

RECONSTRUCT™



User Manual
Version 1.1.0.0

Table of Contents i


Chapter 1: Quick Start	1
Getting Started	1
Main Window.....	2
Create a Series	2
Import Section Images	2
Calibrate Pixel Size	4
Align Sections	5
Step 1.....	5
Step 2.....	5
Step 3.....	5
Step 4.....	6
Step 5.....	6
Step 6.....	6
Step 7.....	6
Step 8.....	6
Step 9.....	6
Trace Profiles	6
View Three-Dimensional Object.....	7



Chapter 2: Software History	9
History	9
Credits.....	10



Chapter 3: Software Overview	12
Series, Sections, Domains, Traces, and Objects	12
Data File Organization.....	13
XML File Formats	13
Image File Formats.....	16
DXF Imports	16
DXF Exports	17
3D Output Formats	18
System Requirements	18
Operating System.....	18
OpenGL.....	18
Hardware Requirements.....	18



Chapter 4: User Interface	20
Windows and Dialogs.....	20
Resizing and Moving Windows.....	20

Switching Between Windows.....	20
Main Window.....	21
Main Menu	22
Tools Window.....	23
Status Bar	24
Section List Window.....	24
Thumbnails Window.....	26
Domain List Window.....	27
Trace List Window.....	28
Trace Palette Window	29
Object List Window	30
3D Scene Window.....	31
Z-Trace List Window	32
Distances List Window	33
Exiting or Aborting.....	33



Chapter 5: Series..... 35

Creating a Series	35
Opening and Closing Series.....	35
Saving a Series	35
Series Options.....	36
General.....	36
Names/Colors	37
Grids	37
3D	38
Lists	38
Wildcards and Displaying a List	39
Movements	39
Proxies.....	39
Thumbnails	40
Autotracing.....	40
Exporting and Importing Images	43
Exporting Trace Lists	43
Exporting Traces as Line Drawings.....	43
Importing Line Drawings.....	43
Importing Series Data.....	43



Chapter 6: Sections..... 45

Section Indexing.....	45
Creating a Section.....	45
Saving a Section	45
Displaying a Section.....	46
Centering.....	46
Zooming and Panning	46

Paging Sections	47
Blending and Flickering Sections	47
Listing all Sections	47
Editing Section Attributes	48
Movement Lock	48
Thickness	48
Section Numbering	49
Deleting Sections	50
The Undo Stack	50
Thumbnails.....	50
Exporting Section Images	52



Chapter 7: Domain Images..... 57

Creating a Domain	57
Importing Multiple Images	57
Montages.....	60
Listing Domains.....	60
Selecting a Domain	61
Domain Boundary	61
Domain Attributes.....	62
Name.....	62
Source	62
Pixel Size.....	63
Contrast and Brightness	63
Color Channels.....	63
Moving and Montaging Domain Images.....	63
Merging a Domain.....	64
When to Merge a Domain	64
Reinitializing a Domain.....	67
Contrast.....	67
Boundary	67
Transform	67
Copying Domains.....	67
Hiding Domains.....	68
Proxies	68
Proxy Rule of Thumb for single domains.....	70



Chapter 8: Traces

Creating a Trace	71
Open Point-by-Point Drawing.....	72
Closed Point-by-Point Drawing	72
Freehand Drawing.....	72
Ellipse Tool.....	72
Rectangle Tool	73

Line Tool	73
Stamping Predefined Shapes.....	73
Creating a Grid of Traces	73
Semi-Automatic Tracing.....	75
Z-trace as a 3-D Ruler.....	76
Setting Trace Attributes for Drawing	76
The Trace Palette.....	77
About Trace Names	78
The Trace List	79
Editing Traces	80
Merge	81
Reverse	82
Simplify.....	82
Smooth	84
Modifying Trace Attributes	85
Trace Name.....	85
Border Color	86
Fill Color	87
Fill Style.....	87
Hidden	87
Simplified.....	88
Closed	89
Adding a Comment.....	89
Trace Clipboard.....	89
Selecting Traces	90
Deselecting Traces.....	90
Zooming to Selected Traces.....	90
Finding Traces.....	91
Tracing with Wildfire	91
Wildfire Editing.....	92
Propagating Wildfires to Adjacent Sections.....	92
Generating Wildfires throughout a Region.....	93



Chapter 9: Calibration and Measurement	94
Units of Measurement	94
Data Precision	94
Image Pixel Size.....	95
Calibration Grids and Scale Bars	96
Calibrating Traces	96
Drawing Calibration Traces	96
Calibrate Selected Traces	97
Resize Traces, Set Pixel Size, or Scale Section.....	97
Apply the Calibration to Other Sections.....	98
Section Thickness	99
Calibrating Section Thickness	99

Z-Distance	100
Status Bar Measurements	102
Domain List Measurements	103
Length	103
Area	103
Midpoint X	103
Midpoint Y	103
Trace List Measurements	104
Length	104
Area	104
Centroid X	104
Centroid Y	105
Min X	105
Min Y	105
Max X	105
Max Y	105
Z Position	105
Thickness	105
Object List Measurements	105
Count	106
Surface Area	106
Flat Area	106
Volume	107
Z-Trace List Measurements	108
Length	108
Distances List Computation	108
Algorithm	108



Chapter 10: Alignments and Movements	113
Alignments	113
Aligning Selected Traces	113
Aligning Entire Sections	114
Evaluating Alignments	114
Undoing and Adjusting Alignments	115
Repeating an Alignment	115
Montaging Domains by Point Correspondences	116
Step 1	117
Step 2	117
Step 3	117
Step 4	117
Step 5	117
Aligning By Correlation	118
Step 1	118
Step 2	118
Step 3	118

Step 4	118
Tips for Aligning Sections.....	119
Section the Material Appropriately.....	119
Beware of Ambiguous Alignments.....	119
Avoid Aligning the Same Features Every Time	120
Don't Over-Align.	121
Avoid Quadratic Aligns.	121
Choose the Starting Point Carefully.....	121
One Strategy	121
Movements.....	121
Transformations.....	121
Keyboard Increments	122
Movement Type-In Dialog	123
Undoing Movements.....	124
Repeating Movements.....	124
Propagate	124
Repeating Movement Sequences.....	125



Chapter 11: Objects.....	126
The Object List	126
Modifying Objects.....	127
Changing Object Attributes.....	127
Renaming Objects	127
Hiding and Unhiding Objects	128
Simplifying Objects	128
Deleting Objects	128
Creating Objects	129
Copying Objects	129
Smoothing Objects	129
Creating Z-traces at Midpoints.....	129
Calculating the Distances Between Objects.....	129
3D Object Representations	130
Traces	131
Trace Slabs	131
Trace Midpoints	132
Box	133
Boissonnat Surface	133
Draw closed traces	134
Use clockwise traces to define outer boundaries.....	134
Use counterclockwise traces to define holes	134
Trace without loops or self-intersections.....	134
Do not intersect traces of the object.....	135
ACOS Domain Error.	135
Cylinder	135
Ellipsoid	135

Sphere	136
3D Parameters	136
Shift whole object by	136
and sections by	136
Normals	137
Facets	138
Size	138
Boissonnat Surfaces	138
Z-Trace List	138
Modifying Z-traces	139
Generating Z-traces	139
Smoother Versions	139
Grid at Z-trace Points	140
Trace Substitutions	140




Chapter 12: The 3D Scene

OpenGL Rendering	141
Composing a Scene	141
Adding Objects	141
Removing Objects	142
Changing Object Appearance	142
Diffuse Color	143
Emissive Color	143
Ambient Intensity	143
Transparency	144
Specular Color	144
Shininess	144
Filled or Wireframe Faces	144
Viewing a Scene	145
Rotating, Panning, and Zooming	146
Spin	146
Background Color	147
Saving a Scene	147
VRML 2.0	147
Traces	148
Trace Slabs	148
Trace Midpoints	148
Box	148
Boissonnat Surface	148
Cylinder	148
Ellipsoid	148
Sphere	148
VRML 1.0	149
DXF	149
Bitmap	149

JPEG	149
360° Bitmaps	149
Scene Information	149

 Bibliography	151
--	-----

 Appendix I	152
RECONSTRUCT™ Frequently Asked Questions	152
General	152
Images	152
Moving and Aligning	152
Traces and Trace Editing	152
Measurements	153
Objects and The Object List	153
Surfacing and 3D Output	153
ANSWERS	154
General	154
Images	156
Moving and Aligning	157
Traces and Trace Editing	158
Measurements	160
Objects and The Object List	161
Surfacing and 3D Output	162
Keyboard Overlay	166



Chapter 1: Quick Start

This chapter gives the steps for reconstructing serial sections. You will learn the basics of using RECONSTRUCT™ to import a series of images, view and align the sections, trace profiles, and generate a 3D view of the traces.

Note: *Spending an hour skimming through the entire Help Manual could save you many hours of frustration.*

Getting Started

The only requirements for getting started are that the RECONSTRUCT™ application is installed and that a set of image files is available for input. To install the software, sign in and download the latest version from the Users Group http://groups.yahoo.com/group/reconstruct_users/files/. Run the downloaded file and a dialog should appear asking you whether you want to install RECONSTRUCT™. Select Yes and proceed through the installation dialogs until the process is complete.

To test your installation, start RECONSTRUCT™ and select *Open...* from the Series menu. If this is the first time you've installed and used RECONSTRUCT™, a series file called *example.ser* should be visible in the Open dialog. Open the example series. One can find the example series in the installation directory as well, normally *C:\Program Files\Reconstruct*. When the series opens you should see some text on the screen. Follow these text instructions to get familiar with the basic operations of RECONSTRUCT™.

This chapter only covers a small fraction of RECONSTRUCT™ functionality. To get a better idea of the full functionality, review the Key and Mouse commands listings available through the Help menu. Then open the *Options...* dialog from the Series menu and review the contents of all the Options tabs. This will give you an idea of the features of the program that can be reconfigured to suit your purpose. Details about these features can be found in subsequent chapters.

This manual and a list of frequently-asked-questions (Appendix 1) were updated in 2009.

Main Window

When RECONSTRUCT™ is started a main window appears along with the status bar, main menu and tools window (Fig 1.0). The main window will appear black until a series is opened or sections are imported.

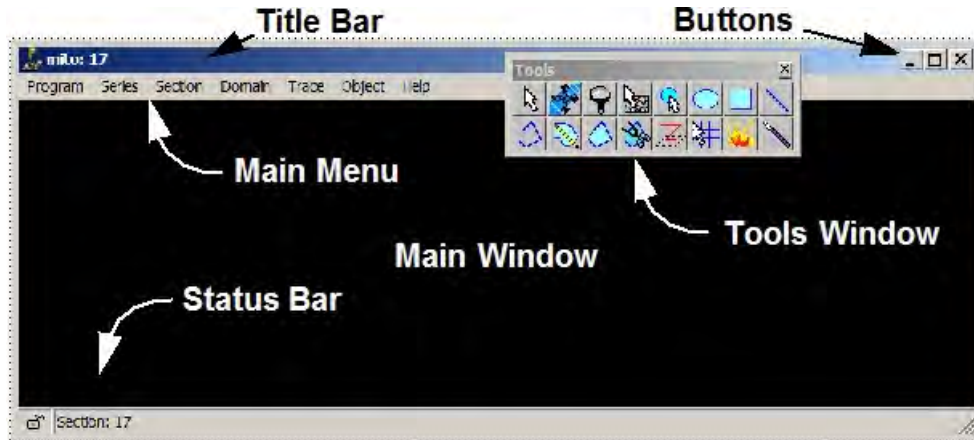


Figure 1.0. RECONSTRUCT™'s window layout. A main window with a narrow border and a black interior is overlaid with the main menu, floating tools window and a status bar.

