CSE 163: Assignment 3—Final Project

Ravi Ramamoorthi

The final project is a chance to showcase what you've learned in the course, to develop a project of your choice. We provide some options here, both well-described project options, and less specified ones. However, note that you can also propose a topic of your choosing, with the instructor's approval. You are also welcome to extend one of your previous assignments in a substantial way. Some of the more free-form options below are deliberately under-specified, to give some flexibility, and to allow you to use your creativity.

As in previous assignments, the submission is to be done by creating a website, including full documentation and examples (or a PDF), and including links to videos of the system running (if a real-time application). You should do this project in groups of two, but as before, this is not strictly required. As always, e-mail the instructor/TA with a link to the website to submit the assignment.

1 Relatively Well-Specified Options

This part of the assignment gives some relatively well-specified (but still flexible) options, related to implementing some of the techniques for real-time and/or image-based rendering. For the real-time parts, you must provide video to enable grading, either on your website or on YouTube or similar medium.

You need only pick one of the choices, based on your interests and expertise. Since many of these are close to the research frontier, I am providing as much guidance as possible, but you will need to be proactive in looking up the original literature and other sources such as online tutorials. No source/skeleton code is available. If you find online source code, you are welcome to use it as a reference, or to use utility functions, but obviously not as a core part of the assignment. If in doubt, please ask the instructor for clarification. Future editions of the course may provide more support, but the point is also to get you to work independently and communicate your results, given that this is an advanced computer graphics course.

1.1 Precomputed Radiance Transfer and Relighting

Recently, a body of work known as precomputed radiance transfer, or precomputation-based relighting, has become a popular technique for real-time rendering in varying illumination, and recently also viewpoint for dynamic scenes. Some variant of this technique is now widely used in movies and games. Good references are the seminal papers by Sloan, Kautz and Snyder in SIGGRAPH 02, and Ng, Ramamoorthi, Hanrahan in SIGGRAPH 03. A good historical source, that also surveys some of the previous work in image relighting, is a survey I recently wrote on precomputation-based rendering (available on both my website, and on the publisher's. I can also give you a spare copy of the book if you ask me quickly enough). In this context, I also encourage you to read some of the earliest work, such as Nimeroff et al. 94, "Efficient Re-Rendering of Naturally Illuminated Environments", EuroGraphics Workshop on Rendering, and Dorsey et al. 95, "Interactive Design of Complex Time-Depending Lighting", IEEE Computer Graphics and Applications, 1995.

The basic idea of precomputation-based methods is based on linearity of light transport. You precompute solutions for different light sources and then linearly combine them in some quick way. Ideally, you would implement a basic PRT system, at least along the lines of Ng et al. 03, and maybe some of the further work like Wang et al. 04 or Ng et al. 04. However, this is difficult in the short time available, so I stratify the various tasks below:

1.1.1 Basic Relighting from a few Point Lights

First, implement a very basic relighting system. Consider a scene with a fixed viewpoint (this is essentially the approach taken in Ng et al. 03). You can use a raytracer if you have one (or any other method, including a publicly available rendering system) to create images from a number of different properly chosen light source locations. For compatibility with later parts of the assignment, choose (say) 20 or so light source directions, which are directional lights, These could be uniformly distributed on the sphere, or chosen in such a way as to create interesting visual effects in your scene. Implement a simple visual system that allows you to change the intensities (and possibly colors) of these lights and re-renders the scene in real-time. This is done simply by linearly combining the source images, based on the light source intensities. This system provides a basic framework for interactive lighting design of images.

Similar to Ng et al. 03, you could also make the computations on vertices of a diffuse mesh, allowing you to visualize a Lambertian model in simple OpenGL, where you can also change viewpoint. Do this last part (vertex computation and meshes) before the other components, only if it is relatively straightforward.

