5 Working with Visualizations

Once visualizations of datasets are up on the screen, they can be modified to bring out features of the data. Modifications include changing the color and surface properties, deforming by a scalar field, and showing axes and bounds. This chapter describes techniques for modifying visualization objects. A later chapter will describe how to add other elements to the visualization, such as annotations, legends, and drawn objects.

Refining visualizations

SciAn uses defaults and heuristics so that it can bring up the first visualization of a dataset quickly and easily. Most of the process of using SciAn involves progressively refining the visualizations to produce the desired result.

When you bring up the control panel of a visualization object, there will be a group of icon buttons at the right. Each icon represents a set of attributes of the visualization object that can be modified through controls. Most modifications to visualizations are done by changing the settings of these controls. It is best to maintain a playful attitude and experiment with the controls to see how they affect the visualization. Often, the most interesting effects are the unexpected result of trying out various combinations of settings.

Working with color

Color is a very important component of an effective visualization. SciAn provides a variety of tools for coloring visualization objects.

Full color and color map modes

Visualization windows generally have two color modes: full color and color map. Which color mode is currently being used is controlled by the setting within the Renderer control panel. Not all workstations can do both color modes. If your workstation cannot do one of the color modes, that button appears gray.

The two color modes differ in the way the workstation hardware is used to produce colored images. Full color mode uses the 24-bit RGB capabilities of the workstation as well as alpha planes, if any. Color map mode uses the color lookup table of the workstation. Which one to use depends on the data you are visualizing.

In full color mode, objects are fully lit and shaded. Full color mode allows the maximum amount of flexibility over the appearance of the visualization. This mode generally requires more hardware bit planes than color map mode.

In color map mode, there is either imitation light shading or none, depending on the settings. Lights are completely ignored in color map mode.

The color mode also affects the way that colors are interpolated. SciAn uses the Gouraud shading capabilities of the hardware to color objects smoothly. In full color mode, colors are interpolated along a straight line in RGB space. In color map mode, colors are interpolated along the color index, which is related to field value. For datasets which change drastically from point to point, color map mode may give a better picture of the data, especially if the color is more important than seeing the shape, such as with 2-D visualizations.

When a window is in color map mode, SciAn tries to allocate colors from the color table as well as it can. However, there are only a finite number of colors available. When viewing many datasets, SciAn can run out of colors on systems which do not have enough color map bit planes.

Experiment with the color mode to find what works best for the data. Most of the time, full RGB mode is the better choice. However, brightly colored flat objects may benefit from color map mode.

Base color and lights

The resulting color of an object is the result of a complex interaction between the base color of a visualization object and the light sources shining upon it. SciAn controls the base color and the light sources, and the workstation hardware calculates the resulting color using Phong lighting and Gouraud shading, treating each RGB color component separately. This process is well explained in Foley [1982].

Because of the differences between models of hardware, images may appear different on different workstations. For example, the Silicon Graphics workstations do automatic gamma correction on the resulting color, but the IBM RS/6000 workstations do not. SciAn corrects base color in an attempt to minimize this effect, but you may get darker images on the IBM, which may require more and brighter light sources to illuminate.

When designing a visualization, keep the interaction of lights and colors in mind. A blue light source, for example, will not illuminate a red object.

Color controls

Most visualization objects can be given color. If a visualization object has color, there will be a Color icon on the right side of its control panel. Click on this icon to bring up the color controls. Figure 5-1 shows an example color control panel.

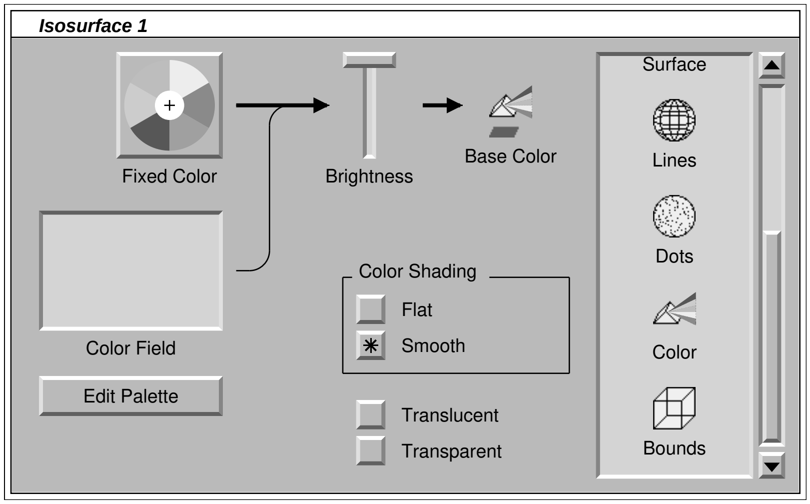

Figure 5-1. *Color control panel*

Visualization objects can be colored either by a fixed color or by a scalar field. The dataflow diagram at the top left of the color control panel shows this process. Either a fixed color, represented by a color wheel, or a color field, shown as an icon in a corral, goes through a brightness control and results in the color of the object.

