Graphics Qualifier, Spring 2016

Strive to be ***c****oncise,* ***c****lear,* ***c****omplete,* ***c****orrect,* ***c****onvincing and* ***c****lever.*

Concise*: skip trivial stuff that a good student can fill in easily.*

Clear*: Clarify your understanding of each problem. Define your terms and notation. Avoid using confusing abstractions or vague statements. Use a figure or metaphor if it helps.*

Complete*: indicate what can go wrong and broadly how you would detect and address such invalid inputs or special cases, but do not explain the details.*

Correct*: Avoid false or controversial statements. Verify your formulae, algorithms, equations. If the questions asks for an exact answer, do not propose a approximation or a sampling. If you can't find the correct answer, propose a guaranteed upper bound.*

Convincing*: Try to find convincing arguments to explain briefly why your solution is complete and correct.*

Clever*: Look for an elegant solution that may be simple to implement or to explain.*

1 General (Answer 4 of 6)

1.1 Cutout

Let M be a triangle mesh (watertight manifold). It is displayed on the screen using some homogeneous transformation H. The user draws a polygonal loop P (with vertices P_i) on the screen over the projection of M with the idea of drawing in 3D **on** the mesh the border of a cutout that she wants to remove from M. Provide a high-level outline and the details of the data structure and algorithms for computing a representation of the result of this drawing and for displaying the modified mesh M (with the **precise** cutout portion removed) when viewing it with another homogeneous transformation H'. Your explanations should suffice to convince the reader that your solution works and to guide a programmer to implement your solution. If your solution may fail in some configurations, specify them precisely and succinctly.

1.2 Ropes

Consider two ropes, P and Q, in 3D. Each one is represented by a polygonal curve. All the edges of both curves have the same length and maintain that length constant throughout animation. You are given two key frames, each showing a state of an animation, where the ropes are entangled. Explain how you would: (1) compute the continuous model (or at least lots of intermediate frames) of the animation and (2) test whether your animation is collision free. Provide the math and the algorithm and a convincing justification.

1.3 Spheres

You are given two spheres of radius r and centers A and B. You are given a triangle mesh (watertight manifold) M that is placed between A and B and has lots of small triangles. We want to test whether M occludes A from B completely (any ray from any point on A to any point on B hits M). If you can, provide an algorithm for verifying this exactly (not by using a set of sample rays). If not, provide an algorithm for a conservative test that can often guarantee total occlusion, even though in some cases when a total occlusion exists your algorithm may say "I am not sure". Include a proof or convincing justification of the correctness of your algorithm.

1.4 Robots

N robots landed on a plane. They walk up at different times in a random order. When a robot wakes up, it takes a single picture to identify the current location of all robots (not their velocity), computes the location X of the common rendezvous point, and starts moving towards X in a straight line. The location X is not given. In the picture, the robots appear as identical dots and are not identifiable, so the location cannot be defined as the place where robot 1 is. Each robot must compute X from the instantaneous locations of the other robots. For example, X cannot be the centroid of the current locations of all the robots, because some of the other robots may already have stated moving, so that centroid would have been different when they wake up. Explain how each robot would compute X and provide a high-level algorithm for doing so. (Hint: There is no trick here. X is defined geometrically be the locations of the unordered set of the points where the robots are now.)

1.5 Stitch

You need to stitch two images (or heightmaps) together as seamlessly as possible. The images are grayscale and the same resolution. The output is the same resolution as the input. Describe algorithms to handle the following situations:

(1) you need to composite the left side of one image and the right side of another image. The composite image must contain the entire leftmost column of image 1 and the entire rightmost column of image 2 but you are free fill the other output pixels from either image.

(2) You need to composite the center of one image with the periphery of another image. The composite image must contain a the central circle with radius half the image height from image 1 and the first and last rows and columns from image 2.

1.6 Soda

You are developing a computer game in which soda cans are crashing into one another. It has been decided that the cans will be represented as perfect cylinders with flat disk end-caps. You have been given the task of determining whether a given pair of cylinders will collide with one another in a given time-step, and if so, exactly when. In any particular time-step, each cylinder is moving along a straight line at constant speed, with NO rotation. Describe an algorithm that determines exactly when there is a collision between a given pair of cylinders in a given time-step. Be sure that you enumerate all of the possible ways that two capped cylinders can strike each other.

2 Animation (Answer 2 of 4)

2.1 The Drone

A camera-mounted drone (Figure 1) can be modelled as an articulated rigid body system with two body nodes: the "drone node" which can translate and rotate in the world space with 6 global degrees of freedom and the "camera node" which connects to the "drone node" by a 3D ball joint. Assume the drone has control over all 9 degrees of freedom in the system, $\mathbf{u} \subseteq \mathbf{R}^9$ with an upper bound $\mathbf{u}_U$ and lower bound $\mathbf{u}_L$.

You are given a task to take photos of an event in a few spots ($\mathbf{x}_1, \ldots, \mathbf{x}_n$) at corresponding time instances ($t_1, \ldots, t_n$), can you design a trajectory that uses the least amount of control force for the drone to accomplish the task?

