

# An Introduction to Ray Tracing

*TR87-038*

*December, 1987*

*Roman Kuchkuda*

The University of North Carolina at Chapel Hill  
Department of Computer Science  
Sitterson Hall, 083A  
Chapel Hill, NC 27599-3175



# An Introduction to Ray Tracing

Roman Kuchkuda  
University of North Carolina  
Chapel Hill

December 20, 1987

## Abstract

This paper is a practical guide to ray tracing for those familiar with graphics. It consists of a conceptual model of ray tracing, C code for a basic system, and an explanation of how and why the code works.

## Key Words & Phrases

Ray Tracing, C code, Conceptual Model, Spheres, Boxes, Triangles, Superquadrics, Reflection, Shadows, Transparency, LEX, YACC

## 1 INTRODUCTION

Ray tracing is one of the most popular algorithms for rendering high quality computer graphics images. It is conceptually simple; rendering is reduced to finding the intersection of a line with an object and then shading the point of intersection. It is powerful; almost any type of object can be rendered. It is extensible; adding a variety of new effects is easy and natural.

For all its conceptual simplicity, it can be a difficult subject to learn. The classic computer graphics texts only mention it. Current technical journals contain articles about advanced features and special extensions. Articles on the basics were printed many years ago and can be difficult to find.

This paper attempts to fill the need for a beginner's guide to ray tracing. It consists of three parts:

- Part 1 presents a conceptual model of ray tracing and discusses the translation of the model into code.
- Part 2 presents C code for a basic ray tracing system. Each section of code is accompanied by an explanation of how and why it works. Figure 1 shows the image generated by the basic system, using the default image parameters.

- Part 3 presents extensions to the basic system. Figures 2 - 4 show images generated by the full system.

All of the code necessary to build a simple ray tracer is included. Figures 1-4 were produced using code which was filtered from the text of this document. Only the standard features of C are used. In some cases, execution speed has been sacrificed in the interest of clarity.

### 1.1 CAUTIONS

The chief drawback of ray tracing is that it is slow. Complicated images take HOURS of CPU time. Be kind to your fellow users; run your ray tracer in the background at low priority.

Image files are large. Keeping many images on disk will quickly use a lot of space.

Finally, beware - ray tracing is addictive. The beauty of your images will entrance you. There is always one more effect that can be added to produce a new and more spectacular picture.

### 1.2 CONCEPTUAL MODEL

The fundamental idea behind ray tracing is to follow rays of light from a light source. Their paths and intensity/color are computed as they reflect off and refract through objects.

Imagine a camera on a tripod sitting in a room filled with a variety of objects. The room is lit by a single light bulb. Rays of light from the light bulb reflect off objects. Some of the light rays travel through the camera's lens and hit the film, creating an image. This is the process which ray tracing attempts to simulate.

Everything in the room can be represented mathematically. The light bulb has an x,y,z position and a brightness. The objects can be modeled as geometric primitives, such

as spheres or polygons, or as collections of primitives. The camera can be described by its position, orientation, and the field of view of the lens.

In the program, light rays are sent out from the light bulb in all directions. The angles of reflection and refraction of the light rays are computed as they bounce off objects. The surface properties (color, shininess, etc.) of the objects change the color and intensity of the rays. Some of the rays eventually strike the film. The color and intensity of those that do are recorded.

To produce a clear image, many rays must strike the film. Because only a small percentage reach the film, a tremendous number of rays must be traced. It is incredibly expensive, but the image should be nearly perfect if the model of light interaction is good.

The key to practical ray tracing is a slight adjustment to this process. Rather than trace rays from the light source, with most of the rays missing the film, rays are traced from points on the film out of the camera into the room. One ray is traced for each spot on the film, with the color of the object hit recorded as the color for that spot. On the graphics display, each pixel corresponds to a spot on the film.

### 1.3 HISTORY

The earliest reference in computer graphics to ray tracing is [Appel 68]. [Goldstein 71] described image production software at MAGI. [Whitted 80] extended ray tracing by adding reflection and refraction. Since this early work, ray tracing has benefitted from the efforts of many researchers. Some areas of development are described in Section 3.9.

## 2 FIRST PROGRAM

The first ray tracing program is designed to allow you to quickly produce some simple images. It also provides a solid base for further development.

A ray tracing system includes routines for many types of object primitives. The first program includes only the code for spheres. In the full system, the code for boxes, triangles, and superquadrics is presented. As you develop your ray tracer, you can add more primitives.

The shading model in the first program is fairly sophisticated. It includes ambient, diffuse, and specular lighting. Several point light sources are allowed. In the full system shadows, reflections and transparency will be added.

In a flexible system many values are specified by the user at run time. The code to obtain user input is fairly long. So,

the values are hard coded in the basic system. In section 3.6, user input routines are presented. For now, run the program and see what the image looks like. Try changing some of the hard coded values and see how the image changes.

### 2.1 MAIN ROUTINE

The main routine determines a color value for each pixel in the image and writes this value to the pixel output file.

Initializations are performed by a call to `setup`. `viewing` is called to calculate the direction and spacing of the rays to be traced. The pixel output file is opened and the image size written to it in `startpic`.

The actual ray trace is then performed. The rays are generated, in order, from left to right, top to bottom, by two nested loops. The outer loop is 1 to  $n$ , where  $n$  is the number of lines in the image. The inner loop is 1 to  $m$ , where  $m$  is the numbers of pixels across each line. Inside the inner loop, a ray is sent into the world by a call to `intersect`.

`intersect` computes the intersections of a ray with all objects. It determines the point of closest intersection and returns the color of the object at that point. This color value is stored. If the ray doesn't hit any objects, the background color is stored.

When the colors of all of the pixels across a line have been determined, the line is output through a call to `linepic`. Every tenth line a status message is sent to `stdout`. When all of the lines in the image have been completed, `endpic` closes the output file.

```
/** main.c **/  
#include <stdio.h>  
#include "typedefs.h"  
#include "maindecl.h"  
  
main()  
{  
    int line_y, pixel_x;  
    t_3d scrnx, scrny, firstray, ray;  
    t_color color;  
    double dis, line[SCREENWIDTH][3];  
  
    setup();  
    viewing(&scrnx, &scrny, &firstray);  
    startpic(outfilename, sizey, sizex);  
  
    for (line_y = 0; line_y < sizey; line_y++)  
    {  
        for (pixel_x = 0; pixel_x < sizex; pixel_x++)  
        {  
            ray.x = firstray.x + pixel_x*scrnx.x -
```

```

        line_y*scrny.x;
ray.y = firstray.y + pixel_x*scrnx.y -
        line_y*scrny.y;
ray.z = firstray.z + pixel_x*scrnx.z -
        line_y*scrny.z;
normalize(&ray);

/* actual ray trace */
dis = intersect(-1,&eyep,&ray,&color);
if (dis > 0) /* ray intersected object */
{
    line[pixel_x][0] = color.r;
    line[pixel_x][1] = color.g;
    line[pixel_x][2] = color.b;
}
else /* use background color */
{
    line[pixel_x][0] = background.r;
    line[pixel_x][1] = background.g;
    line[pixel_x][2] = background.b;
}
}
linepic(line); /* output line of pixels */
if (line_y%10 == 0)
{
    printf("done line %d\n",line_y);
    fflush(stdout);
}
}
endpic(); /* done with picture */
}

```

## 2.2 INCLUDE FILES

Include files define structure types, global variables, constants, and function types. Their use simplifies later modification of the program.

The first include file, `typedefs.h`, defines the structures for the basic ray tracer. Structures allow conceptual grouping of data and simplify information passing between routines. New attributes can easily be added to structures.

Spheres, boxes, triangles, and superquadrics are defined. Each object is described in two structures. The first structure has the fields which are common to all objects. The second has the fields that are unique to a given object.

`t_sphere` describes spheres. It has fields for the sphere's radius and its  $x,y,z$  center position.

The box object type is defined by the `t_box` structure. It has  $x,y,z$  center position, sizes for the  $x$ ,  $y$ , and  $z$  sides and `sidehit`. `Sidehit` is used in the computation of the surface normal of the box.

Triangles are described by the `t_triangle` structure. It

includes the triangle's plane description, three edge vectors, and three edge constants.

The superquadric object structure, `t_superq`, has all of the values of the box structure, plus some new values. These values are described in section 3.5.

The `t_object` structure contains the fields which are common to all objects. It has object id, which is unique for each object, surface type number, which specifies the color of the surface, object type, and object pointer, which points to one of the object type structures defined above.

The `t_light` type defines a point light source. It has an  $x,y,z$  position in space and a brightness value.

The `t_surface` structure describes the properties of an object's surface. The color is specified by red, green, and blue components for each of ambient, diffuse, and specular lighting. The specular coefficient controls the size and intensity of the specular highlight. It is a positive integer and the higher the coefficient, the smaller and brighter the specular highlight. The reflectivity of the surface is specified by a number from zero to one, with zero being no reflection, one being a perfect mirror. Transparency is specified in a similar manner, with zero being opaque, one completely transparent. Reflections and transparency are not included in the basic program; they are added later.

Three-dimensional vectors are represented by the `t_3d` structure. It is used throughout the program for both positions and vectors in 3 space.

