NCSA X DataSlice

for the X Window System™

Version 1.0

September 1989

University of Illinois at Urbana-Champaign

Table of Contents

Introduction

Overview vii

About NCSA X DataSlice vii

Use of This Manual viii

Manual Contents vi

Form of Presentation vii

Suggested Reading ix

Chapter 1 Getting Started

Chapter Overview 1.1

Installing X DataSlice 1.1

Specifying a Display Host 1.1

Identifying the Current Display Host 1.2

Invoking X DataSlice 1.2

Loading a 3D Scientific Dataset 1.3

Windows in NCSA X DataSlice 1.4

Displaying Your Dataset 1.6

Attribute Window 1.6

Cartesian Dicer Display Window 1.7

Working with Palettes 1.9

Changing Palettes in One Raster Window 1.9

Changing Palettes in Multiple Raster Windows 1.10

Saving Generated Images 1.10

Exiting the Application 1.10

Chapter 2 File Formats in NCSA X DataSlice

Chapter Overview 2.1

File Formats 2.1

Raw Raster Files 2.1

Raw Palette/SEQ Files 2.1

HDF Files 2.2

How NCSA X DataSlice Reads and Displays Data Files 2.5

Chapter 3 Planes Along Axes

Chapter Overview 3.1

Viewing Planes Along the Coordinate Axes 3.1

Viewing 2D Planes 3.1

Single Image Window 3.4

Multiple Frame Window 3.5

Unfold Window 3.5

Animation Window 3.5

Viewing 3D Planes 3.8

Visable Volume Rendering 3.8

Setting Parameters 3.9

Defining Substances 3.9

Chapter 4 Arbitrary Planes

Chapter Overview 4.1

Viewing Oblique Planes 4.1

Manipulating the Image 4.1

Viewing the Image 4.3

Animating the Image 4.3

Chapter 5 Cartesian Dicer

Chapter Overview 5.1

Viewing Planes in a Dataset 5.1

Chapter 6 Animate 8bit Raster and Process SDS

Chapter Overview 6.1

Loading 8-bit Raster Images 6.1

Loading the Raster Image 6.1

Animating the Raster Image 6.2

Processing Scientific Datasets 6.3

Figures

Figure 1.1 Control Window 1.2
Figure 1.2 Load SDS File Window 1.3
Figure 1.3 Attributes Window 1.5
Figure 1.4 Empty Cartesian Dicer Window 1.8
Figure 1.5 Cartesian Dicer Window (View Slice) 1.8
Figure 1.6 Load a color map File Window 1.9
Figure 1.7 Save a Dicer Image File Window 1.11
Figure 2.1 HDF Environment 2.2
Figure 3.1 Attributes Window 3.1
Figure 3.2 Cartesian Display Mode Window 3.2
Figure 3.3 Single Frame Window 3.5
Figure 3.4 Multiple Frame Window 3.6
Figure 3.5 Animation Window 3.7
Figure 3.6 Display Mode (View 3D) 3.8
Figure 3.7 V-Buffer 3.10
Figure 4.1 Attributes Window 4.1

Figure 4.2 Arbitrary Display Mode Window 4.2
Figure 4.3 Arbitrary Display View Window 4.4
Figure 4.4 Arbitrary Display Animation Window 4.4
Figure 5.1 Attributes Window 5.1
Figure 5.2 Cartesian Dicer Window 5.2
Figure 5.3 Cartesian Dicer (View Slice) 5.2
Figure 5.4 Cartesian Dicer (Multiple Slices) 5.3
Figure 6.1 Load a RASTER8 set File Window 6.1
Figure 6.2 Animate 8-bit Raster Image Window 6.2
Figure 6.3 Process & save a SDS Window 6.4

Introduction

Overview

This introduction provides an overview of the capabilities and purpose of NCSA X DataSlice and directs you to sources that offer detailed information regarding the X Window System. The organization and purpose of this manual are described and notational conventions explained.

About NCSA X DataSlice

NCSA X DataSlice is a color imaging and data analysis tool based on the X Window System. It was developed for analysis of three-dimensional 32-bit floating-point scientific data stored in NCSA HDF file format. The current release of NCSA X DataSlice features:

- color raster display of 2D slices from 3D datasets
- display of actual data values in spreadsheet form
- synchronized selection between corresponding image and data windows
- continuous and single-step animation of color raster images
- manual frame control, frame skipping, auto-reverse, and asynchronous discontinuation animation capabilities
- tiling of multiple images in a single window
- arbitrary slicing and dicing
- 3D visible volume rendering
- interpolation of color raster images
- support for user-defined, raw palettes, such as those created in NCSA PalEdit

System Requirements

In order to run NCSA X DataSlice, you must have (1) a system and server that support color X Windows—such as Sun 2, 3, and 4 systems, and the IBM RT, and (2) installed X version 11 and server release 2 or later on your system.

Use of This Manual

This section describes the organization of this manual, and the conventions and nomenclature used in developing it. Although the instructions in this manual should enable you to successfully use NCSA X DataSlice, a working knowledge of your window manager is presupposed. For detailed information regarding the X Window System, refer to the guides listed in the section entitled "Suggested Reading."

Manual Contents

This manual is organized into the following chapters:

Chapter 1, "Getting Started," explains how you should install and invoke NCSA X DataSlice. The chapter contains a tutorial that introduces you to the basic procedures involved in using the application, explains similarities among the program's windows, and familiarizes you with the program's user interface:

Chapter 2, "File Formats in NCSA X DataSlice," describes the basic file formats used in NCSA X DataSlice, how the program reads and displays data files, and how to specify the dimensions of raw files. Information is included for obtaining HDF software from NCSA.

Chapter 3, "Planes Along Axes," describes options available in the Planes Along Axes window. It oulines the various options available to manipulate and view both 2D and 3D planes.

Chapter 4, "Arbitrary Planes," outlines the options available in the Arbitrary Planes window which permits you to manipulate, view, and animate 2D oblique slices of data.

Chapter 5, "Cartesian Dicer," outlines the options available in the Cartesian Dicer window.

Chapter 6, "Animate 8bit Raster and Process SDS," describes the use of animation controls and how 3D SDSs may be enlarged by pre-processing.

Form of Presentation

The material in this manual is presented in text, screen displays, or command line notation.

Text

In explaining various features and commands, this manual often presents a word within a paragraph in *italics* to indicate that the word is defined within the paragraph.

Portions of this manual refer to other sections or chapters of the manual where related topics are discussed. These cross references usually indicate the title of sections or chapters enclosed in quotation marks, such as, See Chapter 1, "Getting Started."

Screen Displays

Screen displays in this manual are presented in `courier type`.

Entry Format Notation

Throughout this manual, several explanations instruct you to make entries by typing on the keyboard. These entry instructions are printed in **`courier bold type`** and appear within a screen display, paragraph, or on a separate line.

Figure I.1 shows you how to read and enter a command line.

Figure I.1 Reading and Entering a Command Line

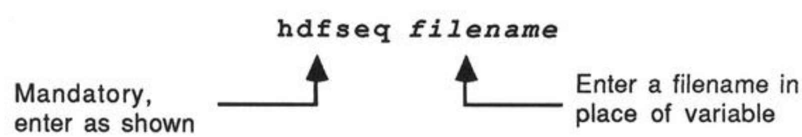

Suggested Reading

This document is not a manual on the X Window System. If you need detailed information regarding X Windows, refer to the following works:

- *The Definitive Guides to the X Window System*, published and copyrighted by O'Reilly & Associates, Inc., 1988.
- *X Window System*, by Robert W. Scheifler, James Gettys, and Ron Newman, published and copyrighted by the Massachusetts Institute of Technology and Digital Equipment Corporation, 1988.
- *Introduction to the Window X System*, by Oliver Jones, published and copyrighted by Prentice Hall, 1988.

READ ME NOW

If you want to see more software like NCSA X DataSlice, please send us a letter, email or U. S. mail, telling us what you are doing with NCSA X DataSlice. We need to know: (1) What science you are working on—an abstract of your work would be fine; and (2) How NCSA X DataSlice has helped you, for example, by increasing your productivity or allowing you to do things you could not do before.

We encourage you to cite the use of NCSA X DataSlice, and any other NCSA software you have used, in your publications. A bibliography of your work would be extremely helpful.

NOTE: This is a new kind of shareware. You share your science and successes with us, and we can get resources to share more software like NCSA X DataSlice with you.

NCSA Contacts

Mail user feedback, bugs, and software and manual suggestions to:

NCSA Software Development
X DataSlice
152 Computing Applications Bldg.
605 E. Springfield Ave.
Champaign, IL 61820

Send communications via electronic mail to one of the following:

Bug Reports
bugs@ncsa.uiuc.edu
bugs@ncsavms.bitnet

All Other Communications
softdev@ncsa.uiuc.edu
softdev@ncsavms.bitnet

Disclaimer

Trademark Acknowledgments

Chapter 1 Getting Started

Chapter Overview

Installing X DataSlice

- Specifying a Display Host
- Identifying the Current Display Host

Invoking X DataSlice

Loading a 3D Scientific Dataset

Windows in NCSA X DataSlice

Displaying Your Dataset

- Attributes Window
- Cartesian Dicer Display Window

Working with Palettes

- Changing Palettes in One Raster Window
- Changing Palettes in Multiple Raster Windows

Saving Generated Images

Exiting the Application

Chapter Overview

This chapter explains how you should install and invoke NCSA X DataSlice. The chapter contains a tutorial that introduces you to the basic procedures involved in using the application, explains similarities among the program's windows, and familiarizes you with the program's user interface:

- installing and invoking the application
- loading a 3D scientific Dataset
- performing basic imaging operations
- loading a palette
- saving generated images
- exiting the application

Installing NCSA X DataSlice

Before using NCSA X DataSlice, you must install the program somewhere within your path on the client system (the machine on which you want to run the application) and specify your display host (the machine from which you wish to control and view the application). Next, define the path environment variable to include the appropriate path, as shown below:

For example, if you were to install NCSA X DataSlice on the CRAY-2 using C shell, you would set the path by entering a command of the form:

```
set path=($path /pathname)
```

where /pathname is the path to NCSA X DataSlice.

Specifying a Display Host

If you are running the program off a remote client machine, e.g. Cray 2, you must first register the client machine on your X-Window host server by typing the following on your local host:

```
xhost clientmachinename
```

where clientmachinename is the name of the remote client.

To specify your display host on your local client machine, enter a command of the following form:

```
setenv DISPLAY hostname:monitornumber
```

where hostname is the name assigned to the desired machine, and monitornumber is the number assigned to the monitor on which you wish to view the program.

For example, to specify a machine named babbage as the display host and use the first of the monitors connected to babbage as the display monitor, you would enter:

```
setenv DISPLAY babbage:0
```

Identifying the Current Display Host

To identify the current display host, enter:

```
printenv $DISPLAY
```

The return specifies the server that is currently set as your display host. For example, the return message

```
babbage:0
```

informs you that NCSA X DataSlice will be displayed on the machine dubbed babbage, using monitor 0—the first of the monitors connected to babbage.

Invoking X DataSlice

To invoke NCSA X DataSlice, enter:

```
xds
```

The control window appears as the first frame (see Figure 1.1). It contains the LogBook and several control buttons you will use to initiate the main operations performed by NCSA X DataSlice: loading a 3D scientific dataset, loading an 8-bit raster image set, processing a scientific dataset, and quitting the application. Thus, you may want to keep it in direct view. You may move the window using mouse movements and operations that are dependent on the type of window manager that you are currently running.

Figure 1.1 Control Window

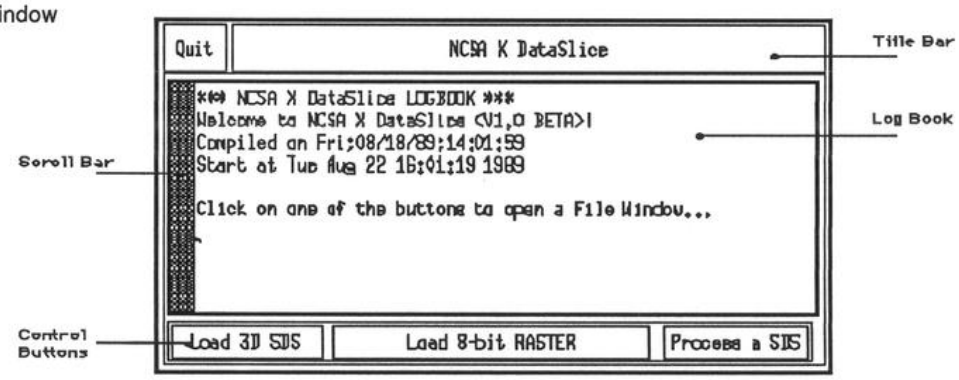

The *LogBook* records and displays every non-trivial action you make while running NCSA X DataSlice. It provides feedback such as instructions, explanations, or warnings regarding your actions, and functions as a notebook in which you may record your observations and comments for future reference.