To change the fixed color, click in the color wheel at the color you desire. If you hold down the Shift key, the colors are constrained to 50% and 100% saturation colors. Double-clicking near one of the six pure colors or their 50% saturation equivalents will snap to that color.

To change the dataset used to colorize the data, drag the icon of the desired dataset into the corral.

To choose whether the color is taken from the fixed color or the scalar field, click on the appropriate portion of the switch between the two and the brightness control.

Raise and lower the brightness slider to change the brightness of the resulting surface.

The colors chosen are taken from the palette of the color field combined with the values in the field. You can show the palette controls of the dataset by clicking on the icon and pressing the Edit Palette button. (See "Color palettes" below.)

Color map mode

When the visualization window is in color map mode, it is not possible to get all of the nuances and subtleties of color shading, but SciAn will try its best. When a fixed color is used, SciAn will produce a ramp palette based on the surface normal to imitate shading. When a color field is used, SciAn will use the colors in the palette directly.

At the bottom right of the control panel are some buttons which control the way color is used. The Color Shading radio button group controls the way the surface is shaded. When the Smooth button is checked, the colors of each polygon in the surface will be smoothly interpolated using the computer's graphics hardware and software. When the Flat button is checked, each polygon will only show a single solid color. These settings interact in interesting ways with the Light Shading radio button group in the Surface control panel, described later. Experiment with different settings of these two groups of controls.

The Translucent check box causes the visualization to be drawn using screen door translucency. This simulation of transparency involves drawing the polygons using a dither pattern onto only half the pixels within the polygon. This capability is available on most supported workstations, although it may slow drawing on some. The Transparent check box causes the visualization object to be drawn with blending transparency. This check box will only appear if the workstation has alpha blending. Both transparency and translucency are useful effects in their own right and can be combined as well.

Hint *When you use blending transparency, you will see the back side of some objects through the front, so objects will appear brighter. Turn down the brightness a bit to bring out details. It is also a good idea to turn down the brightness on brightly colored objects behind the transparent object.*

Color palettes

A **color palette** is a mapping of a range of numeric values onto a range of colors. Color palettes are used to find the base color of a visualization object when it is colored by a field.

Color palettes belong to the scalar field they control, not to the visualization object. One scalar field can only have a single color palette. Using the same color palette for all color operations on a field helps ensure consistency in the visual representation.

Color palettes may have a variable number of colors. They specify minimum and maximum scalar values, which are used to determine which of the colors to use. Three colors are handled specially: the underflow and overflow colors, which are used when the scalar value is below the minimum or above the maximum, and the missing data color, which is used in cases of missing data. The remaining colors are mapped linearly onto the values between the minimum and maximum.

Figure 5-2 shows the Palette control panel. Occupying most of the window is the color bar control. At the bottom is a bar display of the entire color palette showing underflow, overflow, and missing data colors as well as the range of colors mapped on to the data values. Above the display of the full colors are separate bars that display the components of each color.