Lastly, the `t_color` structure specifies a red, green, blue color trio. Color values are passed throughout the program in this form.

```

/** typedefs.h */
typedef struct{double x,y,z;} t_3d;

typedef struct{double r;          /* radius */
               double x,y,z;     /* position */
               }o_sphere ;

typedef struct{int sidehit;      /* side intersected */
               double xs,ys,zs; /* size of sides */
               double x,y,z;     /* center position */
               }o_box;

typedef struct{ t_3d nrm;        /* triangle normal */
               double d;        /* plane constant */
               t_3d e1,e2,e3;    /* edge vectors */
               double d1,d2,d3; /* plane constants */
               } o_triangle ;

typedef struct{int sidehit;      /* side intersected */
               double xs,ys,zs; /* size of sides */

```

```

double x,y,z;      /* center position */
double pow;        /* n in formula */
double a,b,c,r;    /* coefficients */
double err;        /* error measure */
}o_superq;

typedef struct{int id;          /* object number */
int objtyp;          /* object type */
int surfnum;        /* surface number */
union {
o_sphere *p_sphere;
o_box *p_box;
o_triangle *p_triangle;
o_superq *p_superq;
} objpnt; /* object pointer */
}t_object ;

typedef struct{double x,y,z,bright;} t_light;

typedef struct{double ar,ag,ab; /* ambient r,g,b */
double dr,dg,db; /* diffuse r,g,b */
double sr,sg,sb; /* specular r,g,b */
double coef; /* specular coef */
double refl; /* reflection 0-1 */
double transp; /* transparency 0-1 */
}t_surface;

typedef struct {double r,g,b;} t_color;

```

constants.h contains a variety of constant definitions. First, the sizes of the light, object, and surface type arrays are defined. The array sizes must be increased as you render images with more objects (or lights or surface types) in them. Only the main program should require recompilation when these sizes are changed.

Next the screen size is defined. It should not change unless you move to a frame buffer with a different resolution or aspect ratio. The gamma correction constant is likewise frame buffer dependent.

Finally, object type constants are specified. These constants relate the object type to the routines used to find intersections and normals. The constants (0 - n) correspond to the order of the intersection and normal routines in the objint and objnrm arrays in maindecl.h.

```

/** constants.h */
#define LIGHTS 4
#define OBJECTS 50
#define SURFACES 50
#define SCREENWIDTH 512
#define SCREENHEIGHT 512
#define ASPECTRATIO 1.0
#define GAMMA 1.8
#define OTYPSPHERE 0

```

```

#define OTYPBOX 1
#define OTYPTRIANGLE 2
#define OTYPSUPERQ 3

```

maindecl.h contains the *declaration* of the global variables and is included in the main program.

Global information has been minimized. This tends to lessen unforeseen interactions among routines. The global variables consist mainly of values which define the ray tracing world.

First, the object intersection routine and object normal routine pointer arrays are declared. The objint array has pointers to the ray-object intersection routine for each type of object. The objnrm array has pointers to the surface normal routine for each object type.

Next, the arrays for objects, lights, and surface types are declared. A maximum size for each array is kept, along with a count of how many items in each array are currently being used.

The viewing parameters are declared next. They include the size of the image to be produced, the eye point, look point, up vector, and horizontal and vertical fields of view. Section 2.5 describes these parameters.

The current level and maximum level of reflection are used in the full system to limit the number of levels of reflection which are rendered.

The string for the pixel output file name is then declared. Finally, the background color is defined. It is used when a ray does not hit any object.

```

/** maindecl.h */
#include "constants.h"
#include "funcdefs.h"

/* intersection routines */
double (*objint[])() = {intsph,intbox,inttri,intsup};
/* normal routines */
int (*objnrm[])() = {nrmsph,nrmbox,nrmtri,nrmsup};

/* global variables */
int nlight; /* presently in use */
int lightlim = LIGHTS; /* maximum declared */
t_light light[LIGHTS]; /* array of lights */
int nobject; /* presently in use */
int objectlim = OBJECTS; /* maximum declared */
t_object object[OBJECTS]; /* array of objects */
int nsurface; /* presently in use */
int surfacelim = SURFACES; /* maximum declared */
t_surface surface[SURFACES]; /* array of surfaces */
int sizex, sizey; /* image sizes */
t_3d eyep,lookp,up; /* view definition */
double hfov,vfov; /* fields of view */

```

```

int level,maxlevel;          /* reflection level */
char outfile[60];           /* pixel file name */
t_color background;         /* background color */

```

globalvar.h contains the *definition* of the global variables and is included in routines which access the global variables.

```

/** globalvar.h */
extern double (*objint[])( );
extern int (*objnorm[])( );
extern int nlight,lightlim;
extern t_light light[];
extern int nobject,objectlim;
extern t_object object[];
extern int nsurface,surfacelim;
extern t_surface surface[];
extern int sizex,sizey;
extern t_3d eyep,lookp,up;
extern double hfov,vfov;
extern int level,maxlevel;
extern char outfile[];
extern t_color background;

```

funcdefs.h defines the names and types for all of the routines in the ray tracer. It is included in all routines which have calls to other routines.

```

/** funcdefs.h */
double brightness();
int crossp();
double dotp();
int endpic();
int gammacorrect();
double intersect();
int lightray();
int linepic();
double normalize();
int setup();
int shade();
int startpic();
int viewing();
int yyparse();

int maksph();
double intsph();
int nrmsph();
int makbox();
double intbox();
int nrmbbox();
int maktri();
double intrtri();
int nrmttri();
int maksup();
double intsup();
int nrmsup();

```

## 2.3 SETUP ROUTINE

setup initializes global variables and sets defaults for user specified values. yyparse is called to get user input.

```

/** setup.c */
#include "typedefs.h"
#include "globalvar.h"

int setup()
{
    /* set defaults */
    nlight = 0;
    nobject = 0;
    nsurface = 0;
    level = 0;
    sizex = 512;
    sizey = 512;
    hfov = 50;
    vfov = 50;
    eyep.x = 100.0;
    eyep.y = 0.0;
    eyep.z = 0.0;
    lookp.x = 0.0;
    lookp.y = 0.0;
    lookp.z = 0.0;
    up.x = 0.0;
    up.y = 1.0;
    up.z = 0.0;
    strcpy(outfile, "raytrace.pix");
    yyparse(); /* parse user input */
}

```

## 2.4 FIRST PROGRAM STUBS

The basic system contains stubs for routines which will be defined in the full system. yyparse and the intersection and normal routines for boxes, triangles, and superquadrics are stubbed.

The intersection and normal routines are empty, returning zero when called. There is no user input in the basic system; yyparse consists of a series of hard coded assignment statements.

```

/** stubs.c */
#include "typedefs.h"
#include "globalvar.h"

/* intersection and normal routines */
int makbox(){return;}
int maktri(){return;}
int maksup(){return;}
double intbox(){return(0);}
double intrtri(){return(0);}

```

```

double intsup(){return(0);}
int nrmbox(){return;}
int nrmtri(){return;}
int nrmsup(){return;}

int yyparse()
{
    nlight = 1;          /* light source */
    light[0].x = 100;
    light[0].y = 300;
    light[0].z = 0;
    light[0].bright = 1.0;
    /* sphere - surface type 0, radius 40, */
    /* center at 5, 30, 40 */
    maksph(0,40.0,5.0,30.0,40.0);
    nobject++;
    /* sphere - surface type 1, radius 30, */
    /* center at -10, -10, -40 */
    maksph(1,30.0,-10.0,-10.0,-40.0);
    nobject++;
    nsurface = 2;
    surface[0].ar = 30; /* surface 0 - red shiny */
    surface[0].ag = 0;
    surface[0].ab = 0;
    surface[0].dr = 90;
    surface[0].dg = 0;
    surface[0].db = 0;
    surface[0].sr = 190;
    surface[0].sg = 120;
    surface[0].sb = 120;
    surface[0].coef = 15;
    surface[1].ar = 0; /* surface 1 - green, duller */
    surface[1].ag = 50;
    surface[1].ab = 0;
    surface[1].dr = 0;
    surface[1].dg = 100;
    surface[1].db = 0;
    surface[1].sr = 30;
    surface[1].sg = 40;
    surface[1].sb = 30;
    surface[1].coef = 2;
    background.r = 100; /* background color */
    background.g = 200; /* light blue */
    background.b = 250;
    return(0);
}


```

## 2.5 VIEWING ROUTINE

To produce an image, a ray is sent into the world for each pixel. The `viewing` routine calculates the direction and the spacing of the rays. It uses either the viewing parameters specified by the user or the defaults defined in `setup`.

The viewing definition is complicated. It is described using

the camera analogy introduced earlier. The points, vectors, and angles involved are shown in the following diagram.



The eyepoint (E) describes the camera position. The lookpoint (L) describes where the camera is pointing. It is the center of the image plane. Together, the eyepoint and the lookpoint define the gaze vector (G). The up vector (U) describes the orientation of the camera about the gaze vector.

Field of view refers to the length of the camera's lens, from wide angle to telephoto. Both horizontal (H) and vertical (V) fields of view are defined. For the usual rectangular display, the horizontal field of view is somewhat larger than the vertical. For a square display, they are equal.

