of data representations is the marching cubes algorithm a good approach for generating surface meshes.
- In the marching cubes algorithm, what information do you need to know about each voxel in order to render the subsection of the surface mesh that may be contained within the voxel?
- What is the difference between using a meshing algorithm such as marching cubes and a volumetric rendering? In which circumstances would you prefer a meshing algorithm? In which circumstances would you prefer a volumetric shader?
- Give several approaches to identifying contours that can be rendered in a line drawing of an object.
- List some of the aspects of informational drawings (e.g., anatomy books, car repair manuals, and hand-drawn maps) that researchers in NPR have attempted to capture.
- For each item in your list, describe how that effect was captured.
- What is missing from your list? In other words, given what you saw in class, what aspects of such drawings do you believe is missing from the automatically generated results?
- BONUS: Researchers have also studied how to use cross-hatching of varying densities to create an impression of shading, while maintaining the appearance of a line drawing created with pencil or charcoal, for example. How might you write a shader to accomplish this effect?

Differential Equations and Particle Systems

- 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?

Animation and Motion Capture, Cloth and Fluids

- There are three techniques typically used for animation: keyframing, procedural animation (e.g., simulation), 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?
- How can you extend a particle simulator to create a convincing cloth animation? Could you also do character animation using only particles?
- List some animation approaches you saw in class that you would be interested to try.
- What are some of the major differences that must be considered when moving from rigid or deformable bodies (such as cloth) to simulating fluids?

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? How might we use an environment map to act as the major light sources within a scene? What are the advantages of using such a techniques 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 that was present in the real scene captured by the environment map?
- Ignoring sampling issues, will an object lit with a spherical environment map look the same as if lit with a cubic environment map? Why, or why not?
- What is a high dynamic range image? .. a high dynamic range environment map?
- 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 a single photograph capture light from all directions? Almost all directions? Only a hemisphere of directions? Explain by drawing a picture.
- 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?

(BONUS) Photo Clip Art

- Describe the various conditions that contribute to the appearance of a scene, such that the same building, for example, will appear many different ways if viewed at different instances in time. List at least 5 such conditions.
- What information about a scene can reasonably be extracted from an image?
- Referring to the research we discussed in class, can you think of two ways of estimating the illumination in a scene for the purpose of finding close matches to a photograph?
- Why do we want to find close matches to a given photograph?
- Describe some of the flaws that you saw in the edited photographs that we saw in class. What do you think is needed to fix those flaws to create images that are indistinguishable from reality using the Photo Clip Art approach?
- Regarding shadows, what is important to get right when pasting an item from one image to another? What is less important?

Cameras and Displays

- Describe how we can sample a scene in multiple ways and combine the results to obtain a better image
 - For image de-blurring
 - To identify depth discontinuities
 - To display clearly a high dynamic range scene on a low dynamic range device
 - To focus at varying distances
- How can a 3D display be created using water drops? What element of this display is critical for showing detail in multiple layers?
- Explain the principle behind a display that shows a scene depicting the incident illumination of the room in which it is placed (and updates in real time as that illumination changes)
- Explain the principle behind the 3D display shown in class (remember the floating head). How does it capture changes in viewing angle for observers ranged around the display? How does it portray proper parallax for observers at different viewing heights and distances?