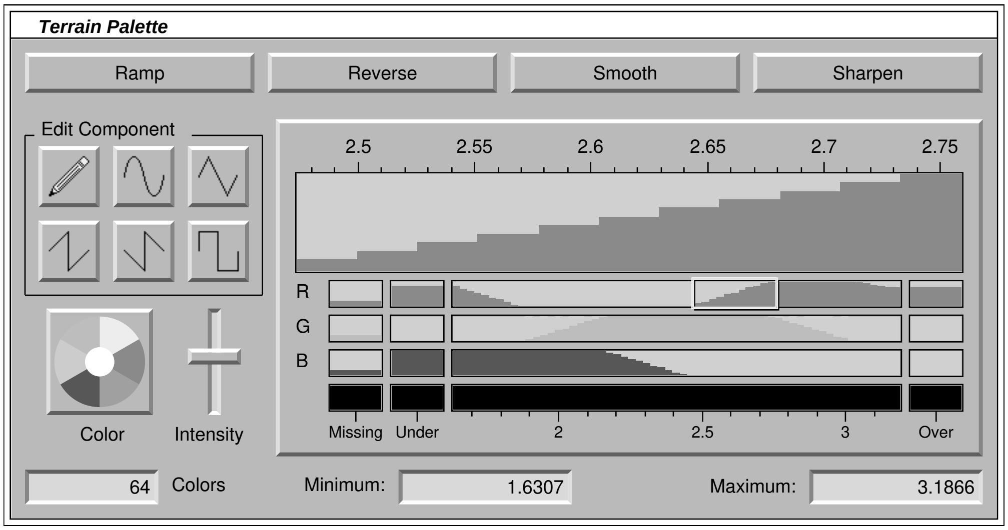

Figure 5-2. *Palette control panel*

You can select a range of colors in any of the four bars by pressing and dragging through the range you want to select. The selected range will be expanded to fill the color readout at the top of the color bar control. This readout shows a more detailed display of the selected range and is also used for the Edit Component tools.

Below the color bar are two text boxes giving the minimum and maximum data values that are mapped onto the colors in the range. You can change either of these values by changing the text and pressing the Enter key.

At the bottom left of the control window is a text box that shows the total number of colors in the control panel. Change this value and press the Enter key to change the number of colors. If you increase the number of colors, it is a good idea to select the entire range and press the Smooth button a few times to smooth the colors out.

At the lower left are a hue/saturation color wheel and an intensity slider. The hue/saturation control will force all of the colors in the selected range to have the hue and saturation you have clicked on. The intensity control will change the intensity of the selected colors. If a range of full colors is selected, this control will change the value without affecting the hue and saturation. If a range of a color component is selected, this will change the amplitude of that component.

At the top are four buttons that do special effects on the selected range. All simple functions work on either a single component or all three components at once, depending on which range is selected. The Ramp button takes the colors defined at the beginning and end of the selected range and generates a smooth increasing or decreasing ramp of colors between them. The Reverse button reverses the selected range of colors. The Smooth button smoothes out abrupt changes in the range of colors. You can press it again and again to smooth even more. The Sharpen button exaggerates changes in the color. This button amplifies sampling variations as well.

Another special effect, Ruffle, is available on the Color menu. This reverses short sections of the selected colors, producing a contour-like effect.

To the left of the color bar control are six icon buttons containing Edit Component tools. These are active whenever a range of a single component is selected and affect the magnified view of the component in the readout at the top of the color bar. The top left component is the free-form edit tool. Click on it and then draw in the component readout with the left mouse button to produce whatever shape of the color function you like.

The other five tools are function generators which create periodic functions in the color range. They are inspired by the ICARE system by Donna Cox [Cox 1988]. By carefully tuning the periodic functions to variations in data, it is possible to bring out variations at a particular scale. By tuning different components to different scales, both large scale and small scale variations can be seen at the same time.

The icons within the tool buttons show the shapes of the waves they create. The top center tool is a sine wave tool, which produces smooth, gradual variations in color. The top right tool is a triangle wave tool, which produces results similar to the sine wave tool with slightly sharper definition. The bottom left and center tools are the sawtooth and reverse sawtooth tools. These produce a contour-like effect similar to the ruffle effect, especially when used on Value in the HSV color model. The bottom right tool is the square wave tool, which produces sharp square waves.

When any of the function generator tools is selected, the component readout is immediately filled with the periodic function. A yellow selection rectangle with four control handles appears around a single period of the function. Click on the top or bottom handles and drag up or down to control the amplitude of the function. Click on the left or right handles and drag side to side to control the period and phase of the function. Hold down the Shift key to constrain motion to a round scalar value. Click in the center of the box and drag side to side to control the phase of the function without changing its period.

Color models

SciAn provides several color models for editing colors. The RGB color model, which represents each color as a linear combination of red, green, and blue, is the most familiar and is used by default in SciAn. The HSV color model represents color with hue, saturation, and "value" or intensity. This is the same approach used by the hue/saturation color wheel and intensity slider. The HLS color model represents color with hue, lightness, and saturation. Hue and saturation are the same as in the HSV model. Lightness is the amount of white or black mixed in with the color. The YIQ color model is the model used by the NTSC television system used in America and Japan.

You can change the color model by using the Color Model submenu in the Color menu.

Color models

Which color model is in use only affects interaction through the color bar. Internally, colors are always kept in the RGB color model, the native color model of most workstations. If you switch from the YIQ model to the RGB model and back again, you may notice that the display in the color bar may not match what it was before. This is because not all possible combinations of Y, I, and Q represent valid colors.

Saving palettes

From within a palette control panel, you can choose Save Palette from the File or Color menu. This will save the settings of the control panel in a special file in the directory of files that contain the dataset. Whenever you open up that dataset from that directory again, the palette will automatically be loaded.