To enter text, position the pointer in the LogBook, click once, and begin to type. Note that once you enter text into the notebook, you are unable to delete it.

The contents of the LogBook are preserved in the file xdsLog, which should be stored in the same directory as the application, even after you exit the program. Consequently, you can review a session with X DataSlice by reading the contents of the file xdsLog, and you can use vi or other text editors to alter the contents of the file.

Loading a 3D Scientific Dataset

To load a three-dimensional scientific dataset stored in an HDF file into NCSA X DataSlice:

1. Click the Load 3D SDS button in the control window. A file window appears, displaying the current directory (see Figure 1.2).

Figure 1.2 Load SDS File Window

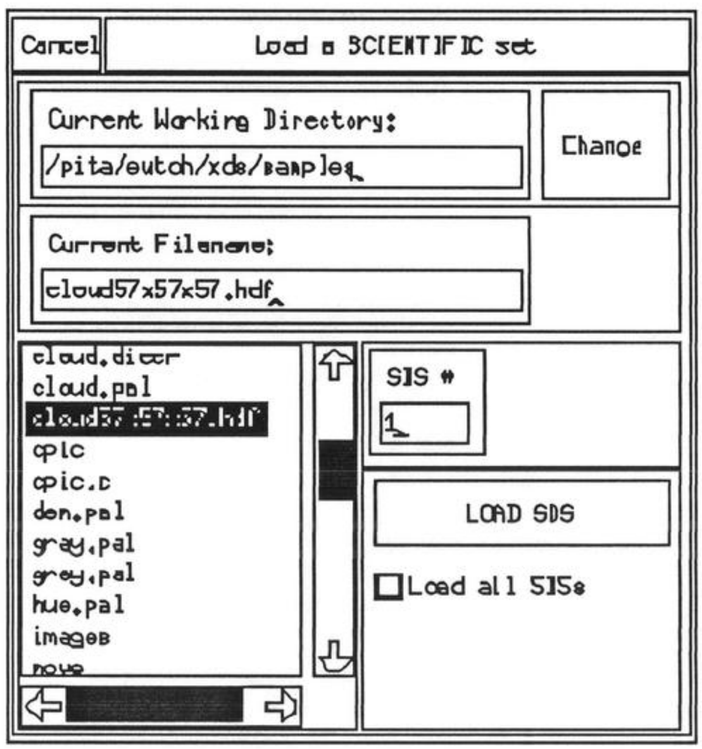

The file window and its contents depend on the particular type of operation you are performing. In brief, if you want to save a file, then the Load buttons are not present, and vice versa. Only one file window may be opened at a time, so you must close the opened window by clicking Cancel if you decide not to use it. The file window automatically closes after each affirmed file operation.

2. Change the current directory, if necessary, by entering the path to a new directory in the text box labeled, Current Directory being displayed. Click Change to access the new directory and its file contents. If an error message appears in the LogBook, the current directory remains unchanged.

3. Choose the desired dataset to be viewed by scrolling up or down the file list window and clicking on its filename. The name appears in the Current Filename text box. The LogBook reports if the file is indeed a 3D SDS HDF file or a Raster8 HDF file. If the file is of HDF format, the message shows the number of scientific datasets contained in the file. For this tutorial, select cloud57x57x57.hdf, which contains three 3D scientific datasets that depict the X, Y, and Z wind velocities in a thunderstorm.

4. Specify which scientific dataset you wish to load into memory from the selected file by entering a number in the text box labeled SDS #. Notice that the text box SDS # reads 1 by default, indicating that the first SDS in the file will be loaded unless you specify otherwise. Be aware that if you enter a value above the designated number of datasets present, the file window closes and an error appears in the LogBook: ***HDF ERROR: Cannot load SDS *4*.

To load all the datasets, check the option, Load all SDSs.

NOTE: A 3D SDS usually requires a few megabytes of memory, so although it is possible to load up to ten SDSs together, it is not advisable to load more that three datasets into memory at any one time.

5. Click Cancel to leave the box without implementing your entries or LOAD SDS to continue with the imaging of your dataset(s). The file window disappears and an Attributes window appears (Figure 1.3).

Windows in NCSA X DataSlice

Most windows in this program, with the exception of the Control and file windows, have certain elements in common. When describing these similarities, this manual will use the Attributes window (Figure 1.3) as a representative for most other windows contained in the X DataSlice program.

In this program you issue commands by clicking buttons in windows, turning toggle buttons on and off, selecting within or moving scrollbars, and entering values.

Figure 1.3 Attributes Window

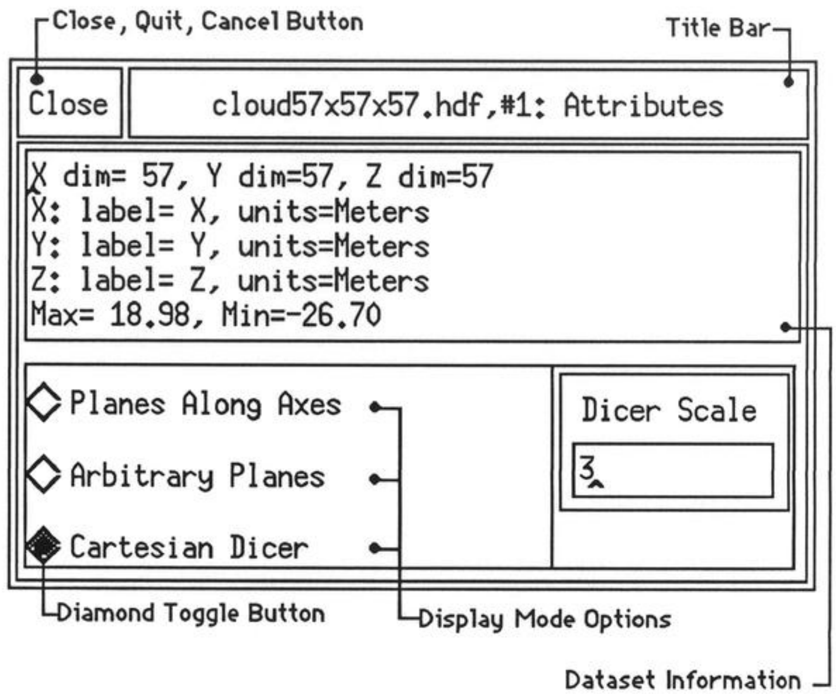

All windows created by X DataSlice are movable under uwm or twm. Each contains a Close, Quit, or Cancel button at the top-left corner for exiting and disposing of the window, and features a title bar that identifies the window according to the following convention:

filename, #n: function

where filename is the name of the HDF file whose 3D scientific dataset is represented in the window, n indicates the position of the SDS or raster image in the file in relation to any other SDSs or raster images the file may contain, and function is the function of the window. For example in the Attributes window, the title bar reads, cloud 57x57x57.hdf, #1: Attributes, which identifies an Attributes window for the first scientific dataset in the HDF file, cloud57x57x57.

All windows other than the control and file windows are spawned in a hierarchical format after the initial loading of a dataset. If you close a parent window, all its children are then automatically closed by NCSA X DataSlice, saving you the trouble of tracking which window belongs to which dataset. Closing a window at the bottom of the hierarchy closes only that particular window.

NOTE: Each time you close an Attributes window, you free up a lot of memory for additional imaging options.

Most image windows in X DataSlice offer buttons for changing palettes and saving images named Palette and Save, respectively.

Refer to the later sections of this chapter for more information on these options.

To activate a window you are not presently in, or to bring a background screen to the foreground, position the arrow in the desired window and click the mouse once. Note that these actions are dependent on the window manager that you are currently running.

Displaying Your Dataset

The Attributes Window

An attributes window (Figure 1.3) appears for each scientific dataset successfully loaded from a specified HDF file and contains information about that particular dataset including: the dimensions of the dataset, the labels and units of measurement of the respective axes, and the maximum and minimum values of the dataset. If you have clicked the option, Load all SDSs, an Attributes window appears for each separate dataset.

Note that when the minimum and maximum values are not specified in the file, NCSA X DataSlice calculates them automatically. The time required for calculation is contingent upon the size of the dataset, so several minutes may elapse before the values are calculated. To avoid this delay, you should pre-calculate and store the values for your dataset, using the program "move" in the samples directory. Instructions as to how move can be used are given in the source file move.c itself.

Besides giving information about the current SDS dataset being examined, the Attributes window also contains a dicer scale and various display modes from which you may choose to view your dataset.

Dicer Scale

The value you enter in the dicer scale text window determines the size of the generated dicer image and its surrounding window. 1 is the smallest number you may enter, and although there is no maximum limit for the scale, 10 is the recommended upper limit. The generation time for the image is contingent upon the dataset and the dicer scale; the larger the dataset and scale, the longer the program takes to generate the image. Note that you should enter the dicer scale value before choosing the Cartesian Dicer display mode. If you fail to enter a number, the scale defaults to 3.

Display Modes

Display modes available to the user are: Planes Along Axes, Arbitrary Planes, and Cartesian Dicer. Planes Along Axes, which portrays 2D slices of data along the three major cartesian coordinate axes, also contains an option to view 3D images generated from a V-Buffer Algorithm. The Arbitrary Planes mode presents arbitrarily selected 2D oblique slices, while the Cartesian Dicer displays a 2D volume slice of data as shown in its relative location in the 3D matrix. For this tutorial, the Cartesian Dicer mode is examined.

To generate a small image with the Cartesian Dicer:

1. Enter a 2 in the text box labeled, Dicer Scale.

2. Click the diamond toggle next to Cartesian Dicer. Notice that all of these diamond-shaped buttons toggle on and off.

The Cartesian Dicer Display window opens.

Cartesian Dicer Display window

The Cartesian Dicer window appears as shown in Figure 1.4. The Cartesian Dicer permits you to view 2D slices of data in their relative positions in the 3D matrix and to manipulate their axes for the particular view you desire. You may also designate the thickness of your slice. To choose a thickness of one, simply click once on a particular point of the designated slider. To choose a thickness greater than one, click on the Frame Control Slider, and drag to the desired position. For instance, if your plane begins at point 25 on the Y axis and ends at point 56, your dataslice will have a thickness of 31. In both instances, the text boxes labeled Start and End reflect the range of your dataslice. For a more detailed discussion on the Cartesian Dicer mode, refer to Chapter 5, "Cartesian Dicer."

For this demonstration, you will manipulate the y slider to identify a data slice with a thickness greater than one.

1. Click and drag to a desired point within the Frame Control Slider text box for the Y axis. Notice the text boxes labeled Start and End reflect the beginning and ending point of your plane, illustrating the thickness of your slice.

2. Click the View button. You must always click View after each slice selection. X DataSlice now generates a new slice visible in the Cartesian Dicer window as shown in Figure 1.5.

