Assignment P2 : RAY CASTING

Computer Graphics 2

Due 16 March 1995 by end of day

For this assignment you will implement a simple ray casting program for spheres and polygons. Ray casting is one-level (non-recursive) ray tracing: at each pixel you determine what surface point is visible and shade it.

Basics

Your program will have the following steps:

1. Read a 3-D scene file. We will provide the scene files and a routine, `scene_read`, to parse scene files. This routine will provide you with lists of spheres, polygons, and lights in the scene, and the 4×4 viewing matrix.
2. Loop over output pixels, for each generating a ray in world space.
3. Test each ray for intersection against all spheres and polygons, and determine the nearest intersection point (if any).
4. If the ray hits a surface, shade the intersection point using Phong's illumination formula (given below). Otherwise, if the ray hits nothing, use the background color. The resulting RGB color is used as the pixel value.
5. Write pixels to a PPM picture file.

Thus, the input of your ray caster is a scene file and the output is a PPM picture file. You can view the picture with `xv`, `xview`, or `xplay`.

The source code and scene files we provide are in `/afs/andrew/scs/cs/15-463/pub/p2`. See the `README_P2` file there. We will provide two scene files describing simple scenes: `min.sc`, a single polygon and a single sphere illuminated by a single light, and `table.sc`, a dodecahedron and sphere on a table. See `www/assts/scene.ps` for documentation on scene file format.

Programming Tips

These files already have camera transformations built into them. If you want to view the scene from different angles, you can split the scene file into two parts as has been done in `min.cam.sc` and `min.nocam.sc`, and modify just the former. To support this, it is suggested that your `raycast` program accept multiple scene files as arguments, e.g. `raycast min.sc` or `raycast min.cam.sc min.nocam.sc`, calling `scene = scene_read(scenefilename, scene)` on each file. On the first call, pass in `scene==NULL`, but on subsequent calls to `scene_read`, pass along the scene pointer returned from the previous call.

After calling `scene_read`, the matrix that transforms world space to screen space (a.k.a. the "viewing matrix") is available in `scene->worldtoscreen`. The requested width and height of the output picture are available in `scene->width` and `scene->height`. See the source files `scene.h` and `scene.c` for more details.

In assignment P3, you will extend your program to do shadows, recursive ray tracing (good for chrome & glass) and antialiasing. Since these modifications imply even more calls to your ray-sphere

and ray-polygon intersection routines, you may want to optimize those routines a bit, especially if you'll be working on a slow machine.

It is suggested that steps 3 & 4 above be modularized into a routine that takes a ray as argument and returns a color, e.g.

```
trace(Ray *ray, Color *col)     /* ray: input; col: returned */
```

Shading

The shading formula used[1] should be Phong's illumination model[2]:

$$\mathbf{I} = \mathbf{C}_{mat}\left[\mathbf{I}_a + \sum_i \mathbf{I}_i k_{dr} \max\{(N \cdot L_i), 0\}\right] + \sum_i \mathbf{I}_i k_{sr} \max\{(N \cdot H_i), 0\}^e$$

where the summation is over all light sources, boldface denotes an RGB triple or color[3], italic capital denotes a 3-D direction vector, lower case denotes a scalar, and "$\cdot$" denotes dot product. General variables are:

$\mathbf{I}$ is the returned radiance of the ray (a color)

N is the surface normal at the intersection point

Light source parameters are:

$\mathbf{I}_a$ is the radiance of ambient light (a color)

$\mathbf{I}_i$ is the radiance of light source$i$, (a color)

L_i is the direction of light i from the intersection point, computed as the difference of the light's position and the intersection position

H_i is the halfway vector for Phong's illumination model:
$H_i = (L_i + V)/|L_i + V|$ where V is direction of the viewer,
i.e. V is the opposite of the incident ray direction

Surface material parameters are:

$\mathbf{C}_{mat}$ is the reflectance of the intersected surface's material (a color)

k_{dr} is the coefficient of diffuse reflection

k_{sr} is the coefficient of specular reflection

e is related to surface roughness and controls highlight size

Light source parameters and material parameters are available in the structures `Light` and `Material` described in `scene.h`.

Shading Notes

All direction vectors should be unit vectors (make them unit length, if you don't know a priori that they are). The max with zero is a backfacing test. The product of $\mathbf{C}_{mat}$ with the radiance quantity

[1] See Foley et. al equations (16.14), (16.20), and discussion of halfway vector H on page 731. To simplify, we're assuming light source attenuation factor = 1 and a white specular color.

[2] Note that Phong illumination is not the same thing as Phong shading. The former is a formula for computing diffuse and specular reflection of any geometric surface type; the latter is a technique for interpolating normal vectors across polygons to make them look curved, not faceted.

[3] "Color" is a vague term, but in computer graphics it usually means spectral samples or color coordinates of something in some color space. Sometimes it specifically means RGB spectral samples of radiance (a.k.a. intensity), sometimes it specifically means RGB spectral samples of reflectance.

$[\mathbf{I}_a + \cdots]$ is a product of two 3-vectors. This product should be done componentwise, e.g. multiply the reds to get the red component, the greens to get the green component, etc. Typically, $k_{dr} + k_{sr}$ and each component of $\mathbf{C}_{mat}$ lie between 0 and 1. If the computed radiance is $\mathbf{I} = (I_r, I_g, I_b)$ then your final pixel values should be $(255 * I_r, 255 * I_g, 255 * I_b)$, but since the radiances ($\mathbf{I}_i$ and $\mathbf{I}$) are not necessarily bounded, you should check each component of the pixel value for overflow, pegging at 255.

What to Submit.

Submit in your `p2` directory:

1. Source code to your `raycast` program, (called "`raycast`", please).
2. A working executable of your `raycast` program that we will be able to run (executable on either a DEC, SGI, or Sun4).
3. PPM picture files for `min.sc` and `table.sc`. (Call them `min.ppm` and `table.ppm`).
4. A modified version of `table.sc` where you have changed the camera and some of the objects (positions, colors, and/or lights). No need to get fancy, but you should make some changes. Call the new scene file `newtable.sc` and create a picture of it with your ray caster, called `newtable.ppm`.
5. A text file `README` where you summarize briefly what you've done, and what each of the files in your directory is. Tell us how to run your program on `min.sc` to regenerate `min.ppm`.

Notes: feel free to modify the data structures in `scene.h`. You will probably want to pre-compute various information (plane equations, etc) for polygons, for instance. You can also modify `scene.c` and `matrix.[ch]` if you like. Optimization is optional, not required.