1) What are the physical constraints? Write down the mathematical formulation and definitions.
2) Please formulate the trajectory design problem into a spacetime optimization with the definitions of free variables, cost function and constraints.
3) How do you plan on solving this optimization efficiently?
4) How would you relax the problem to make it a convex optimization?


(Figure 1)

2.2 The Integrator

(True story) The students in CS4496 were asked to implement the Midpoint Method to integrate a dynamic system. One of the students this year implemented the following to update the state variable from the current state $\mathbf{x}_0$ to the next state $\mathbf{x}_1$. Assume that the time step is h and the differential equations can be expressed as \dot{**x**} = **f**(**x**).

$x_{new} = \mathbf{x}_0 + h\,\mathbf{f}(\mathbf{x}_0)$

$\mathbf{x}_1 = \mathbf{x}_0 + h * ½\,(\mathbf{f}(\mathbf{x}_0) + \mathbf{f}(\mathbf{x}_{new}))$

Is this correct? If the answer is yes, please prove that this formulation is equivalent to the standard Midpoint Method. If the answer is no, please discuss the stability and accuracy analysis for this formula.

2.3 Flesh and Bone

Suppose we want to simulate an animal arm using articulated rigid bodies to represent the bones and the finite elements to represent the flesh around them.

1) Propose a state space and governing differential equations for simulating this system.
2) Describe the constraints used to enforce accurate two-way coupling between the bones and the flesh?
3) How do you make sure that the bones never protrude the flesh, i.e. the bones will always stay inside the boundary of flesh.

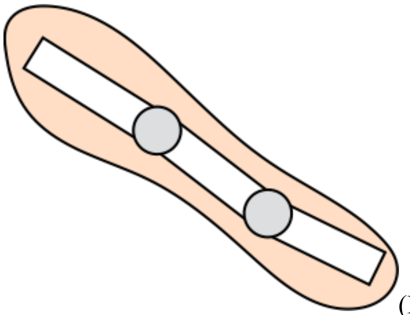
(Figure 2)

2.4 The Screwdriver

Screwing a nail into a wooden plank is a common task that requires coordination between both hands. In this example, you use your left hand to hold down the plank while the right hand is twisting the screwdriver. Your body can be described as an articulated rigid body system with the root body node fixed in the space shown as the figure below. Assume that you are in a sitting position which makes the degrees of freedom in your lower body irrelevant to this manipulation task.

1) How do you compute the joint torque of upper body required to achieve the task of the left hand.
2) How do you compute the joint torque of upper body required to achieve the task of the right hand.
3) The degrees of freedom at torso are needed for both tasks. How do you avoid interference between these two tasks?

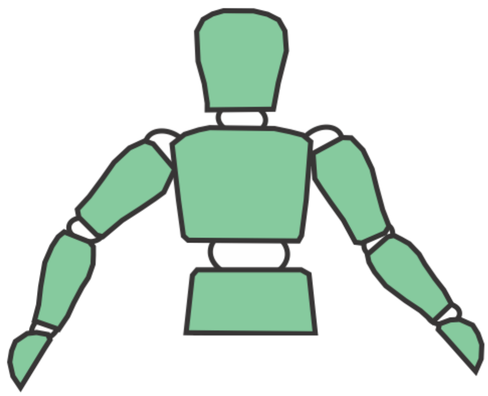

3 Rendering (Answer 2 of 4)

3.1 Surface Reflectances

Different surface properties can require different techniques for rendering. For each of the kinds of surfaces and accompanying effects below, describe the highest quality method that you know for rendering the surfaces. Assume that the lights in the scene are point lights at particular finite 3D positions in the scene.

a) Diffuse surfaces and with diffuse inter-reflection

b) Caustics from shiny surfaces

c) Glossy (slightly blurred) reflections

3.2 Reflection

Explain why shiny metal or mirrors reflect light the way they do without dispersing the directions of the reflected photons. A photon with velocity direction V arrives at a surface that has normal N. Another photon produced as the consequence of this shock leaves the surface with a specific reflected velocity direction R. Why this particular R? Can you think of a reasonable and convincing conjecture explaining this in geometric or physical terms?

3.3 Anti-Aliasing

When considering how to perform anti-aliasing in a ray tracing renderer, there are several important options to decide:

a) super-sampling pattern

b) anti-aliasing pre-filter shape

c) adaptive sampling

For each of the above issues, propose two different methods to use, one "poor" and one "good". For each issue, describe the possible image artifacts of the poor method. Also for each issue, explain why your good method will produce better final images.

3.4 Soft Effects

You have been given the task of creating various "soft" effects using graphics hardware. Your boss tells you that you MUST use rasterization to achieve these effects — no GPU ray tracing is allowed. What is the best method to achieve each of the following effects using graphics hardware? What potential image artifacts may occur in your images? For each method, would ray tracing produce better images, and why?

a) Soft shadows

b) Motion blur

c) Depth-of-field