Chapter 3 Window Types

Chapter Overview

Text Window Actions

- Attributes Dialog and Formats
- Selecting Regions
- Extracting Selection Regions

Imaging and Color Capabilities

- Scaled Color Image
- Interpolated (Smoothed) Image
- Polar Image Generation
- Color Palettes
- How Long Should Image Generation Take?
- Stopping Image Generation

Synchronizing Multiple Datasets across Windows

The Notebook Window

- Notebooks
- Starting and Stopping Calculations
- Built-in Functions
- Constant Result Calculations

Chapter Overview

DataScope has several different types of windows which appear on the screen. They are used for displaying numbers as text, viewing images of different types, and for entering notes and formulas. The three major types are referred to as Text windows, Image windows, and Notebook windows. This chapter describes in detail all of the DataScope operations which are defined for each window type.

Text Window Actions

When you load a dataset, DataScope creates a text window with row and column labels, scroll bars, and a portion of the array printed as floating point numbers. The size and scrolling of this window is controlled with standard Macintosh window actions. Similar in nature to spreadsheet programs, the array of numbers can be scrolled, selected, and copied to the Clipboard.

Attributes Dialog and Formats

Several of the text window characteristics can be controlled with the Attributes command from the Numbers menu. When you Select the Attributes command, a dialog box is displayed with fields that you may use to make alternate settings. The Attributes dialog box is shown in Figure 3.1. If you select the Cancel button, all fields remain unchanged. If you select the OK button, any changes you made take effect.

Figure 3.1 Attributes Dialog Box

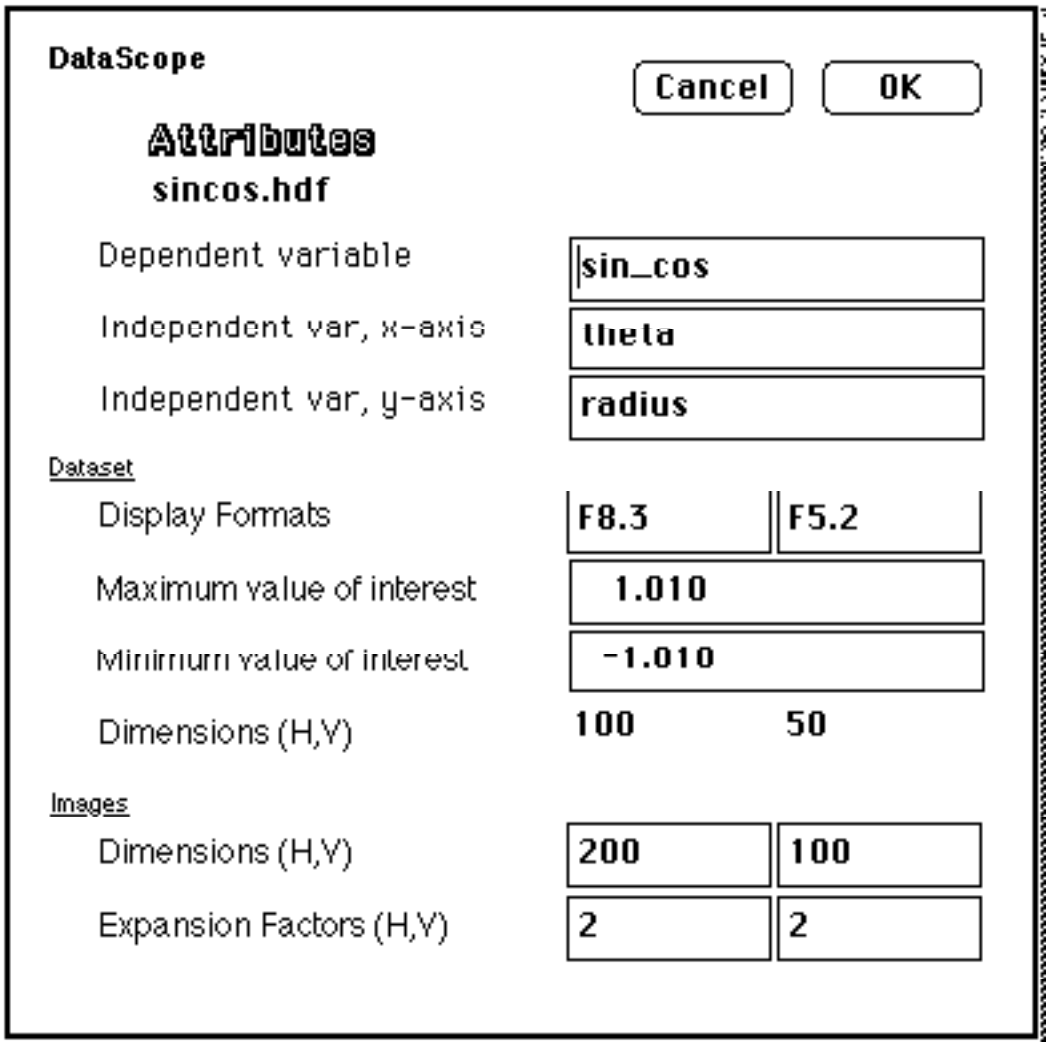

The fields in the attributes dialog box have the following characteristics:

- The filename displayed under the Attributes heading is the original name of the dataset file. This field cannot be changed in this dialog, it changes when the file is saved.
- The first three boxes are associated with the dependent variable and the two axis variables; the user can change the names of these variables. The dependent variable name is used for the text window title and is also used in calculations (see Notebooks section).
- The two Display Format boxes control the printing display format for numbers in the text window. The left box controls the array format, and the right box is used for the column and row labels. The default printing format is E11.4 for the data values and F5.0 for the column and row labels.