After the view has been defined, several values are derived from it. The screen x (X) and screen y (Y) vectors describe the image plane. They are scaled to the width and height of one pixel on the image plane. Firststray (F) describes the ray from the eyepoint to the pixel in the top left corner of the image plane. Screen x, screen y, and firststray are used in the main program to generate the individual rays for each pixel.

```

/** viewing.c */
#include <math.h>
#include "typedefs.h"
#include "funcdefs.h"
#include "globalvar.h"
#define DEGREETORADIAN (3.1415926/180.)

int viewing(scrnx,scrny,firststray)
t_3d *scrnx,*scrny,*firststray;
{
    int i,j;
    t_3d gaze;
    double dist,magnitude;

    gaze.x = lookp.x - eyep.x;
    gaze.y = lookp.y - eyep.y;
    gaze.z = lookp.z - eyep.z;
    dist = normalize(&gaze);
    /* scrnx = gaze cross up */
    crossp(scrnx, &gaze, &up);
}

```

```

/* scrny = scrnx cross gaze */
crossp(scrny, scrnx, &gaze);

dist *= 2.0;
magnitude = dist * tan(hfov*DEGREETORADIAN) / sizex;
scrnx->x *= magnitude;
scrnx->y *= magnitude;
scrnx->z *= magnitude;
magnitude = dist * tan(vfov*DEGREETORADIAN) / sizey;
scrny->x *= magnitude;
scrny->y *= magnitude;
scrny->z *= magnitude;

firstray->x = lookp.x - eyep.x;
firstray->y = lookp.y - eyep.y;
firstray->z = lookp.z - eyep.z;
firstray->x += sizey/2*scrny->x - sizex/2*scrnx->x;
firstray->y += sizey/2*scrny->y - sizex/2*scrnx->y;
firstray->z += sizey/2*scrny->z - sizex/2*scrnx->z;
}

```

## 2.6 INTERSECTION ROUTINE

`intersect` finds the closest intersection between a ray and any of the objects in the world. The source of the ray, the ray vector, and the ray start position are input. The distance to the closest intersection and the color of the object at the point of intersection are returned. If the ray doesn't hit any object, a value of zero is returned.

For each object:

- The intersection routine associated with it is called. The distance to intersection is returned. Zero is returned if the ray misses the object.
- If the intersection is closer than any found previously, the object identity and distance to intersection are saved.

After the closest intersection is found, the point of intersection is computed. Letting  $X_p, Y_p, Z_p$  be the ray origin,  $X_r, Y_r, Z_r$  the ray direction, and  $s$  the distance to the intersection: the parametric equations for the ray give the point of intersection.

$$\begin{aligned} X &= X_p + s * X_r \\ Y &= Y_p + s * Y_r \\ Z &= Z_p + s * Z_r \end{aligned}$$

The normal at the point of intersection is found by calling the normal routine for the object hit by the ray. Then the shading routine is called to calculate the color at the point of intersection.

```
/***/ intersect.c /***/
```

```

#include <math.h>
#include "typedefs.h"
#include "globalvar.h"
#define FAR_AWAY 99.99E+20

double intersect(source, pos, ray, color)
int source;
t_3d *pos, *ray;
t_color *color;
{
    int objhit, objtry;
    double s, ss;
    t_3d hit, normal;

    objhit = -1;
    ss = FAR_AWAY;
    /* check for intersection of ray with all objects */
    for (objtry=0; objtry<nobject; objtry++)
    {
        /* special check used for reflections */
        if (objtry != source) /* don't try source */
        {
            s = (*objint[object[objtry].objtyp])
                (pos, ray, &object[objtry]);
            /* keep track of closest intersection */
            if ((s > 0.0) && (s <= ss))
            {
                objhit = objtry;
                ss = s;
            }
        }
    }

    if (objhit < 0) return(0); /* ray hit no objects */

    /* find point of intersection */
    hit.x = pos->x + ss * ray->x;
    hit.y = pos->y + ss * ray->y;
    hit.z = pos->z + ss * ray->z;
    /* find normal */
    (*objnrm[object[objhit].objtyp])
        (&hit, &object[objhit], &normal);
    /* find color at point of intersection */
    shade(&hit, ray, &normal, &object[objhit], color);
    return(ss);
}

```

## 2.7 SPHERE ROUTINES

The `maksph` routine allocates the space for a sphere description and fills it with the values passed to the routine.

`intsph` computes the intersection between a ray and a sphere. The source position of the ray, the ray vector, and the object, a sphere, are passed to the routine. The distance



to the nearest intersection is returned. A value of zero is returned if no intersection is found.

The equation of a sphere is  $X^2 + Y^2 + Z^2 = R^2$

The equations for the ray in parametric form, where  $X_p, Y_p, Z_p$  is the ray source and  $X_r, Y_r, Z_r$  is the ray vector, are

$$\begin{aligned} X &= X_p + S * X_r \\ Y &= Y_p + S * Y_r \\ Z &= Z_p + S * Z_r \end{aligned}$$

Substituting the parametric form in the sphere equation gives

$$\begin{aligned} S^2 * X_r^2 + 2 * S * X_p * X_r + X_p^2 + \\ S^2 * Y_r^2 + 2 * S * Y_p * Y_r + Y_p^2 + \\ S^2 * Z_r^2 + 2 * S * Z_p * Z_r + Z_p^2 - R^2 = 0 \end{aligned}$$

This is a quadratic of the form  $A * S^2 + B * S + C = 0$

$$\begin{aligned} A &= X_r^2 + Y_r^2 + Z_r^2 \\ B &= 2 * (X_p * X_r + Y_p * Y_r + Z_p * Z_r) \\ C &= X_p^2 + Y_p^2 + Z_p^2 - R^2 \end{aligned}$$

Therefore, there are solutions

$$S = \frac{-B \pm \sqrt{B^2 - 4 * A * C}}{2 * A}$$

Several simplifications can be made.

Since the ray is normalized with length one:

$$A = X_r^2 + Y_r^2 + Z_r^2 = 1$$

Let  $B = 2 * D$  with

$$D = X_p * X_r + Y_p * Y_r + Z_p * Z_r$$

This produces

$$S = \frac{-2 * D \pm \sqrt{4 * D^2 - 4 * C}}{2}$$

Which simplifies to

$$S = -D \pm \sqrt{D^2 - C}$$

This form of the equation is solved in the `intsph` routine. Since there are two solutions; the smallest one greater than zero is returned. If no solution is found, zero is returned.

`nrmsph` finds the surface normal at a point on the sphere's surface. If the point of intersection is P, the sphere center C, and its radius R, then the normal vector is

$$\begin{aligned} N_x &= (P_x - C_x) / R \\ N_y &= (P_y - C_y) / R \\ N_z &= (P_z - C_z) / R \end{aligned}$$

```

/** sphere.c */
#include <math.h>
#include "constants.h"
#include "typedefs.h"
#include "globalvar.h"

/* create sphere object */

```

```

int maksph(surf,r,x,y,z)
    int surf;
    double r,x,y,z;
{
    int size;
    o_sphere *sphere;

    size = sizeof(o_sphere);
    sphere = ((o_sphere *) malloc(size));
    object[nobject].id = nobject;
    object[nobject].objtyp = OTYPSPHERE;
    object[nobject].surfnm = surf;
    object[nobject].objpnt.p_sphere = sphere;
    sphere->r = r;
    sphere->x = x;
    sphere->y = y;
    sphere->z = z;
}

double intsph(pos,ray,obj)
    t_3d *pos; /* origin of ray */
    t_3d *ray; /* ray vector */
    t_object *obj; /* sphere description */
{
    double b,t,s;
    double xadj,yadj,zadj;
    o_sphere *sph;

    sph = obj->objpnt.p_sphere;
    /* translate ray origin to object's space */
    xadj = pos->x - sph->x;
    yadj = pos->y - sph->y;
    zadj = pos->z - sph->z;

    /* solve quadratic equation */
    b = xadj*ray->x + yadj*ray->y + zadj*ray->z;
    t = b*b - xadj*xadj - yadj*yadj - zadj*zadj +
        sph->r*sph->r;
    if (t < 0) return(0.0);
    s = -b - sqrt(t); /* try smaller solution */
    if (s > 0) return(s);
    s = -b + sqrt(t); /* try larger solution */
    if (s > 0) return(s);
    return(0); /* both solutions <= zero */
}

int nrmsph(pos,obj,nrm)
    t_3d *pos; /* point of intersection */
    t_object *obj; /* sphere description */
    t_3d *nrm; /* return surface normal */
{
    o_sphere *sph;

    sph = obj->objpnt.p_sphere;
    nrm->x = (pos->x-sph->x)/sph->r;

```

```

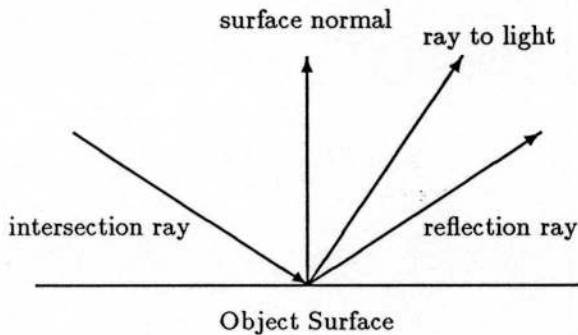
nrm->y = (pos->y-sph->y)/sph->r;
nrm->z = (pos->z-sph->z)/sph->r;
}

```

## 2.8 SHADING ROUTINE

The shading routine calculates the color of a point on an object's surface. It models ambient light as well as diffuse and specular lighting from multiple point light sources. White light is assumed for all light sources.