1.1.2 Image-Based Lighting

The next step is to consider image-based environment lighting, as for example in Sloan et al. 02 and Ng et al. 03. Environment maps are now very common, and high-dynamic range lighting environments are available from a number of websites, like Paul Debevec's. You could also try painting your own. A common representation is a cubemap, where the sphere is projected onto 6 faces of a cube that are unwrapped.

Ng et al. 03 used cubemaps with resolutions up to $6 \times 64 \times 64$. However, the larger resolutions will make your system harder to build, and so I recommend starting small, with say resolutions of $6 \times 16 \times 16$, and perhaps images of size 256×256. Actually, if you implemented the vertex-based method above, you could try very small geometric meshes of about 10,000 triangles. Note that a cubemap resolution of $6 \times 16 \times 16$ is still 1536 images and the rendering precomputation could be slow. You may want to turn off global illumination in your ray tracer, and just do direct lighting, that could be much faster (it effectively reduces to visibility computations by tracing rays). If speed is still a concern, you could also try implementing the hardware rasterization approach in Ng et al. 03, where you go to each vertex or pixel, and rasterize the model to get visibility over the cubemap.

You now have a full light transport matrix. Assuming an image resolution of 256×256 and light resolution of 1536, this is about 100 million entities. This is still small enough that a floating point version could fit in main memory, but it quickly becomes intractable as you scale up cubemap and image resolution. You can see what rendering speeds you get without compression. It will likely not be real-time, but may be a few seconds for modern computers with low resolutions.

Note that you will likely want to implement something to manipulate the environment map lighting, most commonly by rotation, but you could also simply translate it along the cubemap. This will allow new lightings to be created, which shows the power of the method for relighting.

1.1.3 Wavelet and other Transforms

Your next task is to wavelet transform the transport matrix along its rows (i.e. along the lighting cubemap) as in Ng et al. 03. You can use the simple Haar wavelet for this purpose. You can also wavelet transform the lighting and choose only the first 100 or 200 terms, making the transport-lighting dot product tractable. See if you can achieve real-time framerates with this, and do some experimentation. Be careful to program the inner loop in a memory-coherent way for maximal efficiency. Having got this far, you can play with increasing resolutions, and using dithering/quantization, as well as truncating elements of the transport matrix to 0 as in Ng et al. 03.

If you complete the assignment so far, you may also want to try using low-order spherical harmonics instead, and compare wavelets and spherical harmonics. See if you can reproduce the types of graphs and comparisons made in Ng et al. 03. More ambitious groups may want to implement some of the methods in Sloan et al. 02, to produce real-time demos of the type one would want in video games. Even more ambitious extra credits are to do all-frequency lighting and view manipulation along the lines of Wang et al. In any event, if you go beyond the base requirements of the assignment, extra credit will be available.

Note that a key component of this assignment is the final demo of relighting with complex shadowing and/or global illumination effects. Noting which of the stages above you were able to complete, and the quality of the final demo will be important factors in grading.

1.1.4 Grading

The grading scheme is subject to change, but tentatively assigns 7 points for the writeup and video, 12 points for basic relighting, 15 for environment maps, 10 for wavelets, and 6 points for robustness and other issues.

1.2 Basic High Quality Rendering

The purpose of this project is to develop an understanding of the basics of modern high quality rendering. Towards this end, you will implement a basic program, likely using GLSL shaders, and including environment and shadow mapping. The instructions in this assignment are relatively brief; perhaps more detailed step-by-steps will be provided in a future instantiation of the course. The best source of more detailed instructions is probably online tutorials and user manuals (such as for GLSL). Note that you can probably find source code online to implement many of the core components of this (or any other) project. While you may look at that as a reference, do not directly copy it (when in doubt if a code segment is a core component or a utility aspect you can use, please ask). It would also be interesting to do this assignment using the OptiX real-time raytracer, although you are more on your own for support.