A number of keystrokes and mouse actions can be used in the main window (see Chapter 4; Table 4.2). These are mostly for navigating between sections in the series and for creating or manipulating section data. Mouse actions in the main window may depend on which tool has been selected from the tools windows. These actions are discussed in more detail in later chapters. However, moving the mouse wheel in the main window always has the same effect as Page Up and Page Down.

Create a Series

Start the RECONSTRUCT™ program and from the Series menu select the *New...* menu item. A file open dialog will appear. Navigate to a folder on the hard drive where the new series files can be written. If desired you may create a new folder from this dialog. The word "newSeries" should appear in the Filename box near the bottom of the dialog. This will be the name of the series files created to store the series data. Select the Save button in the New Series dialog to create the series.

Import Section Images

Import section images into the series using the *Images...* item from the Import submenu of the Series menu. This will open a dialog for creating the section files from a set of image files (Fig. 1.1). Image files can be in a variety of formats including Windows Bitmaps, GIF, TIFF, and/or JPEG. Each image file should contain a single image (not an image stack) representing one section in a series, and should use a representation of 8 or 24 bits per pixel.

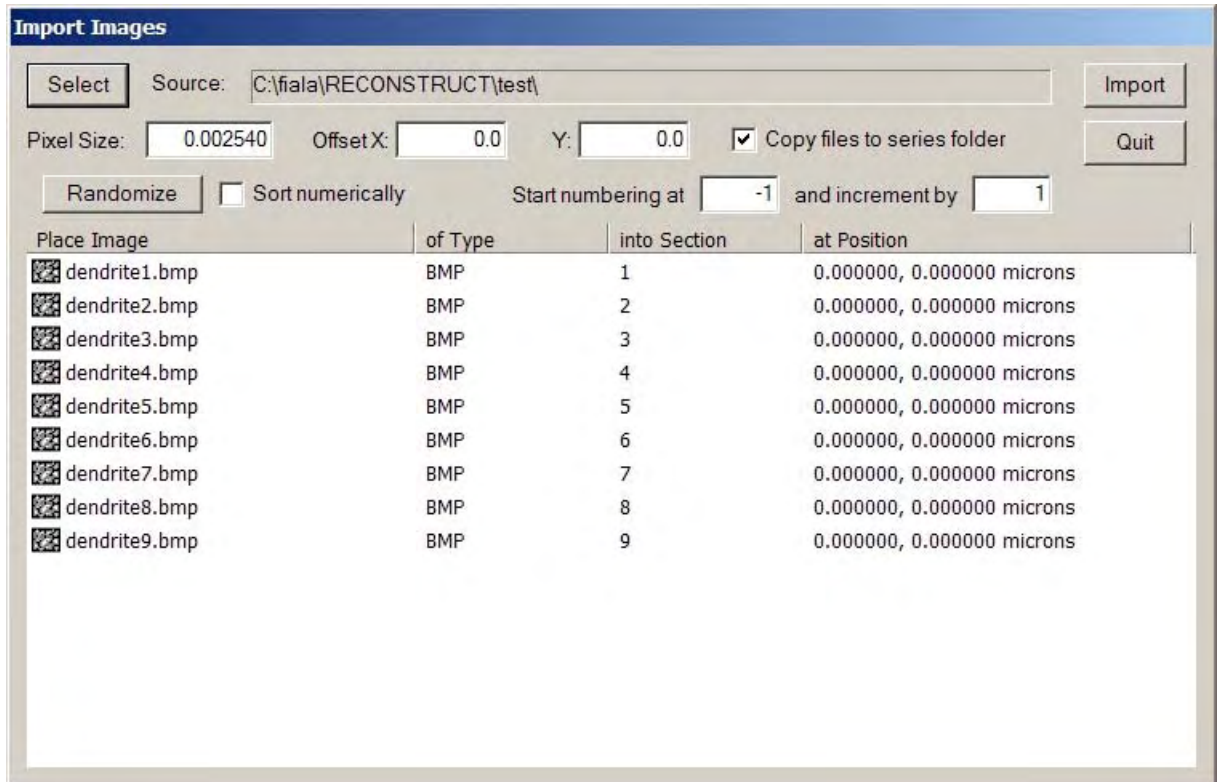


Figure 1.1. The Import Images dialog box after images have been selected using the Select button. Since the file names contain the section index, automatic numbering is realized by having a -1 specified in the First Section box.

Choose image files using the Select button. This button will open a file selection dialog. In this dialog navigate to the folder containing the image files and highlight the image files using the mouse. Multiple files can be highlighted using the Ctrl or Shift keys. The filenames will appear in quotes in the Filename box when multiple files are selected. When all desired images have been selected, use the Open button to transport the filenames into the Import Images list.

Figure 1.1 shows the Import Images dialog after some images have been selected. The source directory is given at the top of the dialog. The filenames of recognized image files and their formats are listed in the left-most columns of the list. The third column gives the section number to which each image will be imported. Set the first section number of the sequence by typing the number, typically 1, in the First Section box. If the sections are not numbered correctly and the filenames contain numbers indicating the correct order, you can use the Sort Numerically checkbox to reorder the files in the list. For finer control of section numbering you may have to import images in stages.

Create the section files by selecting the Import button in the dialog. As the image files are added to the series they will be deleted from the list. After the images have been imported you can repeat the process to add more images to your series. When all images have been added, use the Quit button to close the Import Images dialog.

To view the sections, open the *List sections...* from the Section Menu and double-click on one of the listed sections. This will move the display to that section in the series. Press the Home key. After a brief delay, the image of the first section will be centered in the

main window. Use the Page Up and Page Down keys to advance section by section through the series. Verify that the images were imported as expected with the correct section numbers. The number of each section appears in the title bar of the main window.

Calibrate Pixel Size

An important parameter during image import is the pixel size, which we ignored in the previous section. This parameter specifies the size of image picture elements (pixels) in terms of the units of measure of the series (Chapter 9). If the pixel size in the Import Images dialog was accurate, no further calibration is required. However in most cases the pixel size should be adjusted using the image of a calibration grid or scale that was photographed with the series. This is best done before any alignments or tracing, since changing the pixel size later can invalidate these adjustments.

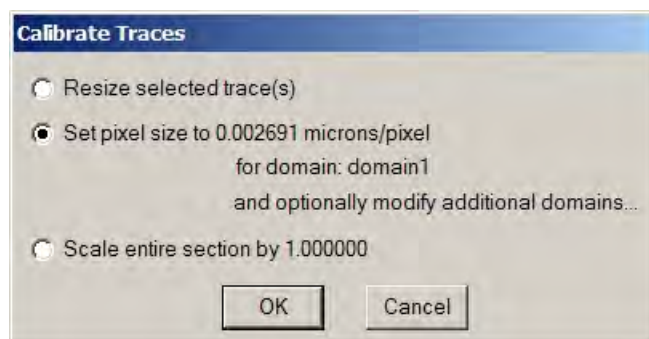
The process of calibration in RECONSTRUCT™ involves drawing traces of known length and then entering that length value into the program. Select the Line Tool (Fig. 1.2) from the Tools window and draw a line trace of known length on a section containing a calibration grid image. (For more details on calibration grids and lines see Chapter 9.)



Figure 1.2. The Tools window. This window should be visible at the top of the RECONSTRUCT™ main window. Use the left mouse button to select a tool. The Line Tool used for calibration lines is the 8th button on the top row. The Pencil Tool is the 10th button. The Wildfire tool button is depressed and so a description of the tool appears in the title bar.

After the calibration trace is drawn, select the *Calibrate...* item from the Trace menu. A dialog box will appear for each calibration trace displaying the current length of the trace. Enter the actual length of the trace in place of the current length.

When the lengths for all selected traces have been entered, a Calibrate Traces dialog will appear (Fig. 1.3) offering three options. Select the second option to set the correct pixel size for the current section. When the "Set Pixel Size on All Domains?" dialog appears, set the First Section to the first section of the series and the Last Section to the last section. Verify the value in the "Set pixel size to..." box and click the OK button. This will set the pixel size of all images in the series to the value determined from the calibration image. The pixel size for any image can be verified by using the Domain List while viewing the section in the main window.



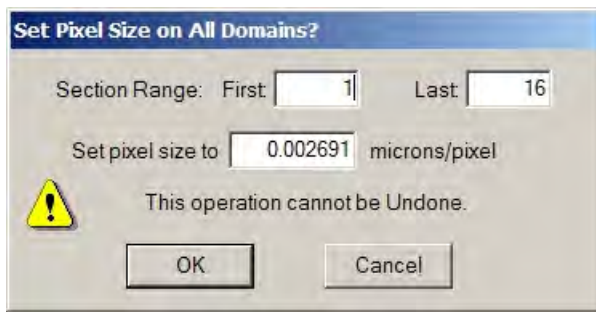


Figure 1.3. The Calibrate dialogs as filled out for a global calibration of the pixel size of images. Top: The first dialog requests that a pixel size calibration be done for the image on which the calibration lines were traced. Bottom: The next dialog allows this same pixel size to be applied to every image in every section of the series.

Note that after calibration of pixel size, the calibration lines are no longer accurately matched to the calibration image on which they were drawn. This is to be expected. Since these calibration lines are no longer useful they can be deleted from the calibration section. An additional check on the calibration would be to redraw one or two calibration lines and verify that their length is now correct.

Align Sections

Tissue that has been physically sectioned and then imaged must be realigned for accurate 3D measurements and visualizations. Section images can be manually moved into relative alignment using keyboard commands (Chapter 10), but it is usually much easier to align sections by specifying correspondence points and then using one of the Align commands from the Trace menu. This option is outlined in the following steps. It will only work if appropriate landmarks are visible in adjacent sections. For best alignment results, read Chapter 10.

Step 1.

Go to a section which will serve as the stationary reference. Then page to the adjacent section that will be aligned to the stationary one. Press the `/' key to switch between the two sections.

Step 2.

In the Trace menu, select *Palette...* to open the Trace Palette window. Click on the first palette button and pick up the Stamp Tool from the Tools window. (It is assumed that this is the default palette for a new series such that the first palette button is an orange circle with the default name "a\$+").

Step 3.

Use the Stamp Tool to place traces at points to be aligned. Use the `/' key to alternate between sections as points are entered. Enter the landmarks in the same order on both sections so that traces of corresponding points will be numbered the same. To delete a point if you make a mistake, use Ctrl-D to deselect all the points, pick up the Arrow Tool (1st tool button) and select the point to be deleted, then use the Delete key to delete it. Use Ctrl-S to reselect all the points to be aligned.

Step 4.

After entering 3 or 4 correspondence points, switch to the section that will be moved. Verify in the title bar that the stationary section is not displayed.

Step 5.

Choose *Linear* from the "Align section" submenu of the Trace menu. When the section is not unlocked this submenu will be titled "Align traces" instead of "Align section". Enter Ctrl-L to unlock section movement. The Status bar will display an open lock image. If traces are accidentally aligned instead of the images, use Ctrl-Z to undo the trace movement. If the message "Alignment needs more traces" is displayed, make sure you have entered at least 3 points on each section and that all points are selected (Ctrl-S) on both sections. If this message still appears make sure the corresponding traces in each section have the same name. In other words, "a1" in section 1 must be matched to a trace "a1" in section 2, "a2" must match "a2", and so on.

Step 6.

After *Linear* has moved the section, use ` / or blend (spacebar) to confirm that all parts of the section are aligned (Chapter 10). If the alignment is acceptable, lock (Ctrl-L) the section to prevent further movements.

Step 7.