Figure 1.4 Empty Cartesian Dicer Window

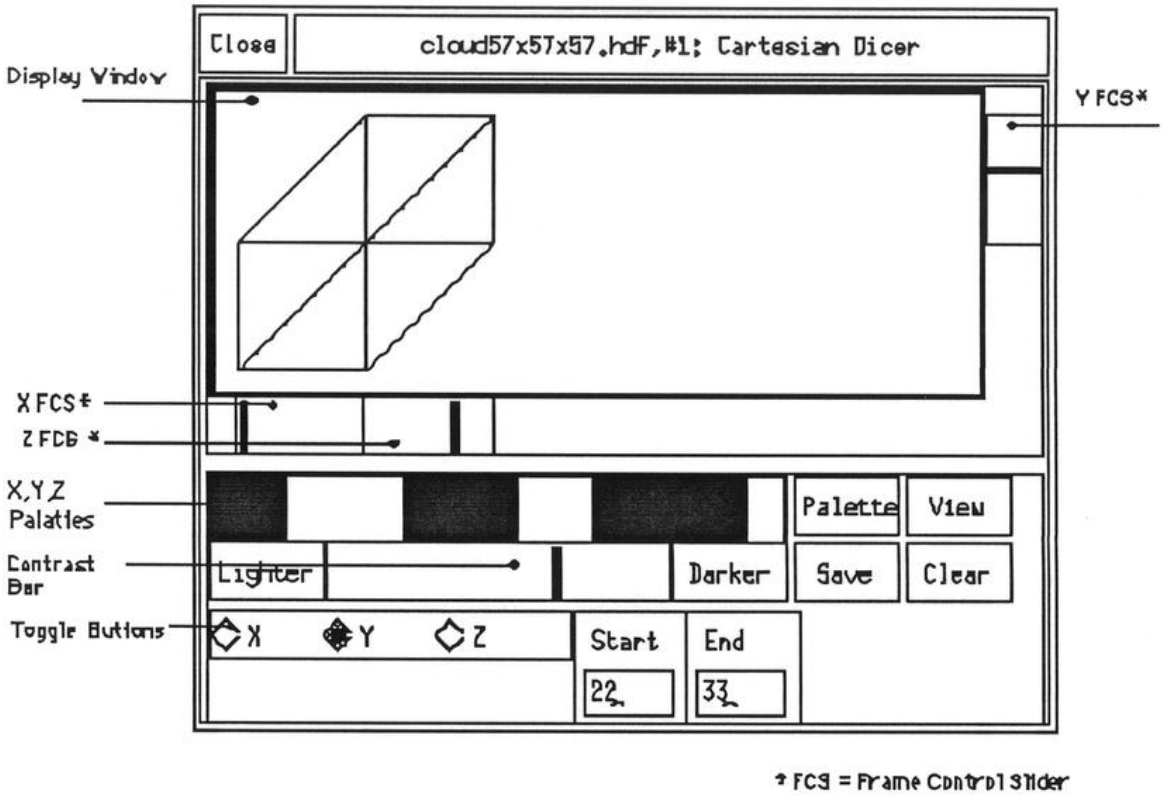

Figure 1.5 Cartesian Dicer (View Slice)

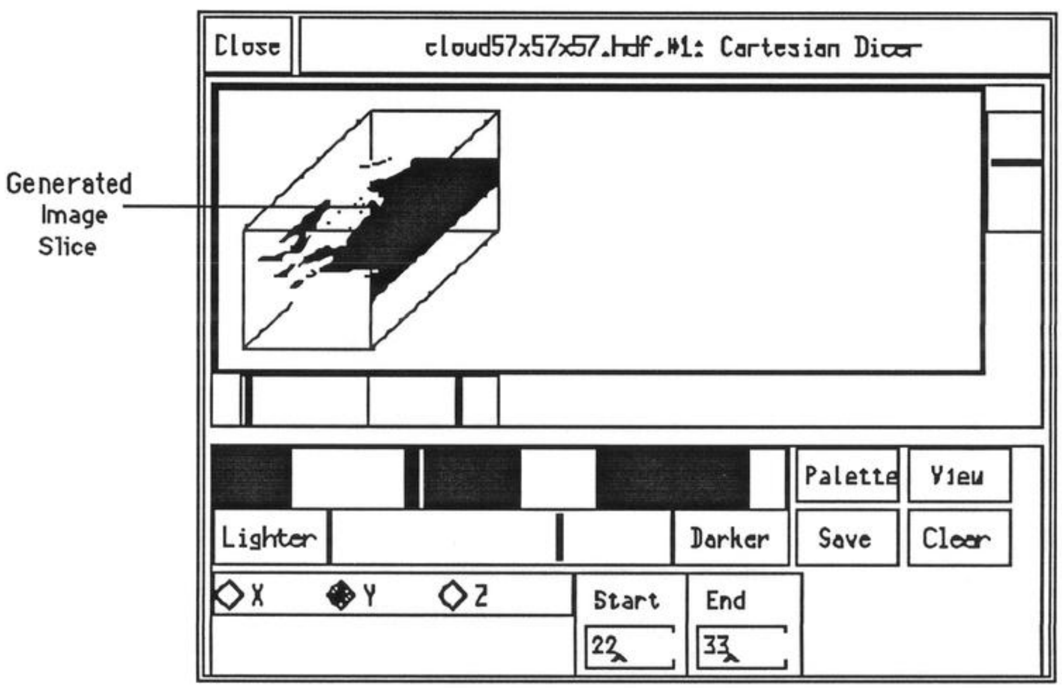

Working with Palettes

You can change the palette in your Dicer window, or in any window which contains a generated dataset image (raster image) by clicking the palette button. A file window labeled, Load a color map, opens.

Changing Palettes in One Raster Window

To change the palette in a single raster window:

1. Change the current directory, if necessary, by entering the path to a new directory in the text box labeled Current Directory being displayed.

2. Click Change to access the new directory and its file contents. If an error message appears in the LogBook, the current directory remains unchanged.

3. Choose the desired palette to be activated by scrolling up or down the file list window and clicking its filename, in this case, carl.pal. The name appears in the Current Filename text box (see Figure 1.6).

Figure 1.6 Load a map File Window

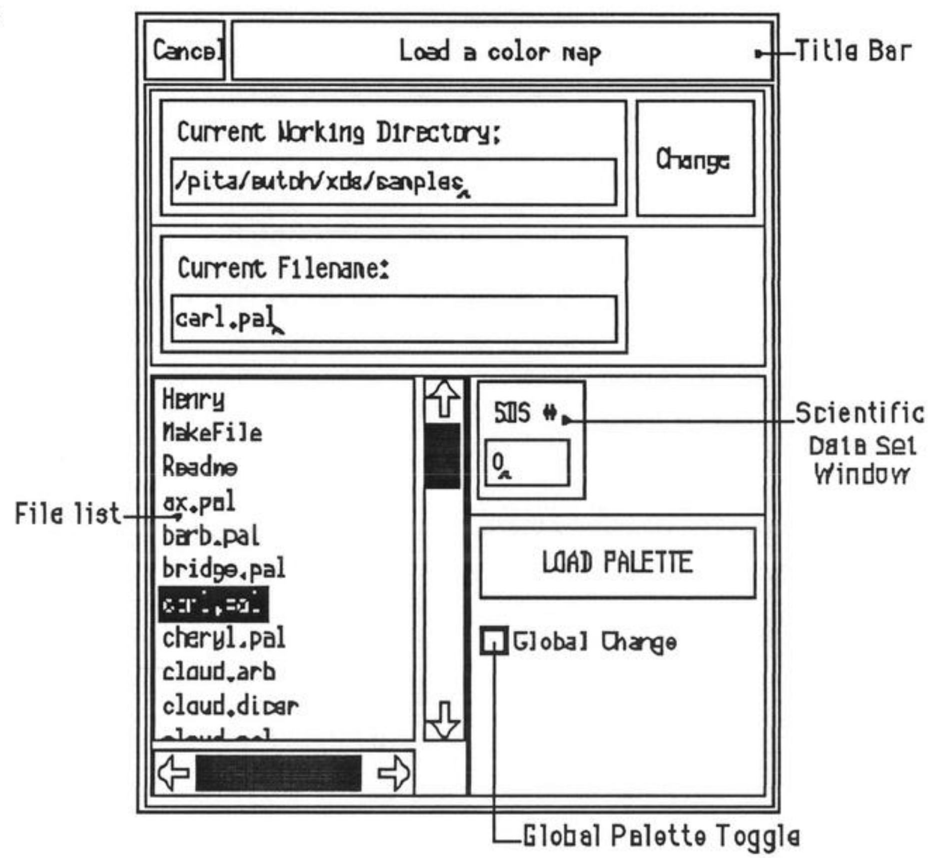

4. Click LOAD PALETTE. When you move the mouse into the raster image window, the former palette changes to the palette you specified.

Note that many image windows contain shading bars with which you may manipulate your palette to create a better view of your image.

Changing Palettes in Multiple Windows

When you change the palette file for a single raster window, X DataSlice provides you the option of changing the palette in every other previously opened raster windows. To do so, activate the Global Change toggle (in the Load a color map file window) by clicking in its box. The designated palette file will now be loaded into all opened raster windows. For more information on raw palette files see Chapter 2, "File Formats in X DataSlice."

Saving Generated Images

To save the Cartesian Dicer image along with its palette, click the Save button located in the lower right-hand corner of the Cartesian Dicer window. A file window, here labeled, Save a Dicer Image, appears. In the text box, Current Filename, enter the HDF filename that will store your Dicer image, and click the Save button to save the image; or, if you want to append an image to the end of an existing file, simply select a filename from the file list (see Figure 1.7), and click Save.

Exiting the Application

To exit NCSA X DataSlice, click Quit in the control window.

Figure 1.7 Save a Dicer Image File Window

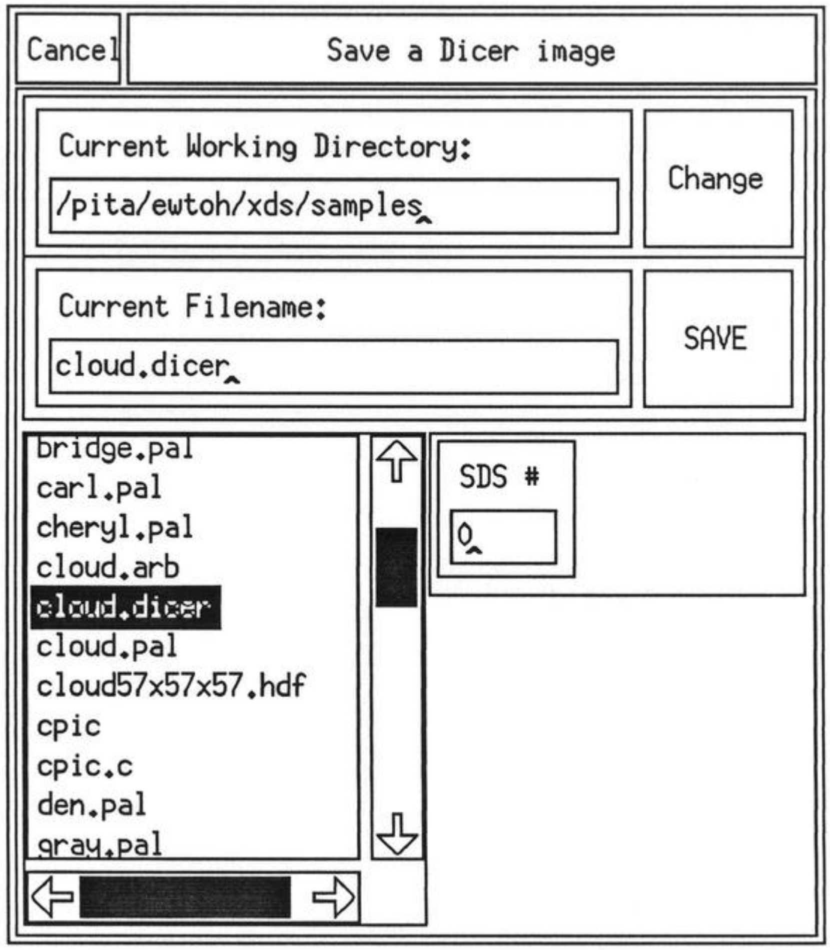

Chapter 2 File Formats in NCSA X DataSlice

Chapter Overview

File Formats

Raw Raster Files

Raw Palette/SEQ Files

HDF Files

How NCSA X DataSlice Reads and Displays Data Files

Chapter Overview

This chapter describes the basic file formats used in NCSA X DataSlice, how the program reads and displays data files, and how to specify the dimensions of raw files. Information is included for obtaining HDF software from NCSA.

File Formats

NCSA X DataSlice reads files containing image data, scientific data, or color palette information. This program can also read in datasets—including user-defined palettes created in applications such as NCSA PalEdit—from raw raster or Hierarchical Data Format (HDF) files.

Raw Raster Files

A raw raster file is a stream of raw, binary, 8-bit raster data in row-major order. Each 8-bit byte corresponds to a pixel in the image. The image is represented in row-major order; that is, the first raster line appears first in the file, succeeded by the next raster line, and so forth. Though raw raster files are easy to create, the raw raster file format is not very flexible and therefore, not highly recommended.

Raw Palette/SEQ Files

A raw palette or SEQ file is a stream of 768 bytes. Raw palette files can store 256 colors, which can be selected from a palette of over 16 million possible colors. Palette files are based on the red, green, and blue representation of color, the RGB color model. The files consist of, in order, 256 bytes of red, 256 bytes of green, and 256 bytes of blue. The 256 color palette entries are calculated by combining the nth element (red), the (n+256)th element (green), and the (n+512)th element (blue) to create the nth RGB component.

In other words, a palette file is a lookup table with 256 entries that tell which color to associate with each of the 256 possible pixel values. Each of the 256 palette entries in the palette is chosen from a master palette of 2^{24} RGB colors. Each palette entry consists of three bytes, one each for red, green, and blue; the first red component, the first green component, and the first blue component, for example, comprise the first palette entry.