The point of intersection, surface normal at the point, intersection ray, distance to intersection, and object pointer are passed to the shading routine.



Using the rays shown above and the object's surface properties, the color of the point being shaded is determined.

The reflection ray from the surface is calculated from the intersection ray and the surface normal. The angle between the reflection ray and the surface normal is equal to the angle between the surface normal and the intersection ray.

The initial color for the point being shaded is the ambient color.

For each light, the diffuse color and the specular highlight color are computed and added to the color of the point being shaded.

- The ray to the light source is calculated by the `lightray` routine. The brightness of the light is found by a call to the `brightness` routine.
- The strength of the diffuse lighting is given by the dot product of the surface normal and the ray to the light. If it is less than zero, the surface is facing away from the light. Thus, the light does not contribute to the color of the surface. If it is greater than zero, the surface faces the light. The dot product is multiplied by both the brightness of the light and the diffuse surface color.
- The specular highlight is computed as the dot product of the reflection ray and the ray to the light. The dot

product, if greater than zero, is raised to the power of `k`, the surface specular coefficient. This result is multiplied by the specular surface color.

At each step, lighting components are added to the color of the surface. The surface becomes brighter as each succeeding term is processed. This matches what you would expect. An object lit only by ambient light is fairly dark and evenly shaded. If the surface is lit from a light source, it is brighter and shows more shape. If it is shiny, bright specular highlights appear.

A more complicated shading model, which includes reflection, transparency and shadows is presented in sections 3.1 and 3.2.

```

/** shade.c */
#include <math.h>
#include "typedefs.h"
#include "funcdefs.h"
#include "globalvar.h"

int shade(pos,ray,nrm,obj,color)
    t_3d *pos,*ray,*nrm;
    t_object *obj;
    t_color *color;
{
    int lnum;
    double k,bright,spec,diffuse;
    t_surface *surf;
    t_3d refl,ltray;

    /* Ambient light contribution */
    surf = &surface[obj->surfnum];
    color->r = surf->ar;
    color->g = surf->ag;
    color->b = surf->ab;

    /* calculate reflected ray */
    k = -2.0 * dotp(ray, nrm);
    refl.x = k*nrm->x + ray->x;
    refl.y = k*nrm->y + ray->y;
    refl.z = k*nrm->z + ray->z;

    for (lnum=0; lnum < nlight; lnum++)
    {
        /* get ray to light */
        lightray(lnum,pos,&ltray);
        diffuse = dotp(nrm, &ltray);
        if (diffuse > 0)
        {
            /* object faces light, add diffuse */
            bright = brightness(obj->id,lnum,pos,&ltray);
            diffuse *= bright;
            color->r += surf->dr * diffuse;

```



```

write(outfile, header, 2*sizeof(int));
}

int linepic(pixels)
double pixels[SCREENWIDTH][3];
{
    unsigned char buffer[SCREENWIDTH][3];
    int i,r,g,b;
    double dr,dg,db;

    for (i=0; i<width; i++)
    {
        r = gammacorrect(pixels[i][0]);
        g = gammacorrect(pixels[i][1]);
        b = gammacorrect(pixels[i][2]);
        buffer[i][0] = r;
        buffer[i][1] = g;
        buffer[i][2] = b;
    }
    write(outfile, buffer, linesize);
}

int endpic()
{
    close(outfile);
}

int gammacorrect(intensity)
double intensity;
{
    int ival;
    double dval;

    /* scale to 0 - 1 range */
    dval = intensity/255.0;
    if (dval > 1.0) dval = 1.0;
    if (dval < 0.001) dval = 0.001;
    /* do gamma correction */
    dval = exp( log(dval) / GAMMA );
    /* convert to integer, range 0-255 */
    dval *= 255.0;
    ival = (int) (dval + 0.5);
    return(ival);
}

```

## 2.11 RAY FUNCTIONS

Several functions operating on rays have been used throughout the basic system.

The **normalize** routine scales a ray to have length one. The ray's original length is also returned. The **dotp** routine returns the dot product of two rays. The **crossp** routine returns the cross product of the two rays.

```

/** raymath.c */
#include <math.h>
#include "typedefs.h"

double normalize(a)
    t_3d *a;
{
    double d;

    d = sqrt(a->x * a->x + a->y * a->y + a->z * a->z);
    a->x /= d;
    a->y /= d;
    a->z /= d;
    return(d);
}

double dotp(a, b)
    t_3d *a, *b;
{
    double d;

    d = (a->x * b->x) + (a->y * b->y) + (a->z * b->z);
    return(d);
}

int crossp(o, a, b)
    t_3d *o, *a, *b;
{
    double d;

    o->x = (a->y * b->z) - (a->z * b->y);
    o->y = (a->z * b->x) - (a->x * b->z);
    o->z = (a->x * b->y) - (a->y * b->x);
    d = sqrt(o->x*o->x + o->y*o->y + o->z*o->z);
    o->x /= d;
    o->y /= d;
    o->z /= d;
}

```

## 3 SYSTEM DEVELOPMENT

By now you should have produced several ray traced images. It's time to extend the capabilities of the first program into a full scale system.

The shading routine was adequate; now it is extended to produce shadows, reflections, and transparency. Shadow generation can be added by itself. It is easiest to add reflections and transparency together because the **shade** routine is modified for each of them.

Spheres are ok for test images. For real images, many different primitives are useful. Each new primitive can be added by itself. For each primitive, four things are needed:

a 'make' routine, an intersection routine, a normal routine, and a structure definition. Routines and structures to handle boxes, triangles, and superquadrics are presented.

### 3.1 SHADOWS

Shadows are an extension to the basic shading routine. They give very strong depth cues and add to the realism of images. In discussing shadows, it is important to differentiate between two terms. Shading is the process of finding the color at a point on an object's surface. Shadowing refers to blocking the light which would have fallen on the object's surface.

Two routines are used to determine whether the point being shaded is in shadow. `lightray` calculates a shadow ray between the point and the light source, saving the distance between the two.

`brightness` returns zero if the point is in shadow. Otherwise, it returns the light's brightness value. `brightness` is very similar to `intersect`, calculating intersections between the shadow ray and all objects. It stops as soon as it finds one object between the point being shaded and the light source.

```

/** light2.c */
#include "typedefs.h"
#include "funcdefs.h"
#include "globalvar.h"

double s_litdis;

int lightray(lnum,objpos,lray)
    int lnum;
    t_3d *objpos, *lray;
{
    lray->x = light[lnum].x - objpos->x;
    lray->y = light[lnum].y - objpos->y;
    lray->z = light[lnum].z - objpos->z;
    s_litdis = normalize(lray);
}

double brightness(source,lnum,pos,ray)
    int source,lnum;
    t_3d *pos,*ray;
{
    int objtry;
    double s;

    for (objtry=0; objtry<nobject; objtry++)
    {
        if (objtry != source) /* don't try source */
        {
            s = (*objint[object[objtry].objtyp])
                (pos,ray,&object[objtry]);
            if ((s > 0.0) && (s <= s_litdis))

```


```

        return(0); /* object in shadow */
    }
}
/* object not in shadow */
return(light[lnum].bright);
}

```

### 3.2 REFLECTION & TRANSPARENCY

Reflection and transparency are extensions to the shading routine. Each adds a new component to the color of the point being shaded.



A reflection ray is sent from the point being shaded by a call to `intersect`. `intersect` finds the closest intersection with another object. Then it calls `shade` to determine the color at the point of intersection and returns the color. This color is combined with the color already calculated for the point being shaded.

$$Color_p = Color_p + R * Color_r$$

where

$Color_p$  is the color already calculated for the point.

$R$  is the reflectivity of the surface being shaded, which ranges from 0.0 for nonreflecting to 1.0 for a perfect mirror.

$Color_r$  is the color of the object hit by the reflection ray. The background color is used if the reflection ray didn't hit any object.

Adding transparency is similar to adding reflection. The transparency ray is sent from the point being shaded by a call to `intersect`. The color returned is mixed with the color already calculated for the point being shaded.

$$Color_p = (1 - T) * Color_p + T * Color_t$$

where

$Color_p$  is the color already calculated for the point, including reflection.

$T$  is the transparency of the surface being shaded, which ranges from 0.0 for opaque to 1.0 for completely transparent.

$Color_t$  is the color of the object hit by the transparency ray. The background color is used if the transparency ray didn't

hit any object.

It is important to understand that shading is now defined recursively. `shade` calls `intersect` which calls `shade` ... A count is used to limit the depth of the recursive calls.