If the alignment is not acceptable, use Ctrl-Z to undo the movement. Check the names and positions of the traces entered and modify or add traces as needed. Go to Step 5. As a last resort, delete all traces by entering Ctrl-S followed by the Delete key. Do this for each of the two sections and then start over with Step 3.

Step 8.

Once the alignment is acceptable you can delete the alignment traces from both sections. If you decide to keep the alignment traces for later editing, you will need to switch the color and name of future correspondence points to avoid confusion with the existing ones. Switching to a new trace name ensures that traces with the same numbering will be created on both sections. To switch correspondence point names, simply select the second button on the Trace Palette. Later, when you repeat this step after aligning the next section, go back to the first palette entry, alternating between trace names after each alignment.

Step 9.

Page Up (or Down) to the next section to be aligned and go to Step 3.

Trace Profiles

Information about sectioned objects is created by tracing the profiles of the objects on the sections. To do this, return to the first section in the series by using Page Down or by selecting the first section in the list generated from the *List sections...* item of the Section menu.

From the Series menu, select *Options...* This opens a dialog that will allow you to configure the default parameters for tracing. Select the Names/Colors options tab at the top the dialog. Put the cursor in the Name edit box and type a name for the object to be traced, such as "myObject". This will be the name given to the profiles you draw. All profiles with the same name are considered part of the same object. Hit Enter or select the OK button to close the Options dialog.

The Tools window (Figure 1.2) should be visible at the top of the main window. Use the cursor to select the tool button that looks like a pencil. The cursor will change to a pencil when the cursor is moved back into the main window.

Place the tip of the pencil on the edge of the profile to be traced as shown in Figure 1.4). Press and hold down the left mouse button. Drag the tip of the pencil around the profile and release the mouse button when the tip gets near the start of the trace. The trace will be completed (closed) and filled with a color.



Figure 1.4. Drawing a trace with the freehand drawing tool. (left) Start of the tracing. The pencil is positioned at the profile to be traced and the left mouse button pressed and held down. (middle) Just before release of the mouse button. The pencil left a trail where it was moved with the mouse from the start position. (right) Right after release of the mouse button. The trace is closed and filled with color.

Press the Page Up key to advance to the next section. The trace on the previous section will be saved automatically as long as the default save options were not changed in the Series Options dialog. Repeat the process of tracing the profile with the mouse. Continue paging and tracing until several sections have been traced. Use the Page Down button to go back and verify that the data was saved correctly.

View Three-Dimensional Object

Select *List objects...* from the Object menu to see the list of objects defined by traces in the series. Use the left mouse button to double-click the object name in the list (or use a single mouse click and select *Add to scene* from the Scene menu at the top of the list window.)

A three-dimensional view of the object traces will appear in a 3D Scene window after a brief delay. You can rotate the object representation by placing the mouse in the 3D Scene window and pressing and dragging with the left mouse button. You can zoom in or stop the object's spin by pressing the right mouse button. Pressing both mouse buttons together allows you to pan the scene.

To surface the object, open the Series Options dialog and select the 3D tab. In the 3D options, change the Generate option from Traces to Boissonnat surface and click OK. Go

back to the Object List window and again double-click on the object name. The object will be surfaced and replace the trace representation in the scene.

Note that traces must conform to certain criteria for Boissonnat surfacing to work correctly. For example, traces of the same object within a section cannot overlap. For a detailed discussion of tracing rules related to Boissonnat surfacing see Chapter 11: Objects.



Chapter 2: Software History

History

The origins of RECONSTRUCT™ began in 1995 with the development of a Windows 3.1 program for tracing serial sections. WinV8 was developed in the Image Graphics Laboratory (Kristen Harris, Director) at Children's Hospital, Boston, as a continuation of the DOS-based V8 program developed by John Davis.

Quickly WinV8 was replaced by a new Win32 application, IGL Trace. Version 1.0a of IGL Trace was completed in August of 1997, with NIH funding to the Mental Retardation Research Center at Children's Hospital. IGL Trace proved useful for identifying, tracing, measuring, and reconstructing 3D objects in serial sections in a number of scientific studies (Fiala, 2005; <http://www3.interscience.wiley.com/cgi-bin/fulltext/118735636/PDFSTART>).

IGL Trace was designed to operate on aligned sections. These can be images obtained from electron or light microscopy. For electron microscopy, where each section is imaged separately, it is necessary to realign the images for accurate reconstructions. IGL Align, also developed in the Image Graphics Laboratory, was a first attempt at aligning grayscale images. Funding from NIH's Human Brain Project to Children's Hospital allowed the development of a more sophisticated alignment program, sEM Align, in 1999. sEM Align was developed to allow on-screen alignment of larger images than was possible through the use of IGL Align. This is accomplished by displaying and moving scaled versions of the original images on screen. Images can be incrementally moved using keyboard input or alignments can be computed from a set of feature correspondences specified by the user. To use the alignments computed by sEM Align, a new set of aligned images are rendered for input into IGL Trace.

Most of the functionality of sEM Align and IGL Trace is incorporated into RECONSTRUCT™. By combining alignment and tracing into one program, RECONSTRUCT™ allows images to be processed more efficiently. Tracing can be done directly on the transformed images and alignments can be easily modified. Converters are available to convert existing sEM Align and IGL Trace series to ones compatible with RECONSTRUCT™. See <http://synapses.clm.utexas.edu/tools/reconstruct/formats.stm> for details.

RECONSTRUCT™ development was funded, through April, 2004, by NIH's Human Brain Project under a grant to Dr. Kristen Harris. A Users Group was created at that time to continue to support the RECONSTRUCT™ user community. The RECONSTRUCT™ Users Group is available to everyone at http://groups.yahoo.com/group/reconstruct_users/. A RECONSTRUCT™ Developers Group has been established at http://groups.yahoo.com/group/reconstruct_developers/ to provide a forum for collaborative, open-source development. Anyone may contribute to the next version of RECONSTRUCT™ by joining the RECONSTRUCT™ Developers Group. Members of the Developers Group gain access to the source code and may participate in developing the next version of RECONSTRUCT™. The developers group also provides additional information and links to the free development tools used for creating the application.

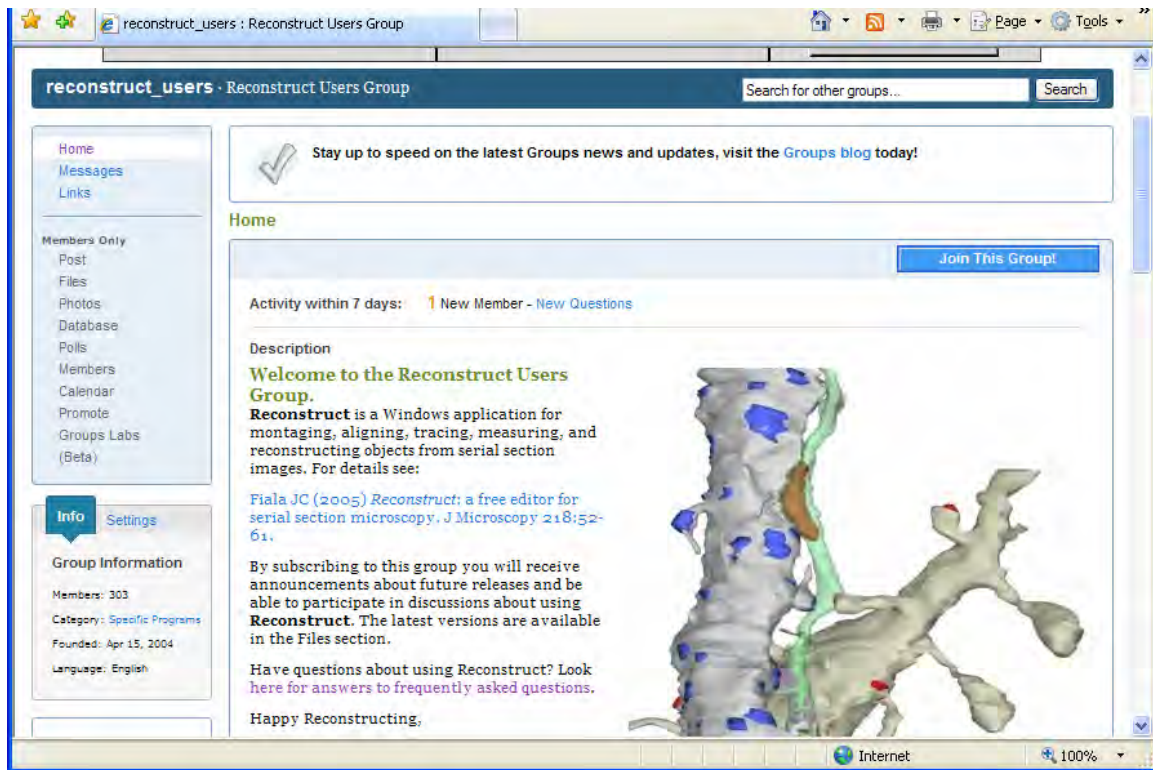


Figure 2.1. The welcome page for the RECONSTRUCT™ Users Group at Yahoo! Groups.

Credits

The development of RECONSTRUCT™ was funded, in part, by the National Institutes of Health under grants (P30 HD 18655, R01 MH/DA 57351, R01 EB 002170, R01 NS024760 and R01 MH057414). A number of different people contributed ideas for the user interface and behavior of the software, including Dr. Harris.

Source Code

The majority of RECONSTRUCT™ was designed and written by **John Fiala**. Other contributors are listed in the License dialog of the Help menu. The contributions of these individuals are as follows:

Andy Key and the Independent JPEG Group

Image input and output was realized by incorporating the Generalized Bitmap Module (GBM) library written by Andy Key (<http://www.nyangau.fsnet.co.uk/>). These libraries incorporated the source code of the Independent JPEG Group (<http://www.ijg.org/>) written by **Thomas G. Lane** for handling the JPEG/JFIF image formats.

Stefan Gustavson

Provided the ANSI C port of the complex FFT that is used for image correlations.

Ju Lu

Implemented the section-to-section propagation of wildfires.

Copyright © 2006 by John C. Fiala
Revised by K. Harris and K. Sorra; 2009



Chapter 3: Software Overview

RECONSTRUCT™ was developed from years of experience working with high magnification serial section images of brain tissue. Since each section is imaged separately in the microscope, the images must be brought back together to understand the structure of the whole tissue. RECONSTRUCT™ aids in the extraction of three-dimensional (3D) information from the serial section images.

Series, Sections, Domains, Traces, and Objects

Serial section microscopy consists of an ordered set of sections, representing the sequential pieces of the tissue block. Data in RECONSTRUCT™ are organized around this concept. An ordered set of sections is called a series. RECONSTRUCT™ operates on one series at a time. Within a series, data are grouped by section and each section is given an index indicating its position within the series. Each section in the series contains the data related to one physical section of the tissue. The section data includes images of the section and traces drawn on these images. The traces are later grouped together across sections to form 3D objects.




Figure 3.1. The mapping of data elements onto the section. The contents of this section are two domains ("domain1" and "domain2") and one trace named "dendrite", composed as shown in the upper left. Each trace and domain is associated with an independent transformation that determines the size and location of the element on the section. In this example each domain has a rectangular boundary that defines the area of the image to be displayed. Reprinted from Fiala (2005).

Each image within a section is called a domain, reflecting that it has a defined boundary and independent position within the section. Multiple image domains can be placed side-by-side within a section to make a montage, which is a composite picture of the section from many small images.

Images are entered into the computer by digitizing photographic prints, negatives, or by direct digital imaging with a digital camera. For example, the negatives on EM film are typically scanned at 1000 dpi using a conventional desktop scanner. These files are saved as Bitmaps or TIFFs from the scanner software, usually with a number in the filename indicating the section number. The image files are then imported into RECONSTRUCT™ as domains within the sections.