When you load a palette file, the hardware color table entries are remapped according to the new palette; that is, each color in the palette file is assigned to its corresponding entry in the hardware color table.

HDF Files

Hierarchical Data Format, or HDF, is a flexible, standard file format designed at NCSA for sharing of graphical and floating-point data among different programs and machines. This general purpose and extensible file format allows you to store raster images, their dimensions, color tables, and annotations all in the same file. In addition, you may store floating-point data, scaling information, dimensions, annotations, and units of measurement in a single file. HDF files are accessible from NCSA software for the Sun Workstation such as NCSA X DataSlice as well as for the Macintosh such as NCSA DataScope, NCSA PalEdit, and NCSA Layout as well as user programs, other NCSA software, and HDF utilities. These files can be used used on such machines as the CRAY X-MP/48, CRAY-2, IBM PC, and Alliant. The portability and usefulness of HDF files is demonstrated in Figure 2.1.

Figure 2.1 HDF Environment

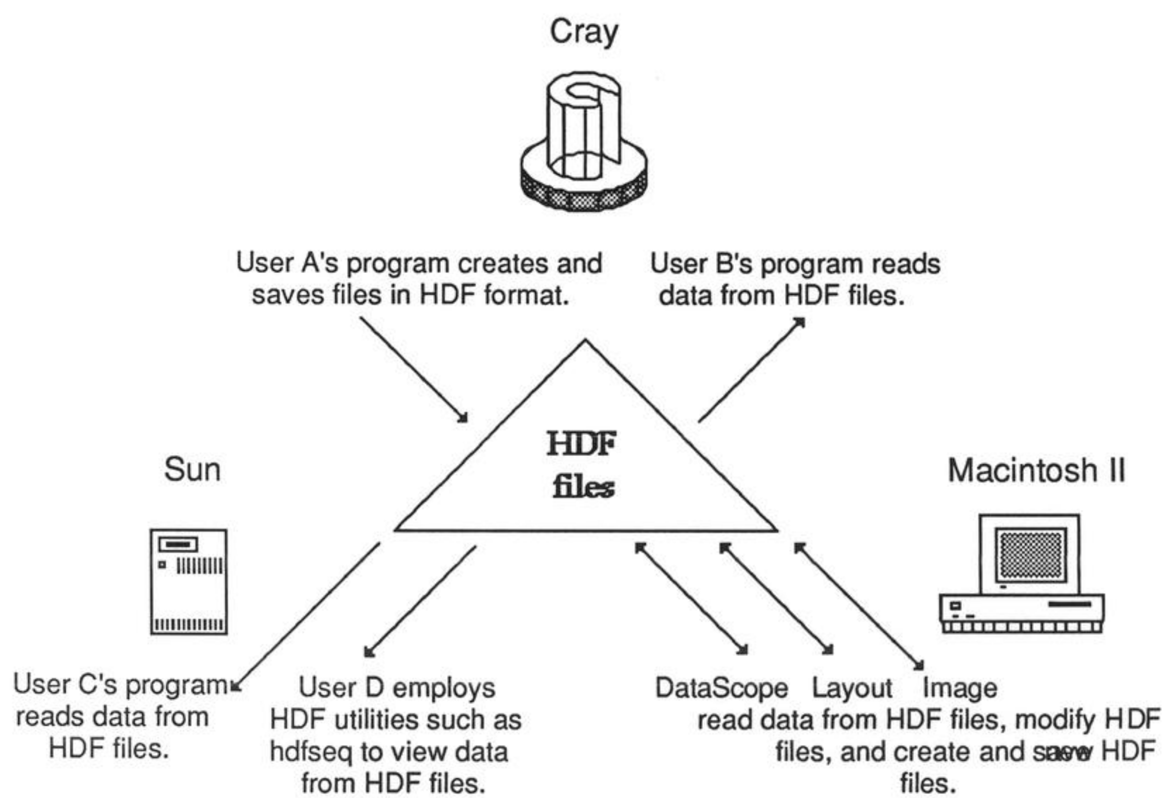

Why Should I Use HDF?

HDF lets you store datasets with extra file information about your data, for example, the dimensions of your image. This makes the files easier to read and manage by programs such as NCSA X DataSlice, and saves you the trouble of tracking this information externally or waiting for delays in image processing.

What Information Goes into an HDF File?

Both raster images and scientific floating-point data can be stored in HDF files for use in NCSA X DataSlice.

A raster image set contains a raster image, together with its dimensions. It may also contain a palette. You may specify that the raster image be stored in compressed or non-compressed form.

NOTE: Images saved by NCSA X DataSlice will be saved in HDF files in compressed form.

A scientific dataset can store scientific data in an array of 32-bit floating-point numbers of any dimension, together with information regarding the rank and size of each dimension. When storing your datasets in HDF files, you have the option of storing other information such as labels for the data and the axes, scales for the axes, and the maximum and minimum values of the data.

How Do I Create an HDF File?

Public domain software is available from NCSA for creating HDF files. NCSA's HDF libraries support both Fortran and C calls on any of the following machines: Cray (UNICOS), Sun (UNIX), Alliant (Concentrix), Macintosh, and IBM PC (MS-DOS).

The best way to store your data in an HDF file is to incorporate calls to the appropriate HDF library in the program that produces your image or scientific data. These calls can store your raw image, palette, scientific data, and other information in an HDF file in proper format.

If you have access to the HDF libraries, you can store floating-point data in your file, and use it directly. NCSA X DataSlice uses floating-point data rather than 8-bit binary data whenever the former is available.

NOTE: With a UNIX-based system, you can use the command line utility called r8tohdf to convert one or more raw raster images and palettes to HDF format.

Where Can I Obtain More Information about HDF?

The 1.0 release of NCSA X DataSlice comes with the source code for HDF Version 2.36. Also, if you are connected to Internet (NSFNET, ARPANET, MILNET, etc.) you can download HDF software and documentation at no charge from an anonymous file transfer protocol (FTP) server at NCSA. The steps you should follow to do so are enumerated below. If you have any questions regarding the connection or procedure, consult your local system administrator or network expert.

1. Log on to a host at your site that is connected to Internet and is running software supporting the FTP command.

2. Invoke FTP on most systems by entering the Internet address of the server:

% ftp ftp.ncsa.uiuc.edu

or

% ftp 128.174.20.50

3. Log in by entering anonymous for the name.

4. Enter your local login name for the password.

5. Enter get README.FIRST to transfer the instructions file (ASCII) to your local host.

6. Enter quit to exit FTP and return to your local host.

7. Review the README.FIRST file for complete instructions concerning the organization of the FTP directories and the procedures you should follow to download the README files specific to the application you want.

Your login session should resemble the following sample, where the remote user's local login name is smith and user entries are indicated in boldface type.

```
harriet_51%  ftp ftp.ncsa.uiuc.edu
Connected to zaphod.
220 zaphod FTP server (Version 4.173 Tue Jan 31 08:29:00
CST 1989) ready.
Name (ftp.ncsa.uiuc.edu: smith):  anonymous
331 Guest login ok, send ident as password.
Password:  smith
230 Guest login ok, access restrictions apply.
ftp> get README.FIRST
200 PORT command successful.
150 Opening ASCII mode data connection for README.FIRST
(10283  bytes).
226 Transfer complete.
local: README.FIRST  remote: README.FIRST
11066 bytes received in .34 seconds (32 Kbytes/s)
ftp>  quit
221 Goodbye.
harriet_52%
```

The README.FIRST file instructs you to copy the HDF README file to your directory and read it before proceeding. Your FTP session should resemble the one listed below:

```
ftp>    cd HDF
```

250 CWD command successful.
ftp> get README
200 PORT command successful.
150 Opening ASCII mode data connection for README (10283 bytes)
226 Transfer complete.
local: README remote: README
2080 bytes received in .14 seconds (15 Kbytes/s)
ftp> quit
221 Goodbye.
harriet_52%

The HDF README file explains how to copy the contents of the HDF directory to your home directory via remote login or anonymous ftp. The precise file transfer procedure varies according to the type of operating system under which you will use HDF—UNICOS or other.

HDF software and manuals are available for purchase—either individually or as part of the anonymous FTP reel or cartridge tapes—through the NCSA Technical Resources Catalog. Orders can only be processed if accompanied by a check in U.S. dollars made out to the University of Illinois. To obtain a catalog, contact:

NCSA Documentation Orders
152 Computing Applications Building
605 East Springfield Avenue
Champaign, IL 61820
(217) 244-0072

You can save all raster plots in HDF single HDF files, i.e a set of animation images are stored in one HDF Raster8 file. You may then reuse the images and palettes in NCSA X DataSlice or any of the other NCSA visualization tools that support HDF such as NCSA X DataSlice or NCSA Layout.

How NCSA X DataSlice Reads and Displays Data Files

NCSA X DataSlice assumes that all data files are to be displayed as two-dimensional data, in 8-bit format, arranged in row-major order, with the origin in the upper-left corner. If the original data was created in column-major order, as in a Fortran program, it must be transposed prior to being read by NCSA X DataSlice.

The term *8-bit format* refers to a process whereby data is scaled onto the numerical values from 0 through 255 and is stored in single bytes, one data element per byte. When data is displayed on the screen, a byte is interpreted as a number from binary 0 through binary 255. The number represents a color from the current palette. For example, a byte that is equivalent to binary

8 is interpreted as the ninth color in the current color palette and is displayed accordingly.

NCSA X DataSlice displays images that have been saved in raw raster, ASCII, or HDF files. To load images from raw files, which do not store information regarding the dimensions of the data, you must specify the dimensions of the data--see the following section.

NOTE: You can scale HDF floating-point data to 8-bit data using NCSA X DataSlice.

Chapter 3 Planes Along Axes

Chapter Overview

Viewing Planes Along the Coordinate Axes

Viewing 2D Planes

- Single Image Window
- Multiple Frame Window
- Unfold Window
- Animation Window

Viewing 3D Planes

- Visable Volume Rendering
- Setting Parameters
- Defining Substances

Using the Animation Window

Chapter Overview

This chapter describes options available in the Planes Along Axes window. It outlines the various options available to manipulate and view both 2D and 3D planes.

Viewing Planes Along the Coordinate Axes

NCSA X DataSlice allows you to view 2D slices of data along the three major cartesian coordinate axes singly, as raster images displayed together with actual data values in a single image window, or multiply, as multiple raster images tiled in a single image window. Multiple images may also be animated.

Viewing 2D Planes

To view 2D slices of data along the three major cartesian coordinate axes:

1. Enable the option, Planes Along Axes, in the Attributes window (Figure 3.1). A Cartesian Display Mode window appears (see Figure 3.2).

Figure 3.1 Attributes Window

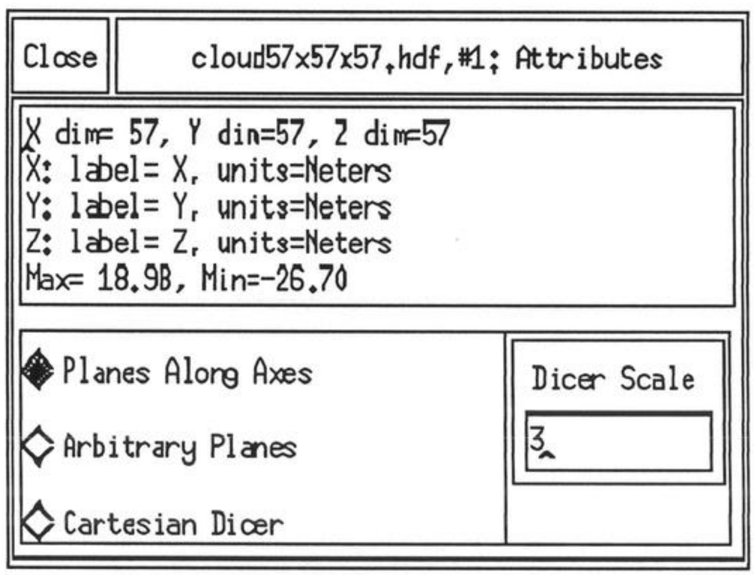

Figure 3.2 Cartesian Display Mode Window

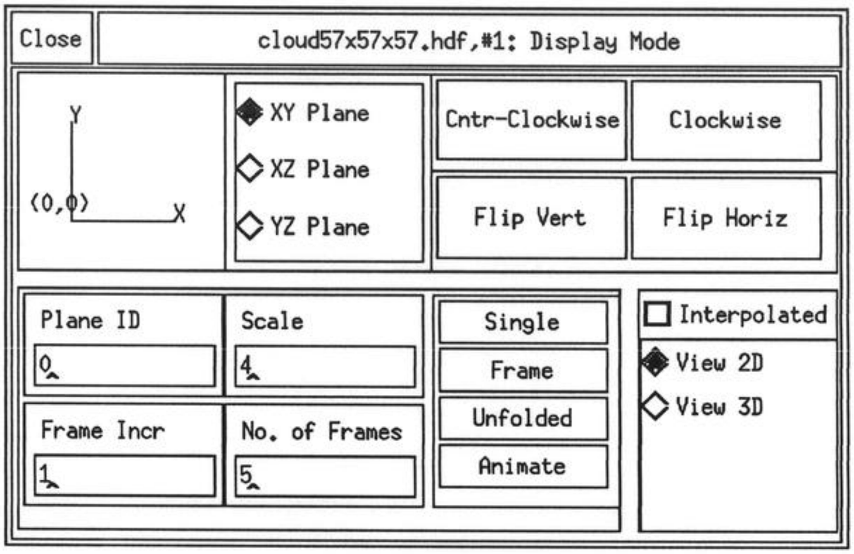

2. Select and orient a plane for viewing data. Click the button labeled XY Plane, XZ Plane, or YZ Plane, respectively, to indicate the coordinate axis of interest.