Gray scale

Gray scale images are sometimes more useful than color images in specific applications. SciAn handles gray as a special case of full color. Simply use the HSV color model, set the saturation (S) to zero, and work with the value (V) exclusively. You can even use a hybrid approach, where most of the values are represented by gray and a specific feature of interest is represented in full color.

Undoing changes

Two items in the Color menu allow you to undo changes. The Revert to Original item reverts the entire palette to its original state. The Keep Changes item keeps the current changes. Subsequent use of Revert to Original will revert to this changed palette.

Surfaces and shading

Most visualization objects produce surfaces, which are shaded based on the color, the surface properties, and the lights in the space. Chapter 4 explained how to use lights. This section explains how to adjust the attributes of surfaces.

Surface controls

All visualizations which have a surface, such as meshes and isosurfaces, have a Surface icon in the control panel. Click on this icon to bring up controls for the surface properties. An example is shown in Figure 5-3.

The Highlights control at the top left corner controls the specular highlights of the object. Move the pointer in the control to the desired combination of brightness and sharpness. The four balls at the corners give a rough idea what the highlights will look like. At the top right, the highlights are bright and sharp. At the top left, the highlights are bright and diffuse. As the pointer is brought down, the highlights become dimmer and dimmer.

Note

The vertical axis controls the specularity from 0 to 1. The horizontal axis controls the Phong exponent from 30 to 128.

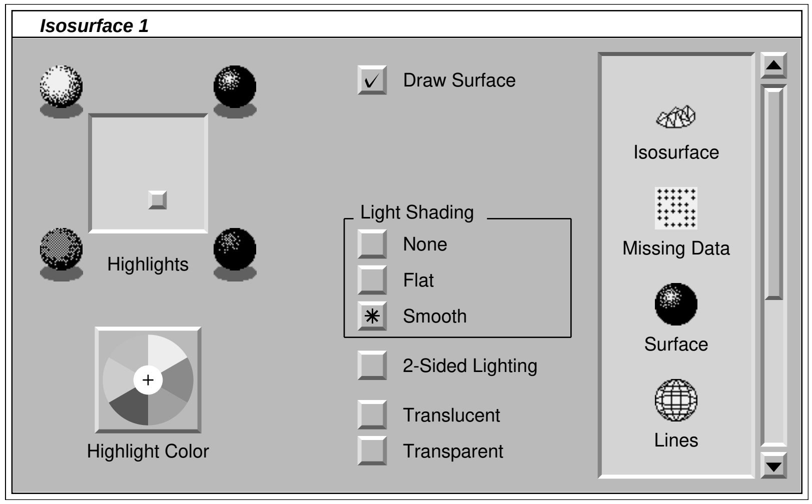

Figure 5-3. *Surface control panel*

The Highlight Color color wheel controls the color of the highlights. White highlights are the most commonly used and usually the most realistic, but other highlight colors can produce interesting effects, especially in colored objects.

At the top right is a Draw Surface check box. When it is checked, the surface is drawn. Unchecking it is an easy way to turn off the surface. When the surface is turned off, lines and dots may still be shown.

Below are three radio buttons that control the light shading used on the object. When the None button is down, the object is not lit at all, and colors appear exactly as the base color of the object. This is useful for planar objects such as meshes where the shape is not as important as the color. When the Flat button is down, the object is lit the same on all vertices of each polygon in the object. When the Smooth button is down, the object is lit smoothly. This produces the most realistic effect.

The 2-Sided Lighting check box determines whether 2-sided lighting is used to illuminate the scene. Normally it is off, and lights from only the top side of surfaces affect the lighting of the surface. When 2-sided lighting is on, each light illuminates the surfaces from above and below. This is an easy way to brighten up an object. When 2-sided lighting is used, twice the number of lights specified are actually used to illuminate the scene.

The Translucent and Transparent check boxes at the bottom right control the translucency and transparency of the object. These are the same controls that appear in the Color controls. For more information, see the earlier section, "Color controls."

Geometric surfaces

Visualization objects are converted into polygons and polygon meshes before they are displayed. The process is automatic and can usually be done without loss in quality. However, in some cases, such as when spheres and cylinders are displayed, polygons can only approximate the shape of the original object. For these objects, there are options to change the quality of the approximation within the Geometry control panel

Every visualization object which produces spheres, cylinders, or conical frusta has a Geometry icon in its control panel. Clicking on this icon brings up a set of controls as shown in Figure 5-4.

