

CHAPTER 2



On to OpenGL and 3D Computer Graphics



The primary goal for this chapter is to get acquainted with OpenGL and begin our journey into computer graphics using OpenGL as our API (Application Programming Interface) of choice. We shall apply an experiment-discuss-repeat approach where we run code and ask questions of what is seen, acquiring thereby an understanding not only of the way the API functions, but underlying CG concepts as well. Particularly, we want to gain insight into:

- (a) The synthetic-camera model to record 3D scenes, which OpenGL implements.
- (b) The approach of approximating curved objects, such as circles and spheres, with the help of straight and flat geometric primitives, such as line segments and triangles, which is fundamental to object design in computer graphics.

We begin in Section 2.1 with our first OpenGL program to draw a square, the computer graphics equivalent of “Hello World”. Simple though it is, with a few careful experiments and their analysis, `square.cpp` yields a surprising amount of information through Sections 2.1-2.3 about orthographic projection, the fixed world coordinate system OpenGL sets up and how the so-called viewing box in which the programmer draws is specified in this system. We gain insight as well into the 3D-to-2D rendering process.

Adding code to `square.cpp` we see in Section 2.4 how parts of objects outside the viewing box are clipped off. Section 2.5 discusses OpenGL as a state machine. We have in this section as well our first glimpse of property





values, such as color, initially specified at the vertices of a primitive, being interpolated throughout its interior.

Next is the very important Section 2.6 where all the drawing primitives of OpenGL are introduced. These are the parts at the application programmer's disposal with which to assemble objects from thumbtacks to spacecrafts.

The first use of straight primitives to approximate a curved object comes in Section 2.7: a curve (a circle) is drawn using straight line segments. To create more interesting and complex objects one must invoke OpenGL's famous three-dimensionality. This involves learning first in Section 2.8 about perspective projection and also hidden surface removal using the depth buffer.

After a bunch of drawing exercises in Section 2.9 for the reader to practice her newly-acquired skills, the topic of approximating curved objects is broached again in Section 2.10, this time to approximate a surface with triangles, rather than a curve with straight segments as in Section 2.7. Section 2.11 is a review of all the syntax that goes into making a complete OpenGL program.

We conclude with a summary, brief notes and suggestions for further reading in Section 2.12.




Figure 2.1: Screenshot of `square.cpp`.

2.1 First Program

Experiment 2.1. Run `square.cpp`.

Note: Visit the book's website www.sumantaguha.com for a guide on how to install OpenGL and run our programs on various platforms.*

In the OpenGL window appears a black square over a white background. Figure 2.1 is an actual screenshot, but we'll draw it as in Figure 2.2, bluish green standing in for white in order to distinguish it from the paper. We are going to understand next how the square is drawn, and gain some insight as well into the workings behind the scene. **End**





Figure 2.2: OpenGL window of `square.cpp` (bluish green pretending to be white).



The following six statements in `square.cpp` create the square:

```
glBegin(GL_POLYGON);
    glVertex3f(20.0, 20.0, 0.0);
    glVertex3f(80.0, 20.0, 0.0);
    glVertex3f(80.0, 80.0, 0.0);
    glVertex3f(20.0, 80.0, 0.0);
glEnd();
```

Important! If, from what you might have seen elsewhere, you have the notion that `glBegin()-glEnd()`, and even `GL_POLYGON`, specifications are

*If your program compiles but doesn't run, this might be because your graphics card doesn't support OpenGL 4.3. See the note following item 2 in Section 2.11 for what to do in this case.





classical and don't belong in the newest version of OpenGL, then you are right insofar as they are not in the core profile of the latter. They are, though, accessible via the compatibility profile which allows for backward compatibility. Moreover, we explain carefully in the book's preface why we don't subscribe to the "shaders from the opening bell and toss everything classical" school of thought as far as *teaching* OpenGL is concerned. Of course, we shall cover thoroughly the most modern – in fact, fourth generation – OpenGL later in the book. If you have not done so yet, we urge you to read about our pedagogical approach in the preface.

The corners of the square are specified by the four vertex declaration statements between `glBegin(GL_POLYGON)` and `glEnd()`. Let's determine how the `glVertex3f()` statements correspond to corners of the square.

If, suppose, the vertices are specified in some coordinate system that is embedded in the OpenGL window – which certainly is plausible – and if we knew the axes of this system, the matter would be simple. For example, *if* the x -axis increased horizontally rightwards and the y -axis vertically downwards, as in Figure 2.3, then `glVertex3f(20.0, 20.0, 0.0)` would correspond to the upper-left corner of the square, `glVertex3f(80.0, 20.0, 0.0)` to the upper-right corner and so on.

However, even assuming that there do exist these invisible axes attached to the OpenGL window, how do we find out where they are or how they are oriented? One way is to "wiggle" the corners of the square! For example, change the first vertex declaration from `glVertex3f(20.0, 20.0, 0.0)` to `glVertex3f(30.0, 20.0, 0.0)` and observe which corner moves. Having determined in this way the correspondence of the corners with the vertex statements, we ask the reader to deduce the orientation of the hypothetical coordinate axes. Decide where the origin is located too.

Well, it seems then that `square.cpp` sets up coordinates in the OpenGL window so that the increasing direction of the x -axis is horizontally rightwards, that of the y -axis vertically upwards and, moreover, the origin seems to correspond to the lower-left corner of the window, as in Figure 2.4. We're making progress but there's more to the story, so read on!

The last of the three parameters of a `glVertex3f(*, *, *)` declaration is evidently the z coordinate. Vertices are specified in *3-dimensional* space (simply called 3-space or, mathematically, \mathbb{R}^3). Indeed, OpenGL allows us to draw in 3-space and create truly 3D scenes, which is its major claim to fame. However, we *perceive* the 3-dimensional scene as a picture *rendered* to a 2-dimensional part of the computer's screen, the rectangular OpenGL window. Shortly we'll see how OpenGL converts a 3D scene to its 2D rendering.

Section 2.1

FIRST PROGRAM




Figure 2.3: The coordinate axes on the OpenGL window of `square.cpp`? *No.*







Figure 2.4: The coordinate axes on the OpenGL window of `square.cpp`? *Almost there ...*





2.2 Orthographic Projection, Viewing Box and World Coordinates

What exactly do the vertex coordinate values mean? For example, is the vertex at (20.0, 20.0, 0.0) of `square.cpp` 20 mm., 20 cm. or 20 pixels away from the origin along both the x -axis and y -axis, or is there some other absolute unit of distance native to OpenGL?

Experiment 2.2. Change the `glutInitWindowSize()` parameter values of `square.cpp`* – first to `glutInitWindowSize(300, 300)` and then `glutInitWindowSize(500, 250)`. The square changes in size, and even shape, with the OpenGL window. Therefore, coordinate values appear not to be in any kind of absolute units on the screen. **End**

Remark 2.1. Of course, you could have reshaped the OpenGL window directly by dragging one of its corners with the mouse, rather than resetting `glutInitWindowSize()` in the program.

Understanding what the coordinates actually represent involves understanding first OpenGL's rendering mechanism, which itself begins with the program's *projection statement*. In the case of `square.cpp` the projection statement is

```
glOrtho(0.0, 100.0, 0.0, 100.0, -1.0, 1.0)
```

in the `resize()` routine, which determines an imaginary *viewing box* inside which the programmer draws. Generally,

```
glOrtho(left, right, bottom, top, near, far)
```



sets up a viewing box, as in Figure 2.5, with corners at the 8 points:

```
(left, bottom, -near), (right, bottom, -near), (left, top, -near),  
(right, top, -near), (left, bottom, -far), (right, bottom, -far),  
(left, top, -far), (right, top, -far)
```

It's a box with sides aligned along the axes, whose span along the x -axis is from *left* to *right*, along the y -axis from *bottom* to *top*, and along the z -axis from *-far* to *-near*. Note the little quirk of OpenGL that the *near* and *far* values are flipped in sign.

The viewing box corresponding to the projection statement `glOrtho(0.0, 100.0, 0.0, 100.0, -1.0, 1.0)` of `square.cpp` is shown in Figure 2.6(a). The reader may wonder at this time how the initial coordinate axes are *themselves* calibrated – e.g., is a unit along an axis one inch, one centimeter or something else – as the size of the viewing box and that of the objects drawn inside it depend on this. The answer will be evident once the rendering process is explained momentarily.

*When we refer to `square.cpp`, or any `program.cpp`, it's always to the original version in a folder in the `ExperimenterSource` directory, so if you've modified the code for an earlier experiment you'll need to copy back the original.








Figure 2.5: Viewing box of `glOrtho(left, right, bottom, top, near, far)`.

Section 2.2
 ORTHOGRAPHIC
 PROJECTION, VIEWING
 BOX AND WORLD
 COORDINATES






Figure 2.6: (a) Viewing box of `square.cpp` (b) With the square drawn inside.

As for drawing now, the vertex declaration `glVertex3f(x, y, z)` corresponds to the point (x, y, z) . For example, the corner of the square declared by `glVertex3f(20.0, 20.0, 0.0)` is at $(20.0, 20.0, 0.0)$. The square of `square.cpp`, then, is as depicted in Figure 2.6(b).

Once the programmer has drawn the entire scene, if the projection statement is `glOrtho()` as in `square.cpp`, then the rendering process is two-step:

1. *Shoot*: First, objects are *projected perpendicularly* onto the front face of the viewing box, i.e., the face on the $z = -near$ plane. For example, the square in Figure 2.7(a) (same as Figure 2.6(b)) is projected as in Figure 2.7(b). The front face of the viewing box is called the *viewing face* and the plane on which it lies the *viewing plane*.

This step is like shooting the scene on film. In fact, one can think of the viewing box as a giant version of those archaic *box cameras* where the photographer ducks behind the film – the viewing face – and covers her head with a black cloth; so big, in fact, that the whole scene is actually inside the box! Moreover, mind that with this analogy there's



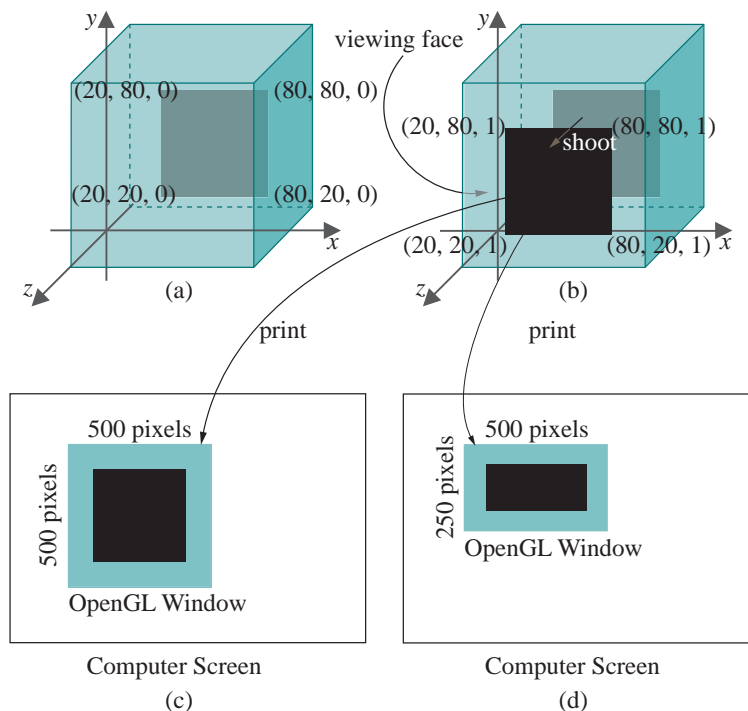




Figure 2.7: Rendering with `glOrtho()`.

no lens, only the film.