NOTE: Each of the three faces is parallel to one cartesian axis and orthogonal to the others. Your 3D SDS, if written in C, should have the implicit ordering of data [x][y][z], or row-major order, in which each entry in the array may also be lexically accessed by:

```
data[x+(y*dimz)+(z*dimy*dimz)]
```

where dimy is the Y dimension of the SDS and dimz is the Z dimension. This order is implicitly assumed when you invoke NCSA HDF to create an SDS file.

Click the buttons Cntr-Clockwise, Clockwise, Flip Vert, and Flip Horiz to change the orientation of the axes view with respect to the origin of the data ([0,0,0] in C programs or [1,1,1] in Fortran).

Note that the default position of the origin is on the bottom left hand corner of the plane with the first literal of the axis plane's label going right and the second literal going up. For example, XY plane toggle has the default of the origin at the left hand corner of the dataset, X as its column axis and Y as its row axis. Click Cntr-Clockwise to rotate the axes plane view 90 degrees counter clockwise, Clockwise to rotate it 90 degrees clockwise, Flip Vert to flip it vertically, and Flip Horiz to flip it horizontally.

3. Set parameters for viewing the 2D plane(s).

NOTE: The View 3D button in the cartesian display mode window, which accesses 3D visible volume rendering features of NCSA X DataSlice, requires different parameters and is discussed in the following subsection.

Enter an integer into the text box labeled, Plane ID, where 0 is the foremost plane in the data matrix, to indicate which plane along the specified coordinate axis you wish X DataSlice to consider the *current* plane. The current plane is displayed if you click the Single button, and appears first if you tile several planes in a single window.

NOTE: If you enter an invalid number, that is, one which does not fall withing the range of the dimensions of the SDS, the nearest available plane (the first or last) is displayed.

Enter an integer in the text box labeled Scale to indicate the magnification value which determines the actual pixel size of every numeric entry in the generated image(s). The default value is 4, because 3D datasets are generally expected to produce very small images.

Enter an integer in the text box labeled Frame Incr, to indicate the intervals by which you wish NCSA X DataSlice to increment the frames when tiling multiple images. The default value is 1.

Enter an integer in the text box labeled, No. of Frames, to specify the number of frames to display when tiling multiple images. The user may open up to a maximum of 15 frames on the Frame Window.The default value is 5.

4. Check the box labeled Interpolated to specify that the raster image(s) generated be interpolated by the factor indicated in the Scale text box. By default, generated raster images are non-interpreted.

NOTE: This is quite a time-consuming operation and is most efficiently done in a batch mode. You can economize on time by calculating several slices as a batch using the Animate or Frame automated buttons.

5. Invoke the image(s) by clicking one of the following buttons:

   Single: Opens a Cartesian Single Plane window that contains both the numeric values and raster display of the current frame.

   Frame: Opens a Cartesian Frame window with the number of frames specified. The frames are tiled beginning with the current plane and incremented by the Increment value specified.

Unfolded: Opens a specialized Cartesian Frame window which tiles the specified number of frames along each coordinate axis. This option creates images for each 3D *box* of data for comparison along the dataset.

Animate: Opens a Cartesian Plane Animation window which creates an animation sequence using the number of frames specified, starting with the current plane and incremented with the specified incremental value.

Each of these windows contains a Palette and a Save button. The Palette button allows you to change the color map, and the Save button allows you to store the raster image in a single HDF file. For more information about these options, see Chapter One, "Getting Started."

The following sections describe each Cartesian Image window and its features in detail. For information on using the animation window, refer to the section of this chapter entitled "Animation Window."

Single Image Window

The Single Image window, opening a slice of data for analysis along a coordinate axis, includes a numeric data display and its corresponding raster image (see Figure 3.3). A palette displays the current range of color, which may be changed by clicking the Palette command button.

The numeric data labels identify the vertical and horizontal cartesian coordinate axis. The horizontal and vertical scales display the units in planar coordinates along each axis, respectively. The two-dimensional data list includes all values in the selected plane. When you select an individual entry with the mouse, the corresponding "pixel" is highlighted in the raster image. Similarly, clicking a specific point in the image highlights a corresponding value in the data list. The list can be scrolled vertically and horizontally using the scroll bars.

You may open single image windows of several other planes along the same axis and synchronize all selections between them by setting their synchronize toggles on or off.

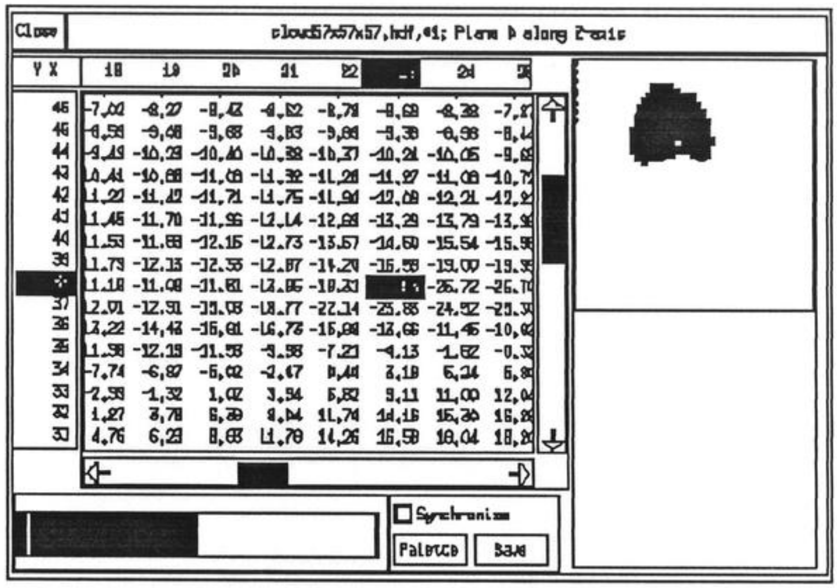

Figure 3.3 Single Frame Window

Multiple Frame Window

The Multiple Frame Window, shown in Figure 3.4, displays a number of raster images, or frames, in increasing frame order and by the increment specified. Each of the frames are labeled by number and axis; the frames are synchronized so that a point selected is highlighted in all the frames. The Palette button allows you to change the color map, and the Save button allows you to store the raster image in a single HDF file.

Unfold Window

The Unfold Window tiles the specified number of frames along each coordinate axis. This option creates images for each 3D *box* of data for comparison along the dataset.

NOTE: No illustration is provided for this window, because of the large size of the window.

Animation Window

The animation window contains an animation mechanism that may be used to quickly display several related raster frames together to give an illusion of motion (see Figure 3.5).Animation adds a 3rd dimension (time) to the 2D slices of 3D data.

To use the animation window:

1. Enter an integer in the Frame text box to indicate the beginning frame of the animation sequence, and click the Set button immediately below. For instance, if the animation sequence contains frames 1 through 10, but you wish to start the animation with frame 2, enter 2.

2. Enter an integer indicating what intervals of frames are to be animated by typing your preference in the text box labeled Skip, and click the Set button immediately below. For

example, if you wish to skip every third frame as your animation is run, you would enter 3.

3. Check the Repeat diamond toggle to specify that the animation be repeated in the same direction as when first started.

4. Check the Reverse diamond toggle to specify that the animation be run continuously from the last image in the sequence to the first, and then first to last.

5. Click the Off diamond toggle to turn the Repeat or Reverse option off.

NOTE: If only one frame exists in your file, the text box below the animation options will contain a solid black bar indicating that the options have no function. No animation can occur for a single frame.

Figure 3.4 Multiple Frame Window

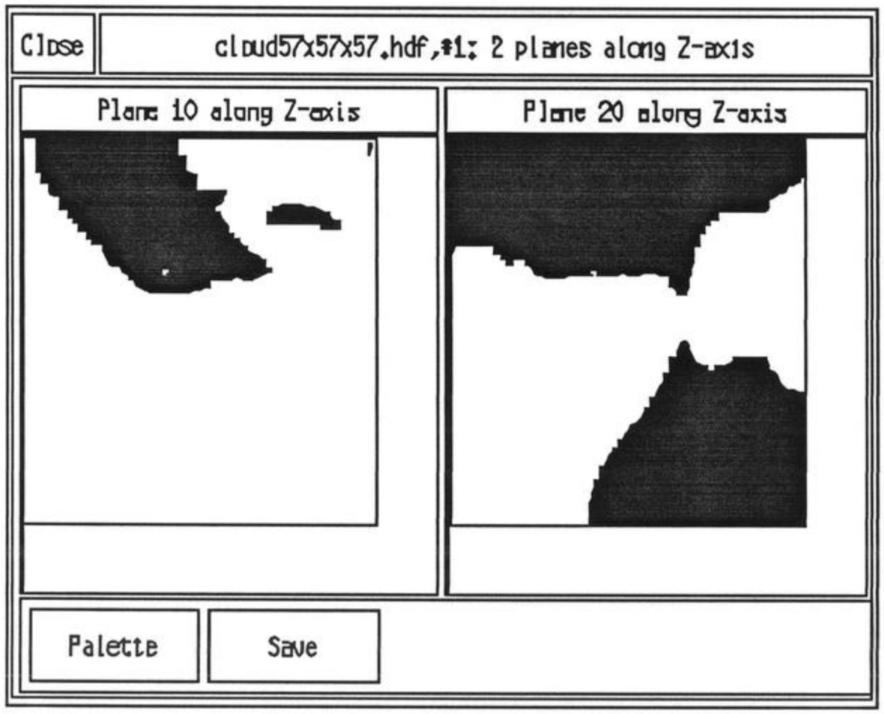

Figure 3.5 Animation Window

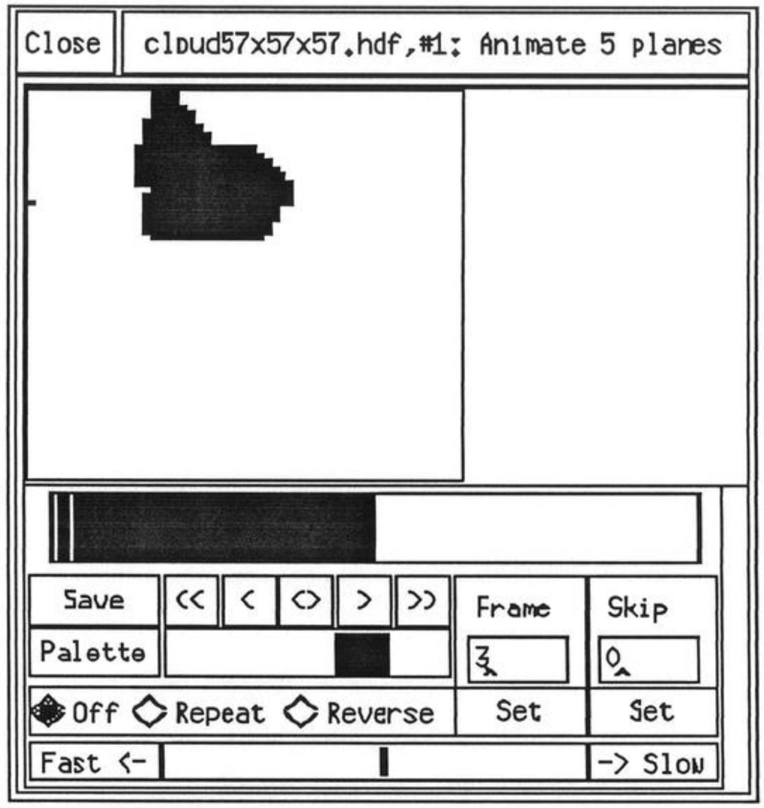

6. Set the animation speed by clicking within the speed bar, located at the bottom of the animation window. You may either click anywhere in the the solid white region of the bar or click on the fast (located on the left end of the bar) or slow button (located on the right end of the bar) before or during animation to set the speed. A slight delay in speed change may occur as X DataSlice adjusts to your preferences.