The net result is that reflections can be seen in reflections.

```
/** shade2.c **/  
#include <math.h>  
#include "typedefs.h"  
#include "funcdefs.h"  
#include "globalvar.h"  
  
int shade(pos,ray,nrm,obj,color)  
  t_3d *pos,*ray,*nrm;  
  t_object *obj;  
  t_color *color;  
{  
  int lnum;  
  double k,dis,bright,spec,diffuse;  
  t_surface *surf;  
  t_3d refl,ltray;  
  t_color newcol;  
  
  /* calculate reflected ray */  
  k = -2.0 * dotp(ray, nrm);  
  refl.x = k*nrm->x + ray->x;  
  refl.y = k*nrm->y + ray->y;  
  refl.z = k*nrm->z + ray->z;  
  
  /* Ambient light contribution */  
  surf = &surface[obj->surfnum];  
  color->r = surf->ar;  
  color->g = surf->ag;  
  color->b = surf->ab;  
  
  for (lnum=0; lnum < nlight; lnum++)  
  {  
    /* get ray to light */  
    lightray(lnum,pos,&ltray);  
    diffuse = dotp(nrm, &ltray);  
    if (diffuse > 0)  
    {  
      /* object faces light, add diffuse*/  
      bright = brightness(obj->id,lnum,pos,&ltray);  
      diffuse *= bright;  
      color->r += surf->dr * diffuse;  
      color->g += surf->dg * diffuse;  
      color->b += surf->db * diffuse;  
  
      spec = dotp(&refl, &ltray);  
      if (spec > 0)  
      {  
        /* highlight is here, add specular */  
        spec = bright * pow(spec,surf->coef);  
        color->r += surf->sr * spec;
```

```
        color->g += surf->sg * spec;  
        color->b += surf->sb * spec;  
      }  
    }  
  }  
  
  /* reflection */  
  k = surf->refl;  
  if ((k > 0) && (level < maxlevel))  
  {  
    level++;  
    dis = intersect(obj->id, pos, &refl, &newcol);  
    if (dis > 0)  
    {  
      color->r += newcol.r * k;  
      color->g += newcol.g * k;  
      color->b += newcol.b * k;  
    }  
    else  
    {  
      color->r += background.r * k;  
      color->g += background.g * k;  
      color->b += background.b * k;  
    }  
    level--;  
  }  
  
  /* transparency */  
  k = surf->transp;  
  if (k > 0)  
  {  
    color->r *= (1-k);  
    color->g *= (1-k);  
    color->b *= (1-k);  
    dis = intersect(obj->id, pos, ray, &newcol);  
    if (dis > 0)  
    {  
      color->r += newcol.r * k;  
      color->g += newcol.g * k;  
      color->b += newcol.b * k;  
    }  
    else  
    {  
      color->r += background.r * k;  
      color->g += background.g * k;  
      color->b += background.b * k;  
    }  
  }  
}
```

### 3.3 BOXES

The `makbox` routine allocates the space for a box description and fills it with the values which were passed to the

routine.

`intbox` computes the intersection between a ray and a box. The source position of the ray, the ray vector, and the object, a box, are passed to the routine. The distance to the nearest intersection is returned. A value of zero is returned if no intersection is found.

The ray can hit any of the six sides of the box. Therefore, the six possible intersections must be checked.

Each intersection calculation has two parts. First, the intersection of the ray with the plane of the box side is found. Then the point of intersection with the plane is checked to determine whether it lies within the box side.

Let  $X_p, Y_p, Z_p$  be the ray origin,  $X_r, Y_r, Z_r$  the ray,  $X_c, Y_c, Z_c$  the box center, and  $X_s, Y_s, Z_s$  the sizes of the box sides. The intersection calculation for the plus X side of the box is

$$S = ((X_c + X_s) - X_p) / X_r$$

This gives the point of intersection with the plane

$$X = X_c + X_s$$

Then, the y and z coordinates of the point are found.

$$Y = Y_p + s * Y_r$$

$$Z = Z_p + s * Z_r$$

If Y is within  $Y_s$  of  $Y_c$  and Z is within  $Z_s$  of  $Z_c$ , the point is on the box.

This calculation is performed for all six sides. The distance to the closest valid intersection is returned. The code is fairly long, but simple.

The normal is found by keeping track of which side of the box is intersected. The normal is 1.0 in the direction of that side of the box.

```

/** box.c */
#include <math.h>
#include "constants.h"
#include "typedefs.h"
#include "globalvar.h"
#define FAR_AWAY 99.99E+20

/* create box object */

int makbox(surf,x,y,z,xs,ys,zs)
    int surf;
    double x,y,z;
    double xs,ys,zs;
{
    int size;
    o_box *box;

    size = sizeof(o_box);
    box = ((o_box *) malloc(size));
    object[nobject].id = nobject;

```

```

    object[nobject].objtyp = OTYPBOX;
    object[nobject].surfnum = surf;
    object[nobject].objpnt.p_box = box;
    box->x = x;
    box->y = y;
    box->z = z;
    box->xs = xs;
    box->ys = ys;
    box->zs = zs;
}

/* intersection calculation for ray and box */

double intbox(pos,ray,obj)
    t_3d *pos; /* origin of ray */
    t_3d *ray; /* ray vector */
    t_object *obj; /* box description */
{
    double s,ss,xhit,yhit,zhit;
    double xadj,yadj,zadj;
    o_box *box;

    box = obj->objpnt.p_box;
    ss = FAR_AWAY;
    /* translate ray origin to objects space */
    xadj = pos->x - box->x;
    yadj = pos->y - box->y;
    zadj = pos->z - box->z;

    if (ray->x != 0) /* check x faces */
    {
        s = (box->xs-xadj)/ray->x;
        if ((s > 0) && (s < ss))
        {
            yhit = fabs(yadj + s * ray->y);
            zhit = fabs(zadj + s * ray->z);
            if ((yhit < box->ys) && (zhit < box->zs))
                { box->sidehit = 0; ss = s;}
        }
        s = (-box->xs-xadj)/ray->x;
        if ((s > 0) && (s < ss))
        {
            yhit = fabs(yadj + s * ray->y);
            zhit = fabs(zadj + s * ray->z);
            if ((yhit < box->ys) && (zhit < box->zs))
                { box->sidehit = 1; ss = s;}
        }
    }

    if (ray->y != 0) /* check y faces */
    {
        s = (box->ys-yadj)/ray->y;
        if ((s > 0) && (s < ss))
        {
            xhit = fabs(xadj + s * ray->x);

```

```

    zhit = fabs(zadj + s * ray->z);
    if ((xhit < box->xs) && (zhit < box->zs))
        { box->sidehit = 2; ss = s;}
}
s = (-box->ys-yadj)/ray->y;
if ((s > 0) && (s < ss))
{
    xhit = fabs(xadj + s * ray->x);
    zhit = fabs(zadj + s * ray->z);
    if ((xhit < box->xs) && (zhit < box->zs))
        { box->sidehit = 3; ss = s;}
}
}

if (ray->z != 0) /* check z faces */
{
    s = (box->zs-zadj)/ray->z;
    if ((s > 0) && (s < ss))
    {
        xhit = fabs(xadj + s * ray->x);
        yhit = fabs(yadj + s * ray->y);
        if ((xhit < box->xs) && (yhit < box->ys))
            { box->sidehit = 4; ss = s;}
    }
    s = (-box->zs-zadj)/ray->z;
    if ((s > 0) && (s < ss))
    {
        xhit = fabs(xadj + s * ray->x);
        yhit = fabs(yadj + s * ray->y);
        if ((xhit < box->xs) && (yhit < box->ys))
            { box->sidehit = 5; ss = s;}
    }
}

if (ss == FAR_AWAY) return(0.0);
return(ss);
}

```

/\* normal calculation for box \*/

```

int nrmbx(pos,obj,nrm)
    t_3d *pos; /* point of intersection */
    t_object *obj; /* box description */
    t_3d *nrm; /* return surface normal */
{
    o_box *box;

    box = obj->objpnt.p_box;
    nrm->x = 0.0;
    nrm->y = 0.0;
    nrm->z = 0.0;
    switch(box->sidehit)
    {
        case(0): nrm->x = 1.0;
                break;

```

```

        case(1): nrm->x = -1.0;
                break;
        case(2): nrm->y = 1.0;
                break;
        case(3): nrm->y = -1.0;
                break;
        case(4): nrm->z = 1.0;
                break;
        case(5): nrm->z = -1.0;
                break;
    }
    return;
}

```

### 3.4 TRIANGLES

The **maktri** routine allocates the space for a triangle structure and fills the fields. The three corner points of the triangle are passed into **maktri**. In order to make the ray-triangle intersection calculation fast, some values are precomputed. A plane is defined by a normal vector and a plane constant. The normal vector and plane constant for the triangle are computed and stored. Each edge of the triangle can be defined by a plane which is perpendicular to the plane of the triangle. The vectors and constants for the triangle's three edges are computed and stored.

The normal vector for the triangle is the cross product of the vectors from the first point to the second and the second point to the third. The plane constant is the dot product of the normal vector and any one of the triangle's points.

The normal for an edge is defined by the cross product of the vector from its start point to its end point and the triangle's normal vector. The plane constant for an edge is the dot product of the edge vector and one of the points in the edge.