2. *Print*: Next, the viewing face is *proportionately scaled* to fit the rectangular OpenGL window. This step is like printing the film on paper. In the case of `square.cpp`, printing takes us from Figure 2.7(b) to (c).

If, say, the window size of `square.cpp` were changed to one of *aspect ratio* (= width/height) of 2, by replacing `glutInitWindowSize(500, 500)` with `glutInitWindowSize(500, 250)`, printing would take us from Figure 2.7(b) to (d) (which actually distorts the square into a rectangle).

The answer to the earlier question of how to calibrate the coordinate axes of the space in which the viewing box is created should be clear now: the 2D rendering finally displayed is the same *no matter* how they are calibrated, because of the proportionate scaling of the viewing face of the box to fit the OpenGL window. So it does not matter what unit we use, be it an inch, millimeter, mile, ...! Here's a partly-solved exercise to drive home the point.



Exercise 2.1.

- (a) Suppose the viewing box of `square.cpp` is set up in a coordinate system where one unit along each axis is 1 cm. Assuming pixels to be 0.2 mm. \times 0.2 mm. squares, compute the size and location of the square rendered by shoot-and-print to a 500 pixel \times 500 pixel OpenGL window.
- (b) Suppose next that the coordinate system is re-calibrated so that a unit along each axis is 1 meter instead of 1 cm., everything else remaining same. What then are the size and location of the rendered square in the OpenGL window?
- (c) What is rendered if, additionally, the size of the OpenGL window is changed to 500 pixel \times 250 pixel?

Section 2.2
 ORTHOGRAPHIC
 PROJECTION, VIEWING
 BOX AND WORLD
 COORDINATES

Part answer:






Figure 2.8: The viewing face for `square.cpp`, given that one unit along each coordinate axis is 1 cm., scaled to a 500 pixel \times 500 pixel OpenGL window.

- (a) Figure 2.8 on the left shows the square projected to the viewing face, which is 100 cm. square. The viewing face is then scaled to the OpenGL window on the right, which is a square of sides 500 pixels = 500 \times 0.2 mm. = 100 mm. The scaling from face to the window, therefore, is a factor of 1/10 in both dimensions. It follows that the rendered square is 60 mm. \times 60 mm., with its lower-left corner located both 20 mm. above and to the right of the lower-left corner of the window.
- (b) Exactly the same as in part (a) because, while the viewing box and viewing face are now 10 times larger, the scaling from face to window is now a factor of 1/100, rather than 1/10.



We conclude that the size and location of the rendering in each coordinate direction are independent of how the axes are calibrated, but determined rather by the *ratio* of the original object's size to that of the viewing box in that direction.

Although the calibration of the axes doesn't matter, nevertheless, we'll make the sensible assumption that all three are calibrated *identically*, i.e., one unit along each axis is of equal length (yes, oddly enough, we could make them different and still the rendering would not change, which you can verify yourself by re-doing Exercise 2.1(a), after assuming that one unit along the x -axis is 1 cm. and along the other two 1 meter). The only other assumptions about the initial coordinate system that we make are conventional ones:

- (a) It is *rectangular*, i.e., the three axes are mutually perpendicular.
- (b) The x -, y - and z -axes in that order form a *right-handed* system in the following sense: the rotation of the x -axis 90° about the origin so that its positive direction matches with that of the y -axis appears *counter-clockwise* to a viewer located on the positive side of the z -axis (Figure 2.9).




Figure 2.9: The x -, y - and z -axes are rectangular and form a (a) right-handed system (b) left-handed system.



Fixed World System

To summarize, set up an initial rectangular right-handed coordinate system located wherever you like in space, but with axes all calibrated identically. Call a unit along each axis just "a unit" as we know it doesn't matter what the unit is. Then leave it *fixed forever* – imagine it screwed to the top of your desk!



Figure 2.10: A dedicated 3D graphics programmer in a world all her own.

See Figure 2.10. This system coordinatizes *world space* and, in fact, we shall refer to it as the *world coordinate system*. All subsequent objects,



including the viewing box and those that we create ourselves, inhabit world space and are specified in world coordinates. These are all virtual objects, of course!

Remark 2.2. Because it's occupied by user-defined objects, world space is sometimes called *object space*.

Experiment 2.3. Change only the viewing box of `square.cpp` by replacing `glOrtho(0.0, 100.0, 0.0, 100.0, -1.0, 1.0)` with `glOrtho(-100, 100.0, -100.0, 100.0, -1.0, 1.0)`. The location of the square in the new viewing box is different and, so as well, the result of shoot-and-print. Figure 2.11 explains how. **End**

Section 2.2
ORTHOGRAPHIC
PROJECTION, VIEWING
BOX AND WORLD
COORDINATES

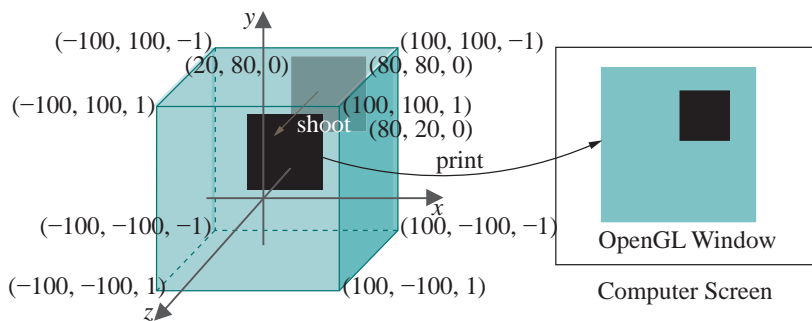




Figure 2.11: The viewing box of `square.cpp` defined by `glOrtho(-100, 100.0, -100.0, 100.0, -1.0, 1.0)`.



Exercise 2.2. (Programming) Change the viewing box of `square.cpp` by replacing `glOrtho(0.0, 100.0, 0.0, 100.0, -1.0, 1.0)` successively with the following, in each case trying to predict the output before running:

- (a) `glOrtho(0.0, 200.0, 0.0, 200.0, -1.0, 1.0)`
- (b) `glOrtho(20.0, 80.0, 20.0, 80.0, -1.0, 1.0)`
- (c) `glOrtho(0.0, 100.0, 0.0, 100.0, -2.0, 5.0)`

Exercise 2.3. If the viewing box of `square.cpp` is changed by replacing `glOrtho(0.0, 100.0, 0.0, 100.0, -1.0, 1.0)` with `glOrtho(-100.0, 100.0, -100.0, 100.0, -1.0, 1.0)`, and the OpenGL window size changed replacing `glutInitWindowSize(500, 500)` with `glutInitWindowSize(500, 250)`, then calculate the area (in *number of pixels*) of the image of the square.

Exercise 2.4. (Programming) We saw earlier that, as a result of the print step, replacing `glutInitWindowSize(500, 500)` with `glutInitWindowSize(500, 250)` in `square.cpp` causes the square to be distorted





into a rectangle. By changing *only one* numerical parameter elsewhere in the program, eliminate the distortion to make it appear square again.

Incidentally, it's clear now that our working hypothesis after the first experiment in Section 2.1, that the OpenGL window comes with axes fixed to it, though not unreasonable, was not accurate either. The OpenGL window it turns out is simply an empty target rectangle on which the front face of the viewing box is printed. This rectangle is called *screen space*.

So there are two spaces we'll be interacting with: world and screen. The former is a virtual 3D space in which we create our scenes, while the latter is a real 2D space where images concocted from our scenes by shoot-and-print are rendered for viewing.

Exercise 2.5. (Programming) Alter the z coordinates of each vertex of the “square” – we should really call it a polygon if we do this – of `square.cpp` as follows (Block 1*):

```
glBegin(GL_POLYGON);  
    glVertex3f(20.0, 20.0, 0.5);  
    glVertex3f(80.0, 20.0, -0.5);  
    glVertex3f(80.0, 80.0, 0.1);  
    glVertex3f(20.0, 80.0, 0.2);  
glEnd();
```

The rendering does not change. Why?

Remark 2.3. Always set the parameters of `glOrtho(left, right, bottom, top, near, far)` so that $left < right$, $bottom < top$, and $near < far$.



Remark 2.4. The aspect ratio (= width/height) of the viewing box should be set same as that of the OpenGL window or the scene will be distorted by the print step.



Remark 2.5. The perpendicular projection onto the viewing plane corresponding to a `glOrtho()` call is also called *orthographic projection* or *orthogonal projection* (hence the name of the call). Yet another term is *parallel projection* as the lines of projection from points in the viewing box to the viewing plane are all parallel.

2.3 The OpenGL Window and Screen Coordinates

We've already had occasion to use the `glutInitWindowSize(w , h)` command which sets the size of the OpenGL window to width w and height h measured in pixels. A companion command is `glutInitWindowPosition(x , y)` to specify the location (x, y) of the upper-left corner of the OpenGL window on the computer screen.

*To cut-and-paste you can find the block in text format in the file `chap2codeModifications.txt` in the directory `ExperimenterSource/CodeModifications`.





Experiment 2.4. Change the parameters of `glutInitWindowPosition(x, y)` in `square.cpp` from the current (100, 100) to a few different values to determine the location of the origin (0, 0) of the computer screen, as well as the orientation of the screen's own x -axis and y -axis. **End**

Section 2.4
CLIPPING

The origin (0,0) of the screen it turns out is at its upper-left corner, while the increasing direction of its x -axis is horizontally rightwards and that of its y -axis vertically downwards; moreover, one unit along either axis is *absolute* and represents a pixel. See Figure 2.12, which shows as well the coordinates of the corners of the OpenGL window initialized by `square.cpp`.




Figure 2.12: The screen's coordinate system: a unit along either axis is the pitch of a pixel.



Note the inconsistency between the orientation of the screen's y -axis and the y -axis of the world coordinate system, the latter being directed *up* the OpenGL window (after being projected there). One needs to be careful about this, especially when coding programs where data is read from one system and used in the other.



2.4 Clipping

A question may have come to the reader's mind about objects which happen to be drawn outside the viewing box. Here are a few experiments to clarify how they are processed.

Experiment 2.5. Add another square by inserting the following right after the code for the original square in `square.cpp` (Block 2):

```
glBegin(GL_POLYGON);
    glVertex3f(120.0, 120.0, 0.0);
    glVertex3f(180.0, 120.0, 0.0);
    glVertex3f(180.0, 180.0, 0.0);
    glVertex3f(120.0, 180.0, 0.0);
glEnd();
```





From the value of its vertex coordinates the second square evidently lies entirely outside the viewing box.

If you run now there's no sign of the second square in the OpenGL window! This is because OpenGL *clips* the scene to within the viewing box before rendering, so that objects or parts of objects drawn outside are not rendered. Clipping is a stage in the graphics pipeline. We'll not worry about its implementation at this time, only the effect it has. **End**

Exercise 2.6. (Programming) In the preceding experiment can you redefine the viewing box by changing the parameters of `glOrtho()` so that both squares are visible?



Figure 2.13: Screenshot of a triangle.

Experiment 2.6. For a more dramatic illustration of clipping, first replace the square in the original `square.cpp` with a triangle; in particular, replace the polygon code with the following (Block 3):

```
glBegin(GL_POLYGON);
  glVertex3f(20.0, 20.0, 0.0);
  glVertex3f(80.0, 20.0, 0.0);
  glVertex3f(80.0, 80.0, 0.0);
glEnd();
```

See Figure 2.13. Next, lift the first vertex up the z -axis by changing it to `glVertex3f(20.0, 20.0, 0.5)`; lift it further by changing its z -value to 1.5 (Figure 2.14 is a screenshot), then 2.5 and, finally, 10.0. Make sure you believe that what you see in the last three cases is indeed a triangle clipped to within the viewing box – Figure 2.15 may be helpful. **End**





Figure 2.14: Screenshot of the triangle clipped to a quadrilateral.