7. Click the < to single-step backwards or the > button to single-step forward through the animation.

8. Click the << button to move continously backwards or the >> button to move continously forward through the animation.

9. Click the <> button to halt a continous animation. If your animation is running at a fast speed, a slight delay may occur before X DataSlice recognized the interrupt signal.

10. Click the Palette button to change the palette of your image. For more information on changing palettes, see Chapter One, "Getting Started." Notice that the palette bar located directly below the raster image represents the current palette spectrum of the displayed image.

11. Click Save to save your image animation and its current palette. For more information on saving images, see Chapter One, "Getting Started."

Note that the speed of any display is dependent on the size of the images contained within the file.

Viewing 3D Planes

Visable Volume Rendering

The V-Buffer Algorithm generates 3D images, from a surface representation to an armorphous wisp of form, simply by changing the opacity and other parameters that control the rendering.

To access the V-Buffer, you must first load a 3D scientific dataset if you have not done so already. For more information on loading datasets, see Chapter One, "Getting Started." For this discussion, use the sample spine64x64x64.hdf.

In the display mode window, click the appropriate plane, XY, XZ or YZ, and click View 3D. Note that you can only have views parallel to the cartesian axes. The parameters for visible volume rendering appear in the bottom left part of the screen, replacing the View 2D options (see Figure 3.6).

Figure 3.6 Display Mode--View 3D

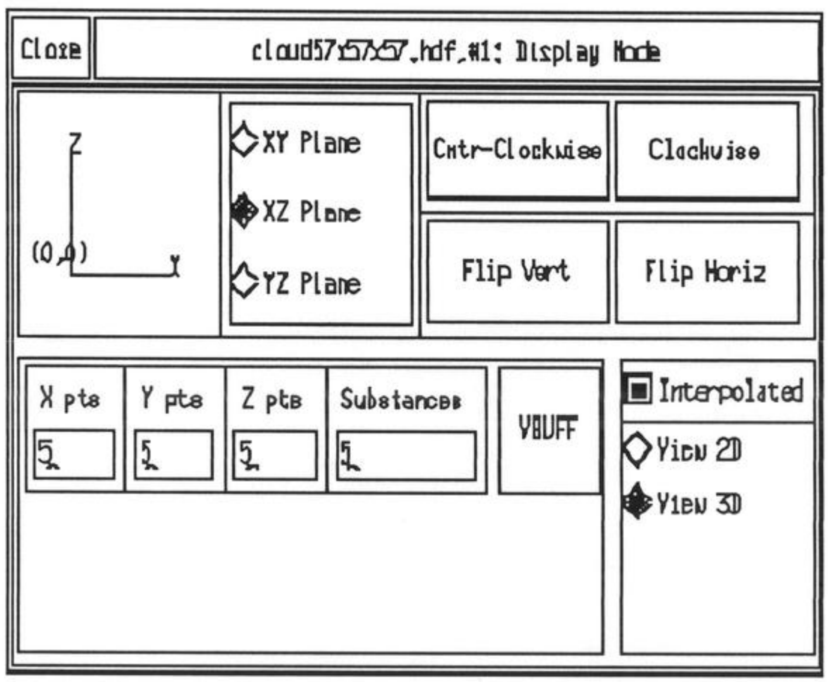

Setting Parameters

The X pts, Y pts, and Z pts are scales that control the number of pixels per side of the computational cube along each axis. The X and Y values control the horizontal and vertical dimensions of the output, while the Z value has no effect on the dimensions of the final image.

Also, indicate the number of substances you want to view in your dataset. Note that the application provides default values. Clicking on Interpolate smooths out the resulting image.

Click V-Buff to set the remaining parameters and to generate the image. See Figure 3.7.

Defining Substances

A substance is defined as a range of values in the 3D dataset. For each substance included in the dataset, you must set lower and upper percentage values and an opacity value. These selections will largely determine the quality of your output. After entering values for the first substance, click Set. Then, pull down the scroll bar to move to the remaining substances. For each substance in the dataset, you must enter its values and click Set.

NOTE: Make sure all values are nomalized between 0 and 1.

For example, imagine you want to view two substances in the spine dataset: the bone and the tissue. Suppose the tissue lies between 23% and 36% of the range of values in the dataset. For the tissue, substance 1, you would set the lower value to 0.23 and the upper value to 0.36 Because the tissue surrounds the bone, you would not want the tissue to be so opaque that it would occlude

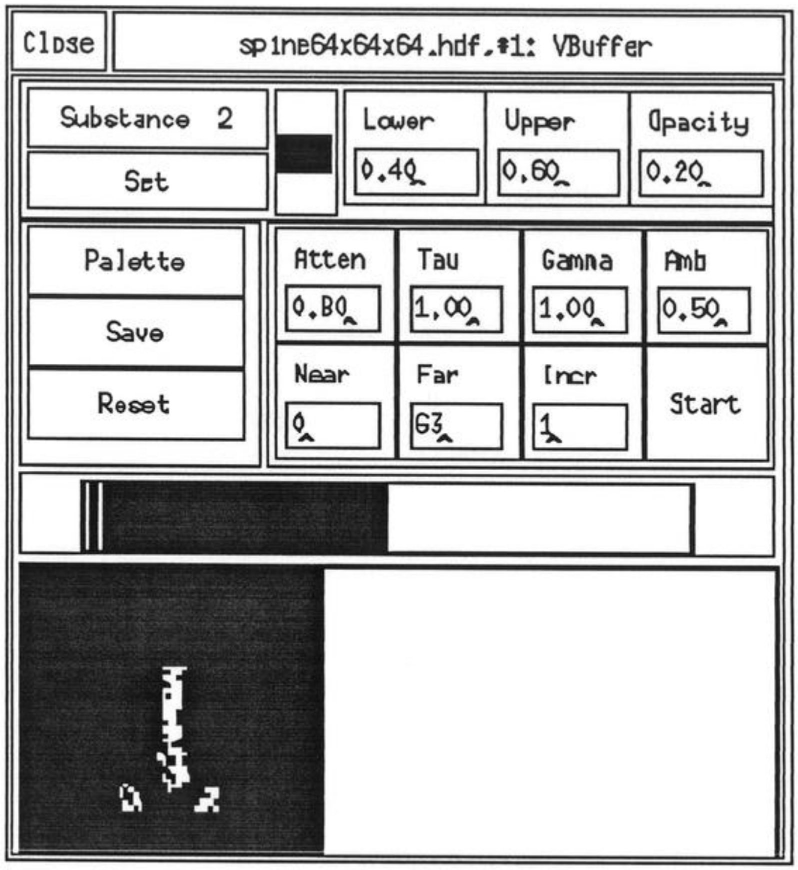

Figure 3.7 V-Buffer

the bone material. Hence you would assign an opacity of 0.1 to the tissue. Thus, to define substance 2 as the bone, set the lower value to 0.48 and the upper value to 0.63 with a higher opacity of 0.3. You may obtain these values by using the other tools supplied by X DataSlice to locate the data values.

Once you have set all substance values, you can fine-tune the image presentation using the following viewing parameters which have been preset with some common default values:

Attenuate [Atten]: Controls the amount of depth cueing of the view. It lies between 0 and 1, and best results occur between 0.5 to 0.7. An attenuaion of 0.7 implies that the intensity attributed to the last plane of data is 70% of what it should be without attenuation (i.e., Atten = 1.0).

NOTE: Do not input either 0 or 1 as the ATTENUATE factor.

Gamma: Controls the spread of density (intensity) values. The higher the gamma entry, the greater the bias is towards larger values in the range of data values.

Ambience [Amb]: Controls the amount of ambient light falling on the dataset. If a view is too dark in general, this parameter may be increased to brighten the image. Usually it ranges from 0 to 1, because the actual intensity of the image also falls within this range.

Enter range values for the Near and Far planes of the image (for example, 0 to 63). If you wish to see the image backwards, enter the high value in Near and the low value in Far (for example, 64 to 1). Finally, enter the increments (Incr) at which you want to process planes of the image.

With all parameters entered, click Start to generate the 3D image. The Control window will indicate each plane being generated together with the maximum intensity accumulated thus far. As soon as out Max becomes non-zero, it implies that a substance has been found in the dataset. The Pause button (which toggles with Start) allows you to stop image generation in order to see the image at any particular point. Click Cont to continue generation, or if necessary, click the Reset button and change parameters.

Clicking Palette and Save, respectively, allows you to load a new palette and to save the V-Buffer image.

Chapter 4 Arbitrary Planes

Chapter Overview

Viewing Oblique Planes

- Manipulating the Image
- Viewing the Image
- Animating the Image

Chapter Overview

This chapter outlines the options available in the Arbitrary Planes window which permits you to manipulate, view, and animate 2D oblique slices of data.

Viewing Oblique Planes

Manipulating the Image

To view arbitrarily selected 2D oblique slices of data:

1. Enable the option Arbitrary Planes in the Attributes window (Figure 4.1). An Arbitrary Display Mode window appears (see Figure 4.2). This window allows you to specify an angle for the slicing plane so that NCSA X DataSlice slices 2D planes from the 3D volume in an arbitrary fashion.

Figure 4.1 Attributes Window

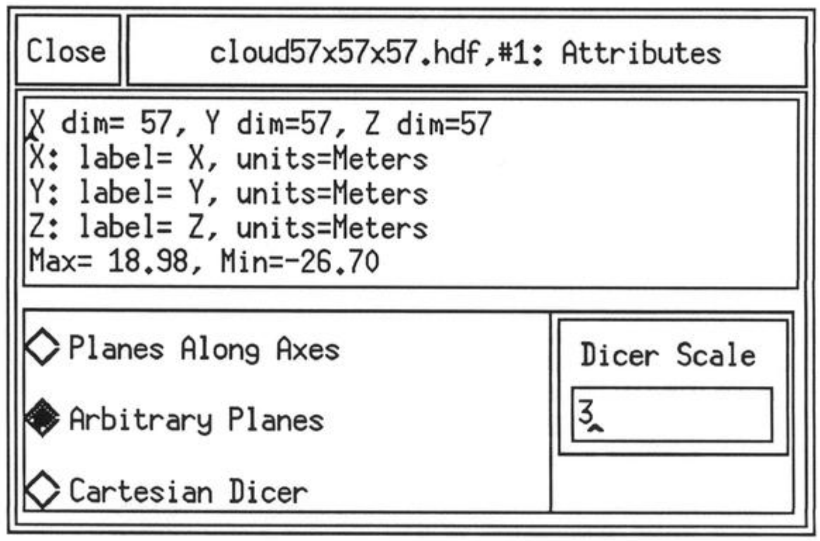

Figure 4.2 Arbitrary Display Mode

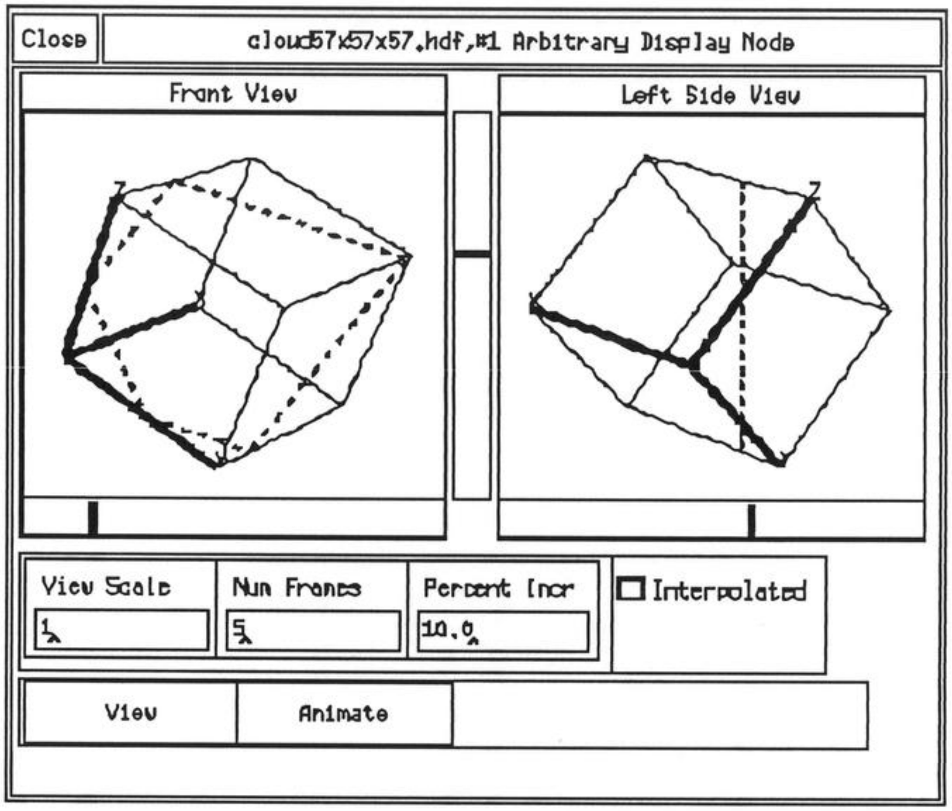

2. Drag the Yaw Slider beneath the Front View drawing or the Pitch Slider beneath the Left Side View drawing to rotate the drawings horizontally or vertically, respectively. The Front View drawing depicts the orthogonal view of the 3D volume, whereas the Left Side View shows a left side view of the volume.