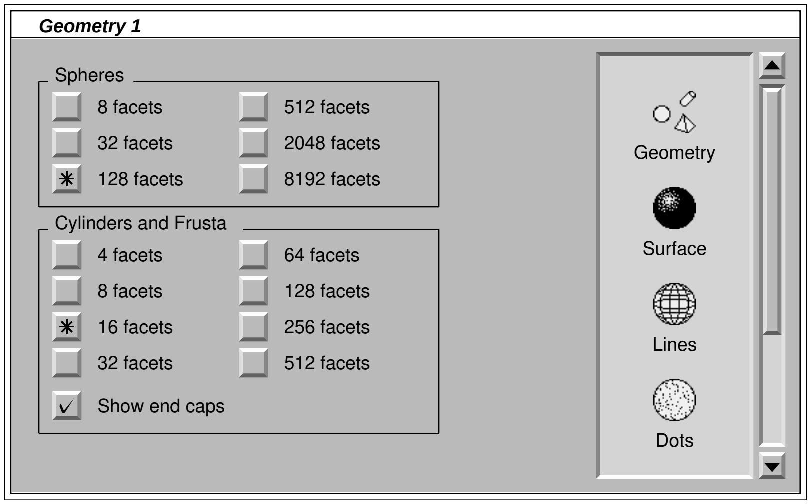

Figure 5-4. *Geometry control panel*

The control panel contains two sets of controls, one for spheres, and one for both cylinders and conical frusta. In each set, a group of radio buttons controls the number of facets, or polygons, used to approximate the object. The more facets used, the higher the quality and the slower the speed.

The sphere approximation is very sensitive to the number of facets. 128 facets provides and adequate approximation for interactive work, but the results are rough. At 2048 facets, the sphere looks fairly smooth, but there are still artifacts with specular reflection. 8192 facets provides a very nice result, but it is slow.

Because they only curve around one dimension, not two, cylinders and frusta are much more forgiving than spheres. As few as 16 facets provide a convincing imitation. 256 facets are enough for high-quality work.

Deformation

Some visualization objects, such as meshes and contours, can be deformed. The deformation works by pushing each vertex of the object a certain amount in the direction of its surface normal. After the object is deformed, surface normals are calculated again.

An example of this is the Mesh visualization object. This object produces a color mesh from a dataset with two topological dimensions. If the dataset has only two spatial dimensions, the mesh is calculated flat, spreading out in the X and Y axes with

Z equal to 0. The surface normal is then 1.0 in the positive Z direction. Deforming the object by the field used to generate the mesh deforms each point by the field at that point in the Z direction. The net result, a mesh or carpet plot where the height of the mesh is the value of the scalar field, is very good for displaying fields such as terrain elevation.

All objects which can be deformed have a Deform icon in their control panel. Clicking on it brings up a window like Figure 5-5.

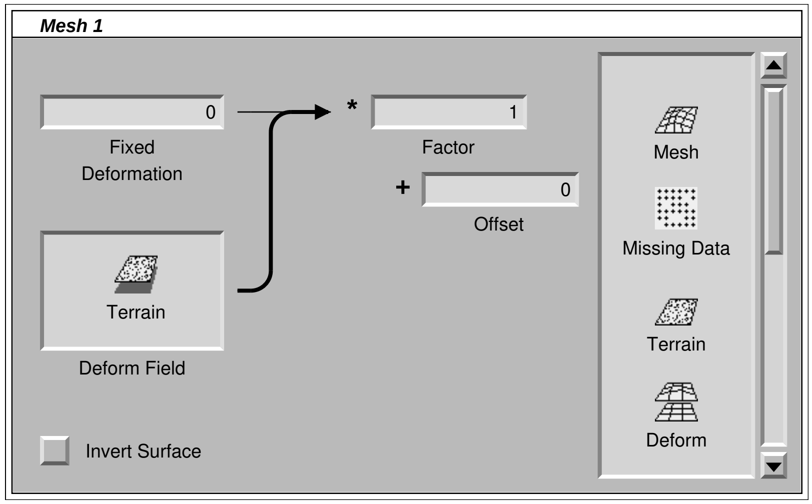

Figure 5-5. *Deformation control panel*

As in the color control panel, at the left is a switch which can either be set to a fixed deformation or to a field for deforming. The deformation is multiplied by the factor and added to the offset to give the total deformation. All text boxes can be edited to change parameters. The Invert Surface check box inverts the resulting surface when it is checked, including inverting the surface normals.

Lines

In addition to shaded surfaces, some visualization objects can draw lines. Some objects, such as contours, only have lines. Others, such as isosurfaces, allow a wire frame of their surfaces to be drawn.

Like surfaces, the lines are drawn in the color defined by the Color controls; however, they are drawn as bright luminous wires without light shading. When both the surface and lines are drawn, the lines usually appear as slightly brighter lines on the surface, which is quite an interesting effect. Another useful technique is to have two visualizations (use Duplicate from the Object menu to get the second one), one showing only the surface and one showing only the lines, each with different colors.