Exercise 2.7. (Programming) A triangle was clipped to a quadrilateral in the viewing box in the preceding experiment. What is the maximum number of sides of a figure to which you can clip a triangle in the box (quadrilateral, pentagon, hexagon, ...)? Code and show.

Exercise 2.8. Use pencil and paper to guess the output if the polygon declaration part of `square.cpp` is replaced with the following (Block 4):

```
glBegin(GL_POLYGON);
  glVertex3f(-20.0, -20.0, 0.0);
  glVertex3f(80.0, 20.0, 0.0);
  glVertex3f(120.0, 120.0, 0.0);
  glVertex3f(20.0, 80.0, 0.0);
glEnd();
```

The viewing box has six faces that lie on different six planes and, effectively, OpenGL clips *off* the scene on one side of each of these six planes, accordingly called *clipping planes*. Imagine a knife slicing down each plane as in Figure 2.15. Specifically, in the case of the viewing box set up by `glOrtho(left, right, bottom, top, near, far)`, clipped off is to the left of the plane $x = left$, to the right of the plane $x = right$ and so on.








Figure 2.15: Six clipping planes of the `glOrtho(left, right, bottom, top, near, far)` viewing box. The lightly shaded part of the triangle sticking out of the box is clipped by a “clipping knife”.

Remark 2.6. As we shall see in Chapter 3, the programmer can define clipping planes in addition to the six that bound the viewing box.

We’ll leave this section with a rather curious phenomenon for the reader to explain.

Exercise 2.9. (Programming) Raising the first vertex of (the original) `square.cpp` from `glVertex3f(20.0, 20.0, 0.0)` to `glVertex3f(20.0, 20.0, 1.5)` causes the square – actually, the new figure which is not a square any more – to be clipped. If, instead, the second vertex is raised from `glVertex3f(80.0, 20.0, 0.0)` to `glVertex3f(80.0, 20.0, 1.5)`, then the figure is clipped too, *but* very differently from when the first vertex is raised. Why? Should not the results be similar by symmetry?

Hint: OpenGL draws polygons after triangulating them as so-called *triangle fans* with the first vertex of the polygon the center of the fan. For example, the fan in Figure 2.16 consists of three triangles around vertex v_0 .








Figure 2.16: A triangle fan.

2.5 Color, OpenGL State Machine and Interpolation

Experiment 2.7. The color of the square in `square.cpp` is specified by the three parameters of the `glColor3f(0.0, 0.0, 0.0)` statement in the `drawScene()` routine, each of which gives the value of one of the three primary components, *blue*, *green* and *red*.

Determine which of the three parameters of `glColor3f()` specifies the blue, green and red components by setting in turn each to 1.0 and the others to 0.0. In fact, verify the following table:





Chapter 2
ON TO OpenGL AND
3D COMPUTER
GRAPHICS

Call	Color
<code>glColor3f(0.0, 0.0, 0.0)</code>	Black
<code>glColor3f(1.0, 0.0, 0.0)</code>	Red
<code>glColor3f(0.0, 1.0, 0.0)</code>	Green
<code>glColor3f(0.0, 0.0, 1.0)</code>	Blue
<code>glColor3f(1.0, 1.0, 0.0)</code>	Yellow
<code>glColor3f(1.0, 0.0, 1.0)</code>	Magenta
<code>glColor3f(0.0, 1.0, 1.0)</code>	Cyan
<code>glColor3f(1.0, 1.0, 1.0)</code>	White

End

Generally, the `glColor3f(red, green, blue)` call specifies the *foreground color*, or *drawing color*, which is the color applied to objects being drawn. The value of each color component, which ought to be a number between 0.0 and 1.0, determines its *intensity*. For example, `glColor3f(1.0, 1.0, 0.0)` is the brightest yellow while `glColor3f(0.5, 0.5, 0.0)` is a weaker yellow

Remark 2.7. The color values are each *clamped* to the range [0, 1]. This means that, if a value happens to be set greater than 1, then it's taken to be 1; if less than 0, it's taken to be 0.

Exercise 2.10. (Programming) Both `glColor3f(0.2, 0.2, 0.2)` and `glColor3f(0.8, 0.8, 0.8)` should be grays, having equal red, green and blue intensities. Guess which is the darker of the two. Verify by changing the foreground color of `square.cpp`.



The call `glClearColor(1.0, 1.0, 1.0, 0.0)` in the `setup()` routine specifies the *background color*, or *clearing color*. Ignore for now the fourth parameter, which is the *alpha* value. The statement `glClear(GL_COLOR_BUFFER_BIT)` in `drawScene()` actually clears the window to the specified background color, which means that every pixel in the color buffer is set to that color.



Experiment 2.8. Add the additional color declaration statement `glColor3f(1.0, 0.0, 0.0)` just after the existing one `glColor3f(0.0, 0.0, 0.0)` in the drawing routine of `square.cpp` so that the foreground color block becomes

```
// Set foreground (or drawing) color.
glColor3f(0.0, 0.0, 0.0);
glColor3f(1.0, 0.0, 0.0);
```

The square is drawn red because the *current* value of the foreground color is red when each of its vertices is specified. **End**

Foreground color is one of a collection of variables, called *state variables*, which determine the state of OpenGL. Among other state variables are point





size, line width, line stipple, material properties, etc. We'll meet several as we go along or you can refer to the red book* for a full list. OpenGL remains and functions in its current state until a declaration is made changing a state variable. For this reason, OpenGL is often called a *state machine*. The following simple experiment illustrates a couple of important points about how state variables control rendering.

Section 2.5
COLOR, OpenGL STATE
MACHINE AND
INTERPOLATION

Experiment 2.9. Replace the polygon declaration part of `square.cpp` with the following to draw two squares (Block 5):

```
glColor3f(1.0, 0.0, 0.0);
glBegin(GL_POLYGON);
    glVertex3f(20.0, 20.0, 0.0);
    glVertex3f(80.0, 20.0, 0.0);
    glVertex3f(80.0, 80.0, 0.0);
    glVertex3f(20.0, 80.0, 0.0);
glEnd();

glColor3f(0.0, 1.0, 0.0);
glBegin(GL_POLYGON);
    glVertex3f(40.0, 40.0, 0.0);
    glVertex3f(60.0, 40.0, 0.0);
    glVertex3f(60.0, 60.0, 0.0);
    glVertex3f(40.0, 60.0, 0.0);
glEnd();
```

A small green square appears inside a larger red one (Figure 2.17). Obviously, this is because the foreground color is red for the first square, but green for the second. One says that the color red *binds* to the first square – or, more precisely, to each of its four specified vertices – and green to the second square. These bound values specify the color *attribute* of either square. Generally, the values of those state variables which determine how it is rendered collectively form a primitive's attribute set.

Flip the order in which the two squares appear in the code by cutting the seven statements that specify the red square and pasting them after those to do with the green one. The green square is overwritten by the red one and no longer visible because OpenGL draws in *code order*.






Figure 2.17: Screenshot of a green square drawn in the code after a red square.



End

Experiment 2.10. Replace the polygon declaration part of `square.cpp` with (Block 6):

```
glBegin(GL_POLYGON);
    glColor3f(1.0, 0.0, 0.0);
```

*The *OpenGL Programming Guide* [104] and its companion volume, the *OpenGL Reference Manual* [105], are the canonical references for the OpenGL API and affectionately referred to as the red book and blue book, respectively. Note that the on-line reference docs at www.opengl.org pretty much cover all that is in the blue book.





Chapter 2
ON TO OPENGL AND
3D COMPUTER
GRAPHICS



Figure 2.18: Screenshot of a square with differently colored vertices.

```
glVertex3f(20.0, 20.0, 0.0);
glColor3f(0.0, 1.0, 0.0);
glVertex3f(80.0, 20.0, 0.0);
glColor3f(0.0, 0.0, 1.0);
glVertex3f(80.0, 80.0, 0.0);
glColor3f(1.0, 1.0, 0.0);
glVertex3f(20.0, 80.0, 0.0);
glEnd();
```

The different color values bound to the four vertices of the square are evidently *interpolated* over the rest of the square as you can see in Figure 2.18. In fact, this is most often the case with OpenGL: numerical attribute values specified at the vertices of a primitive are interpolated throughout its interior. In a later chapter we'll see exactly what it means to interpolate and how OpenGL goes about the task. **End**

2.6 OpenGL Geometric Primitives

The geometric primitives – also called drawing primitives or, simply, primitives – of OpenGL are the parts that programmers use in Lego-like manner to create mundane objects like balls and boxes, as well as elaborate spacecrafts, the worlds to which they travel, and pretty much everything in between. The only one we've seen so far is the polygon. It's time to get acquainted with the whole family.



Experiment 2.11. Replace `glBegin(GL_POLYGON)` with `glBegin(GL_POINTS)` in `square.cpp` and make the point size bigger with a call to `glPointSize(5.0)`, so that the part drawing the polygon is now



```
glPointSize(5.0); // Set point size.
glBegin(GL_POINTS);
glVertex3f(20.0, 20.0, 0.0);
glVertex3f(80.0, 20.0, 0.0);
glVertex3f(80.0, 80.0, 0.0);
glVertex3f(20.0, 80.0, 0.0);
glEnd();
```

End

Experiment 2.12. Continue, replacing `GL_POINTS` with `GL_LINES`, `GL_LINE_STRIP` and, finally, `GL_LINE_LOOP`. **End**

In the explanation that follows of how OpenGL draws each primitive, assume that the n vertices declared in the code between `glBegin(primitive)` and `glEnd()` are v_0, v_1, \dots, v_{n-1} in that order, i.e., the declaration of the primitive is of the form:





```

glBegin(primitive);
  glVertex3f(*, *, *); // v0
  glVertex3f(*, *, *); // v1
  ...
  glVertex3f(*, *, *); // v_{n-1}
glEnd();

```

Section 2.6
OPENGL GEOMETRIC
PRIMITIVES

Refer to Figure 2.19 as you read.

GL_POINTS draws a point at each vertex

$$v_0, v_1, \dots, v_{n-1}$$

GL_LINES draws a *disconnected* sequence of straight line segments (henceforth, we'll simply use the term “segment”) between the vertices, taken two at a time. In particular, it draws the segments

$$v_0v_1, v_2v_3, \dots, v_{n-2}v_{n-1}$$

if n is even. If n is not even then the last vertex v_{n-1} is simply ignored.

GL_LINE_STRIP draws the *connected* sequence of segments

$$v_0v_1, v_1v_2, \dots, v_{n-2}v_{n-1}$$

Such a sequence is called a *polygonal line* or *polyline*.

GL_LINE_LOOP is the same as **GL_LINE_STRIP**, *except* that an additional segment $v_{n-1}v_0$ is drawn to complete a loop:

$$v_0v_1, v_1v_2, \dots, v_{n-2}v_{n-1}, v_{n-1}v_0$$

Such a segment sequence is called a *polygonal line loop*.

The thickness of lines can be set by a `glLineWidth(width)` call.



Remark 2.8. In world space points have zero dimension and lines zero width; values specified by `glPointSize()` and `glLineWidth()` are used only for rendering. Otherwise, it would be rather hard to see a point actually of zero dimension or a line of zero width!