RECONSTRUCT™ is designed to facilitate image cropping, scaling and alignment. Cropping is achieved by drawing a new domain boundary that specifies the region to be displayed on the section. Image scaling is specified by a pixel size (magnification) parameter. Alignment is realized by a non-linear transformation associated with the image. All this information is stored within the section.

In addition to the images and transformations, lines within the section specify object profiles on the image data. These traces are added to the sections by drawing with a mouse (or a pen on a Tablet PC). All the traces that share the same name belong to the same 3D object. This convention allows the 3D objects to be extracted from the section files by grouping together all the traces with the same name. Objects are thus represented by the traces in the sections. Objects can be split apart or combined together simply by renaming the traces.

Data File Organization

The data for each section (images and traces) is saved to a file indexed by the section number as the filename extension. A series with 100 sections would have 100 section data files, numbered sequentially 1 through 100. For a series named "myseries", the section data is stored in files named "myseries.1", "myseries.2", "myseries.3", and so on to "myseries.100".

All the section data files are stored in the same folder or directory. Information about the whole series (e.g. option settings) is stored in a separate series data file in this same folder. The series data file has a .ser file extension, e.g. "myseries.ser". Thus, the data files associated with a series reside in the same folder and share the same filename. Only the file extension is different, .ser for the series data file and a positive integer for each section data files.

The image data itself is not stored in the section data file. For each domain in the section there is a reference to the corresponding image file. The actual image data can be stored anywhere on the computer, but generally it is best to store the image files in the same folder as the section data files. This allows the entire series to be easily off-loaded to CD-ROM or other archival media, and restored at a later date.

XML File Formats

RECONSTRUCT™ uses XML (eXtensible Markup Language, www.xml.com), for representing series and section data. XML files are in plain ASCII text and are therefore readable by text editors. XML can also be validated, allowing errors in the files to be automatically detected. A Document Type Definition (DTD) specifies the format for a valid

file. For RECONSTRUCT™ there are two DTDs, one for the section data file and one for the series data file. By including a reference to the DTD in the header of the XML data file, an XML validator can determine whether the syntax of the data file is correct.

The section DTD (Fig 3.2) specifies the XML file format for all section data files. The basic data types are defined by the ENTITY descriptions, referenced later by the % sign. SFBool defines a single field Boolean, either "true" or "false". SFColor is an RGB color defined by 3 floats where 0 0 0 is black and 1 1 1 is white. SFFloat is a single 32-bit floating point value. MFFloat is an array of floating point values separated by white space. SFInt32 is a single 32-bit integer value. SFString is a string of characters excluding the characters , " = / < and >. MFVec2f is an array of pairs of floats specifying (x,y) coordinates on the section. The rest of the DTD defines the basic ELEMENTS of a section and their attributes.

```

<!ENTITY % SFBool      "(true|false)">
<!ENTITY % SFColor     "CDATA">
<!ENTITY % SFFloat     "CDATA">
<!ENTITY % MFFloat     "CDATA">
<!ENTITY % SFInt32     "CDATA">
<!ENTITY % SFString    "CDATA">
<!ENTITY % MFVec2f     "CDATA">

<!ELEMENT Contour EMPTY>
<!ATTLIST Contour
  name      %SFString;    "unknown"
  hidden    %SFBool;     "false"
  closed    %SFBool;     "true"
  simplified %SFBool;     "false"
  border    %SFColor;    "1 0 1"
  fill      %SFColor;    "1 0 1"
  mode      %SFInt32;    "9"
  comment   %SFString;   #IMPLIED
  points    %MFVec2f;    #IMPLIED>

<!ELEMENT Image EMPTY>
<!ATTLIST Image
  mag      %SFFloat;    "1.0"
  contrast %SFFloat;    "1"
  brightness %SFFloat;  "0"
  red      %SFBool;     "true"
  green    %SFBool;     "true"
  blue     %SFBool;     "true"
  src      %SFString;   ""
  proxy_src %SFString;  ""
  proxy_scale %SFFloat; "1.0">

<!ELEMENT Transform ((Image,Contour)|Contour+) >
<!ATTLIST Transform
  dim      %SFInt32;    "6"
  xcoef    %MFFloat;    "0 1 0 0 0 0"
  ycoef    %MFFloat;    "0 0 1 0 0 0">

<!ELEMENT Section (Transform+)>
<!ATTLIST Section
  index    %SFInt32;    "-1"

```

```

thickness    %SFFloat;    "0.05"
alignLocked  %SFBool;    "false">

```

Figure 3.2. The Document Type Definition of the section data file (section.dtd).

A section contains one Section data element. The Section has three attributes, an *index* specifying the position of the section in the series, a *thickness* specifying the section thickness, and an *alignLocked* flag that specifies whether the entire section can be moved. The Section may contain one or more Transform elements.

The Transform element contains the coefficients of the linear combination of the bivariate polynomial basis, and a *dim* parameter specifying how many of these terms are to be used. Each Transform element contains either an Image element and boundary Contour, or a set of Contour elements. For more information, transformation equations are provided in Fiala and Harris, 2001 (<http://www.jamia.org/cgi/content/abstract/8/1/1>).

The Image element is the reference to an image file. Image attributes are *mag*, the width of a square pixel in series units; *src*, the path to the file containing the image data; *contrast* and *brightness* operate on a virtual intensity range of [0,1]; *proxy_src*, the path to a file containing image data at reduced dimensions; *proxy_scale*, the dimension of the proxy image divided by the dimension of src image. The Contour element paired with an Image element is the set of points defining the boundary of the image domain.

A Contour element is the set of points defining a trace on the section. The Contour element contains no other elements but has many attributes: *name*, the identifier of object to which the trace belongs; *hidden*, whether trace is displayed; *closed*, whether the trace closes on itself; *simplified*, whether trace points have been simplified; *border*, the displayed color of the polygonal representation; *fill*, the color of the interior of the trace when filled; *mode*, the type of filling operation used; *comment*, a string of user information; *points*, the sequence of (x,y) coordinates of the trace.

Fig 3.3 shows a valid XML section data file based on the DTD. Note that a similar DTD is defined for the series data file (series.dtd), but it is not expounded here because the large number of series options makes the file very long. To read the DTD format, open the series.dtd file using any text editor.

```

<?xml version="1.0"?>
<!DOCTYPE Section SYSTEM "section.dtd">

<Section index="7" thickness="0.05"
alignLocked="true">
<Transform dim="6"
  xcoef=" 0.240299 0.999416 -0.0182671 0 0 0"
  ycoef=" 0.0104287 0.0297386 0.997894 0 0 0">
<Image mag="0.002691" contrast="1" brightness="0"
  src="_unaligned.7" />
<Contour name="domain1" hidden="false" closed="true"
simplified="false" border="1 0 1" fill="1 0 1"
mode="11"
  points="0 0,
          768 0,
          768 768,
          0 768,
          "/>

```



```

</Transform>

<Transform dim="0"
  xcoef=" 0 1 0 0 0 0"
  ycoef=" 0 0 1 0 0 0">
<Contour name="a4" hidden="false" closed="true"
simplified="false" border="1 0.5 0" fill="1 0.5 0"
mode="13"
  points="0.377196 0.55957,
          0.377196 0.555605,
          0.381161 0.55164,
          0.385125 0.55164,
          0.38909 0.555605,
          0.38909 0.55957,
          0.385125 0.563534,
          0.381161 0.563534,
          "/>
</Transform>

</Section>

```

Figure 3.3. An example XML section data file. Section number 7 consists of one image domain "domain1" and one trace "a4".

Image File Formats

RECONSTRUCT™ currently supports several common single image file formats including Windows Bitmaps, JPEG, TIFF, and GIF. The complete set of acceptable file formats is determined by the image input/output library developed by Andy Key (see Credits, below). Acceptable formats include Windows and OS/2 bitmaps (.BMP .VGA .BGA .RLE .DIB .RL4 .RL8) except for CCITT G3 MH Fax compressed and 24 bit RLE files; CompuServe GIF files (.GIF); JPEG File Interchange Format files (.JPG .JPEG .JPE); TIFF files (.TIF .TIFF) with 1, 4, 8 and 24-bits per pixel; Truevision Targa/Vista bitmaps (.TGA .VST .AFI); ZSoft PC Paintbrush Image format (.PCX .PCC); GEM Raster (.IMG); Portable Greyscale-map (binary P5 type) (.PGM); and Portable Pixel-map (binary P6 type) (.PPM). Convert image files as needed into a format that RECONSTRUCT™ will recognize. ImageJ or Adobe Photoshop can be useful tools for adjusting image file formats.

For most formats it is required that the image filename extension reflect the file format used. For example, a portable pixel-map file could be named "myimage.ppm" but not "myimage.junk". Exceptions to this rule are the common formats: Windows Bitmap, JPEG, TIFF, and GIF. Filenames for these formats are arbitrary because the program can identify the file format from header information in the file.

DXF Imports

Line drawing information can be imported into a series using the ASCII Drawing Interchange File (DXF) format (www.autodesk.com). This DXF file is a sequence of code-value pairs in text format where each integer code appears on a separate line, followed by the value string on the next line. RECONSTRUCT™ scans the file for zero codes signifying new entities. Only LINE, POLYLINE, VERTEX, CIRCLE, and SOLID entities are recognized. Each such entity found creates a corresponding trace on the section. The trace name reflects the entity type and the color code (62) value of the entity. Entities with

different colors in the DXF file will be given different names, while entities of the same type and color will have the same name.

DXF Exports

Trace data can also be exported using the ASCII Drawing Interchange File (DXF) format (www.autodesk.com) by using the Lines... command from the Series Export menu. A .dxf is created for each section that contains all the non-hidden traces as POLYLINES. The trace name is written after group code 3. The approximate trace color is converted to an 8-bit value that is output in code 62. Whether the trace is closed or not is output as a boolean value after code 70. The vertices of the trace are output in series units with the precision specified in the General tab of the Series Options. Each VERTEX contains the X and Y positions on the section and the z-distance for the section. The following example demonstrates the DXF output for a section containing a single rectangular trace called "test1".

```
0
SECTION
2
ENTITIES
0
POLYLINE
62
227
66
1
70
1
3
test1
0
VERTEX
10
74.7673
20
75.2601
30
0.05
0
VERTEX
10
181.965
20
75.2601
30
0.05
0
VERTEX
10
181.965
20
185.168
30
0.05
0
```

```
VERTEX
10
74.7673
20
185.168
30
0.05
0
SEQEND
0
ENDSEC
0
EOF
```

3D Output Formats

In addition to saving a reconstructed scene as an image file, the scene can be exported as a 3D model in one of three commonly used formats: VRML 1.0, VRML 2.0, and DXF. For more information on Virtual Reality Modeling Language (VRML) and 3D output DXF file formats, see Chapter 12: The 3D Scene.

System Requirements

RECONSTRUCT™ was written and tested under Microsoft Windows XP on a Dell Optiplex GX260 Pentium 4 (2 GHz) computer with 512 Mb of RAM. The RECONSTRUCT™ binary executable is compiled for Pentium-compatible processors, and the application relies entirely on the Win32 libraries for interacting with the computer's user interface. RECONSTRUCT™ will therefore run on any Pentium-compatible computer that supports Win32.

Operating System

RECONSTRUCT™ will run on all Microsoft Windows operating systems: Windows 95, Windows 98, Windows 2000, Windows NT, Windows XP and Vista. RECONSTRUCT™ will run under other operating systems provided a Win32 emulator/translator is installed. A Win32 translation layer for X Windows, such as WINE (www.winehq.com), translates the Win32 API calls from RECONSTRUCT™ into the equivalent X Windows operations. WINE is available for a variety of operating systems, including GNU/LINUX, FreeBSD, and Solaris. Users have verified WINE for running RECONSTRUCT™ under Mandriva Linux w/KDE 3.3.2. In addition, the RECONSTRUCT™ binary may run on Virtual PC, Apple's Pentium/Windows emulator.