Every object which can draw lines has a Lines button in its control panel. Pressing this button brings up a control panel like Figure 5-6.

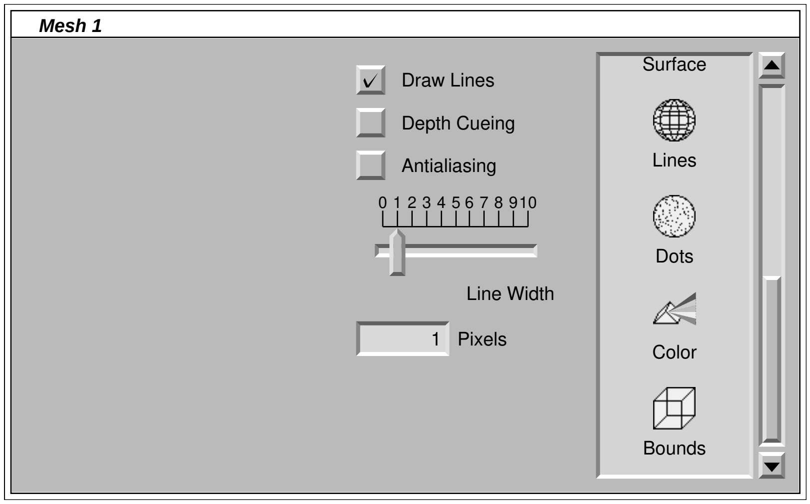

Figure 5-6. *Lines control panel*

At the top right of the panel is a Draw Lines check box. When this check box is on, the lines are drawn. The check box provides an easy way to turn just the lines of the visualization object on and off.

Below the Draw Lines check box is a Depth Cueing check box. Because lines lack the lighting effects of surfaces, it is sometimes harder to see the 3-D shape of an object drawn only in lines. Depth cueing is a technique which draws lines brighter or dimmer depending on how far they are from the observer. This is an approximation of the effect of a local light source, which illuminates distant objects less than nearby ones. Depth cueing will only work and provide reasonable results on some workstations with enough physical bitplanes to represent the entire range of brightness values. To get a good range of brightness values, it may be necessary to adjust the near and far clipping planes within the Observer control panel.

Below this is an Antialiasing check box. Normally, lines are simply drawn on the screen in the fastest available way, which involves turning the screen pixels on when the line goes over that pixel. Because of aliasing, this can produce a rough and jagged appearance. Antialiasing is a technique to provide smoother looking lines by coloring pixels according to the percentage they are covered by the line. Antialiasing does not work on all workstations and may not work for all line widths.

Below the Antialiasing check box is a Line Width slider. This slider controls the width of the lines in pixels. The width of the line also appears as text in the Pixels readout below, allowing you to enter the width as text.

Dots

In addition to shaded surfaces and lines, some visualization objects can draw dots. Some objects, such as point clouds, only have dots. Others may allow dots to be drawn at the vertices of their lines or surfaces.

Like lines, dots are drawn in the color defined by the Color controls as bright luminous points without light shading. When both the surface and dots are drawn, the dots usually appear as slightly brighter dots on the surface.

Dots are drawn very rapidly. A visualization that shows only dots can show the basic shape of a visualization object while still drawing in a short amount of time for quick interaction.

Every object which can draw dots has a Dots button in its control panel. Pressing this button brings up a control panel like Figure 5-7.

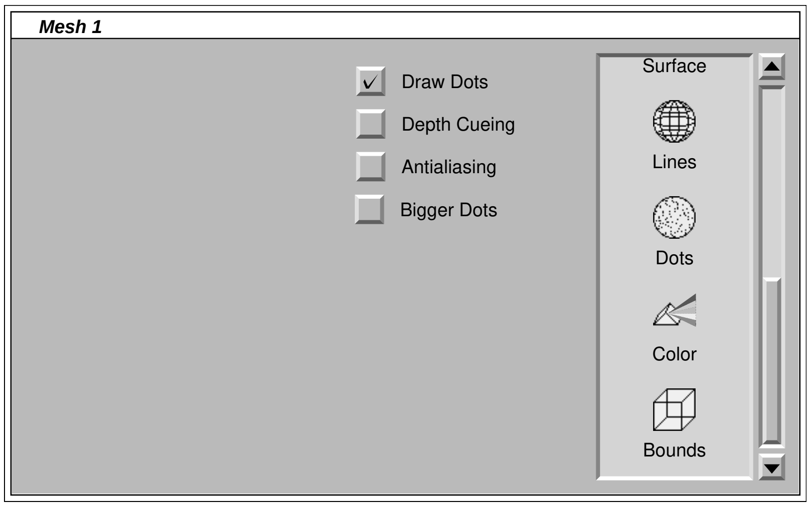

Figure 5-7. *Dots control panel*

At the top right of the panel is a Draw Dots check box. When this check box is on, the dots are drawn. The check box provides an easy way to turn just the dots of a visualization object on and off.