Why does OpenGL provide separate primitives to draw polygonal lines and line loops when both can be viewed as a collection of segments and drawn using **GL_LINES**? For example,

```

glBegin(GL_LINE_STRIP);
  v0;
  v1;
  v2;
  ...
glEnd();

```



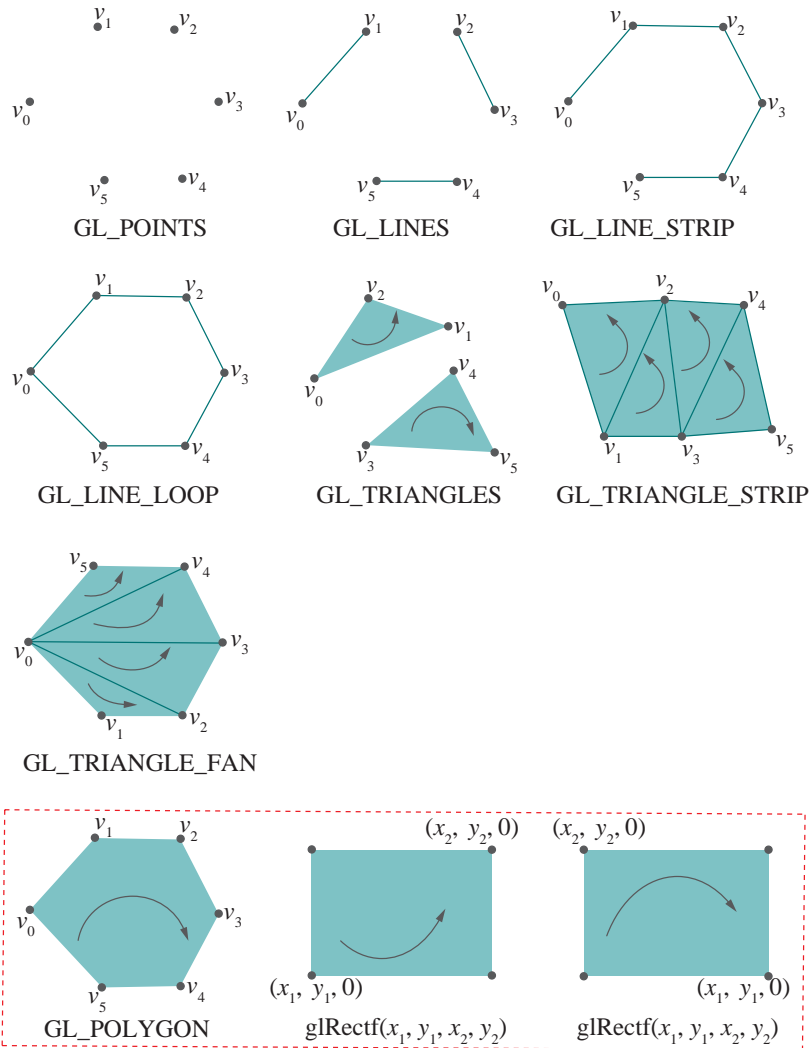




Figure 2.19: OpenGL's geometric primitives. Vertex order is indicated by a curved arrow. Primitives inside the red rectangle have been discarded from the core profile of later versions of OpenGL, e.g., 4.3; however, they are accessible via the compatibility profile.

is equivalent to

```
glBegin(GL_LINES);
v0;
v1
```



```
v1;  
v2  
v2;  
...  
glEnd();
```

Section 2.6
OPENGL GEOMETRIC
PRIMITIVES

The answer is first to avoid redundancy in vertex data. Secondly, possible rendering error is avoided as well because OpenGL does not know that the two `v1`s in the `GL_LINES` specification above are supposed to represent the same vertex, and may render them at slightly different locations because of differences in floating point round-offs.

Exercise 2.11. (Programming) This relates to the brief discussion on interpolation at the end of Section 2.5. Replace the polygon declaration part of `square.cpp` with (Block 7):

```
glLineWidth(5.0);  
glBegin(GL_LINES);  
glColor3f(1.0, 0.0, 0.0);  
glVertex3f(20.0, 20.0, 0.0);  
glColor3f(0.0, 1.0, 0.0);  
glVertex3f(80.0, 20.0, 0.0);  
glEnd();
```

Can you say what the color values should be at the midpoint (50.0, 20.0, 0.0) of the segment drawn? Check your answer by drawing a point with those color values just above the midpoint, say at (50.0, 22.0, 0.0), and comparing.

Experiment 2.13. Replace the polygon declaration part of `square.cpp` with (Block 8):

```
glBegin(GL_TRIANGLES);  
glVertex3f(10.0, 90.0, 0.0);  
glVertex3f(10.0, 10.0, 0.0);  
glVertex3f(35.0, 75.0, 0.0);  
glVertex3f(30.0, 20.0, 0.0);  
glVertex3f(90.0, 90.0, 0.0);  
glVertex3f(80.0, 40.0, 0.0);  
glEnd();
```



End



`GL_TRIANGLES` draws a sequence of triangles using the vertices three at a time. In particular, the triangles are

$$v_0v_1v_2, v_3v_4v_5, \dots, v_{n-3}v_{n-2}v_{n-1}$$

if n is a multiple of 3; if isn't, the last one, or two, vertices are ignored.

The given order of the vertices for each triangle, in particular, v_0, v_1, v_2 for the first, v_3, v_4, v_5 for the second and so on, determines its *orientation* as perceived by a viewer. Figure 2.19 indicates orientation with curved arrows.





Orientation is important because it enables OpenGL to decide which side of a primitive, front or back, the viewer sees. We'll deal with this topic separately in Chapter 9. Till then disregard orientation when drawing, listing the vertices of a primitive in any order you like.

`GL_TRIANGLES` is a *2-dimensional* primitive and, by default, triangles are drawn filled. However, one can choose a different drawing mode by applying the `glPolygonMode(face, mode)` command where *face* may be one of `GL_FRONT`, `GL_BACK` or `GL_FRONT_AND_BACK`, and *mode* one of `GL_FILL`, `GL_LINE` or `GL_POINT`. Whether a primitive is front-facing or back-facing depends, as said above, on its orientation. To keep matters simple for now, though, we'll use only `GL_FRONT_AND_BACK` in a `glPolygonMode()` call, which applies the given drawing mode to a primitive regardless of which face is visible. The `GL_FILL` option is, of course, the default filled option for 2D primitives, while `GL_LINE` draws the primitive in *outline* (or *wireframe* as it's also called), and `GL_POINT` only the vertices.

Experiment 2.14. In fact, it's often easier to decipher a 2D primitive by viewing it in outline. Accordingly, continue the preceding experiment by inserting the call `glPolygonMode(GL_FRONT_AND_BACK, GL_LINE)` in the drawing routine and, further, replacing `GL_TRIANGLES` with `GL_TRIANGLE_STRIP`. The relevant part of the display routine then is as below:

```
// Set polygon mode.
glPolygonMode(GL_FRONT_AND_BACK, GL_LINE);

// Draw a triangle strip.
glBegin(GL_TRIANGLE_STRIP);
  glVertex3f(10.0, 90.0, 0.0);
  glVertex3f(10.0, 10.0, 0.0);
  glVertex3f(35.0, 75.0, 0.0);
  glVertex3f(30.0, 20.0, 0.0);
  glVertex3f(90.0, 90.0, 0.0);
  glVertex3f(80.0, 40.0, 0.0);
glEnd();
```

End



`GL_TRIANGLE_STRIP` draws a sequence of triangles – called a *triangle strip* – as follows: the first triangle is $v_0v_1v_2$, the second $v_1v_3v_2$ (v_0 is dropped and v_3 brought in), the third $v_2v_3v_4$ (v_1 dropped and v_4 brought in), and so on. Formally, the triangles in the strip are



$$v_0v_1v_2, v_1v_3v_2, v_2v_3v_4, \dots, v_{n-3}v_{n-2}v_{n-1} \quad (\text{if } n \text{ is odd})$$

or

$$v_0v_1v_2, v_1v_3v_2, v_2v_3v_4, \dots, v_{n-3}v_{n-1}v_{n-2} \quad (\text{if } n \text{ is even})$$

Exercise 2.12. (Programming) Create a square annulus as in Figure 2.20(a) using a *single* triangle strip. You may first want to sketch the





annulus on graph paper to determine the coordinates of its eight corners. The figure depicts one possible *triangulation* – division into triangles – of the annulus.

Hint: A solution is available in `squareAnnulus1.cpp` of Chapter 3.

Exercise 2.13. (Programming) Create the shape of Figure 2.20(b) using a single triangle strip. A partial triangulation is indicated.

Experiment 2.15. Replace the polygon declaration part of `square.cpp` with (Block 9):

```
glBegin(GL_TRIANGLE_FAN);
  glVertex3f(10.0, 10.0, 0.0);
  glVertex3f(15.0, 90.0, 0.0);
  glVertex3f(55.0, 75.0, 0.0);
  glVertex3f(80.0, 30.0, 0.0);
  glVertex3f(90.0, 10.0, 0.0);
glEnd();
```

Apply both the filled and outlined drawing modes.

End

`GL_TRIANGLE_FAN` draws a sequence of triangles – called a *triangle fan* – around the first vertex as follows: the first triangle is $v_0v_1v_2$, the second $v_0v_2v_3$ and so on. The full sequence is

$$v_0v_1v_2, v_0v_2v_3, \dots, v_0v_{n-2}v_{n-1}$$

Exercise 2.14. (Programming) Create a square annulus using *two* triangle fans. First sketch a triangulation different from that in Figure 2.20(a).

`GL_POLYGON` draws a polygon with the vertex sequence

$$v_0 v_1 \dots v_{n-1}$$


(n must be at least 3 for anything to be drawn).

Finally:


`glRectf`($x1, y1, x2, y2$) draws a rectangle lying on the $z = 0$ plane with sides parallel to the x - and y -axes. In particular, the rectangle has diagonally opposite corners at $(x1, y1, 0)$ and $(x2, y2, 0)$. The full list of four vertices is $(x1, y1, 0)$, $(x2, y1, 0)$, $(x2, y2, 0)$ and $(x1, y2, 0)$. The rectangle created is 2-dimensional and its vertex order depends on the situation of the two vertices $(x1, y1, 0)$ and $(x2, y2, 0)$ with respect to each other, as indicated by the two drawings at the lower right of Figure 2.19.

Note that `glRectf`() is a stand-alone call; it is not a parameter to `glBegin`() like the other primitives.

Section 2.6 OPENGL GEOMETRIC PRIMITIVES







(a)



(b)

Figure 2.20: (a) Square annulus – the region between two bounding squares – and a possible triangulation (b) A partially triangulated shape.





Important: The preceding two, `GL_POLYGON` and `glRectf()`, have both been discarded from the core profile of later versions of OpenGL, e.g., the one we are going to study ourselves later in the book, namely, fourth generation 4.3; however, they are accessible via the compatibility profile.

The reason that polygons and rectangles have been discarded is not hard to understand: both can be made from triangles, so are really redundant. The reason we do use them in the first part of this book is because they afford an easily understood way to make objects – e.g., a polygon is certainly more intuitive for a beginner than a triangle strip.

However, when drawing a polygon one must be careful in ensuring that it is a *plane convex* figure, i.e., it lies on one plane and has no “bays” or “inlets” (see Figure 2.21); otherwise, rendering is unpredictable. So, even though we draw them occasionally for convenience, we recommend that the reader, in order to avoid rendering issues and to prepare for the fourth generation of OpenGL, *avoid polygons and rectangles altogether* in her own projects, and, instead, *use exclusively triangles*.




Figure 2.21: OpenGL polygons should be planar and convex.

In fact, following are a couple of experiments, the second one showing how polygon rendering can behave oddly indeed if one is not careful.

Experiment 2.16. Replace the polygon declaration of `square.cpp` with (Block 10):



```
glBegin(GL_POLYGON);
  glVertex3f(20.0, 20.0, 0.0);
  glVertex3f(50.0, 20.0, 0.0);
  glVertex3f(80.0, 50.0, 0.0);
  glVertex3f(80.0, 80.0, 0.0);
  glVertex3f(20.0, 80.0, 0.0);
glEnd();
```




You see a convex 5-sided polygon (Figure 2.22(a)).