OpenGL

OpenGL libraries are required for previewing (and saving) 3D objects. The OpenGL libraries provide a software interface to the computer's graphics hardware. OpenGL is shipped standard with every version of Windows (95/98/2000/NT/XP/Vista) and should work with RECONSTRUCT™ without any additional configuration. For more information on OpenGL see Chapter 12: The 3D Scene.

Hardware Requirements

Performance of RECONSTRUCT™ is determined by three main hardware components: the graphics display adaptor, the amount of RAM, and the speed of the hard drive. A high-quality graphics display adaptor that accelerates OpenGL is highly desirable for 3D Scene operations. Most modern game adaptors are satisfactory for this purpose. The software was tested on a Radeon VE adaptor with 32 Mb of memory. This device gave excellent OpenGL performance.

RECONSTRUCT™ is not a very demanding application from a memory standpoint because the data from only two sections is held in memory at one time. Thus, a very long series of hundreds of sections will not require any more memory than a short one. A minimum requirement for good performance is free RAM equal to at least 3 times the total amount of data on a section. Thus, if a section contains 40 Mb of image data, $40+40+40=120$ Mb of RAM would be needed for running the program. Allow for an additional 128Mb for the Windows XP operating system and desktop, along with 20-40 Mb for each major application (Acrobat, Outlook, Word, Excel, etc.) used simultaneously with RECONSTRUCT™. 512Mb of memory was adequate for all data used to test the software.

The greatest performance enhancer is a fast hard drive. Note that hard drive performance is not necessarily the transfer rate given in the hard drive specifications. This Data Transfer Rate is usually a burst rate for the interface to the system memory, not the actual speed of data access from the drive. Look for drives with the highest speed of rotation (RPM) and the most sectors per track. Real sustained data transfer rates are roughly proportional the product of rotational speed and sectors per track. For example, a 15,000 RPM SCSI drive will outperform a 5,400 RPM IDE/ATA drive, resulting in faster loading and paging of sections in RECONSTRUCT™. A 7,200 RPM IDE/ATA drive was used during development and testing of the software. A 7200 RPM drive gives satisfactory performance on images less than 15 Mb.



Chapter 4: User Interface

The user of RECONSTRUCT™ interacts with a set of screen elements using the mouse (or pen) and keyboard. This chapter describes the basics of the graphical user interface.

Windows and Dialogs

Two kinds of screen elements will be distinguished: windows and dialogs. Dialogs must be completed prior to returning to the main window, while windows release control to each other. The exception to this is that some message dialogs (such as the Version dialog from the Help menu) release control back to the main window. Dialogs cannot be resized but can be moved. Windows float on top of the application's main window and can be resized. A special case is the status bar that remains docked at the bottom of the main window.

Resizing and Moving Windows

Windows can be resized by placing the pointer on the narrow border of the window, pressing the left mouse button, and dragging the border. The pointer should change to resize arrows when over the border of a realizable window. The main window can also be resized by dragging from the resize area in the lower right-hand corner of the status bar. Both windows and dialogs can be moved by dragging the title bar with the left mouse button.

Under some operating systems the windows can be maximized or minimized, but after this operation they cannot be restored to their original size and position except by closing and reopening the window. To maximize or minimize a window, place the pointer on the window's title bar and click the right mouse button to open a popup menu of window commands.

Windows that display an icon in the title bar, namely the main window and the Series Options dialog, can be moved with the keyboard. Selecting the icon with left mouse button opens the system menu. To resize or move the window select the menu item and then use the arrows keys to adjust the position of the window. Holding the Ctrl key down allows finer control of the movements. Release the keyboard mode by hitting Enter or by pressing the left mouse button again. Changes can be canceled with the Esc key.

Switching Between Windows

The keyboard focus can be switched between windows of the application by using Ctrl-Tab, much in the same way as the Alt-Tab command is used to switch between active applications on the desktop. The window that has the keyboard focus is called the active window. An active window usually has a colored title bar while inactive windows often have grayed title bars.

TABLE 4.1. Basic Keyboard and Mouse Actions for All Windows

Keystroke	Action
Alt	Access the menu of the currently active window
Alt-Tab	Switch between applications
Ctrl-Tab	Switch between floating windows of the application
Left Mouse Drag	Action
on title bar	Position the window
on border	Size the window

Main Window

When the application is started the RECONSTRUCT™ main window appears along with two of the sub-windows, the tools window and the status bar (Fig 4.1). The main window consists of a title bar at the top, a set of menus beneath the title bar, and a rendering area (a.k.a. client area) for displaying data.

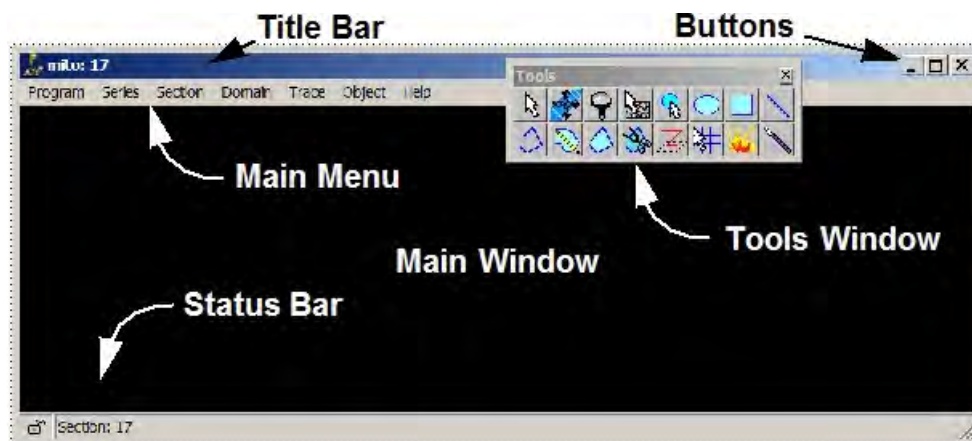


Figure 4.1. RECONSTRUCT™'s initial window layout. A main window with a narrow border and a black interior is overlaid with a floating tools window and a status bar.

At the left end of the title bar is the RECONSTRUCT™ icon. When a series is open the series name is displayed in the title bar along with the current section number. On the right end of the title bar are the usual minimize, maximize, and close buttons. Selecting the minimize button, by clicking the left mouse button when the pointer is over the button, hides the window. The hidden window can be reopened by selecting it on the task bar or by using Alt-Tab (pressing the Tab key while holding the Alt key down.) Alt-Tab can also be used to rapidly switch between RECONSTRUCT™ and other applications running on the desktop. The maximize button resizes the main window to fill the screen. The main window cannot be manually resized when it is maximized. Unmaximize the window by the selecting the maximize button again.

A number of keystrokes and mouse actions can be used in the main window (Table 4.2). These are mostly for navigating between sections in the series and creating or manipulating section data. Mouse actions in the main window may depend on which tool has been selected from the tools windows. These actions are discussed in more detail in

later chapters. However, moving the mouse wheel in the main window always has the same effect as Page Up and Page Down.

TABLE 4.2. Keyboard and Mouse Actions in the Main Windows

Keystroke	Action
Page Up	Display next (higher) section in series
Page Down	Display preceding lower section in series
/	Switch front and back sections. Hold to flicker.
Backspace	Delete last trace point entered when drawing
Esc	Abandon current drawing or tool operation
0, 1, ... 9	Switch default trace attributes to palette entry
Ctrl-0, ctrl-1, ... ctrl-9	Set the corresponding palette entry to the current default attributes and stamp tool shape
F1-F12 keys	Section, domain and trace movements
Arrow keys	Section, domain and trace movements
-	Decrease brightness of selected domain
=	Increase brightness of selected domain
[Decrease contrast of selected domain
]	Increase contrast of selected domain
Mouse input	May depend on the currently active tool

Main Menu

The main menu consists of Program, Series, Section, Domain, Trace, Object, and Help menus (Fig 4.1). Submenus of these menus are indicated by an arrowhead on the right-hand side of the item. Items that open a dialog requiring further user input are indicated by a trailing "...". These operations can generally be canceled in the dialog. Other menu items execute immediately.

The mouse pointer can be used to select items in the main menu, but menu items can also be selected with the keyboard. Keyboard control of the menu is activated by pressing and releasing the Alt key. When the menu is under keyboard control certain letters in the menu items will be highlighted. These items can be selected by tapping the key for the highlighted letter. For example, pressing and releasing the Alt key and then the T key, will open the Trace menu. The L key will then open the Trace List. Alternatively, the arrow keys can be used to browse through the menus. The highlighted menu item can be selected by using the Enter key. The menu selection process can be aborted by using the Esc key.

Some menu items have special keystrokes associated with them that allow them to be rapidly executed without using the menu selection process (Table 4.3). For example, pressing the S key while holding down the Ctrl key when the main window is active will select all the traces in the section.

Many of the menu items may be grayed or disabled if they are not currently appropriate. For example, if no series is open then the menus for manipulating sections, domains, and traces are not applicable. Likewise the domain and trace menus are mostly disabled

when no domain or trace has been selected, since most operations apply to selected data elements.

TABLE 4.3. Keyboard Accelerators for the Main Menu

Keystroke	Action
Space	Blend the front and back section images
Esc	Abandon current menu selection activity
Ctrl-\	Move front section By Correlation with back section
Ctrl-A	Edit the attributes of selected traces
Ctrl-B	Paste attributes from clipboard traces to selected traces
Ctrl-C	Copy the selected traces
Ctrl-D	Deselect all the selected traces
Ctrl-H	Hide the selected traces
Ctrl-L	Toggle the movement lock on the section
Ctrl-M	Merge the selected traces
Ctrl-O	Open the Series Options dialog
Ctrl-P	Switch to/from precision cursor when drawing
Ctrl-R	Repeat last movement
Ctrl-S	Select all traces on section
Ctrl-U	Undo all current changes to the section
Ctrl-V	Paste traces into section
Ctrl-X	Cut selected traces from section for later pasting
Ctrl-Z	Undo section or domain changes
Home	Center section image in view
Ctrl-Home	Zoom and center selected traces in view
End	Undo/redo last zoom
Insert	Paste traces into section
Delete	Delete selected traces or domain

Tools Window

The Tools window (Fig 4.2) is a floating toolbar that allows the user to select different modes of mouse input in the main window. The toolbar can be opened or closed by selecting the *Tools window* item from the Program menu. The Tools window can also be closed by selecting the close box on the right-hand side of the title bar.



Figure 4.2. The Tools window. This window should be visible at the top of the RECONSTRUCT™ main window. Use the left mouse button to select a tool. In this

example, the Wildfire tool button is depressed and so a description of the tool appears in the title bar.

A description of the tool appears in the title bar of the Tools window whenever a tool button is pressed. Pressing the tool button also selects the tool for use in the main window. The keyboard (Table 4.4) can be used to change tools without moving the mouse pointer. To do this, give the Tools window the focus using Ctrl-Tab. Then use the Tab key to move through the tool buttons. The highlighted button can be selected with the Enter key.

The Tools window can be resized and placed anywhere on the screen. The tool buttons will automatically reorganize to fit within the new window frame.

TABLE 4.4. Keyboard and Mouse Actions in Tools Window

Command	Action
Tab	Move to next tool button
Enter	Select highlighted tool
Left mouse click	Select the tool under the pointer

Status Bar

The Status bar (Fig 4.1) can be opened or closed by selecting the *Status bar* item from the *Windows* submenu in the Program Menu.