Below the Draw Dots check box is a Depth Cueing check box and an Antialiasing check box. Depth cueing and antialiasing work on dots the same way they do on lines.

Below the Antialiasing check box is a Bigger Dots check box. Normally, dots are drawn as one screen pixel. This is the fastest and is quite satisfactory for dense datasets. When this check box is on, slightly bigger dots will be drawn if the workstation hardware supports it. This may be preferable for sparse datasets with fewer than one thousand points.

Bounds and axes

Every visualization object which has clearly defined bounds in 3-space has a Bounds icon in its control panel. Clicking on this icon brings up a control panel as shown in Figure 5-8.

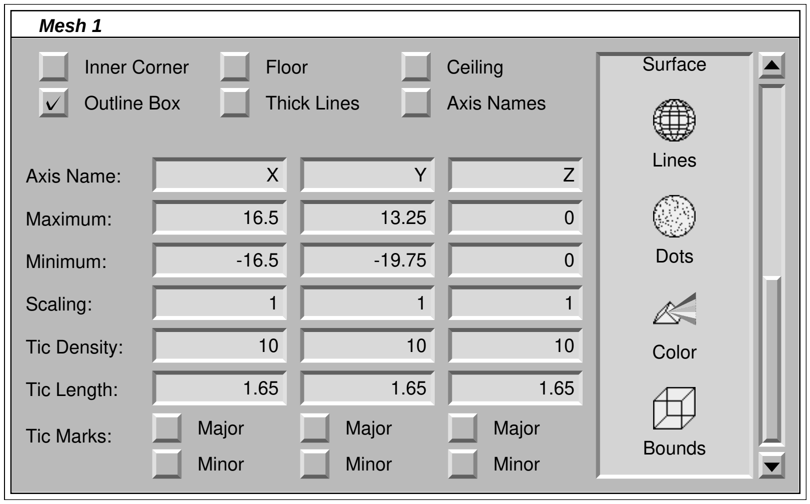

Figure 5-8. *Bounds control panel*

At the top are six check boxes which allow the bounds enclosing the visualization objects to be displayed in a variety of ways:

Inner Corner
When this check box is on, a shaded inner corner of the outline box is shown. This works better than Outline Box for video. To show the entire box, check Inner Corner, Floor, and Ceiling.

Outline Box
When this check box is on, a box will be drawn around the entire visualization object. This works well for interactive work, but the thin lines are not the best for recording onto video.

Floor
When this check box is on, the shaded floor of the outline box is shown. (The floor is transparent from below, however, so it won't obscure the visualization.)

Ceiling
When this check box is on, the shaded ceiling of the outline box is shown. (The ceiling is transparent from above, however, so it won't obscure the visualization.)

Thick Lines
When this check box is on, lines used to draw the outline box are a little thicker than usual. This is better for video.

Axis Names
When this button is checked, axis names are shown around the outline.

Below the check boxes is a series of text boxes and buttons which control if and how the axes are displayed. Each of three columns controls a different axis, which start out named X, Y, and Z. You can change the names of any of the axes by editing the text box. The Maximum and Minimum text boxes control the maximum and minimum values of the visualization object. They start off just big enough to enclose the object. The Scaling controls the factor by which each axis is scaled before drawing. This is normally 1 but can be changed for special effects. One application might be to increase the scaling of the Z axis for meteorological data. The remaining two text boxes control the drawing of tic marks. The Tic Length text box gives the length of the major tic marks in the coordinates used by the visualization object. The Tic Density text box controls how densely the tic marks are drawn. Higher numbers specify denser tic marks. At the bottom are two check boxes per axis: a Major check box and a Minor check box. These control whether major and/or minor tic marks are displayed.

Dataset modifier controls

Dataset modifiers are created automatically when a visualization produced by Visualize As needs to modify the data before visualizing it. When the main dataset of a visualization needs to be modified, one or more control panels will appear in the control window to adjust the modification. This section explains these controls.

Orthogonal Slicer

When a visualization object is used to visualize a dataset with more topological dimensions than it usually requires, such as when a 3-D scalar field is visualized with the 2-D contour visualization object, the dataset is sliced along an orthogonal axis. The slice can be changed through the control panel, such as shown in Figure 5-9.