**inttri** computes the intersection between a ray and a triangle in two steps. First, it finds the intersection of the ray with the plane of the triangle. Next, it determines whether the point of intersection is inside the triangle. Plane equations are used for each step. Note that the points of the triangle are either in clockwise or counterclockwise order. You must be consistent, because the direction of the triangle's normal is determined by the order of the points. The convention used is that as you look at a triangle, its normal points towards you, and its points are in counterclockwise order.

```

/** triangle.c */
#include "constants.h"
#include "typedefs.h"
#include "globalvar.h"
#include "funcdefs.h"

```



```

int maktri(surf,p1,p2,p3)
    int surf;
    t_3d *p1,*p2,*p3;
{
    int size;
    o_triangle *triangle;
    t_3d vc1,vc2,vc3;

    size = sizeof(o_triangle);
    triangle = ((o_triangle *) malloc(size));
    object[nobject].id = nobject;
    object[nobject].objtyp = OTYPTRIANGLE;
    object[nobject].surfnum = surf;
    object[nobject].objpnt.p_triangle = triangle;
    vc1.x = p2->x - p1->x;      /* edge vectors */
    vc1.y = p2->y - p1->y;
    vc1.z = p2->z - p1->z;
    vc2.x = p3->x - p2->x;
    vc2.y = p3->y - p2->y;
    vc2.z = p3->z - p2->z;
    vc3.x = p1->x - p3->x;
    vc3.y = p1->y - p3->y;
    vc3.z = p1->z - p3->z;
    /* plane of triangle */
    crossp(&triangle->nrm,&vc1,&vc2);
    triangle->d = dotp(&triangle->nrm,p1);
    /* edge planes */
    crossp(&triangle->e1,&triangle->nrm,&vc1);
    triangle->d1 = dotp(&triangle->e1,p1);
    crossp(&triangle->e2,&triangle->nrm,&vc2);
    triangle->d2 = dotp(&triangle->e2,p2);
    crossp(&triangle->e3,&triangle->nrm,&vc3);
    triangle->d3 = dotp(&triangle->e3,p3);
}

double intri(pos,ray,obj)
    t_3d *pos,*ray;
    t_object *obj;
{
    double s,k;
    t_3d point;
    o_triangle *triangle;

    triangle = obj->objpnt.p_triangle;
    /* plane intersection */
    k = dotp(&triangle->nrm,ray);
    if (k == 0) return(0);
    s = (triangle->d - dotp(&triangle->nrm,pos)) / k;
    if (s <= 0) return(0.0);

    point.x = pos->x + ray->x * s;
    point.y = pos->y + ray->y * s;
    point.z = pos->z + ray->z * s;
}

```

```

/* edge checks */
k = dotp(&triangle->e1, &point) - triangle->d1;
if (k < 0) return(0);
k = dotp(&triangle->e2, &point) - triangle->d2;
if (k < 0) return(0);
k = dotp(&triangle->e3, &point) - triangle->d3;
if (k < 0) return(0);

return(s);
}

int nrmtri(pos,obj,nrm)
    t_3d *pos;
    t_object *obj;
    t_3d *nrm;
{
    o_triangle *triangle;

    triangle = obj->objpnt.p_triangle;
    nrm->x = triangle->nrm.x;
    nrm->y = triangle->nrm.y;
    nrm->z = triangle->nrm.z;
}

```

### 3.5 SUPERQUADRICS

Superquadrics can be described as boxes with rounded corners. They are specified by equations of the form

$$a * |x|^n + b * |y|^n + c * |z|^n = r^n$$

The **maksup** routine allocates the space for a superquadric description and fills it with the values which were passed to the routine and values which it calculates.

The superquadric structure has all of the fields in the box structure, along with some new fields. The values which correspond to the box structure fields are passed into the routine. The new fields are the A, B, C, R, and N values in the superquadric equation, along with an error metric. The N value is passed into the routine. The other values are derived from the box values. R is the size of the largest box side. A, B, and C are found by dividing the box sizes by R. The error metric is used in the iterative solution of the ray-object intersection. The solution is accurate enough when the measured error is less than this value.

The intersection of a ray with a superquadric cannot be calculated directly because the solution of an order N polynomial is required. Therefore, an iterative technique is needed. The secant method is used because it is very simple.

The superquadric surface lies entirely within the box used to define it. Any ray which will strike the superquadric will first strike the box. So, the ray-box intersection is used as the initial value for the ray-superquadric intersection. This

also gives an inexpensive rejection criteria for most of the rays which would not hit the superquadric.

There is one special case. If the ray starts within the surrounding box, it can hit the superquadric without hitting the box.

```

/**/ superquadric.c /**/
#include <math.h>
#include "constants.h"
#include "typedefs.h"
#include "funcdefs.h"
#include "globalvar.h"
#define ERRCONST 1.05

int maksup(surf,x,y,z,xs,ys,zs,power)
    int surf;
    double x,y,z;
    double xs,ys,zs;
    double power;
{
    int size;
    double max;
    o_superq *super;

    size = sizeof(o_superq);
    super = ((o_superq *) malloc(size));
    object[nobject].id = nobject;
    object[nobject].objtyp = OTYPSUPERQ;
    object[nobject].surfnum = surf;
    object[nobject].objpnt.p_superq = super;
    super->x = x;
    super->y = y;
    super->z = z;
    super->xs = xs;
    super->ys = ys;
    super->zs = zs;
    super->pow = power;
    max = xs;
    if (ys > max) max = ys;
    if (zs > max) max = zs;
    super->a = xs/max;
    super->b = ys/max;
    super->c = zs/max;
    super->r = pow(max,power);
    super->err = pow((ERRCONST * max),power) - super->r;
}

double intsup(pos,ray,obj)
    t_3d *pos,*ray;
    t_object *obj;
{
    double xsiz,ysiz,zsiz;
    double s,s1,t;
    double xadj,yadj,zadj;
    double old,result,p;

```

```

o_superq *super;

/* find box intersection */
s = intbox(pos,ray,obj);
if (s == 0) return(0);
super = obj->objpnt.p_superq;
xadj = pos->x - super->x;
yadj = pos->y - super->y;
zadj = pos->z - super->z;

/* special case - ray origin within box */
if ((fabs(xadj) < super->xs) &&
    (fabs(yadj) < super->ys) &&
    (fabs(zadj) < super->zs)) s = 0;

/* initial solution */
p = super->pow;
result = pow(fabs((xadj + ray->x*s) / super->a), p)
    + pow(fabs((yadj + ray->y*s) / super->b), p)
    + pow(fabs((zadj + ray->z*s) / super->c), p)
    - super->r;

if (result < super->err) return(s);

s1 = s;
s = s + 0.001;
/* iterative refinement */
while (result > super->err)
{
    old = result;
    result = pow(fabs((xadj + ray->x*s) / super->a), p)
        + pow(fabs((yadj + ray->y*s) / super->b), p)
        + pow(fabs((zadj + ray->z*s) / super->c), p)
        - super->r;

    if (result >= old) return(0.0);
    t = (result*(s-s1))/(result-old);
    s1 = s;
    s -= t;
}

return(s);
}

int nrmsup(pos,obj,nrm)
    t_3d *pos,*nrm;
    t_object *obj;
{
    double k;
    o_superq *super;

    super = obj->objpnt.p_superq;
    nrm->x = (pos->x-super->x) / super->a;
    nrm->y = (pos->y-super->y) / super->b;
    nrm->z = (pos->z-super->z) / super->c;

```

```

k = super->pow - 1;
if (nrm->x > 0) nrm->x = pow(nrm->x,k);
else nrm->x = -pow(-nrm->x,k);
if (nrm->y > 0) nrm->y = pow(nrm->y,k);
else nrm->y = -pow(-nrm->y,k);
if (nrm->z > 0) nrm->z = pow(nrm->z,k);
else nrm->z = -pow(-nrm->z,k);
normalize(nrm);
}

```

### 3.6 INPUT PROCESSING

Writing routines to parse and process user input is a large task. In a ray tracing system, many things should be under user control. It should be easy to specify viewing parameters. Large files of object descriptions should be easy to process. Files containing surface type descriptions should be produced once and used for many different images.

The many and varied requirements of the user interface make it quite complicated. Additionally, it is very nice to be able to change the interface easily. One of the cleanest solutions is to use parsing tools, in particular, the UNIX tools LEX and YACC. They are simple to use and because they generate C code, are easily added to the ray tracer.

LEX handles character by character input processing. It recognizes key words and can read numbers in just about any form. It produces tokens which are used by YACC.

```

/**/ input_lex.l /**/
%{
#include <stdio.h>
#include "y.tab.h"
%}
alpha [a-zA-Z]
special [\.\\_]
digit [0-9]
exp [Ee][+]?{digit}+
string {alpha}({alpha}|{digit}|{special})*
%start COMMENT
%%
" " ;
\t ;
\n ;
"/*" {BEGIN COMMENT;}
<COMMENT>[^\n]*\n ;
<COMMENT>[^\n]+"/" {if (yytext[yytext-2] == '*')
    BEGIN 0;}
eyep {return(TEYEP);}
lookp {return(TLOOKP);}
up {return(TUP);}
fov {return(TFOV);}
screen {return(TSCREEN);}
light {return(TLIGHT);}