The Status bar displays information about the data under the pointer when a series is open. The left-most element on the Status bar is a lock/unlock icon that indicates whether the current section can be moved or aligned. The Status bar also contains the number of the section being displayed, the name of the trace nearest to the mouse pointer, and the position of the pointer on the section in series units. On a domain, the Status bar displays the position of the pointer in image pixels. When drawing, the Status bar displays the name of trace being drawing and either position, length, or area information for the drawing increment.

Section List Window

The Section List window is opened by selecting the *List sections...* item from the Section menu (Fig 4.3). The Section List displays the sections of the series in order of ascending section number. The two attributes of each section can be displayed: the lock/unlock status and the section thickness in series units. The columns of the list can be resized by dragging the column borders with the left mouse.

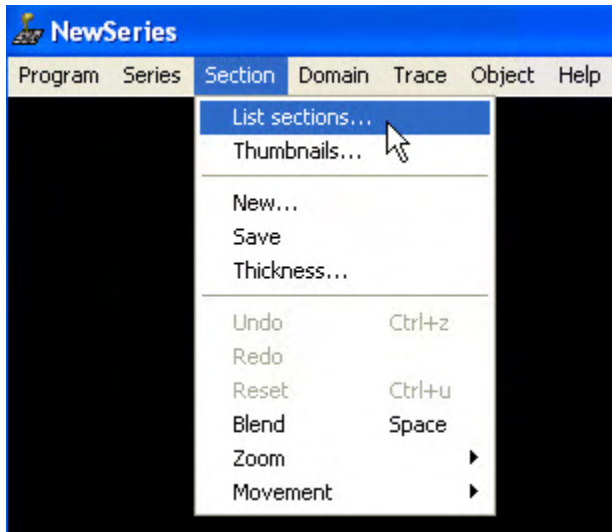


Figure 4.3. The *List sections...* item from the *Section* menu in the main *RECONSTRUCT™* window.

When the Section List is open the current section is highlighted. A different section can be displayed in the main window by double-clicking the new section number in the Section List window. The Section List will be closed once a new section is selected. The paging operation can also be executed by selecting the *Page To* menu item at the top of the Section List window, in which case the Section List is not closed. The Section List window can also be closed by selecting the *List sections...* item from the Section menu in the main window.

Navigate the Section List with the keyboard (Table 4.5) using the up and down arrow keys, the Page Up and Page Down keys, or the Home and End keys. Use Enter to execute the *Page To* operation for the first highlighted section.

TABLE 4.5. Keyboard and Mouse Actions in Section List Window

Command	Action
Up & down arrows	Move to next/previous section in list
Page Up & Down	Move through the list by the number of visible sections
Home	Move to the first section in the list
End	Move to the last section in the list
Delete	Delete the highlighted sections from the series
Enter	Page to the first highlighted section
Ctrl or Shift	Highlight multiple sections for Modify operations
Left Mouse Click	Highlight the section under the pointer
Left Mouse Dbl-Click	Page to the section under the pointer

Multiple sections can be selected in the section list by holding the Ctrl or Shift key down while selecting sections in the list with the mouse. The highlighted sections can be changed simultaneously using the Modify menu at the top of the Section List window. The Modify operations allow highlighted sections to be locked/unlocked, renumbered, deleted, or given particular pixel size or thickness attributes.

Each list window has a List menu that contains operations on the list itself rather than the data elements in the list. The number of items in the list can be displayed by selecting the *Info...* item. The *Save...* item allows the list to be saved as a Comma Separated Values (.csv) file. This file is an ASCII text file that can be readily imported into spreadsheet programs such as Excel. The *Save...* dialog allows the filename and folder for the .csv file to be specified. The list columns and rows will be saved in the file just as they appear in the window.

The information in a list may occasionally need refreshing after changes to the sections or traces of the series. Any list can be regenerated by closing it and reopening it, or by selecting the *Refresh* item from the List menu. *Refresh* will erase the current list and refill it based on the current series data.

Thumbnails Window

Thumbnail images of the sections can be viewed by selecting *Thumbnails...* from the Section menu (Fig. 4.4). Each thumbnail image is placed on a button in the Thumbnails window and the buttons are displayed in one of two ways. In the default display mode the thumbnail buttons are arranged to fill the window starting in the upper left-hand corner. Resizing the window results in rearrangement of the buttons to fill the window.

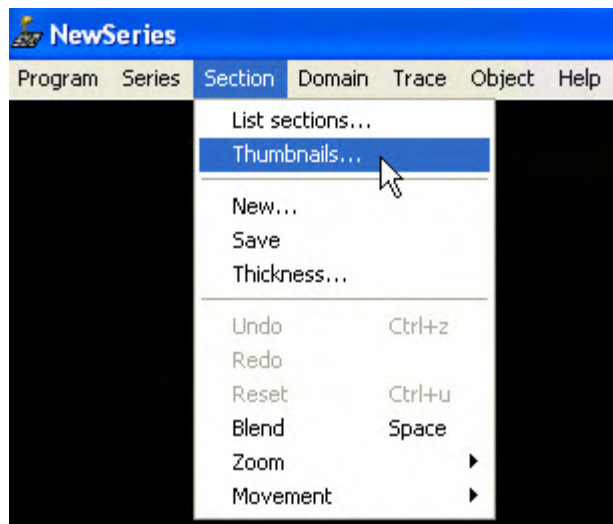


Figure 4.4. The *Thumbnails...* item from the Section menu in the main RECONSTRUCT™ window.

Putting the pointer over a thumbnail button for a certain amount of time results in display of the section number in the title bar when the Thumbnails window is active. Selecting a thumbnail with the left mouse button results in the main window being activated and paging to the section. The thumbnail buttons cannot be navigated with the keyboard but can be scrolled through using the scroll button on the mouse (Table 4.6).

TABLE 4.6. Keyboard and Mouse Actions in Thumbnails Window

Command	Action
Tab	In flipbook mode, stop/start shuttle play
Enter	In flipbook mode, page to the section shown
Left Mouse Click	Page to the section under the pointer
Mouse Scroll	In flipbook mode, move through the thumbnails

For more on the configuration of thumbnails see Chapter 6: Sections.

Domain List Window

The Domain List window is opened using the *List image domains...* item from the Domain menu (Fig 4.5). The window displays a list of all the domains within the section ordered from back to front. The domain name is displayed along with the source path to the image data, the pixel size, and the centroid position of the domain on the section. The columns of the list can be resized by dragging the borders of the column headers.

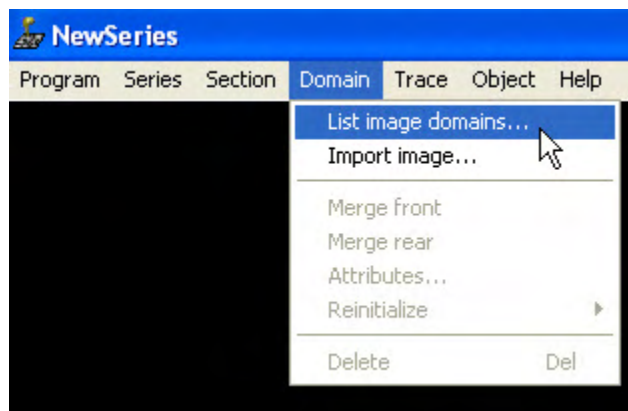


Figure 4.5. The Domain List...item from the Domain menu in the main RECONSTRUCT™ window.

A domain can be highlighted in the list with the left mouse button or by using the arrow keys. The attributes of the highlighted domain can be edited by selecting the *Attributes...* item from the Modify menu. This action does not close the domain list. Double-clicking the left mouse button or hitting Enter will cause the *Select* menu item to be executed on the highlighted domain (Table 4.7). In this case the domain is selected in the main window and the Domain List window is closed. See Chapter 7: Domain Images for more on working with selected domains.

TABLE 4.7. Keyboard and Mouse Actions in Domain List Window

Command	Action
Up & down arrows	Move to next/previous domain in list
Page Up & Down	Move through the list by the number of visible domains
Home	Move to the first domain in the list
End	Move to the last domain in the list
Delete	Remove the domain from the section
Enter	Select the domain from the section
Left Mouse Click	Highlight the domain under the pointer
Left Mouse Dbl-Click	Select the domain under the pointer

Trace List Window

The Trace List window is opened using the *List traces...* item from the Trace menu (Fig 4.6). The window displays a list of all the traces within the section ordered alphabetically. The trace name is displayed along with a set of measurements in series units. The columns of the list can be resized by dragging the borders of the column headers. The set of displayed columns can be specified from the Lists tab in the Series Options dialog. The entire Trace List can be saved to a file by choosing the *Save...* item from the List menu.

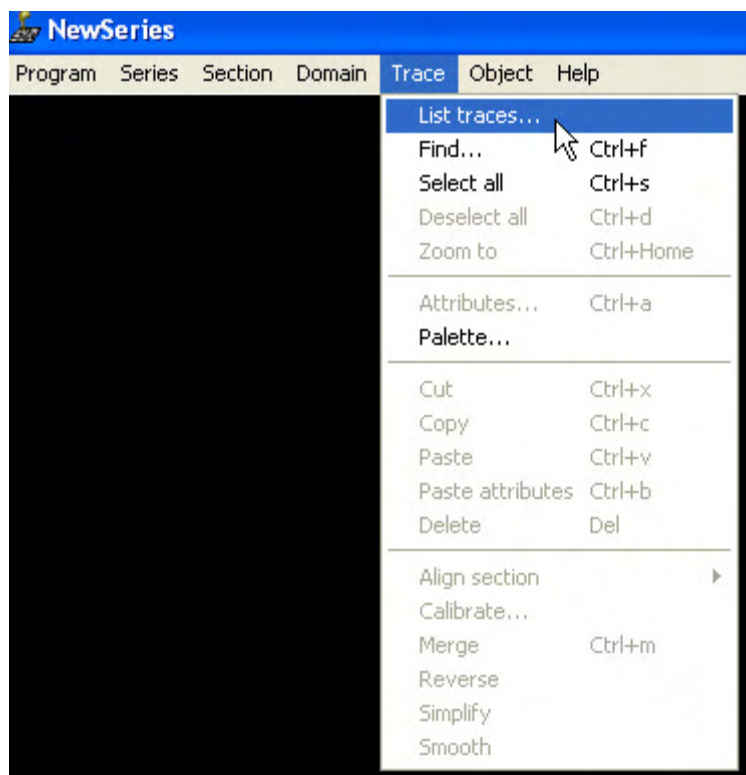


Figure 4.6. The *List Traces...* item from the Trace menu in the main RECONSTRUCT™ window.

A trace in the list can be highlighted with the left mouse button or by using the arrow keys. Multiple traces can be highlighted by holding down the Shift or Ctrl keys. Pressing Enter or using the *Select* menu item will cause the highlighted traces to be selected and the keyboard focus to be shifted to the main window. Any traces that were previously selected in the main window are unselected.

TABLE 4.8. Keyboard and Mouse Actions in Trace List Window

Command	Action
Up & down arrows	Highlight next/previous item in list
Page Up & Down	Move through the list by the number of visible traces
Home	Move to the first trace in the list
End	Move to the last trace in the list
Delete	Delete the highlighted traces.
Enter	Select the trace in the main window
Ctrl or Shift	Highlight multiple traces
Left Mouse Click	Highlight the trace under the pointer
Left Mouse Dbl-Click	Select the trace under the pointer

Trace Palette Window

The Trace Palette window can be opened or closed using the *Palette...* item from the Trace menu (Fig. 4.7). The window displays the set of predefined trace attributes as a set of buttons. Each palette entry represents the trace name, border color, fill color, and fill mode for drawing new traces. This is displayed on the button by a trace drawn in the shape of the stamp trace also defined in the palette entry. For more on palette entries see Chapter 8: Traces.

