Transfer functions emphasize regions in the volume by assigning color and opacity to data values. Histograms are useful for analyzing which ranges of values are important in the data. In general, histograms show the distribution of data values and other related data measures. A 1D histogram is created by dividing up the value range into a number of bins. Each bin contains the number of voxels within the lower and upper bounds assigned to the bin. By examining the histogram, one can see which values are frequent in the data. Histograms, however, do not show the spatial distribution of the samples in the volume.

The output of the data-processing step is a set of textures that are downloaded to the GPU in a later stage. It is sometimes more efficient to combine several textures into a single texture. For example, to reduce the cost of texture lookup and interpolation, the value and normalized gradient textures are usually stored and used together in a single RGBA texture.

39.4.2 Proxy Geometry

During the rendering stage, images of the volume are created by drawing the proxy geometry in sorted order. When the data set is stored in a 3D texture, view-aligned planes are used for slicing the bounding box, resulting in a set of polygons for sampling the volume. Algorithm 39-2 computes the proxy geometry in view space by using the modelview matrix for transforming vertices between the object and view coordinate systems. Proxy polygons are tessellated into triangles, and the resulting vertices are stored in a vertex array for more efficient rendering.

Figure 39-6 illustrates Algorithm 39-2 with two slice polygons. The first polygon contains three vertices, the second is composed of six vertices and is tessellated into six triangles.