1.2.1 GLSL Shaders

Modern graphics cards include the capability of developing vertex and fragment shaders, that provide control over various aspects of the graphics pipeline. They are therefore a key element to implement various high quality rendering effects. Shader development has undergone several iterations, including Cg and HLSL. However, they are now a core component of OpenGL 2.0, 3.0 etc. documented in the OpenGL red book with a GLSL (GL shading language) available. Your first task is to get basic GLSL programmable shaders working, based on the documentation in the red book, the GLSL manual, and other online sources. Demonstrate a simple basic program that uses shaders to do something interesting, that may be more difficult or impossible in the standard OpenGL pipeline.

1.2.2 Shadow Mapping

Next, you will implement a common shadowing technique known as shadow mapping. The scene is first rendered from the light source, with depth values stored. There are special OpenGL $GL\_DEPTH\_COMPONENT$ specifications for these depths. This scene is then rendered from the camera, with depth values projected back to the light and compared to determine shadowing. The exact details of the technique can be found in many textbooks and online tutorials. The key issues are to set up the projection matrices correctly to project back to the light, and execute the right programs. You can implement this with GLSL shaders. Try to implement percentage-closer filtering for antialiasing (the basic idea dates back to Reeves Salesin Cook from SIGGRAPH 87). By applying this properly, you can also simulate penumbrae from soft shadows.

1.2.3 Environment Mapping

The other technique you need to implement is environment mapping. This gives the effects of mirror reflections from image-based lighting. For each point on the surface, you look up the surface normal and view direction, reflect the view about the normal, and use that to index into the lighting. The indexing is simply a texture lookup into the environment cube map. This should be relatively straightforward to program using GLSL shaders, although there are also techniques to do so using standard OpenGL. With shadow and environment mapping, you should be able to make nice images with complex lighting and shadows. Note that environment-mapped shadows are a challenging problem, so you may want to pick the brightest lights in the environment for shadow mapping.

1.2.4   Add-Ons

Besides the basic requirements above, you need to implement at least two add-ons (use your imagination). Since this project is somewhat easier than the others, we reserve the right to consider the complexity of the add-ons before providing full credit. The most natural example could be in environment mapping, where besides mirror reflections, you consider diffuse and Phong Shading. You may also want to look at Heidrich et al. SIGGRAPH 99, "Realistic Hardware-Accelerated Shading and Lighting" for some ideas on interesting things you could do (many of the methods there could be implemented more simply than in the paper, using programmable shaders).

As described by Heidrich et al., and earlier in Miller and Hoffman, diffuse shading simply involves convolving your environment map with the Lambertian shading function (a cosine over the upper hemisphere). This diffuse map can then be looked up directly using the surface normal. For an easier procedural technique, read Ramamoorthi and Hanrahan, SIGGRAPH 01, "An Efficient Representation for Irradiance Environment Maps". For the specular, you instead convolve with the Phong filter, and look up based on the reflected direction. More ambitious projects may want to implement general BRDFs, using Cabral et al. 99, "Reflection Space Image Based Rendering" or Ramamoorthi and Hanrahan 02, "Frequency Space Environment Map Rendering".

With shadow maps, there is wide literature on various enhancements possible, including perspective shadow maps, adaptive shadow maps and so on. There are also various soft shadow mapping techniques that can "fake" shadows from an area light. If you are interested, talk to the instructor.

A number of other possibilities for add-ons also exist. In any event, the final goal is a high-quality real-time rendering demo that shows object(s) with a variety of visual effects being created.

1.2.5   Grading

The grading scheme is subject to change but tentatively has 13 points for basic real-time rendering (a solid video and demo is needed), and 10 points each for shadows, environment maps and add-ons. 7 points will be given for documentation and scene/video selection.

1.3   Image-Based Rendering: Light Field Viewer

Image-Based Rendering is an exciting newer area of computer graphics. Instead of rendering synthetic images from "scratch", one uses acquired or rendered images as inputs, interpolating between them to render intermediate views. The recent commercial development of light field cameras and focusing and view changes after the fact has brought this area to the average consumer; please ask the instructor if you want to do a project closer to the current state of the art of light field cameras, rather than what is recommended below.