3. Enter an integer in the text box labeled, View Scale, to indicate the desired magnification value. This value is used to determine the actual pixel size of the generated image(s). Note that the speed of image generation is directly dependent on the size of the dataset image and the view scale window.

4. Enter an integer in the text box labeled, No. of Frames, to specify the number of frames NCSA X DataSlice should slice along the orthogonal axis during batch processing (animation).

5. Enter an integer in the text box labeled, Percent Incr, to indicate the intervals at which you wish NCSA X DataSlice to slice the frames (along the orthogonal axis) for batch processing (animation sequences). For example if you were to enter 10, multiple slices, each separated by a 10% distance, would appear within your image.

6. Check the box labeled, Interpolated, to specify that the generated raster image(s) be interpolated trilinearly by the factor indicated in the View Scale text box. If this option is not checked, the default slicing uses the floor of the data values during fractional slicing.

7. Click the View or Animate button to specify that the program either generate an image for the data slice or animate the sequence of 2D data slices, respectively.

Viewing the Image

Clicking View in the Arbitrary Display Mode window produces a generated image of your data slice. Three options common to most image windows in X DataSlice, Close, Palette, and Save, exist for you to manipulate your object (Figure 4.3). For more information on these buttons, refer to Chapter One, "Getting Started."

Animating the Image

Clicking Animate in the arbitrary display mode window opens an animation window (see Figure 4.4). To manipulate your image using the available options, refer to the section, "Animation Window," in Chapter Three, "Planes Along Axes."

Figure 4.3 Arbitrary Display View Window

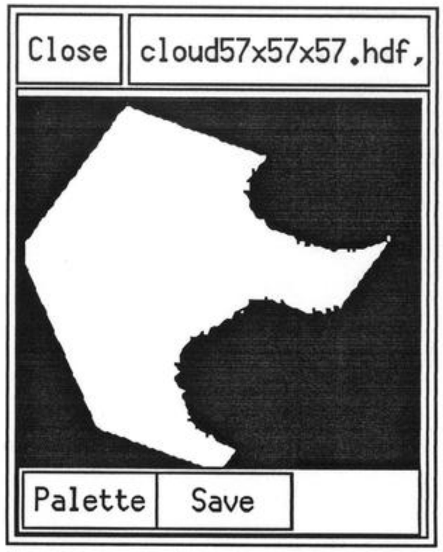

Figure 4.4 Arbitrary Display Animation Window

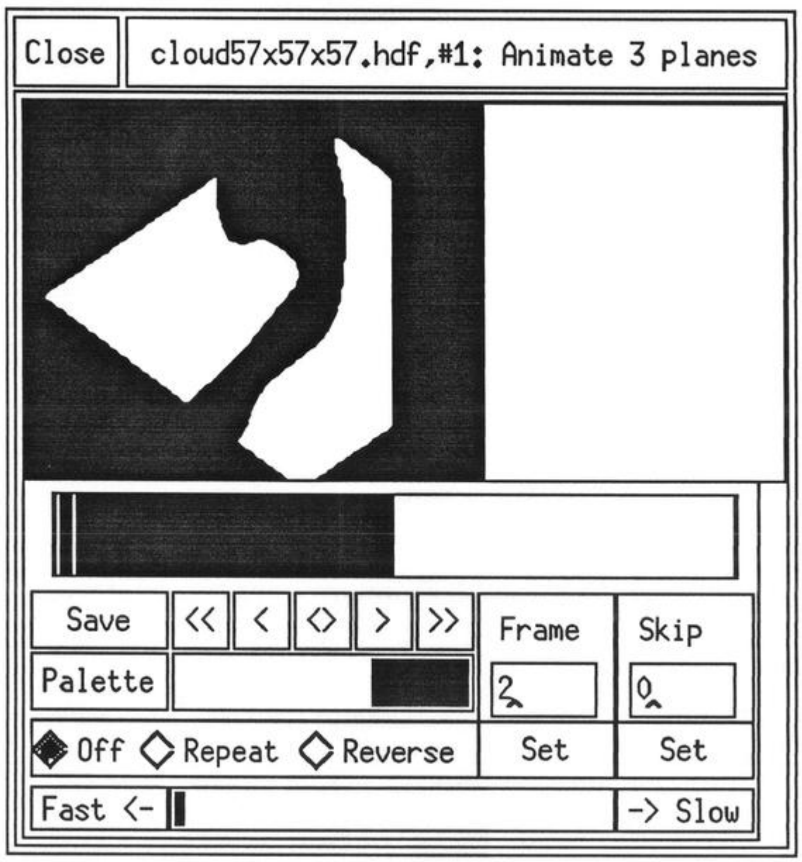

Chapter 5 Cartesian Dicer

Chapter Overview

Viewing Planes of a Dataset

Chapter Overview

This chapter outlines the options available in the Cartesian Dicer window.

Viewing Planes of a Dataset

To view 2D slices of data in their relative positions in the 3D matrix:

1. Enter an integer in the text box labeled, Dicer Scale, in the Attributes window (Figure 5.1). This value indicates indicate the desired magnification value to be used in determining the actual pixel size of the generated image.

Figure 5.1 Attributes Window

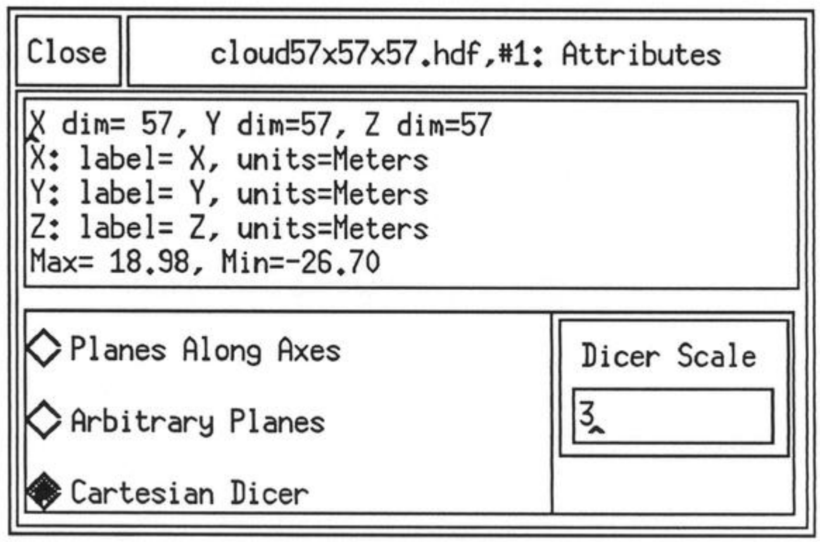

2. Enable the option, Cartesian Dicer, in the Attributes window. The Cartesian Dicer window appears (see Figure 5.2).

3. Click the X,Y, or Z sliders to identify the data slice you wish to depict. The text box labeled, Start, reflects which starting slice (an outline depicts the selected slice) on the respective axis you have chosen. To choose a thickness greater than one, click and drag the selected slider. When you release the mouse button, the Start and End text boxes reflect the range of your slice. Alternatively, you may select the X, Y, and Z toggles, and type the starting and ending planes in the respective dialogs. Click View and the slice appears in the image window.

4. Click the View button after each slice choice to generate the image. The selected slice is represented as a color raster image within the drawing (see Figure 5.3).

Figure 5.2 Cartesiar Dicer Window

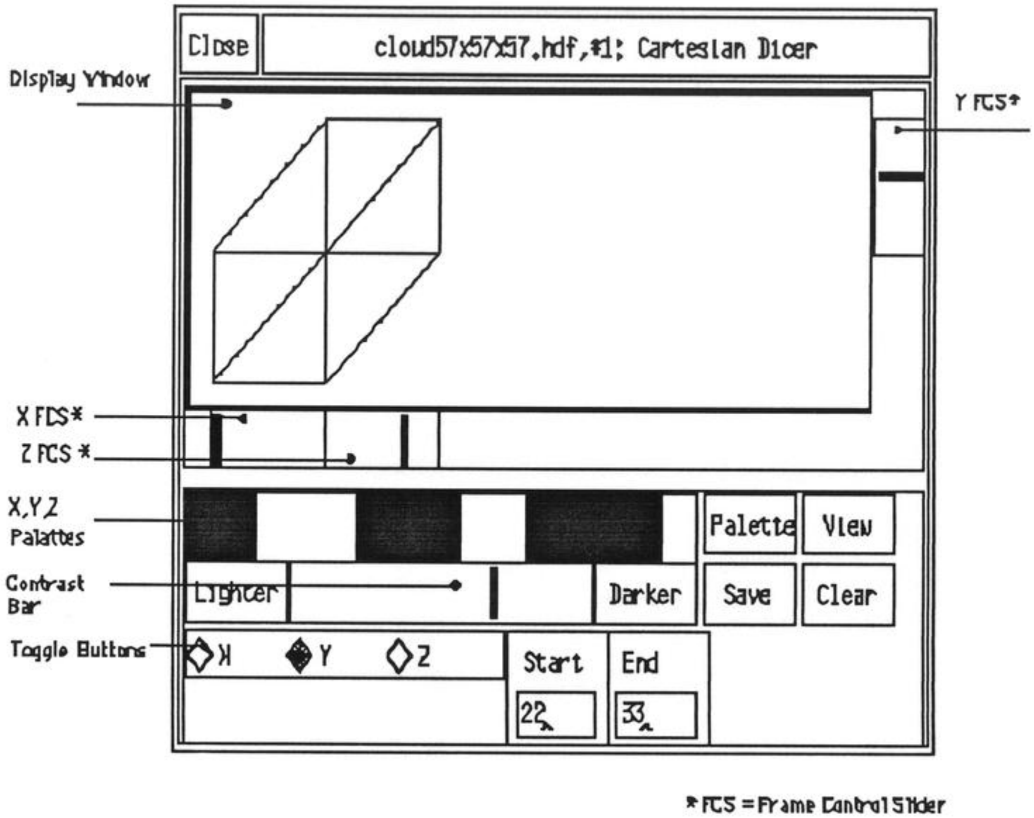

Figure 5.3 Cartesiar Dicer (View Slice)