End

Experiment 2.17. Replace the polygon declaration of `square.cpp` with (Block 11):

```
glBegin(GL_POLYGON);
```





Section 2.6
OPENGL GEOMETRIC
PRIMITIVES

Figure 2.22: Outputs: (a) Experiment 2.16 (b) Experiment 2.17 (c) Experiment 2.17, vertices cycled.

```

glVertex3f(20.0, 20.0, 0.0);
glVertex3f(80.0, 20.0, 0.0);
glVertex3f(40.0, 40.0, 0.0);
glVertex3f(20.0, 80.0, 0.0);
glEnd();

```

Display it *both* filled and outlined using appropriate `glPolygonMode()` calls. A non-convex quadrilateral is drawn in either case (Figure 2.22(b)). Next, keeping the same *cycle* of vertices as above, list them starting with `glVertex3f(80.0, 20.0, 0.0)` instead (Block 12):

```

glBegin(GL_POLYGON);
glVertex3f(80.0, 20.0, 0.0);
glVertex3f(40.0, 40.0, 0.0);
glVertex3f(20.0, 80.0, 0.0);
glVertex3f(20.0, 20.0, 0.0);
glEnd();

```

Make sure to display it both filled and outlined. When filled it's a triangle, while outlined it's a non-convex quadrilateral (Figure 2.22(c)) identical to the one output earlier! Because the cyclic order of the vertices is unchanged, shouldn't it be as in Figure 2.22(b) both filled and outlined? **End**

We'll leave the apparent anomaly* of this experiment as a mystery to be resolved in Chapter 8 on triangulation. But, if you are impatient then refer to the hint provided with Exercise 2.9.

Exercise 2.15. (Programming) Verify, by cycling the vertices, that no such anomaly arises in the case of the convex polygon of Experiment 2.16.

Exercise 2.16. (Programming) Draw the double annulus (a figure '8') shown in Figure 2.23 using as few triangle strips as possible. Introduce additional vertices on the three boundary components if you need to (in addition to the original twelve).








Figure 2.23: Double annulus.

Note: Such additional vertices are called *Steiner vertices*.

*The rendering depends on the particular OpenGL implementation. However, all implementations that we are aware of show identical behavior.





Remark 2.9. Here’s an interesting semi-philosophical question. OpenGL claims to be a 3D drawing API. Yet, why does it not have a single 3D drawing primitive, e.g., cube, tetrahedron or such? All its primitives are 0-dimensional (GL_POINTS), 1-dimensional (GL_LINE*) or 2-dimensional (GL_TRIANGLE*, GL_QUAD*, GL_POLYGON, glRectf()).

The answer lies in how we humans (the regular ones that is and not supers with X-ray vision) perceive 3D objects such as cubes, tetrahedrons, chairs and spacecraft: *we see only the surface, which is two-dimensional*. It makes sense for a 3D API, therefore, to draw only as much as can be seen.

2.7 Approximating Curved Objects

Looking back at Figure 2.19 we see that the OpenGL geometric primitives are composed of points, straight line segments and flat pieces, the latter being triangles, rectangles and polygons. How, then, to draw curved objects such as discs, ellipses, spirals, beer cans and flying saucers? The answer is to *approximate* them with straight and flat OpenGL primitives well enough that the viewer cannot tell the difference. As a wag once put it, “Sincerity is a very important human quality. If you don’t have it, you *gotta* fake it!” In the next experiment we fake a circle.



Figure 2.24: Screenshot of circle.cpp.

Experiment 2.18. Run circle.cpp. Increase the number of vertices in the line loop



```
glBegin(GL_LINE_LOOP);
  for(i = 0; i < numVertices; ++i)
  {
    glColor3f((float)rand()/(float)RAND_MAX,
              (float)rand()/(float)RAND_MAX,
              (float)rand()/(float)RAND_MAX);
    glVertex3f(X + R * cos(t), Y + R * sin(t), 0.0);
    t += 2 * PI / numVertices;
  }
glEnd();
```

by pressing ‘+’ till it “becomes” a circle, as in the screenshot of Figure 2.24. Press ‘-’ to decrease the number of vertices. The randomized colors are just a bit of eye candy. **End**

The vertices of the loop of circle.cpp, which lie evenly spaced on the circle, are collectively called a *sample of points* or, simply, *sample* from the circle. See Figure 2.25(a). The denser the sample evidently the better the approximation.

The parametric equations of the circle implemented are

$$x = X + R \cos t, \quad y = Y + R \sin t, \quad z = 0, \quad 0 \leq t \leq 2\pi \quad (2.1)$$








Figure 2.25: (a) A line loop joining a sample of points from a circle (b) Parametric equations for a circle.

Section 2.7
APPROXIMATING
CURVED OBJECTS

where $(X, Y, 0)$ is the center and R the radius of the circle. See Figure 2.25(b). A `numVertices` number of sample points equally spaced apart is generated by starting with the angle $t = 0$ and then incrementing it successively by $2\pi/\text{numVertices}$.

Observe that the vertex specifications occur within a loop construct, which is pretty much mandatory if there is a large number of vertices.

Incidentally, the program `circle.cpp` also demonstrates output to the command window, as well as non-trivial user interaction via the keyboard. The routine `keyInput()` is registered as the key handling routine in `main()` by the `glutKeyboardFunc(keyInput)` statement. Note the calls to `glutPostRedisplay()` in `keyInput()` asking the display to be redrawn after each update of `numVertices`.

Follow these conventions when writing OpenGL code:

1. Program the “Esc” key to exit the program.
2. Describe user interaction at two places:
 - (a) The command window using `cout()`.
 - (b) Comments at the *top* of the source code.

Here’s a parabola.

Experiment 2.19. Run `parabola.cpp`. Press ‘+/-’ to increase/decrease the number of vertices of the approximating line strip. Figure 2.26 is a screenshot with enough vertices to make a smooth-looking parabola.



The vertices are equally spaced along the x -direction. The parametric equations implemented are

$$x = 50 + 50t, \quad y = 100t^2, \quad z = 0, \quad -1 \leq t \leq 1$$

the constants being chosen so that the parabola is centered in the OpenGL window. **End**



Figure 2.26: Screenshot of `parabola.cpp`.








Figure 2.27: Flat spiral.



Figure 2.28: Flat leaf.



Figure 2.29: Screenshot of `circularAnnuluses.cpp`.

Exercise 2.17. (Programming) Modify `circle.cpp` to draw a flat 3-turn spiral as in the screenshot of Figure 2.27.

Hint: Extending the range of t to 0 to 6π gives 3 full turns. At the same time the “radius” R should decrease with each iteration of the `for` loop, e.g., `R-=20.0/numVertices` decreases it linearly from 40 to 20.

Exercise 2.18. (Programming) Modify `circle.cpp` to draw a disc (a filled circle) by way of (a) a polygon and (b) a triangle fan.

Exercise 2.19. (Programming) Draw a flat leaf like the one in Figure 2.28.

Exercise 2.20. (Programming) Modify `circle.cpp` to draw a circular annulus, like one of those shown in Figure 2.29, using a triangle strip. Don’t look at the program `circularAnnuluses.cpp`!

We’ll be returning shortly to the topic of approximating curved objects, but it’s on to 3D next.

2.8 Three Dimensions, the Depth Buffer and Perspective Projection

The reader by now may be getting impatient to move on from the plane (pun intended) and simple to full 3D. Okay then, let’s get off to an easy start in 3-space by making use of the third dimension to fake a circular annulus. Don’t worry, we’ll be doing fancier stuff soon enough!



Experiment 2.20. Run `circularAnnuluses.cpp`. Three identical-looking red circular annuluses (Figure 2.29) are drawn in three *different* ways:



- i) Upper-left: There is not a real hole. The white disc *overwrites* the red disc as it appears later in the code.

```
glColor3f(1.0, 0.0, 0.0);
drawDisc(20.0, 25.0, 75.0, 0.0);
glColor3f(1.0, 1.0, 1.0);
drawDisc(10.0, 25.0, 75.0, 0.0);
```

Note: The first parameter of `drawDisc()` is the radius and the remaining three the coordinates of the center.

- ii) Upper-right: There is not a real hole either. A white disc is *drawn closer* to the viewer than the red disc thus blocking it out.





```

glEnable(GL_DEPTH_TEST);
glColor3f(1.0, 0.0, 0.0);
drawDisc(20.0, 75.0, 75.0, 0.0);
glColor3f(1.0, 1.0, 1.0);
drawDisc(10.0, 75.0, 75.0, 0.5);
glDisable(GL_DEPTH_TEST);

```

Section 2.8
THREE DIMENSIONS,
THE DEPTH BUFFER
AND PERSPECTIVE
PROJECTION

Observe that the z -value of the white disc's center is greater than the red disc's. We'll discuss the mechanics of one primitive blocking out another momentarily.

iii) Lower: A true circular annulus with a real hole.

```

if (isWire) glPolygonMode(GL_FRONT, GL_LINE);
else glPolygonMode(GL_FRONT, GL_FILL);
glColor3f(1.0, 0.0, 0.0);
glBegin(GL_TRIANGLE_STRIP);
...
glEnd();

```

Press the space bar to see the wireframe of a triangle strip. **End**

Exercise 2.21. (Programming) Interchange in `circularAnnuluses.cpp` the drawing orders of the red and white discs – i.e., the order in which they appear in the code – in either of the top two annuluses. Which one is affected? (*Only the first!*) Why?



Remark 2.10. Note the use of a text-drawing routine in `circularAnnuluses.cpp`. OpenGL offers only rudimentary text-drawing capability but it often comes in handy, especially for annotation. We'll discuss text-drawing in fair detail in Chapter 3.

By far the most important aspect of `circularAnnuluses.cpp` is its use of the *depth buffer* to draw the upper-right annulus. Following is an introduction to this critical utility which enables realistic rendering of 3D scenes.

2.8.1 A Vital 3D Utility: The Depth Buffer

Enabling the depth buffer, also called the *z-buffer*, causes OpenGL to eliminate, prior to rendering, parts of objects that are *obscured* (or, *occluded*) by others.

Precisely, a point of an object is not drawn if its projection – think of a ray from that point – toward the viewing face is obstructed by another object. See Figure 2.30(a) for the making of the upper-right annulus of `circularAnnuluses.cpp`: the white disc obscures the part of the red disc behind it (because the projection is orthogonal, the obscured part is exactly






Figure 2.30: (a) The front white disc obscures part of the red one (b) The point A with largest z -value is projected onto the viewing plane so P is red.

the same shape and size as the white disc). This process is called *hidden surface removal* or *depth testing* or *visibility determination*.



Stated mathematically, the result of hidden surface removal in case of orthographic projection is as follows. Suppose that the set of points belonging to drawn objects in the viewing box, with their first two coordinate values particularly equal to X and Y , respectively, is $S = \{(X, Y, z)\}$, where z varies. In other words, S is the set of drawn points lying on the straight line through $(X, Y, 0)$ parallel to the z -axis.

Then only the point (X, Y, Z) of S , with the *largest* z -value, say, Z , lends its color attributes to their shared projection $(X, Y, -near)$ on the viewing face. The implication is that only (X, Y, Z) is drawn of the points in S , the rest obscured.

For example, in Figure 2.30(b), the three points A , B and C , colored red, green and blue, respectively, share the same first two coordinate values, namely, $x = 30$ and $y = 20$. So, all three project to the same point P on the viewing face. As A has the largest z coordinate of the three, it obscures the other two and P , therefore, is drawn red.