In this project, you will implement a viewer for one of the first image-based rendering techniques, Light Field Rendering by Levoy and Hanrahan in SIGGRAPH 96. This paper is very similar to The Lumigraph by Gortler et al., also in SIGGRAPH 96, but the former paper is conceptually cleaner. The idea is to write a light field viewer for light field datasets. Useful resources will be the datasets available at the Stanford graphics lab website (look at both the old and new data), as well as in other places. Note that there are some issues with the exact format of the Stanford data, and you may need to do some reverse engineering to get it to work (this was somewhat frustrating for students in the past). You may want to start with synthetic datasets (see below).

You can also use a raytracer (including PBRT or other online program) to render synthetic datasets, which can then be more tightly coupled to your viewer. In the initial phase, this may be the preferred method of choice, since the Stanford lif files may require VQ decompression to be implemented in order to read. It is easy to simply render images from multiple viewpoints one at a time. Of course, eventually you will want to use datasets with acquired light field data, captured as actual images. (note that rendering images from multiple views as a precomputation is very similar to the multiple light images that will need to be rendered for the precomputed radiance transfer project. Indeed, similar compression methods can also be used).

Textual material of interest is the original paper, and the image-based modeling and rendering course notes from SIGGRAPH over the years. You may also look for a reference at the light field viewers and code

at the Stanford graphics lab website, though of course you cannot copy directly.

1.3.1 Basics

In its simplest 2-plane parameterization form, a "ray" in a light field is given by its intersection on two planes. Therefore, to render an image, for each ray, one needs to look up the appropriate light field element, based on intersections of the ray with the two planes. There are of course more effective hardware texture-mapping based approaches, some of which are discussed in the original papers, that you may want to think about (but this is not required).

When you find the intersection of the ray with the two planes, make sure to account for if the ray intersects outside the valid regions or behind the viewer (in these cases, the source light field has no data, and the pixel is simply set to the background color). Once you find the intersection, experiment with finding the closest sample, as well as bilinear and quadrilinear interpolation. Which works best?

As a simple basic test, you may want to render your own light fields at low resolution, and compare with direct hardware rendering of objects, to make sure the basic mechanism is working.

1.3.2 Input Data and Light Field Viewer

Your basic light field viewer will read a source input file and draw the light field. At this time, we don't specify a particular file format for the inputs, but you may want to be compatible with common datasets such as the Stanford light field archive. The elements of the parameters file are:

- *Resolution:* The first line could have the resolution in u v s t, for example 16 16 256 256.

- *Area of light field slice:* This should tell you the locations of where the uv and st planes are respectively. The line could have the form: $uv_x^1 \ uv_y^1 \ uv_z^1 \ uv_w^1 \ uv_p^1 \ uv_q^1$, where the superscript 1 stands for the first corner of the uv plane, and subscripts stand for x,y,z and possibly w homogeneous coordinate. p and q are the corresponding texture coordinates (in the $[0,1]$ range). You would have a similar line for each of the 4 corners of the uv plane and then 4 corners of the st plane. For example, a line of the form $-2 \ 2 \ -2 \ 1 \ 0 \ 0$ defines the "bottom left" corner with texture coordinates $(0,0)$ and spatial location $(-2,2,-2)$. This is similar to the way vertex and texture coordinates are specified in OpenGL

- *Data:* Finally, the data must be given. This could be in a separate file. It could simply be the raw uvst data in RGB triples (the size will be the product of resolutions of u,v,s,t times 3).

Finally, you need to implement the viewer. For this purpose, develop some kind of mouse-controlled view manipulation, such a crystal ball interface. Then, for each pixel, generate a ray as in standard ray-tracing, except you actually intersect with the light field planes, and use that to look up the color in the light field data. Finally, you can draw the image on the screen using standard OpenGL (for example, glDrawPixels()).