```

```

surface {return(TSURFACE);}
background {return(TBACKGROUND);}
maxlevel {return(TMAXLEVEL);}
sphere {return(TSPHERE);}
box {return(TBOX);}
triangle {return(TTRIANGLE);}
superq {return(TSUPERQ);}
outfile {return(TOUTFILE);}
{string} {yyval.c=yytext;
    return(TSTRING);}
[+]?{digit}+ {sscanf(yytext,"%d",&yyval.i);
    return(TINT);}

[+]?{digit}+."{digit}*({exp})? |
[+]?{digit}*"."{digit}*({exp})? |
[+]?{digit}+{exp} {sscanf(yytext,"%F",&yyval.d);
    return(TFLOAT);}

%%
yywrap() {return(1);}

YACC recognizes patterns of tokens and performs user
specified actions based on them.

For example, LEX recognizes the letters 'e' 'y' 'e' 'p' as the
string "eyep". It recognizes the characters '1' '0' ' ' '0' ' ' '0'
as the numbers 10 0 0. For each string of characters, it sends
a token to YACC. In this case it sends TEYEP TINT TINT
TINT, with each TINT having a numerical value associated
with it. YACC recognizes that this sequence of tokens spec-
ifies an eyepoint position. User written C code actions then
process the eyepoint position.

/**/ input_yacc.y /**/
%{
#include <stdio.h>
#include "typedefs.h"
#include "globalvar.h"
int number;
t_3d p1,p2,p3;
%}
%union {
    char *c;
    int i;
    double d;
}
%token <i> TINT
%token <d> TFLOAT
%token <c> TSTRING
%token TEYEP TLOOKP TUP TFOV TSCREEN TMAXLEVEL
%token TLIGHT TSURFACE TSPHERE TBOX TTRIANGLE TSUPERQ
%token TBACKGROUND TOUTFILE
%type <d> Fnumber
%type <c> String
%%

```

```

File      : File Item
          | Item
          ;
Item      : Eyep
          | Lookp
          | Up
          | Fov
          | Screen
          | Maxlevel
          | Background
          | Light
          | Surface
          | Sphere
          | Box
          | Triangle
          | Superq
          | Outfile
          ;
Eyep      : TEYEP Fnumber Fnumber Fnumber
          { eyep.x = $2;
            eyep.y = $3;
            eyep.z = $4;
          }
          ;
Lookp     : TLOOKP Fnumber Fnumber Fnumber
          { lookp.x = $2;
            lookp.y = $3;
            lookp.z = $4;
          }
          ;
Up        : TUP Fnumber Fnumber Fnumber
          { up.x = $2;
            up.y = $3;
            up.z = $4;
          }
          ;
Fov       : TFOV Fnumber Fnumber
          { hfov = $2; vfov = $3; }
          ;
Screen    : TSCREEN TINT TINT
          { sizey = $2; sizex = $3; }
          ;
Maxlevel  : TMAXLEVEL TINT
          { maxlevel = $2; }
          ;
Background : TBACKGROUND Fnumber Fnumber Fnumber
          { background.r = $2;
            background.g = $3;
            background.b = $4;
          }
          ;
Light     : TLIGHT Fnumber
          Fnumber Fnumber Fnumber
          { if (nlight == lightlim)
            yyerror("too many lights");
          }

```

```

light[nlight].bright = $2;
light[nlight].x = $3;
light[nlight].y = $4;
light[nlight].z = $5;
nlight++;
}
;
Surface   : TSURFACE TINT
          Fnumber Fnumber Fnumber
          Fnumber Fnumber Fnumber
          Fnumber Fnumber Fnumber
          Fnumber Fnumber Fnumber
          { number = $2;
            if (number >= surfacelim)
              yyerror("surface # too big");
            surface[number].ar = $3;
            surface[number].ag = $4;
            surface[number].ab = $5;
            surface[number].dr = $6;
            surface[number].dg = $7;
            surface[number].db = $8;
            surface[number].sr = $9;
            surface[number].sg = $10;
            surface[number].sb = $11;
            surface[number].coef = $12;
            surface[number].refl = $13;
            surface[number].transp = $14;
            nsurface++;
          }
          ;
Sphere    : TSPHERE TINT Fnumber
          Fnumber Fnumber Fnumber
          { if (nobject == objectlim)
            yyerror("too many objects");
            maksph($2,$3,$4,$5,$6);
            nobject++;
          }
          ;
Box       : TBOX TINT
          Fnumber Fnumber Fnumber
          Fnumber Fnumber Fnumber
          { if (nobject == objectlim)
            yyerror("too many objects");
            makbox($2,$3,$4,$5,$6,$7,$8);
            nobject++;
          }
          ;
Triangle  : TTRIANGLE TINT
          Fnumber Fnumber Fnumber
          Fnumber Fnumber Fnumber
          Fnumber Fnumber Fnumber
          { if (nobject == objectlim)
            yyerror("too many objects");
            p1.x = $3; p1.y = $4; p1.z = $5;
            p2.x = $6; p2.y = $7; p2.z = $8;
          }

```

```

        p3.x = $9; p3.y = $10; p3.z = $11;
        maktri($2,&p1,&p2,&p3);
        nobject++;
    }
    ;
Superq      : TSUPERQ TINT
            Fnumber Fnumber Fnumber
            Fnumber Fnumber Fnumber
            Fnumber
{ if (nobject == objectlim)
    yyerror("too many objects");
  maksup($2,$3,$4,$5,$6,$7,$8,$9);
  nobject++;
}
;
Outfile    : TOUTFILE String
            { strcpy(outfilename, $2); }
;
Fnumber    : TFLOAT
            { $$=$1; }
            | TINT
            { $$=$1; }
;
String     : TSTRING
            { $$ = $1; }
;
%
yyerror(s)
  char *s;
  {
    fprintf(stderr,"%s\n",s);
  }

```

### 3.7 SAMPLE INPUT FILES

An image description consists of a series of keywords and parameters read from standard input. The keywords can occur in any order with spaces and C-like comments interspersed.

`basic.pds` contains the image description to create the default picture produced by the basic system (figure 1). The command `tracer2.go < basic.pds` runs the full system, to produce this image.

```

/*** basic.pds ***/
screen 512 512
fov 50 50
eyep 100 0 0
lookp 0 0 0
up 0 1 0
maxlevel 0
outfile basic.pix
light /* brightness */ 1.0
      /* position */ 100.0 300.0 0.0
sphere /* color */ 0

```

```

/* radius */ 40
/* center */ 5 30 40
sphere /* color */ 1
      /* radius */ 30
      /* center */ -10 -10 -40
surface 0
  30 0 0 /* ambient */
  90 0 0 /* diffuse */
  190 120 120 /* specular */
  15 0.0 0.0 /* coef,transp,refl */
surface 1
  0 50 0 /* ambient */
  0 100 0 /* diffuse */
  30 40 30 /* specular */
  2 0.0 0.5 /* coef,transp,refl */
background /* color */ 100 200 250

```

`all.pds` contains the image description for figure 2. Each of the four object types as well as reflection and transparency is shown in this image.

```

/*** all.pds ***/
screen 512 512
fov 20 20
eyep 100 150 90
lookp 0 0 0
up 0 1 0
maxlevel 1
outfile all.pix
light /* brightness */ 1.0
      /* position */ 100.0 300.0 0.0
superq /* color */ 0
      /* center */ 20 10 10
      /* size */ 10 20 5
      /* power */ 5
sphere /* color */ 1
      /* radius */ 10
      /* center */ -10 25 -20
triangle /* color */ 2
      /* point */ -30 30 0
      /* point */ -34 15 19
      /* point */ -34 -12 -9
box /* color */ 3
    /* center */ -17 -10 7
    /* size */ 12 5 25
surface 0
  0 50 0 /* ambient */
  0 100 0 /* diffuse */
  30 60 0 /* specular */
  6 0.0 0.2 /* coef,transp,refl */
surface 1
  90 90 0 /* ambient */
  230 230 0 /* diffuse */
  110 90 30 /* specular */
  2 0.0 0.0 /* coef,transp,refl */

```

```

surface 2
  0 90 90 /* ambient */
  0 230 230 /* diffuse */
  30 110 90 /* specular */
  8 0.0 0.0 /* coef,transp,refl */
surface 3
  60 0 90 /* ambient */
  210 0 230 /* diffuse */
  90 30 110 /* specular */
  15 0.2 0.0 /* coef,transp,refl */
background /* color */ 0 0 250

```

### 3.8 UNIX MAKEFILE

Unix provides a very convenient tool for building large programming systems: the makefile. It specifies the routines needed to build the system, dependencies between the routines, and the process for building each routine.

If you are not on a Unix system, you should be able to build command files to compile and link the basic system and the full system.

```

#::: makefile :::
# C compile command for all routines #
.c.o::cc -O -c $*.c

# build basic system #
basic: main.o setup.o viewing.o intersect.o shade.o \
  light.o raymath.o stubs.o sphere.o outputp.o
cc -O -o tracer1.go \
  main.o setup.o viewing.o intersect.o shade.o \
  light.o raymath.o stubs.o sphere.o outputp.o -lm

# build full system #
system: main.o setup.o viewing.o intersect.o shade2.o \
  light2.o raymath.o input_lex.o input_yacc.o \
  sphere.o box.o triangle.o superquadric.o outputp.o
cc -O -o tracer2.go \
  main.o setup.o viewing.o intersect.o shade2.o \
  light2.o raymath.o input_lex.o input_yacc.o \
  sphere.o box.o triangle.o superquadric.o outputp.o \
  -lm

# YACC and LEX for full system #
input_yacc.c: input_yacc.y
yacc -d input_yacc.y
mv y.tab.c input_yacc.c

input_lex.c: input_lex.l input_yacc.c
lex -t input_lex.l > input_lex.c

```

### 3.9 FURTHER SUGGESTIONS

This article described only a small subset of what can be accomplished with ray tracing. A few important topics are listed below.