- The maximum and minimum values of interest fields define the region of interest. In many cases, the maximum and minimum values are calculated for you and placed here. If you want to eliminate outliers, or confine the number range which is used to scale colors during image generation, set these values to your region of interest. Outliers are not actually removed from the dataset, they are just ignored during image generation.

- The remaining boxes control the image generation size. The three image generation commands in the Image menu obtain their target image size from these values. For example, if the dataset is of the size 100 by 50, and the expansion factor is set to 3, the resulting image is exactly 300 pixels wide and 150 pixels tall. Only integer expansion sizes are allowed. All settings are rounded to the nearest integer expansion size. The row and column expansion sizes do not have to be equal.

Selecting Regions

All of the display windows in DataScope support the selection of a region of data values. The most common way to select these regions, according to Macintosh standards, is to press the mouse button down in one corner of the region and drag to the opposite corner of the desired region before releasing the button. This works in both the text window and the image windows.

A key feature of DataScope is the synchronization of selection regions between windows. When you select a region in the text window, the corresponding region in each associated image window is outlined with a rectangle. Associated images are those which were generated from the dataset. Conversely, selecting a region in the image window, by dragging out a rectangle, causes the text window to highlight the selection region for the corresponding data values. The text window also scrolls and centers the selection region in the window.

Extracting Selection Regions

One way to create a new dataset is to extract it from another dataset. To do this, first select the region of interest in the source dataset. Then choose the Extract Selection command from the Numbers menu. This pulls the data from the current selection region of the current dataset and creates a new dataset with those values. A new text window with the extracted data then appears on the screen.

NOTE: Unless you select the entire window as the selection region, this new subset will have smaller row and column dimensions than the original. The necessary values for the row and column scales are extracted from the original row and column scales.

Imaging and Color Capabilities

With an array in memory, DataScope can calculate and display new windows with images generated from the dataset. The three supported types of image generation are described in the following section. They are as follows:

- simple, scaled color image
- interpolated (smoothed) image
- polar image

Scaled Color Image

The simplest and quickest imaging operation is the basic scaled color image. For any dataset that is already loaded, select the Generate Image command from the Image menu.

The algorithm is simple. Calculating the difference between the maximum and minimum values of interest yields the range of values of interest. Scaling each value of the dependent variable with the following formula produces a linearly scaled set of numbers between 0 and 240. For example,

colorvalue = 240 * (value - minimum) / range

Using the Macintosh color palette manager, DataScope assigns the values from 0 to 240 to a range of hardware colors. When displayed, this produces a grey-scale or color raster image of the scaled data.

Display of the data is a straightforward duplication of pixels. The one-to-one case, where the number of pixels used to display the image matches the number of data points in the array, is easy. Each pixel is set to the color for its corresponding data point.

When the data array has fewer points than the number of pixels to display, each data point is blown up to a rectangle of pixels. For example, one data value may become a square of 3x3 pixels which are all set to the same color. This creates a "chunky" or "jagged" effect. To get smoother enlargements, use the interpolation method discussed next.

Interpolated (Smoothed) Image

The Interpolated Image command from the Image menu utilizes the same linear color scaling as the Generate Image command. For one-to-one display, when the number of pixels matches the number of data points, the images produced by the two methods are identical. But when the image is larger, each pixel of the image must be interpolated between multiple data values before it is scaled on the 0 to 240 color scale.

Stretch grids, those grids without regularly spaced row and column scales, should be imaged with the interpolation method. This

method takes into account the X and Y distances between points when performing the interpolation.

Polar Image Generation

Again, the color scaling for this type of image generation is on a linear 0 to 240 scale. In this case, the pixels are placed into the image according to a polar interpretation of the row and column scaling information. The row scale is treated as a set of radius values; and the column scale, as a set of angle values in the range 0 to 2*PI radians. The zero angle corresponds to the right side of the horizontal axis with the origin at the center of the window.

Color Palettes

The term *color palette* refers to the colors that are assigned to the 0 to 240 scale of values which DataScope creates for the pixels. Generally, the hardware used supports 256 colors from a total selection of over 16 million colors. The assignment of the scaled data values, 0 to 240, to the colors that you see on the screen is the color palette.

DataScope displays images with a grey-scale palette by default. The level of grey, or intensity of white, in the image shows the relative magnitude of the data value. This is the simplest mapping of colors (levels of grey in this case) to the data values, but many datasets can benefit from alternative mappings – a spectrum from blue to red, for example. The Load Palette command from the Image menu is used to assign a new set of colors to the current image window.