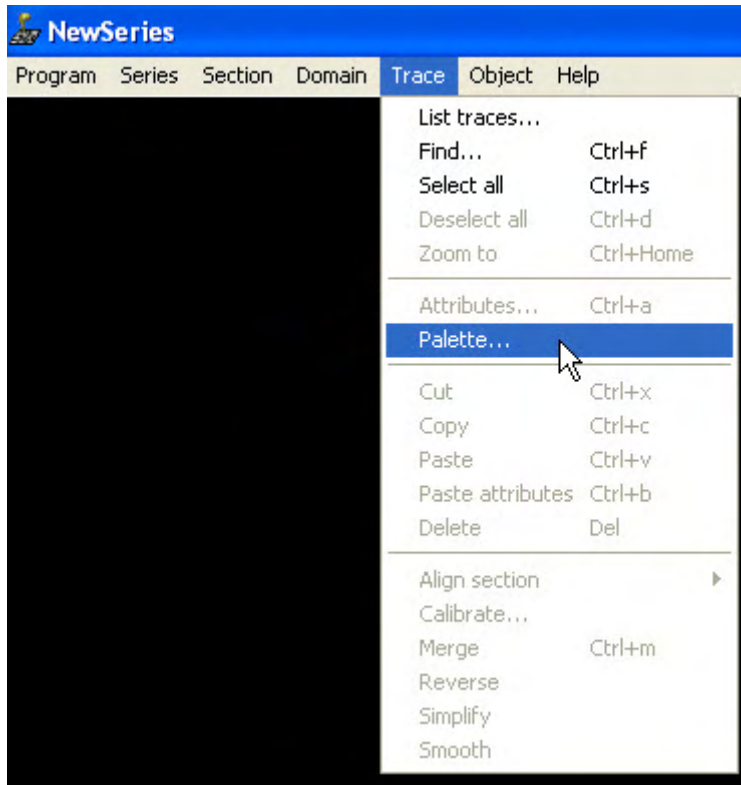


Figure 4.7. The Trace Palette...item from the Trace menu in the main RECONSTRUCT™ window.

The Trace Palette window, like the Tools window, can be resized and placed arbitrarily on the screen. Default trace attributes are selected with a left mouse click (Table 4.9). The Tab key can be used to move through the palette buttons. The highlighted button can be selected with the Enter key. When a palette button is pressed the default trace name of the palette entry is displayed on the title bar.

TABLE 4.9. Keyboard and Mouse Actions in Trace Palette Window

Command	Action
Tab	Move to next palette entry button
Enter	Select highlighted palette entry
Left mouse click	Select the palette entry under the pointer

Object List Window

The Object List window is opened or closed by the *List objects...* item of the Object menu (Fig. 4.8). Each time the list is regenerated the sections are scanned for traces and these are grouped by name into objects. If there are large numbers of traces and sections, filling the Object List may be slow. Each unique trace name is displayed in the list as an object. The list may also display the section range in which the trace name is found, a count of the number of component traces that make up the object, and the object's surface area and volume.

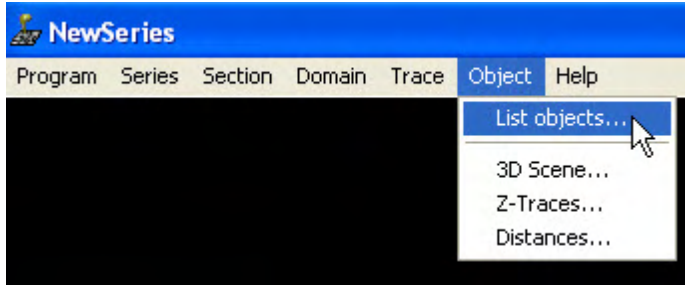


Figure 4.8. The *List objects...* item from the Object menu in the main RECONSTRUCT™ window.

Like the other list windows, the object list can be navigated by mouse or keyboard (Table 4.10). Columns can be resized at the column headings. The window can be resized by dragging the narrow border of the window.

Multiple objects can be highlighted by holding down the Shift or Ctrl keys. Hitting Enter or selecting the *Add to scene* item in the Scene menu will cause a 3D representation of the objects to be generated and added to the 3D Scene. If the 3D Scene window is not open, the window is created and the 3D Scene is displayed.

The object list can be saved to a .csv file using the *Save...* item from the List menu. The attributes of the traces of objects can be modified by selecting the *Attributes...* item from the Modify menu. Similarly, highlighted objects can be deleted from the series with the *Delete...* operation. For more on the Object List see Chapter 11: Objects.

TABLE 4.10. Mouse and Keyboard Actions in Object List Window

Command	Action
Up & down arrows	Highlight next/previous object in list
Page Up & Down	Move through the list by the number of visible objects
Home	Move to the first object in the list
End	Move to the last object in the list
Delete	Delete all traces of highlighted objects from the series.
Enter	Generate the 3D objects and add to the scene
Shift	Highlight multiple objects
Left Mouse Click	Highlight the object under the pointer
Left Mouse Dbl-Click	Generate and add object under the pointer

3D Scene Window

The 3D Scene window displays the 3D representations of objects using OpenGL. OpenGL libraries must be available in order to view this window. The window can be opened or closed from the Object menu (Fig. 4.9), or by creating a 3D object representation from the Object List window.

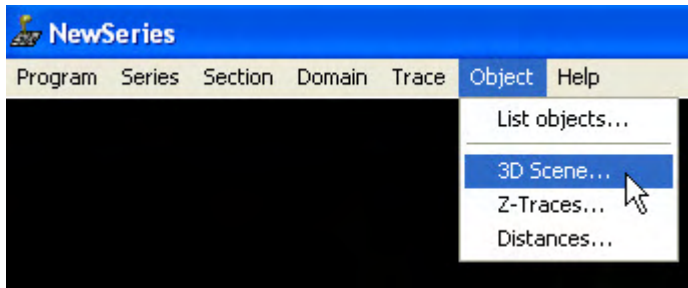


Figure 4.9. The 3D Scene...item from the Object menu in the main RECONSTRUCT™ window.

Details of the 3D scene manipulations using the mouse are described in Chapter 12: The 3D Scene, but are briefly summarized in Table 4.11. If the scene or menus flicker in the 3D Scene window, unchecking the *Spin* item in the View menu may remedy the problem. The animation performance may also be affected by having other windows open in the application or on the desktop.

TABLE 4.11. Mouse Actions in 3D Scene Window

Mouse	Action
Left mouse drag	Rotate scene.
Right mouse* drag	Zoom scene.
Left and right* drag	Pan scene.

*Note: Pressing the right mouse button will pause any on-going animation.

Z-Trace List Window

The Z-Trace List window is opened from the Object menu (Fig. 4.10). The window displays a list of all the z-traces in the section ordered alphabetically. The trace name is displayed along with the section range and the length in series units. An additional column can be used to make a notation about the trace. The columns of the list can be resized by dragging the borders of the column headers. The entire list can be saved to a .csv file by choosing the Save... item from the List menu.

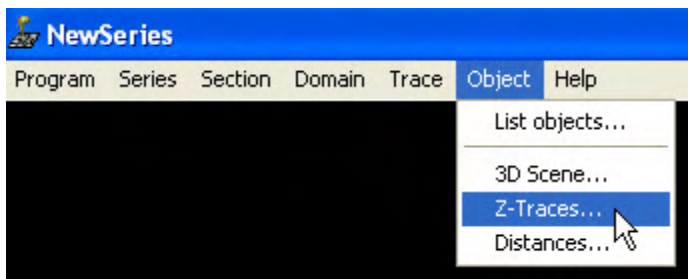


Figure 4.10. The Z-Traces...item from the Object menu in the main RECONSTRUCT™ window.

A trace in the list can be highlighted with the left mouse button or by using the arrow keys (Table 4.12). Multiple traces can be highlighted by holding down the Shift or Ctrl keys.

Pressing Enter or using the *Add to scene* item from the Scene menu will cause the highlighted traces to be represented as 3D objects and added to the 3D Scene.

TABLE 4.12. Keyboard and Mouse Actions in Z-Trace List Window

Command	Action
Up & down arrows	Highlight next/previous trace in list
Page Up & Down	Move through the list by the number of visible traces
Home	Move to the first trace in the list
End	Move to the last trace in the list
Delete	Delete the highlighted traces from the series
Enter	Put the trace into the 3D scene
Shift	Highlight multiple traces
Left Mouse Click	Highlight the trace under the pointer
Left Mouse Dbl-Click	Put the trace under the pointer into the scene

Distances List Window

The Distances List window (Fig 4.11) displays a table of distances between objects in the series. (For more information on the distance computation see Chapter 9: Calibration and Measurement.) The object names and calculated distance (in series units) are displayed for each pair of objects that matches the list's limit strings. To limit the number of distances calculated and displayed, set the limit strings in the Series Options BEFORE selecting *Distances...* from the Object menu. The list display may seem to hang for 65,535 or more distances, but displaying even a few thousand distances may be slow. If more than 65,535 distances are calculated, a warning message offers the opportunity to cancel the action before the list is displayed.

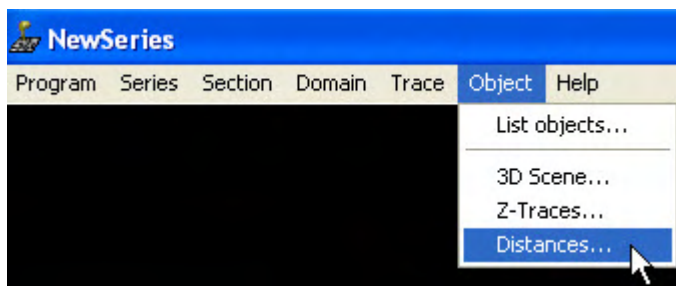


Figure 4.11. The Distances...item from the Object menu in the main RECONSTRUCT™ window.

There are no special keyboard or mouse operations on the Distances List. As with all list views, the columns of the list can be resized by dragging the borders of the column headers. The entire list can be saved to a .csv file by choosing the *Save...* item from the List menu.

Exiting or Aborting

Most windows in RECONSTRUCT™ can be closed by choosing the "x" in the upper right-hand corner of the window. This also typically aborts the operation being done by the

window. For example, if too many distances are being computed in the Distances List, the Distances calculation can usually be aborted without adverse consequences by closing the window. However, if a list view attempts to display too many items, it may make the entire program unresponsive so that the window cannot be closed. When this happens, you may need to abort the entire program to stop the operation.

RECONSTRUCT™ can usually be terminated by closing the main window. If a series is open, it will be closed prior to exiting the program. One exception is when adding objects to the 3D Scene from the Objects List. When 3D representations are being generated, the current series should not be closed nor should the program be terminated. This will cause a fatal error that may result in the loss of unsaved changes to the series. During 3D object generation, use the Abort button provided and wait for the operation to terminate on its own before closing the series or exiting the program.

In the event that RECONSTRUCT™ becomes unresponsive, such as when trying to fill a Distances List with 65,000 measurements, you may need to abort the program from the outside. To do this, open the task manager by using Ctrl-Alt-Delete or by selecting *Run...* from the Start menu and entering "Taskmgr". In the Task Manager, select the Applications tab and find the RECONSTRUCT™ icon beside the name of the series. If the Status is listed as "Not Responding", highlight this entry with the mouse and then click on the End Task button.

Warning: Unsaved changes, to the current and previous sections or to the series options, will be lost by aborting RECONSTRUCT™ from the Task Manager!