**Supersampling & Antialiasing** Sending only one ray per pixel leads to 'jaggies' along the edges of objects. Supersampling is the process of sending multiple rays per pixel. Antialiasing deals with which rays to send for a pixel and how to combine the resulting colors from the rays. [Whitted 80] suggests sending rays at the corners of each pixel. Adapting the number of rays sent into a pixel based on the complexity of the image is suggested in [Lee 85].

**Object Types** Many different types of objects have been ray traced. Ray tracing of prisms, surfaces of revolution, and fractals is presented in [Kajiya 83].

**Bounding Volumes** Checking every ray for intersection with every object becomes prohibitively expensive for large databases. One solution [Rubin 80] is to place bounding volumes around complex objects. Simple objects such as spheres or boxes are used as bounding volumes. Intersection with the complex object is only checked after a ray strikes the bounding object.

**Octree & Other Space Subdivision** Algorithms for automatic generation of bounding volumes led to the idea of space subdivision. [Glassner 84] used octrees to subdivide space; [Kay 86] and [Fujimoto 86] present other subdivision algorithms. Objects are checked for intersection with a ray only when the ray and the object are in the same region of space.

**Shading Models** Many different shading models have been proposed. The goal of each is to more accurately describe the real world. [Cook 81] and [Hall 84] present two particularly good models.

**Patterns & Textures** Adding patterns and textures is an efficient way to increase the realism of images. [Blinn 76] and [Blinn 78] discuss patterns and textures in general. Procedurally defined marble, soap films, water droplets and waves are presented in [Perlin 85]. Procedurally defined wood texture is discussed in [Peachey 85].

**Distributed Ray Tracing** [Cook 84] and [Dippe 85] introduced extensions to ray tracing which produce antialiased images with depth of field, motion blur, surface gloss, shadow penumbra, and other wonderful effects.

### References

[Appel 68] Appel, Arthur Some Techniques for Shading Ma-

- chine Renderings of Solids *AFIPS Spring Joint Computer Conference* 32 (1968) 37-45.
- [Blinn 76] Blinn, James F. Texture and Reflection in Computer Generated Images *Communications of the ACM* 19,10 (October 1976) 542-547.
- [Blinn 78] Blinn, James F. Simulation of Wrinkled Surfaces *Communications of the ACM* 12,3 (August 1976) 286-292.
- [Cook 81] Cook, Robert L. and Torrance, Kenneth E. A Reflectance Model for Computer Graphics *Computer Graphics* 15,3 (August 1981) 307-316.
- [Cook 84] Cook, Robert L., Porter, Thomas, Carpenter, Loren Distributed Ray Tracing *Computer Graphics* 18,3 (July 1984) 137-145.
- [Dippe 85] Dippe, Mark A. and Wold, Erling Henry Antialiasing Through Stochastic Sampling *Computer Graphics* 19,3 (July 1985) 69-78.
- [Foley 82] Foley, James D. and Van Dam, Andries *Fundamentals of Interactive Computer Graphics* Addison-Wesley, Reading, Mass. 1982
- [Fujimoto 86] Fujimoto, Akira, Tanaka, Takyuki, Iwata, Kansei ARTS: Accelerated Ray Tracing System *IEEE Computer Graphics and Applications* 6,4 (April 1986) 16-26.
- [Glassner 84] Glassner, Andrew S. Space Subdivision for Fast Ray Tracing *IEEE Computer Graphics and Applications* 4,10 (October 1984) 15-22.
- [Goldstein 71] Goldstein, Robert A. and Nagel, Roger 3-D Visual Simulation *Simulation* 16 (January 1971) 25-31.
- [Hall 84] Hall, Roy A. A Testbed for Realistic Image Synthesis *IEEE Computer Graphics and Applications* 3,8 (November 1983) 10-20.
- [Kajiya 83] Kajiya, James T. New Techniques for Ray Tracing Procedurally Defined Objects *Computer Graphics* 17,3 (July 1983) 91-102.
- [Kay 86] Kay, Timothy L. and Kajiya, James T. Ray Tracing Complex Scenes *Computer Graphics* 20,4 (July 1986) 269-278.
- [Lee 85] Lee, Mark E., Redner, Richard A., Uselton, Samuel P. Statistically Optimized Sampling for Distributed Ray Tracing *Computer Graphics* 19,3 (July 1985) 61-67.
- [Newman 73] Newman, William M. and Sproul, Robert F. *Principles Of Interactive Computer Graphics* McGraw-Hill, New York 1973
- [Peachey 85] Peachey, Darwyn R. Solid Texturing of Complex Surfaces *Computer Graphics* 19,3 (July 1985) 279-286.
- [Perlin 85] Perlin, Ken An Image Synthesizer *Computer Graphics* 19,3 (July 1985) 287-296.
- [Roth 82] Roth, Scott D. Ray Casting for Modeling Solids *Computer Graphics and Image Processing* 18 (1982), 109-144.
- [Rubin 80] Rubin, Steven M. and Whitted, Turner A 3-Dimensional Representation for Fast Rendering of Complex Scenes *Computer Graphics* 14,3 (June 1980) 110-116.
- [Whitted 80] Whitted, Turner An Improved Illumination Model for Shaded Display *Communications of the ACM* 23,6 (June 1980) 343-349.

## The Author

Roman Kuchkuda is a graduate student in Computer Science at the University of North Carolina at Chapel Hill. He is on fellowship from Megatek Corp, where he is a software engineer in CAD systems development. He previously worked at Evans & Sutherland on flight simulator software. His research interest is graphics software for image synthesis and animation. Current projects involve procedural and real 3D textures. He received a BS in Computer Science from Rensselaer Polytechnic Institute in 1982. He is a member of IEEE and ACM SIGGRAPH.

## APPENDIX 1: RUNCODE OUTPUT

The pixel output routine presented earlier is very simple and general. It stores R, G, and B values for every pixel. Often in ray traced images many adjacent pixels are the same color.

A method which takes advantage of this is run-length encoding. Strings of pixels which have the same color are only written once. The image file is written line by line, as in pixel output. Each line is composed of runcodes. Each runcode consists of r, g, b and length, with length being the number of successive pixels of this color.

Buffering is also used in the runcode output routine to reduce the number of output operations performed.

The interface to the ray tracer is the same as pixel output. To get runcode output you simply link with the file "outputr.o" instead of with "outputp.o".

```
/** outputr.c **/  
/* runcode output routines */  
#include <stdio.h>  
#include <math.h>  
#include "constants.h"  
#include "typedefs.h"  
#include "funcdefs.h"  
  
#define MAGIC 0x22873  
#define MAXBUF 1024  
  
/* TYPES FOR RUNCODE OUTPUT */  
typedef struct {  
    unsigned char count;  
    unsigned char red;  
    unsigned char green;  
    unsigned char blue;  
} CODE;  
  
typedef struct {  
    int size;  
    int magic;  
} HEADER;  
  
/* STATIC VARIABLES */  
CODE codebuf [MAXBUF];  
CODE *code;  
int wrote;  
int width;  
int outfile;  
int full;  
  
int startpic(fname,y,x)  
char *fname;  
int y,x;  
{  
    HEADER header;
```

```
CODE *getcode();  
width = x;  
full = 0;  
wrote = 0;  
  
outfile = creat(fname, 0666);  
if (outfile == -1)  
{ fprintf(stderr, "ERROR CREATING RUNCODE FILE\n");  
  exit(1);  
}  
header.magic = MAGIC;  
header.size = width;  
write(outfile, &header, sizeof(header));  
code = getcode();  
}  
  
int linepic(pixels)  
double pixels[SCREENWIDTH][3];  
{  
    CODE *getcode();  
    int j,length;  
    int r,g,b;  
    int or,og,ob;  
  
    or = gammacorrect(pixels[0][0]);  
    og = gammacorrect(pixels[0][1]);  
    ob = gammacorrect(pixels[0][2]);  
    length = 0;  
    for (j=1; j<width; j++)  
    { r = gammacorrect(pixels[j][0]);  
      g = gammacorrect(pixels[j][1]);  
      b = gammacorrect(pixels[j][2]);  
      if ((r == or) && (g == og) && (b == ob) &&  
          (length < 255)) length++;  
      else  
      { code->red = or;  
        code->green = og;  
        code->blue = ob;  
        code->count = length;  
        code = getcode();  
        length = 0;  
        or = r;  
        og = g;  
        ob = b;  
      }  
    }  
    code->red = or;  
    code->green = og;  
    code->blue = ob;  
    code->count = length;  
    code = getcode();  
}
```



```

int endpic()
{
    flushr();
    close(outfile);
    printf("\nEND runcodes %d\n",wrote);
}

CODE *getcode()
{
    if (full == MAXBUF)
    { flushr();
      full = 0;
    }
    return (&codebuf[full++]);
}

int flushr()
{
    if (write (outfile, codebuf, full*sizeof(CODE))
        < full*sizeof(CODE))
    { fprintf (stderr, "BAD WRITE\n"); exit (1); }
    wrote += full;
}

int gammacorrect(intensity)
double intensity;
{
    int ival;
    double dval;

    /* scale to 0 - 1 range */
    dval = intensity/255.0;
    if (dval > 1.0) dval = 1.0;
    if (dval < 0.0001) dval = 0.0001;

    /* do gamma correction */
    dval = exp( log(dval) / GAMMA );

    /* convert to integer, range 0-255 */
    dval *= 255.0;
    ival = (int) (dval + 0.5);
    return(ival);
}

```