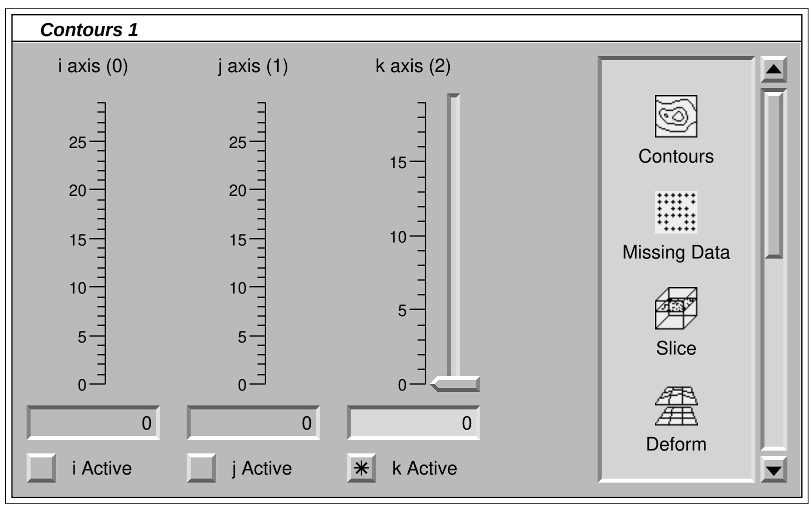

Figure 5-9. *Orthogonal Slicer control panel*

The control panel contains a series of sliders, each of which controls the position of the slice over one topological axis of the dataset. The text boxes at the bottom of each slider allow you to enter the position of the slice as text. Press the Return key after entering the number. The check boxes to the right of the slider control which topological axis the slider controls. Because the orthogonal slicer works on topological, not spatial, dimensions, the axes are denoted by **i**, **j**, **k**, etc. instead of by **x**, **y**, and **z**. Each slider must control a separate axis. If you try to set the radio buttons of two sliders to the same axis, an error dialog will appear.

If there are more sliders than can be shown in the window, a scroll bar at the bottom of the panel allows you to scroll through them.

Sliced objects are created automatically by the Visualize As button or menu item when the type of visualization object takes a dataset of smaller dimensionality than the dataset. For example, if you try to visualize a 3-D scalar field using a contour visualization object, it will automatically become sliced.

Sliced datasets can also be produced with Modify. However, at the present time, this doesn't really gain you anything, as all modifications appropriate to a visualization object are done automatically.

Missing Data

Datasets can contain missing data points. Visualization objects deal with missing data points in a variety of ways, usually by not drawing the visualization where there are missing data points. Sometimes, it is useful to treat missing data points as if they had a certain value. For example, with radar reflectivity data, it is usually best to treat missing data points as having a reflectivity of 0.

The Missing Data control panel, shown in Figure 5-10, allows you to control this behavior. At the top there are two radio buttons and a text box containing a number. If the Treat As Missing radio button is on, missing data points are passed to the visualization object as missing. If the Treat As: button is on, missing data points are treated as if they were the same as the number in the text box to the right. You can enter any real number into the text box, including Infinity and -Inifinity.

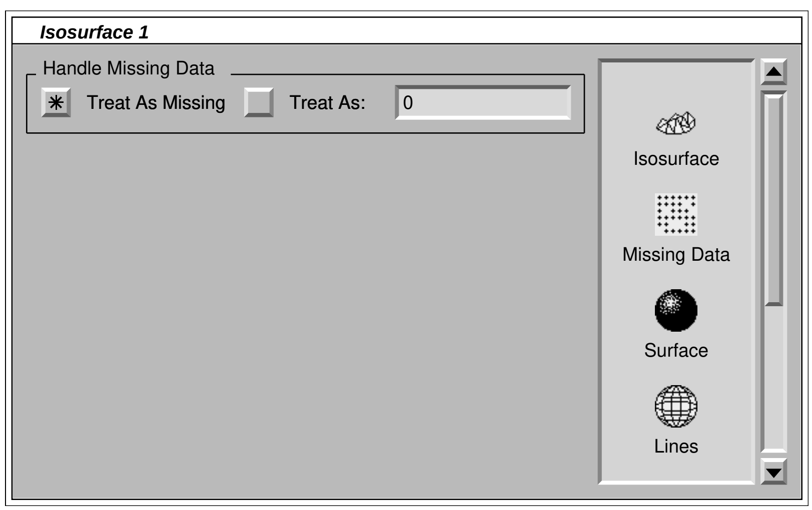

Figure 5-10. *Missing Data control panel*