The z -buffer itself is a block of memory containing z -values, one per pixel. If depth testing is enabled, then, as a primitive is processed for rendering, the z -value of each of its points – or, more accurately, each of its pixels – is compared with that of the one with the same (x, y) -values currently resident in the z -buffer. If an incoming pixel's z -value is greater, then its RGB attributes and z -value replace those of the current one; if not, the incoming pixel's data is discarded. For example, if the order in which the points of Figure 2.30(b) happen to appear in the code is C , A and B , here's how the color and z -buffer values at the pixel corresponding to P change:

```
draw C; // Pixel corresponding to P gets color blue
```



```

// and z-value -0.5.
draw A; // Pixel corresponding to P gets color red
// and z-value 0.3: A's values overwrite C's.
draw B; // Pixel corresponding to P retains color red
// and z-value 0.3: B is discarded.

```

Next, note in `circularAnnuluses.cpp` the enabling syntax of hidden surface removal so that you can begin to use it:

1. The `GL_DEPTH_BUFFER_BIT` parameter of `glClear(GL_COLOR_BUFFER_BIT | GL_DEPTH_BUFFER_BIT)` in the `drawScene()` routine causes the depth buffer to be cleared.
2. The command `glEnable(GL_DEPTH_TEST)` in the `drawScene()` routine turns hidden surface removal on. The complementary command is `glDisable(GL_DEPTH_TEST)`.
3. The `GLUT_DEPTH` parameter of `glutInitDisplayMode(GLUT_SINGLE | GLUT_RGB | GLUT_DEPTH)` in `main()` causes the depth buffer to be initialized.

Exercise 2.22. (Programming) Draw a bull's eye target as in Figure 2.31 by means of five discs of different colors, sizes and depths.

Section 2.8
THREE DIMENSIONS,
THE DEPTH BUFFER
AND PERSPECTIVE
PROJECTION




Figure 2.31: Bull's eye target.

2.8.2 A Helix and Perspective Projection

We get more seriously 3D next by drawing a spiral or, more scientifically, a helix. A helix, though itself 1-dimensional – drawn as a line strip actually – can be made authentically only in 3-space.

Open `helix.cpp` but don't run it as yet! The parametric equations implemented are

$$x = R \cos t, y = R \sin t, z = t - 60.0, \quad -10\pi \leq t \leq 10\pi \quad (2.2)$$

See Figure 2.32. Compare these with Equation (2.1) for a circle centered at $(0, 0, 0)$, putting $X = 0$ and $Y = 0$ in that earlier equation. The difference is that the helix climbs up the z -axis *simultaneously* as it rotates circularly with increasing t (so, effectively, it coils around the z -axis). Typically, one writes simply $z = t$ for the last coordinate; however, we tack on “-60.0” to push the helix far enough down the z -axis so that it's contained entirely in the viewing box.

Exercise 2.23. Even before viewing the helix, can you say from Equation (2.2) how many times it is supposed to coil around the z -axis, i.e., how many full turns it is supposed to make?

Hint: One full turn corresponds to an interval of 2π along t .








Figure 2.32: Parametric equations for a helix.





Experiment 2.21. Okay, run `helix.cpp` now. All we see is a circle as in Figure 2.33(a)! There’s no sign of any coiling up or down. The reason, of course, is that the orthographic projection onto the viewing face flattens the helix. Let’s see if it makes a difference to turn the helix upright, in particular, so that it coils around the y -axis. Accordingly, replace the statement

```
glVertex3f(R * cos(t), R * sin(t), t - 60.0);
```

in the drawing routine with

```
glVertex3f(R * cos(t), t, R * sin(t) - 60.0);
```

Hmm, not a lot better (Figure 2.33(b))!

End

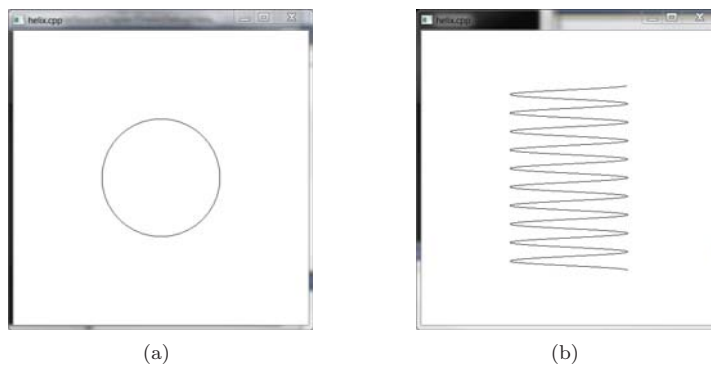







Figure 2.33: Screenshots of `helix.cpp` using orthographic projection with the helix coiling around the: (a) z -axis (b) y -axis.

Because it squashes a dimension, typically, orthographic projection is not suitable for 3D scenes. OpenGL, in fact, provides another kind of projection, called *perspective projection*, more appropriate for most 3D applications. Perspective projection is implemented with a `glFrustum()` call.

Instead of a viewing box, a `glFrustum(left, right, bottom, top, near, far)` call sets up a *viewing frustum* – a frustum is a *truncated pyramid* whose top has been cut off by a plane parallel to its base – in the following manner (see Figure 2.34):

The apex of the pyramid is at the origin. The front face, called the *viewing face*, of the frustum is the rectangle, lying on the plane $z = -near$, whose corners are $(left, bottom, -near)$, $(right, bottom, -near)$, $(left, top, -near)$, and $(right, top, -near)$. The plane $z = -near$ is called the *viewing plane*. The four edges of the pyramid emanating from the apex pass through the four corners of the viewing face. The base of the frustum, which is also the base of the pyramid, is the rectangle whose vertices are precisely where the pyramid’s four edges intersect the $z = -far$ plane. By proportionality with the front vertices, the coordinates of the base vertices are:





Section 2.8
 THREE DIMENSIONS,
 THE DEPTH BUFFER
 AND PERSPECTIVE
 PROJECTION



Figure 2.34: Rendering with `glFrustum()`.



$((far/near) \text{ left}, (far/near) \text{ bottom}, -far),$
 $((far/near) \text{ right}, (far/near) \text{ bottom}, -far),$
 $((far/near) \text{ left}, (far/near) \text{ top}, -far),$
 $((far/near) \text{ right}, (far/near) \text{ top}, -far)$

Values of the `glFrustum()` parameters are typically set so that the frustum lies symmetrically about the z -axis; in particular, *right* and *top* are chosen to be positive, and *left* and *bottom* their respective negatives. The parameters *near* and *far* should both be positive and $near < far$.

Example 2.1. Determine the corners of the viewing frustum created by the call `glFrustum(-15.0, 15.0, -10.0, 10.0, 5.0, 50.0)`.

Answer: By definition, the corners of the front face are $(-15.0, -10.0, -5.0)$, $(15.0, -10.0, -5.0)$, $(-15.0, 10.0, -5.0)$ and $(15.0, 10.0, -5.0)$. The x - and y -values of the vertices of the base (or back face) are scaled from those on the front by a factor of 10 (because $far/near = 50/5 = 10$). The base vertices are, therefore, $(-150.0, -100.0, -50.0)$, $(150.0, -100.0, -50.0)$, $(-150.0, 100.0, -50.0)$ and $(150.0, 100.0, -50.0)$.





Exercise 2.24. Determine the corners of the viewing frustum created by the call `glFrustum(-5.0, 5.0, -5.0, 5.0, 5.0, 100.0)`.

The rendering sequence in the case of perspective projection is a two-step shoot-and-print, similar to orthographic projection. The shooting step again consists of projecting objects within the viewing frustum onto the viewing face, *except that the projection is no longer perpendicular*. Instead, each point is projected along the line joining it to the apex, as depicted by the black dashed lines from the bottom and top of the man in Figure 2.34. Perspective projection causes *foreshortening* because objects farther away from the apex appear smaller (a phenomenon also called *perspective transformation*). For example, see Figure 2.35 where *A* and *B* are of the same height, but the projection *pA* is shorter than the projection *pB*.




Figure 2.35: Section of the viewing frustum showing foreshortening.

Time now to see perspective projection turn on its magic!

Experiment 2.22. Fire up the original `helix.cpp` program. Replace orthographic projection with perspective projection; in particular, replace the projection statement

```
glOrtho(-50.0, 50.0, -50.0, 50.0, 0.0, 100.0);
```

with



```
glFrustum(-5.0, 5.0, -5.0, 5.0, 5.0, 100.0);
```

You can see a real spiral now (Figure 2.36(a)). View the upright version as well (Figure 2.36(b)), replacing

```
glVertex3f(R * cos(t), R * sin(t), t - 60.0);
```

with

```
glVertex3f(R * cos(t), t, R * sin(t) - 60.0);
```



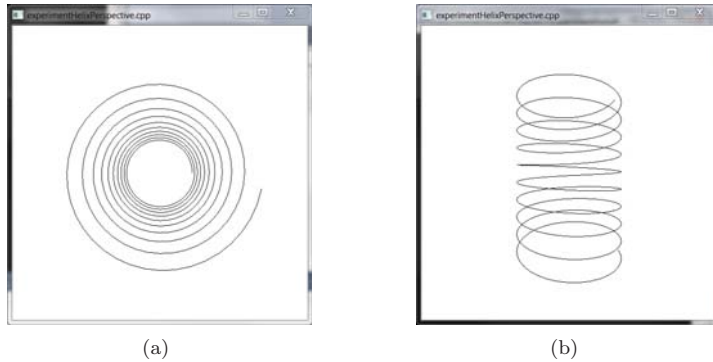




Figure 2.36: Screenshots of `helix.cpp` using perspective projection with the helix coiling up the (a) z -axis (b) y -axis.

Section 2.8
THREE DIMENSIONS,
THE DEPTH BUFFER
AND PERSPECTIVE
PROJECTION

Perspective projection is more realistic than orthographic projection as it mimics how images are formed on the retina of the eye by light rays traveling toward a fixed point. And, in fact, it's precisely foreshortening that cues us humans to the distance of an object.

Remark 2.11. One can think of the apex of the frustum as the location of a *point camera* and the viewing face as its film.



The second rendering step where the viewing face is proportionately scaled to fit onto the OpenGL window is exactly as for orthographic projection. Similar to orthographic projection as well, the scene is clipped to within the viewing frustum by the 6 planes that bound the latter.



Remark 2.12. One might think of orthographic and perspective projections *both* as being along lines of projection convergent to a single point, the center of projection (COP). In the case of perspective projection, this is a regular point with finite coordinates; however, for orthographic projection the COP is a “point at infinity” – i.e., infinitely far away – so that lines toward it are parallel.

Remark 2.13. There do exist 3D applications, e.g., in architectural design, where foreshortening amounts to distortion, so, in fact, orthographic projection is preferred.

Remark 2.14. It's because it captures the image of an object by intersecting rays projected from the object – either orthographically or perspectively – with a plane, which is similar to how a real camera works, that OpenGL is said to implement the *synthetic-camera* model.

Exercise 2.25. (Programming) Continuing from where we were at the end of the preceding experiment, successively replace the `glFrustum()` call as follows, trying in each case to predict the change in the display before running the code:





- (a) `glFrustum(-5.0, 5.0, -5.0, 5.0, 5.0, 120.0)`
- (b) `glFrustum(-5.0, 5.0, -5.0, 5.0, 10.0, 100.0)`
- (c) `glFrustum(-5.0, 5.0, -5.0, 5.0, 2.5, 100.0)`
- (d) `glFrustum(-10.0, 10.0, -10.0, 10.0, 5.0, 100.0)`

Parts (b) and (c) show, particularly, how moving the film forward and back causes the camera to “zoom” in and out, respectively.