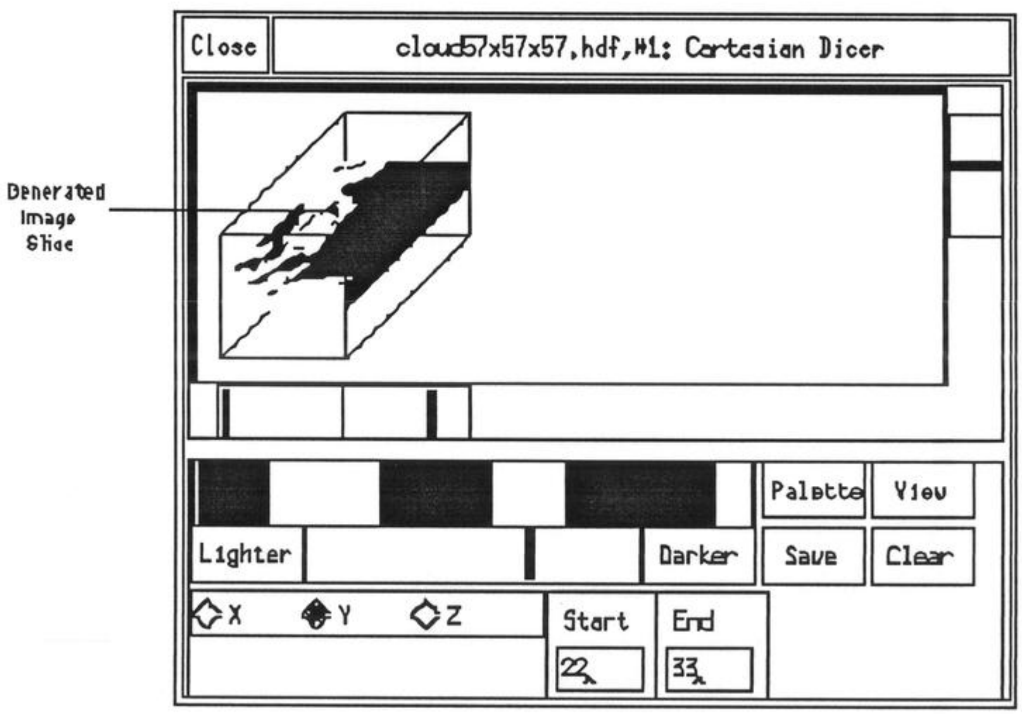

5. Repeat Steps 3 and 4 for each slice you wish to view. Figure 5.4 illustrates multiple slice selections.

Figure 5.4 Cartesian Dicer (Multiple Slices)

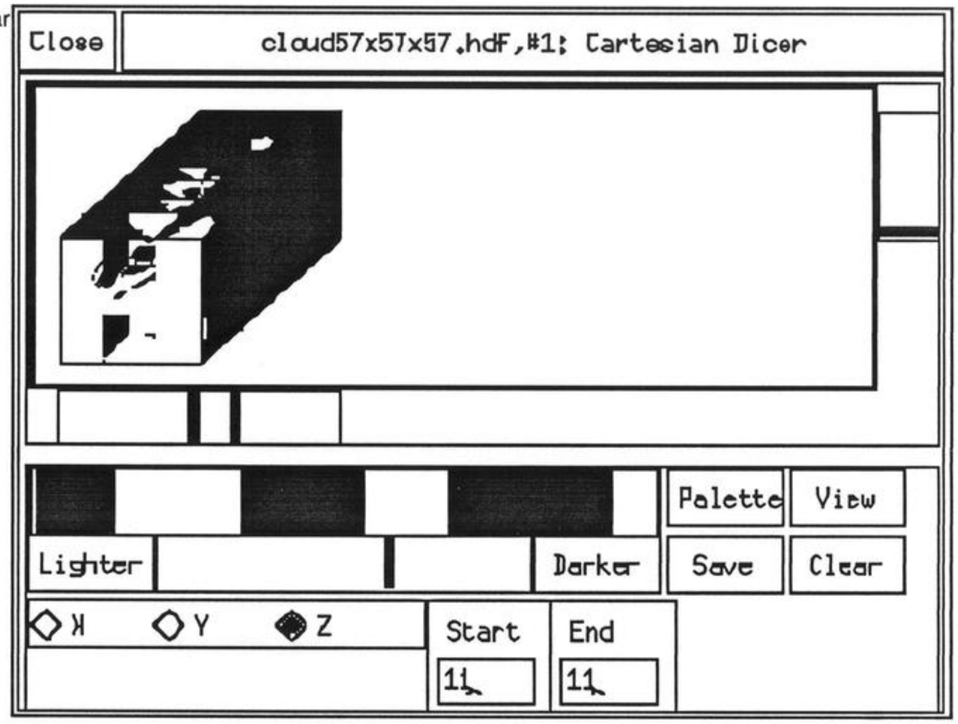

6. Drag the contrast slider a number of places to increase the contrast between the images. The three sections of the palette represent the X, Z, and Y axis, respectively, and reflect the corresponding changes in intensity. The front plane becomes lighter and the more distant planes become darker as you drag the palette slider to the right.

7. Click the Palette button to load a palette that better represents the data. A file window appears so that you may select and load the desired palette. For information on loading palettes, see Chapter One, "Getting Started."

8. Click the Clear button to clear all images from the Cartesian Dicer window.

9. Click Save to save the current palette in a file together with the raster images currently depicted in the cartesian dicer window. A file window appears so that you may save the information into a specific file and directory. For information on saving files using NCSA X DataSlice, see Chapter One, "Getting Started."

Chapter 6 Animate 8-bit Raster and Process SDS

Chapter Overview

Loading 8-bit Raster Images

Loading the Raster Image

Animating the Raster Image

Processing Scientific Datasets

Chapter Overview

This chapter explains the procedures for loading and animating a saved 8-bit Raster image and processing a scientific dataset.

Loading an 8-bit Raster Image

Loading the Raster Image

To load a saved 8-bit raster image, click the Load 8-bit Raster button located in the control window. A file window labeled, Load a RASTER8 set, appears on the screen (see Figure 6.1).

Figure 6.1 Load a RASTER8 set File Window

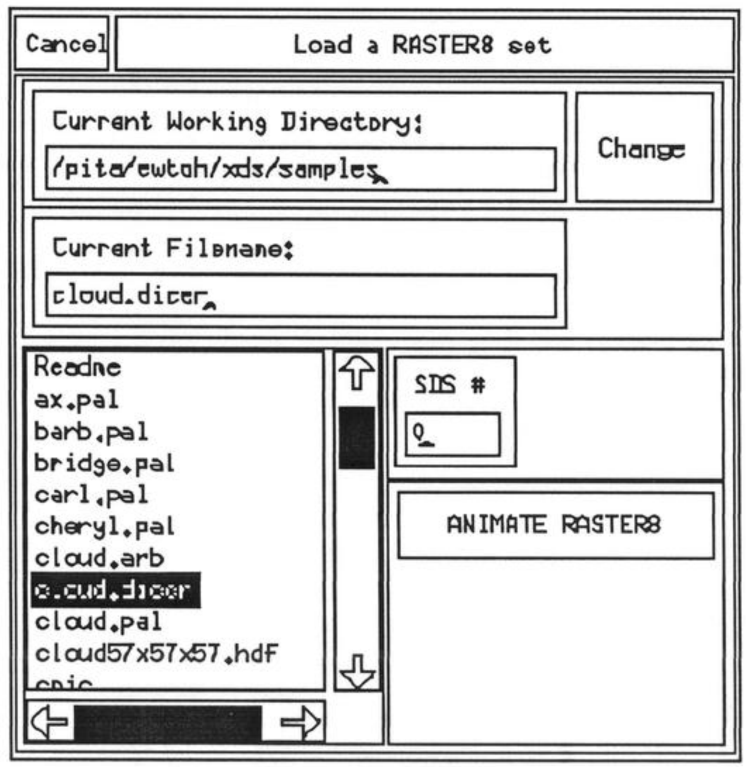

1. Select your image in the current directory, or in another by entering the path to a new directory in the text box labeled, Current Directory being displayed. Click Change to access the new directory and its file contents. If an error message appears in the LogBook, the current directory remains unchanged.

2. Choose the desired Raster8 file to be loaded by scrolling up or down the file list window and clicking on its filename. The name appears in the Current Filename text box. The LogBook reports if the file is indeed a 3D SDS HDF file or a Raster8 HDF file. If the file is of HDF format, the message shows the number of scientific datasets contained in the file.

3. Specify which scientific dataset you wish to load into memory from the selected file by entering its number in the text box labeled SDS #. Notice that the text box SDS # reads 1 by default, indicating that the first SDS in the file will be loaded unless you specify otherwise. Be aware that if you enter a value above the designated number of datasets present, the file window closes and an error appears in the LogBook:
   ***HDF ERROR: Cannot load SDS *4*.

4. Click Cancel to leave the box without loading your entry or Animate Raster8 to load all the frames into memory for animation. The file window disappears and an Animation window appears (see Figure 6.2).

Figure 6.2 Animate 8-bit Raster Image Window

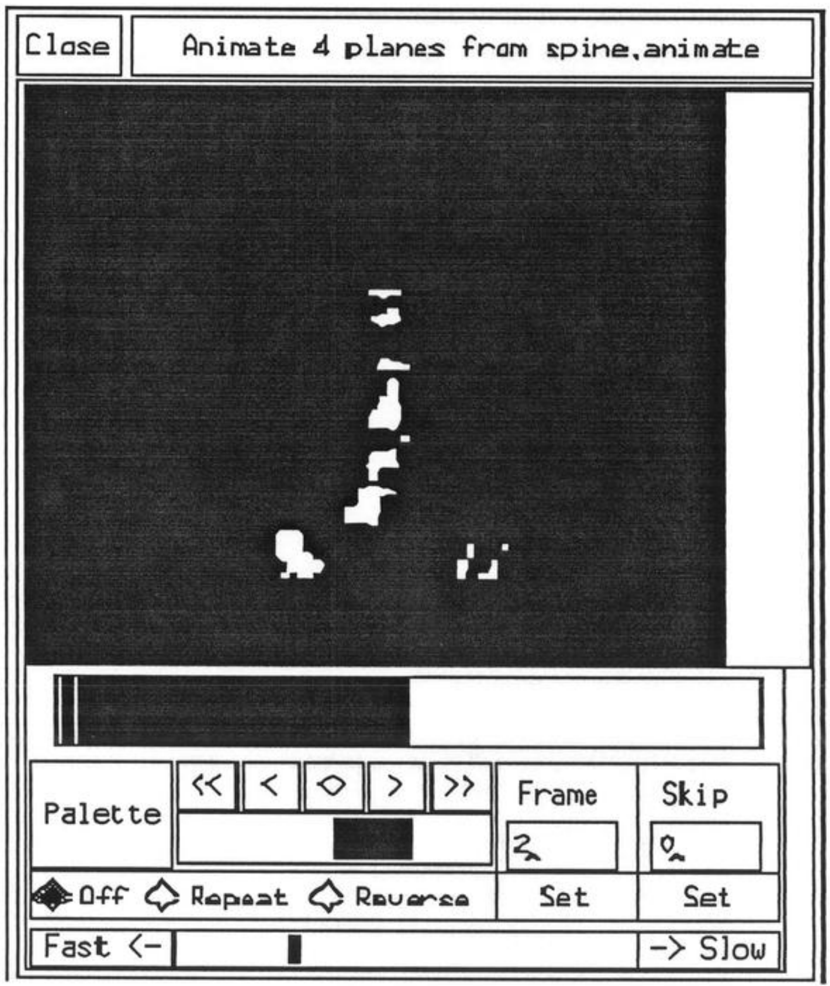

Animating the Raster Image

The animation window contains an animation mechanism that may be used to quickly display several related raster frames together to give an illusion of motion. For more information on animating your raster image using the animation window, refer to Chapter 3, "Planes Along Axes."

Processing Datasets

By selecting the Process a SDS from the Control window, you have the option of expanding a small dataset by trilinear interpolation During interpolation, each subvolume bounded by every eight adjacent data points is assumed to contain a scalar function linearly dependent on all the eight values, so that trilinear interpolation can be performed. The processed datasets, once they have been saved, may be reused by NCSA X DataSlice or any of the other NCSA visualization tools that support HDF such as NCSA Layout and NCSA Image.

To process or interpolate a scientific dataset:

1. Click Process a SDS located in the Control window. A File window labeled, Process & save a SDS (see Figure 6.3), appears on the screen.

2. Select your image in the current directory, or in another by entering the path to a new directory in the text box labeled, Current Directory being displayed. Click Change to access the new directory and its file contents. If an error message appears in the LogBook, the current directory remains unchanged.

3. Select the 3D HDF scientific dataset file containing the dataset to be interpolated and input the filename of the HDF to which the enlarged dataset is to be stored.

4. Enter the number of the desired SDS contained within the chosen file.

5. Enter an integer in the Scale box to dictate the magnification value determining the actual pixel size of the generated image(s). The default value is 3, because 3D datasets are generally expected to produce very small images. Remember that the larger your dataset and bigger your scale value, the longer the dataset will take to process.

6. Click Cancel to abort the File window or click Save to start the interpolation process. Your dataset is automatically saved after the interpolation.

When NCSA X DataSlice has finished interpolating and saving the dataset, the file window disappears from the screen and the LogBook reports on the completed process.

Figure 6.3 Process & save a SDS Window