The Load Palette command can retrieve a palette from any HDF file which contains a palette. It always pulls the first palette in the file. Several sample palettes are provided with DataScope, including a spectrum sample. The HDF documentation should be referred to for information on how to create your own palettes.

How Long Should Image Generation Take?

A 400x400 interpolated image might take 2 to 5 minutes to complete, whereas the non-interpolated image only takes 15 seconds. The standard image generation is dependent only on the size of the original data; the interpolated image, on the other hand, is more dependent on the size of image being generated, for every displayed pixel is separately calculated. Polar image generation primarily depends on the source data size, but the target image size is also a factor.

In all cases, first an empty window and then a watch cursor appears while the creation of the image takes place. The window is periodically updated as the image is generated, displaying the portion of the image that has been calculated so far.

Stopping Image Generation

All three methods of image generation update the screen with partial results as they are calculated. To stop the image calculation before it is finished, hold down the **COMMAND** (⌘) key and type a period (.). When you enter ⌘-PERIOD, the calculations stop and the window remains in its current state with a partially calculated image. Leaving a window in this state does not affect the operation of other actions; it is possible to work with only a portion of an image if you do not want to wait for the whole image to generate.

Synchronizing Multiple Datasets across Windows

One of the most powerful features of DataScope is its ability to synchronize multiple datasets across windows. These might be two, three, or four variables from the same simulation; or, multiple time-steps from a simulation run. The only requirement is that the number of rows and the number of columns of the datasets are identical.

To invoke the synchronization of selection regions, select the Synchronize command from the Numbers menu to turn on the menu checkmark. A selection in one window produces matching selections in all text window datasets that have the same row and column dimensions and all image windows that were generated from those datasets. When you are done, select Synchronize again to turn the feature off and remove the checkmark.

NOTE: While Synchronize is on, any selection activity is reflected in all synchronized windows.

Image windows are always synchronized with their originating text windows regardless of the synchronize setting. The Synchronize command adds to this feature by synchronizing separate datasets, also.

The Notebook Window

Notebooks

For DataScope, every dataset in an HDF file has a *notebook*. Initially, this notebook is empty. You may enter any text in the notebook window of DataScope and save that information with the dataset. If the dataset was previously saved with a notebook, then the notebook window appears automatically when you re-open the HDF file. If there is no notebook window or if it has been removed from the screen, select the See Notebook command from the Numbers menu. The notebook for the current dataset will appear on the screen. If it is already on the screen, it will be brought to the front.

The notebook window is a standard Macintosh text window, so edit text as you normally would for a Macintosh. You can copy and

paste between notebook windows and other text programs under MultiFinder. When you save the dataset, the entire contents of the notebook window are saved with it.

As described in the following section, the notebook window can be used to enter formulas and perform calculations. These formulas can be saved in the notebook window for later use.

Starting and Stopping Calculations

Calculations are performed in the notebook window. Enter a calculation as a FORTRAN-like assignment statement, using the variable names for the currently loaded datasets that appear as the window titles. Then, select the statement with the mouse. Now, select the Calculate From Notes command from the Numbers menu. The assigned variable in the formula will be created from the calculation. Below are examples of these assignment statements.

```
logden = log(density)
Ftemp = (Ctemp*5.0/9.0) + 32.0
ke = 1/2 * mass * vel * vel
c = 3.1416*sin(.65)
```

These statements may contain the operators *, / , +, -, and unary -. They may also contain parentheses and a selection of functions. The supported functions are described in the next section, Built-in Functions.

All arrays used in the calculation must have the same row and column dimensions. The new variable created will have the same dimensions as the arrays that were used to produce it. For example, adding the two arrays means that each element from the first array is added to its corresponding element in the second array. This produces one resulting value at each position that is placed into the array for the new variable.

Attributes for the newly created variable are copied from the first source variable in the expression. Obviously, the newly created variable, then, has the same dimensions as the source variable. A text window is created and appears on the screen with the new variable name from the formula. The maximum and minimum values that define the region of interest are automatically calculated by finding the maximum and minimum values of the resulting array. You may change these in the Attributes Dialog Box (see Figure 3.1) to reflect a more accurate region of interest.

Built-in Functions

DataScope contains very few built-in functions at this time. DataScope will develop and improve in this respect as functions are added. If you need functions that are not provided, please request them, so that they may be added to future releases. You may reach us via US mail or electronic mail at the addresses listed on the copyright page at the beginning of this manual. The currently supported list of DataScopes built-in functions is as follows:

log()	natural logarithm	
sin()	sine, angle measured in radians	
cos()	cosine, angle measured in radians	
ddx()	discrete difference in the X direction, taken from left to	right
ddy()	discrete difference in the Y direction, taken from top to	bottom

Constant Result Calculations

As shown in a previous example, the calculation does not need to contain any array variables. It may contain only constants; therefore, the formula may yield a constant value result. In this case, a text window with an array is not required. Instead, the value is returned to the notebook window. After selecting the Calculate From Notes command, the answer appears in the notebook window.

```
x = 1.23456*1.00
***Result: 1.23456
```

The result information is placed into the notebook directly after the selected region. This value is the result of the calculation. The assignment variable is ignored.