Exercise 2.26. Formulate mathematically how hidden surface removal should work in the case of perspective projection, as we did in Section 2.8.1 for orthographic projection.



Figure 2.37: Screenshot of `moveSphere.cpp`.

Experiment 2.23. Run `moveSphere.cpp`, which simply draws a movable sphere in the OpenGL window. Press the left, right, up and down arrow keys to move the sphere, the space bar to rotate it and ‘r’ to reset.

The sphere appears distorted as it nears the periphery of the window, as you can see from the screenshot in Figure 2.37. Can you guess why? Ignore the code, especially unfamiliar commands such as `glTranslatef()` and `glRotatef()`, except for the fact that the projection is perspective.

This kind of *peripheral distortion* of a 3D object is unavoidable in any viewing system which implements the synthetic-camera model. It happens with a real camera as well, but we don’t notice it as much because the field of view when snapping pictures is usually quite large and objects of interest tend to be centered. End

2.9 Drawing Projects



Here are a few exercises to stretch your drawing muscles. The objects may look rather different from what we have drawn so far, but as programming projects aren’t really. In fact, you can probably cannibalize a fair amount of code from earlier programs.

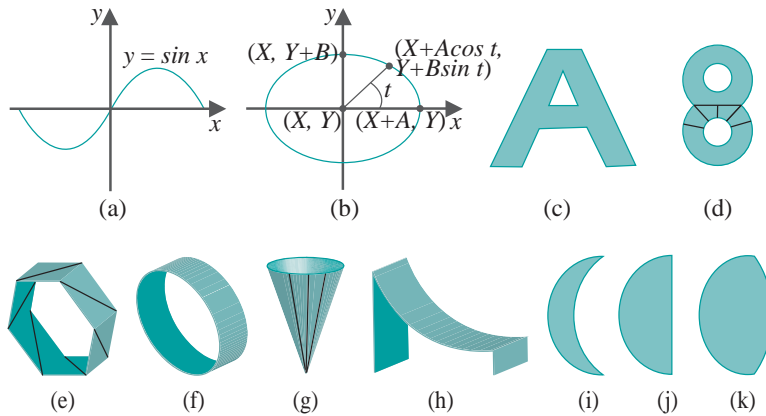


Exercise 2.27. (Programming) Draw a sine curve between $x = -\pi$ and $x = \pi$ (Figure 2.38(a)). Follow the strategy of `circle.cpp` to draw a polyline through a sample from the sine curve.

Exercise 2.28. (Programming) Draw an ellipse. Recall the parametric equations for an ellipse on the xy -plane, centered at (X, Y) , with semi-major axis of length A and semi-minor axis of length B (Figure 2.38(b)):

$$x = X + A \cos t, \quad y = Y + B \sin t, \quad z = 0, \quad 0 \leq t \leq 2\pi$$

Again, `circle.cpp` is the template to use.





Section 2.9
DRAWING PROJECTS

Figure 2.38: Draw these!

Exercise 2.29. (Programming) Draw the letter ‘A’ as a *two-dimensional* figure like the shaded region in Figure 2.38(c). It might be helpful to triangulate it first on graph paper.



Allow the user to toggle between filled and wireframe *a la* the bottom annulus of `circularAnnuluses.cpp`.


Exercise 2.30. (Programming) Draw the number ‘8’ as the 2D object in Figure 2.38(d). Do this in two different ways: (i) drawing 4 discs and using the *z*-buffer and (ii) as a true triangulation, allowing the user to toggle between filled and wireframe. For (ii), a method of dividing the ‘8’ into two triangle strips is suggested in Figure 2.38(d).

Exercise 2.31. (Programming) Draw a ring with cross-section a regular (equal-sided) polygon as in Figure 2.38(e), where a scheme to triangulate the ring in one triangle strip is indicated. Allow the user to change the number of sides of the cross-section. Increasing the number of sides sufficiently should make the ring appear cylindrical as in Figure 2.38(f). Use perspective projection and draw in wireframe.

Exercise 2.32. (Programming) Draw a cone as in Figure 2.38(g) where a possible triangulation is indicated. Draw in wireframe and use perspective projection.

Exercise 2.33. (Programming) Draw a children’s slide as in Figure 2.38(h). Choose an appropriate equation for the cross-section of the curved surface – part of a parabola, maybe – and then “extrude” it as a triangle strip. (If you did Exercise 2.31 then you’ve already extruded a polygon.) Draw in wireframe and use perspective projection.





Remark 2.15. Your output from Exercises 2.31-2.33 may look a bit “funny”, especially viewed from certain angles. For example, the ring viewed head-on down its axis may appear as two concentric circles on a single plane. This problem can be alleviated by drawing the object with a different alignment or, equivalently, changing the viewpoint. In Experiment 2.24, coming up shortly, we’ll learn code for the user to change her viewpoint.

Exercise 2.34. (Programming) Draw in a single scene a crescent moon, a half-moon and a three-quarter moon (Figures 2.38(i)-(k)). Each should be a true triangulation. Label each as well using text-drawing.

2.10 Approximating Curved Objects Once More

Our next 3-space drawing project is a bit more challenging: a hemisphere, which is a 2-dimensional object. We’ll get in place, as well, certain design principles which will be expanded in Chapter 10 dedicated to drawing (no harm starting early).

Remark 2.16. A hemisphere is a 2-dimensional object because it is a surface. Recall that a helix is 1-dimensional because it’s line-like. Now, both hemisphere and helix need 3-space to “sit in”; they cannot do with less. For example, you could sketch either on a piece of paper (2-space) but it would not be the real thing. On the other hand, a circle – another 1D object – does sit happily in 2-space.



Consider a hemisphere of radius R , centered at the origin O , with its circular base lying on the xz -plane. Suppose the spherical coordinates of a point P on this hemisphere are a longitude of θ (measured clockwise from the x -axis when looking from the plus side of the y -axis) and a latitude of ϕ (measured from the xz -plane toward the plus side of the y -axis). See Figure 2.39(a). The Cartesian coordinates of P are by elementary trigonometry




$$(R \cos \phi \cos \theta, R \sin \phi, R \cos \phi \sin \theta)$$

The range of θ is $0 \leq \theta \leq 2\pi$ and of ϕ is $0 \leq \phi \leq \pi/2$.

Exercise 2.35. Verify that the Cartesian coordinates of P are as claimed. *Suggested approach:* From the right-angled triangle OPP' one has $|PP'| = R \sin \phi$ and $|OP'| = R \cos \phi$. $|PP'|$ is the y -value of P . Next, from right-angled triangle $OP'P''$ find $|OP''|$ and $|P'P''|$, the x - and z -values of P , respectively, in terms of $|OP'|$ and θ .

Sample the hemisphere at a mesh of $(p + 1)(q + 1)$ points P_{ij} , $0 \leq i \leq p$, $0 \leq j \leq q$, where the longitude of P_{ij} is $(i/p) * 2\pi$ and its latitude $(j/q) * \pi/2$. In other words, $p + 1$ longitudinally equally-spaced points are chosen along each of $q + 1$ equally-spaced latitudes. See Figure 2.39(b), where $p = 10$ and





Section 2.10
APPROXIMATING
CURVED OBJECTS ONCE
MORE

Figure 2.39: (a) Spherical and Cartesian coordinates on a hemisphere (b) Approximating a hemisphere with latitudinal triangle strips.

$q = 4$. The sample points P_{ij} are not all distinct. In fact, $P_{0j} = P_{pj}$, for all j , as the same point has longitude both 0 and 2π ; and, the point P_{iq} , for all i , is identical to the north pole, which has latitude $\pi/2$ and arbitrary longitude.

The plan now is to draw one triangle strip with vertices at

$$P_{0,j+1}, P_{0j}, P_{1,j+1}, P_{1j}, \dots, P_{p,j+1}, P_{pj}$$

for each j , $0 \leq j \leq q - 1$, for a total of q triangle strips. In other words, each triangle strip takes its vertices alternately from a pair of adjacent latitudes and, therefore, approximates the circular band between them. Figure 2.39(b) shows one such strip. The stack of all q triangle strips approximates the hemisphere itself.

Experiment 2.24. Run `hemisphere.cpp`, which implements exactly the strategy just described. You can verify this from the snippet that draws the hemisphere:

```
for(j = 0; j < q; j++)
{
    // One latitudinal triangle strip.
    glBegin(GL_TRIANGLE_STRIP);
    for(i = 0; i <= p; i++)
    {
        glVertex3f(R * cos((float)(j+1)/q * PI/2.0) *
            cos(2.0 * (float)i/p * PI),
            R * sin((float)(j+1)/q * PI/2.0),
            R * cos((float)(j+1)/q * PI/2.0) *
            sin(2.0 * (float)i/p * PI));
        glVertex3f(R * cos((float)j/q * PI/2.0) *
            cos(2.0 * (float)i/p * PI),
            R * sin((float)j/q * PI/2.0),
            R * cos((float)j/q * PI/2.0) *
            sin(2.0 * (float)i/p * PI));
    }
}
```








Figure 2.40: Screenshot of `hemisphere.cpp`.





```

        sin(2.0 *(float)i/p * PI));
    }
    glEnd();
}

```

Increase/decrease the number of longitudinal slices by pressing ‘P/p’. Increase/decrease the number of latitudinal slices by pressing ‘Q/q’. Turn the hemisphere about the axes by pressing ‘x’, ‘X’, ‘y’, ‘Y’, ‘z’ and ‘Z’. See Figure 2.40 for a screenshot. **End**

Experiment 2.25. Playing around a bit with the code will help clarify the construction of the hemisphere:

- (a) Change the range of the hemisphere’s outer loop from

```
for(j = 0; j < q; j++)
```

to

```
for(j = 0; j < 1; j++)
```

Only the bottom strip is drawn. The keys ‘P/p’ and ‘Q/q’ still work.

- (b) Change it again to

```
for(j = 0; j < 2; j++)
```

Now, the bottom two strips are drawn.

- (c) Reduce the range of both loops:

```

for(j = 0; j < 1; j++)
...
    for(i = 0; i <= 1; i++)
...

```

The first two triangles of the bottom strip are drawn.



- (d) Increase the range of the inner loop by 1:



```

for(j = 0; j < 1; j++)
...
    for(i = 0; i <= 2; i++)
...

```

The first four triangles of the bottom strip are drawn. **End**





There's syntax in `hemisphere.cpp` – none to do with the actual making of the hemisphere – which you may be seeing for the first time. The command `glTranslatef(0.0, 0.0, -10.0)` is used to move the hemisphere, drawn initially centered at the origin, into the viewing frustum, while the `glRotatef()` commands turn it. We'll explain these so-called *modeling transformations* in Chapter 4 but you are encouraged to experiment with them even now as the syntax is fairly intuitive. The set of three `glRotatef()`s, particularly, comes in handy to re-align a scene.

Section 2.10
APPROXIMATING
CURVED OBJECTS ONCE
MORE

Exercise 2.36. (Programming) Modify `hemisphere.cpp` to draw:

- (a) the bottom half of a hemisphere (Figure 2.41(a)).
- (b) a 30° slice of a hemispherical cake (Figure 2.41(b)). Note that simply reducing the range of the inner loop of `hemisphere.cpp` produces a slice of cake without two sides and bottom, so these have to be added in separately to close up the slice.

Make sure the 'P/p/Q/q' keys still work.

Exercise 2.37. (Programming) Just to get you thinking about animation, which we'll be studying in depth soon enough, guess the effect of replacing `glTranslatef(0.0, 0.0, -10.0)` with `glTranslatef(0.0, 0.0, -20.0)` in `hemisphere.cpp`. Verify.

And, here are some more things to draw.