1.3.3 Add-Ons

You should implement at least one add-on. One simple idea is to extend the single-slab light field to multiple slabs, so that you have a closed view of the object. Another idea is to implement VQ or another compression/decompression method (VQ decompression was used in the original light field paper, but some of the light transport theory for precomputed radiance transfer is also relevant). Another simple test is to implement various combinations of nearest, bilinear and quadrilinear interpolation, and compare their relative benefits versus cost. Finally, for the more ambitious, you could use partial geometry like in the lumigraph paper to improve the quality of the light field interpolation and rendering.

1.3.4 Grading

The tentative grading scheme provides 13 points for a basic working demo (with video), 8 points for good documentation and scene selection, 10 points each for the viewer and suitable interpolation method, and 10 points for the add-ons.

2   Less Well-Specified Projects

What follows are some projects that are specified in less detail, that give you more flexibility. Note that you have the three options above, all of the options below, and you can also propose a project of your own.

Subdivision Surfaces: This project aligns most closely with the geometric modeling aspect of homework 2. Subdivision is a new and very popular paradigm for shapes with complex topology. Given a control mesh, it creates a smooth limit surface. The basic idea is to implement Catmull Clark Subdivision with exact evaluation of the limit surface per Jos Stam's Siggraph 98 paper on "Exact Evaluation of Catmull-Clark Subdivision Surfaces at Arbitrary Parameter Values."

The first part is to write a basic subdivision routine. As with all subdivision schemes, you should operate on arbitrary input meshes with complex topology. You may want to start with basic Catmull Clark. Thereafter, include ways of "pushing the surface to the limit" after some number of subdivisions, per Stam's work. Then, develop code to compute tangents, normals and curvatures (Gauss and Mean) at the vertices. See if you can include some color coded plots, as in his paper. The input-output interface is up to you, but you could for example, read OBJ or OFF files for geometry. You could use as examples many of the meshes in assignment 2.

For extra credit, we ask you to write code to move specified points of the control mesh to make the surface as smooth as possible—this typically involves minimizing surface area or integral squared curvature, but define a suitable error metric. You can use any suitable optimization method, and can use off the shelf code for this purpose (if in doubt write your own conjugate gradient descent solver from Numerical Recipes or Shewchuk's website). Your optimization can make use of discrete approximations to the minimized function but your visualizations (part 1 or the paragraph above) should be exact. Show some interesting surfaces that you make by optimizing various energy functions, including combinations of surface area and curvature.

You will need to carefully demonstrate this assignment, including example images with different parameters, shaded and wireframe examples, the control mesh and a few steps of subdivision, animations and (if you do the extra credit) still images for results of optimization and visualizations for gradients and curvatures.

Animation: The goal is to build a basic inverse kinematics system. You may want to start with the analytic formulae for a two-joint arm, but you would ideally like to have a general solver for at least four joints with different lengths. The root should be stationary (attached to an immobile base). There should be some interactive or offline mechanism to specify the goal point in 3D (or create a procedural animation), and the IK solver should select joint angles to reach the desired goal. For submission, make a video that clearly shows your system working (and any problems if relevant), where the goal path is something complicated, that tests the full range of motion (a circle or large arc or figure-eight, not just a line). It would also be nice if you gracefully handled configurations that cannot be achieved (are out of reach). Make sure to try to cover the full range of the arm's operating space in terms of joint angles. Clearly, a full video documenting everything that was done is crucial. Note that the rendering can be very simple, just in OpenGL with simple shading, or any other relevant code (including software online); we want to focus on the animation.

Imaging: Just like physical simulation, image acquisition and manipulation, and computational photography, are areas of growing interest in computer graphics. In this project, you should implement some of the basic imaging algorithms. As with physical simulation this is largely open-ended, and grading will be based on documentation and apparent effort. One suggestion is to do the HDR imaging method from Debevec 97 (if you don't want to take your own photographs, you can probably find some images on the web), and create a high-dynamic range image that can be displayed with any of the recent tonemapping algorithms. Another suggestion is to focus on texture synthesis, implementing at least the basic Efros and Freeman 01 method for image quilting, and at least one major extension.