Exercise 2.38. (Programming) Draw the objects shown in Figure 2.42. Give the user an option to toggle between filled and wireframe renderings.




Figure 2.41: (a) Half a hemisphere (b) Slice of a hemisphere.








Figure 2.42: More things to draw.

A suggestion for the football, or ellipsoid, is to modify `hemisphere.cpp` to make half of an ellipsoid (a hemi-ellipsoid?). Two hemi-ellipsoids back to back would then give a whole ellipsoid.

Remark 2.17. Filled renderings of 3D scenes, even with color, rarely look pleasant in the absence of lighting. See for yourself by applying color to 3D objects you have drawn so far (remember to invoke a `glPolygonMode(*, GL_FILL)` call). For this reason, we'll draw mostly wireframe till Chapter 11,





which is all about lighting. You'll have to bear with this. Wireframe, however, fully exposes the geometry of an object, which is not a bad thing when one is learning object design.

2.11 An OpenGL Program End to End

Of `square.cpp`, in particular, we have touched on almost every command which is functional from a graphics points of view. However, let's run over the whole program to see all that goes into making OpenGL code tick.

We start with `main()`:



1. `glutInit(&argc, argv)` initializes the FreeGLUT library. FreeGLUT, successor to GLUT (OpenGL Utility Toolkit), is a library of calls to manage a window holding OpenGL contexts (the reason such a separate library is needed is that OpenGL itself is only a library of graphics calls).



2. `glutInitContextVersion(4, 3)`
`glutInitContextProfile(GLUT_COMPATIBILITY_PROFILE)`

ask FreeGLUT to provide an OpenGL 4.3 context which is backward-compatible in that legacy commands are implemented. This, for example, allows us to draw with the `glBegin()-glEnd()` operations from OpenGL 2.1, which do not belong in the core profile of OpenGL 4.3.

Important: If your graphics card doesn't support OpenGL 4.3 then the program may compile but not run as the system is unable to provide the context asked. What you might do in this case is thin the context by replacing the first line above with `glutInitContextVersion(3, 3)`, or even `glutInitContextVersion(2, 1)`, instead. Of course, then, programs using later-generation calls will not run, but you should be fine early on in the book.

3. `glutInitDisplayMode(GLUT_SINGLE | GLUT_RGBA)` asks the OpenGL context to support a single-buffered frame, each pixel having red, green, blue and alpha values.
4. `glutInitWindowSize(500, 500)`
`glutInitWindowPosition(100, 100)`
as we have already seen, set the size of the OpenGL window and the location of its top left corner on the computer screen.
5. `glutCreateWindow("square.cpp")` creates the window (precisely, the rendering context) with the specified string parameter as title.
6. `glutDisplayFunc(drawScene)`
`glutReshapeFunc(resize)`





```
glutKeyboardFunc(keyInput)
```

register the routines to call – so-called *callback routines* – when the OpenGL window is to be drawn, when it is resized (and first created), and when keyboard input is received, respectively.

7. `glewExperimental = GL_TRUE`
`glewInit()`
initializes GLEW (the OpenGL Extension Wrangler Library) which handles the loading of OpenGL extensions, with the switch set so that extensions implemented in even pre-release drivers are exposed.
8. `setup()` invokes the initialization routine.
9. `glutMainLoop` begins the event-processing loop, calling registered callback routines as needed.

We have already seen that the only command in the initialization routine `setup()`, namely, `glClearColor(1.0, 1.0, 1.0, 0.0)`, specifies the clearing color of the OpenGL window.

The callback routine to draw the OpenGL window is:

```
void drawScene(void)
{
    glClearColor(GL_COLOR_BUFFER_BIT);

    glColor3f(0.0, 0.0, 0.0);

    // Draw a polygon with specified vertices.
    glBegin(GL_POLYGON);
    ...
    glEnd();

    glFlush();
}
```



The first command clears the OpenGL window to the specified clearing color, in other words, paints in the background color. The next command `glColor3f()` sets the foreground, or drawing, color, which is used to draw the polygon specified within the `glBegin()-glEnd()` pair (we have already examined this polygon carefully). Finally, `glFlush()` forces the prior draw calls to actually execute, which, in this case, means the polygon is drawn.



The callback routine when the OpenGL window is resized, and first created, is `void resize(int w, int h)`. The window manager supplies the width `w` and height `h` of the resized OpenGL window (or, initial window, when it is first created) as parameters to the `resize` routine.

The first command

```
glViewport(0, 0, w, h);
```

Section 2.11
AN OpenGL PROGRAM
END TO END





Chapter 2
ON TO OPENGL AND
3D COMPUTER
GRAPHICS

of `square.cpp`'s `resize` routine specifies the rectangular part of the OpenGL window in which actual drawing is to take place; with the given parameters it is the entire window. We'll be looking more carefully into `glViewport()` and its applications in the next chapter.

The next three commands

```
glMatrixMode(GL_PROJECTION);  
glLoadIdentity();  
glOrtho(0.0, 100.0, 0.0, 100.0, -1.0, 1.0);
```

activate the projection matrix stack, place the identity matrix at the top of this stack, and then multiply the identity matrix by the matrix corresponding to the final `glOrtho()` command, effectively setting up the viewing box of `square.cpp` described in Section 2.2.

The final two commands

```
glMatrixMode(GL_MODELVIEW);  
glLoadIdentity();
```

of the `resize` routine activate the modelview matrix stack and place the identity matrix at the top in readiness for modelview transformation commands in the drawing routine, of which there happen to be none in `square.cpp`. We'll be learning much more about OpenGL's matrix stacks and how they are at the heart of its operation in Chapters 4 and 5.



The callback routine to handle ASCII keys is `keyInput(unsigned char key, int x, int y)`. When an ASCII key is pressed it is passed in the parameter `char` to this callback routine, as is the location of the mouse in the parameters `x` and `y`. As is easy to see, all that `keyInput` of `square.cpp` does is terminate the program when the escape key is pressed. In the next chapter we'll see callback routines to handle non-ASCII keys, as well as interaction via the mouse.



As the reader might well guess, the guts of an OpenGL program are in its drawing routine. Interestingly, the initialization routine often pulls a fair load, too, because one would want to locate there tasks that need to be done once at start-up, e.g., setting up data structures. In fact, it's a common beginner's mistake to place initialization chores in the drawing routine, as the latter is invoked repeatedly if there is animation, leading to inefficiency.

The other routines, such as `main()` and the interactivity and reshape callbacks, are either standard and pretty much transferable from program to program (e.g., `main()`), or easy to code.

2.12 Summary, Notes and More Reading

In this chapter we began the study of 3D CG, looking at it through the "eyes" of OpenGL. OpenGL itself was presented to the extent that the reader acquires functional literacy in this particular API. The drawing primitives were probably the most important part of the API's vernacular.







We discovered as well how OpenGL functions as a state machine, attributes such as color defining the current state. Moreover, we learned that quantifiable attribute values, e.g., RGB color, are typically interpolated from the vertices of a primitive throughout its interior. We saw that OpenGL clips whatever the programmer draws to within a viewing volume, either a box or frustum.



Beyond acquaintance with the language, we were introduced as well to the synthetic-camera model of 3D graphics, which OpenGL implements via two kinds of projection: orthographic and perspective. This included insights into the world coordinate system, the viewing volume – box or frustum – which is the stage in which all drawings are made, the shoot-and-print rendering process to map a 3D scene to a 2D window, as well as hidden surface removal. We made a first acquaintance as well with another cornerstone of 3D graphics: the technique of simulating curved objects using straight and flat primitives like line segments and triangles.

Historically, OpenGL evolved from SGI's IRIS GL API, which popularized the approach to creating 3D scenes by drawing objects in actual 3-space and then rendering them to a 2D window by means of a synthetic camera. IRIS GL's efficiently pipelined architecture enabled high-speed rendering of animated 3D graphics and, consequently, made possible as well real-time interactive 3D. The ensuing demand from application developers for an open and portable (therefore, platform-independent) version of their API spurred SGI to create the first OpenGL specification in 1992, as well as a sample implementation. Soon after, the OpenGL ARB (Architecture Review Board), a consortium composed of a few leading companies in the graphics industry, was established to oversee the development of the API. Stewardship of the OpenGL specification passed in 2006 to the Khronos Group, a member-funded industry consortium dedicated to the creation of open-standard royalty-free API's. (That no *one* owns OpenGL is a good thing.) The canonical, and very useful, source for information about OpenGL is its own home page [103].

Microsoft has a non-open Windows-specific 3D API – Direct3D [90, 141] – which is popular among game programmers as it allows optimization for the pervasive Windows platform. However, outside of the games industry, where it nonetheless competes with Direct3D, and leaving aside particular application domains with such high-quality rendering requirements that ray tracers are preferred, by far the dominant graphics API is OpenGL.

It's safe to say that OpenGL is the de facto standard 3D graphics API. A primary reason for this, other than the extremely well-thought-out design which it had from inception – initial versions of Direct3D in contrast were notoriously buggy and hard to use – is OpenGL's portability. With their recent versions, though, OpenGL and Direct3D seem to be converging, at least in functionality (read an interesting comparison in Wikipedia [27]). It's worth knowing as well that, despite its intended portability, OpenGL can take advantage of platform-specific and card-specific capabilities via





Chapter 2
ON TO OPENGL AND
3D COMPUTER
GRAPHICS

so-called extensions, at the cost of clumsier code.

An unofficial clone of OpenGL, Mesa 3D [93], which uses the same syntax, was originally developed by Brian Paul for the Unix/X11 platform, but there are ports now to other platforms as well.

Perhaps the best reason for OpenGL to be *the* API of choice for students of 3D computer graphics is – and this is a consequence of its almost universal adoption by the academic, engineering and scientific communities – the sheer volume of learning resources available. Not least among these is the number of textbooks that teach computer graphics with the help of OpenGL. Search amazon.com with the keywords “computer graphics opengl” and you’ll see what we mean. Angel [2], Buss [22], Govil-Pai [60], Hearn & Baker [69], Hill & Kelley [72] and McReynolds & Blythe [92] are some introductions to computer graphics via OpenGL that the author has learned much from.

In case the reader prefers not to be distracted by code, here are a few API-independent introductions: Akenine-Möller, Haines & Hoffinan [1], Foley et al. [47, 48] (the latter being an abridgment of the former), Shirley & Marschner [130], Watt [147] and Xiang & Plastock [154]. Keeping different books handy in the beginning is a good idea as, often, when you happen to be confused by one author’s presentation of a topic, simply turning to another for help on just that may clear the way.

With regard to the code which comes with this book, we don’t make much use of OpenGL-defined data types, which are prefixed with `GL`, e.g., `GLsizei`, `GLint`, etc., though the red book advocates doing so in order to avoid type mismatch issues when porting. Fortunately, we have not yet encountered a problem in any implementation of OpenGL that we’ve tried.

In addition to the code that comes with this book, the reader should try to acquire OpenGL programs from as many other sources as possible, as an efficient way to learn the language – any language as a matter of fact – is by modifying live code. Among the numerous sources on the web – there are pointers to several coding tutorials at the OpenGL site [103] – special mention must be made of Jeff Molofee’s excellent tutorials at NeHe Productions [102], covering a broad spectrum of OpenGL topics, each illustrated with a well-documented program. The book by Wright, Lipchak & Haemel [129] is specifically about programming OpenGL and has numerous example programs. The red book comes with example code as well. Incidentally, in addition to the somewhat bulky red and blue books, a handy reference manual for OpenGL is Angel’s succinct primer [3].

Hard-earned wisdom: Write experiments of your own to clarify ideas. Even if you are sure in your mind that you understand something, do write a few lines of code in verification. As the author has repeatedly been, you too might be surprised!