Chapter 5: Series

A series consists of set of files in a single, writable folder, although the image data can reside in read-only files in another folder or drive. This chapter describes how to create and manipulate a series. Manipulation of sections within the series is described in the next chapter.

Creating a Series

A new series can be created by selecting *New...* from the Series menu. The New Series dialog will appear. Navigate to an existing folder or create a new one for the series. Type a name for the series in the Filename box. Each series in a folder must be given a unique name following the usual rules for filenames given by the operating system. The maximum length for the series name is 255 characters. The extension *.ser* will be appended to the filename if it is not typed in the box. After selecting the Save button, the *.ser* file will be created with default settings.

WARNING: Any existing series file with the same name will be overwritten.

The next step is to add sections to the series. This can be done by using the Import Images menu item (Chapter 7: Domain Images) or by creating new sections from the Section menu (Chapter 6: Sections).

Opening and Closing Series

An existing series can be opened by using the *Open...* item from the Series menu. Use the Open Series dialog to navigate to an existing folder and select the *.ser* file for the series. Alternatively, a series can be opened by dragging an *.ser* file onto the RECONSTRUCT™ application icon or filename.

An open series can be closed by selecting the Close menu item from the Series menu, or by exiting the application. Within RECONSTRUCT™ only one series is open at a time. Opening or creating another series will close the current one. Multiple series can be opened on the desktop by running multiple copies of the application.

Saving a Series

A series can be saved by selecting the *Save* item from the Series menu. This saves only the *.ser* file, not the section data files. The series file stores the current view position in the series, z-traces, palette entries, and the series options described below (Table 5.1).

The series file is saved automatically whenever a new series is closed. To turn off automatic series saving, change the option in the Series Options dialog and then select *Save* from the Series menu to put the change into the *.ser* file. Similarly, a manual *Save* is needed to turn this option back on.

Series Options

The series file contains a set of attributes that remember the user-defined options (Table 5.1). These are accessed through the *Options...* dialog from the Series menu. The Options dialog (Fig. 5.1) contains tabbed pages that separate the options into categories. To change options, set the desired values in the pages and then select the OK button at the bottom of the dialog. Selecting Cancel will cancel the changes in most cases. The options are saved when the series is saved.

General

The General tab contains parameters for controlling how data is saved and displayed. The units description string that specifies what units are used for quantitative values such as measurements. All quantities within a series are expressed in these units. All non-image data is saved as text (versus binary format). The data precision parameter controls how many significant digits are written to represent each number. The General tab also contains the default section thickness, given to new sections, and an option for changing how the section thickness is used in calculating distances perpendicular to the sections. Details on the use of these options can be found in Chapter 9: Calibration and Measurement. The hide and unhide options in the General tab determine whether traces and domains (images) will be displayed on-screen when a section is accessed, and also whether the traces and domains will be marked as hidden when a section is saved. Consequently, turning on the "Hide all traces..." option with the "Automatically save sections..." option will hide traces as sections are accessed and these traces will remain hidden even after the "Hide all traces..." option is turned off. To make the traces visible again the "Unhide all traces..." option should be turned on.

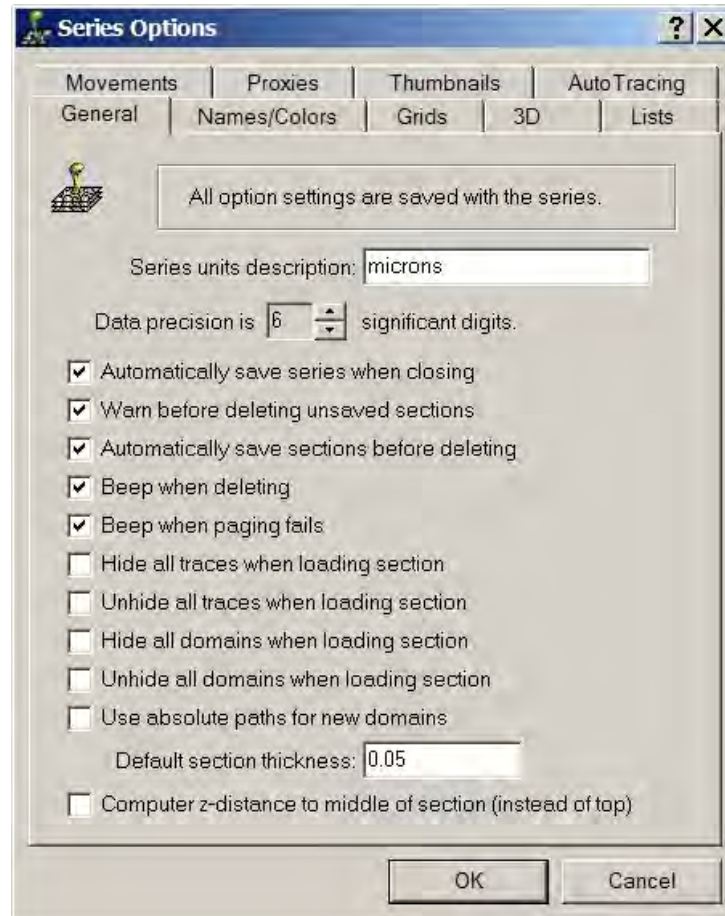


Figure 5.1. The Series Options Dialog. The General tab has options for automatically saving the series and sections files, and the units description and default section thickness. Any changes will be accepted with the OK button and rejected with the Cancel button. To ensure that the change persist after the series is closed, save the series before closing it. Shown are the default settings for a new series.

Names/Colors

The Names/Colors tab defines the default trace parameters, i.e. the name, color, and other attributes that will be used whenever a new trace is drawn. Use this tab to set the values before starting to draw. See Chapter 8: Traces for details on trace names and other attributes. The Names/Colors tab also allows the trace palette entries to be edited. The trace palette can be accessed from the Trace menu to quickly switch between trace names and attributes while drawing. Note that changes to palette entries are carried out immediately rather than waiting for an OK in the Series Options dialog. Such changes cannot be cancelled using the Cancel button.

Grids

The trace parameters for the Grid Tool are specified in the Grids tab. Grid traces are predefined shapes that are drawn onto the section in repetitive patterns or arrays. For more details on the types of grid patterns that can be generated using these options see Chapter 8: Traces.

3D

The 3D tab specifies how three-dimensional representations of objects are generated from traces (Chapter 11: Objects). These representations determine the size, position, and appearance of objects in the 3D Scene. The 3D tab also controls the section number range from which such representations are created. When generating objects from the traces make sure the sections are included in the section number range in this tab.

Lists

The options in the Lists tab (Fig. 5.2) determine what data values are displayed in the data columns of the list windows. The columns of the lists are specified by checkboxes for each list. These settings are saved with the series. The rows of a list are determined by the limit string in the Limit tool box. The limit strings are reset to the default value whenever a series is opened. The rows of this list will be all elements that match the limit string. A limit string can be a specific value such as "mito", which would display only rows named "mito".

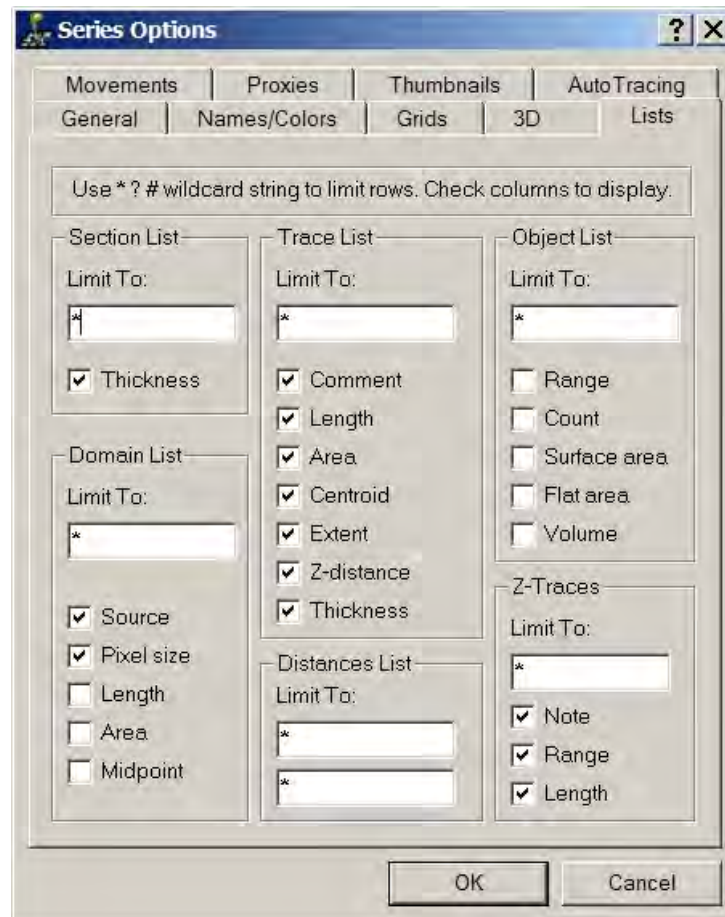


Figure 5.2. The Lists tab of the Series Options Dialog. The settings in the Lists tab determine what rows and columns will be displayed in the list windows. Any changes will be accepted with the OK button and rejected with the Cancel button. Shown are the default settings for a new series.

Wildcards and Displaying a List

Wildcard characters are asterisk (*), question mark (?), and pound sign (#). These characters allow a range of values to be displayed in a list when needed. An asterisk matches any string of characters. A question mark matches any single character. The pound sign matches any single number. For example, "D*" will limit the list rows to those that begin with a capital D. The limit string "D##" will list rows "D01", "D02" and "D99", but not with other characters such as "Da1", "D100", or "Dendrite".

It is possible in RECONSTRUCT™ to limit viewing structures of interest by manipulating object names and limit strings in the Lists tab of the Series Options. For example, if objects are named as follows:

```
N01D01
N01D01Synapse01
N01D01Synapse02

N01D01Synapse22
N01D02
N01D02Synapse01

N01D03
N01D03Synapse01

N02D01

N02D02

N03D01
```

Then, with the Limit string for the Object List set to "N01*", the Object List will only show the objects associated with neuron N01. To see all synapses on all neurons, one can use the limit string "N##D##Synapse*". To see all dendrites (D) on Neuron N02, use limit string "N02D##" and so on.

Movements

The Movements tab contains the options related to keyboard movements described in Chapter 10: Movements and Alignments. The initial values contained in this tab may not be appropriate for a given data set since this depends on the magnification used. For example, the initial value for translational movements (using the arrow keys) is one full series unit. So if your data is examined at high magnification 1 unit may move more than the whole screen width, resulting in objects "disappearing" whenever an arrow key is pressed. Adjust the size of the arrow key movements by editing the values in the "Translate X" and "Translate Y" boxes.

Proxies

The Proxies tab contains options for the use, creation, and deletion of proxy images on the section, described in detail in Chapter 7: Domain Images. When this tab is opened, the program searches through all the sections to determine how many domains have image proxies. This can be slow if there are a lot of sections, during which time the

program might become unresponsive. When the process is complete the percent of domain images that have proxies is displayed. When creating or deleting proxies the program again may become unresponsive while it processes all the section. The Cancel button is ineffective in terminating these operations.

Thumbnails

The Thumbnails options apply to the Thumbnail window, described in Chapter 6: Sections. The Thumbnail window displays a view of the sections in pre-rendered form. These section thumbnails can be quickly browsed and animated. (When viewing sections in the main window, each section is redrawn when it is accessed. This process is slow because each image is transformed into position using complex calculations.) Checking the box next to "Overlay thumbnails to make a flipbook" causes the Thumbnail window to display sections on top of each other, creating a fast version of the normally slow Page Up and Page Down operation.

Autotracing

The final tab contains the options related to automatic simplification, smoothing, and region growing of traces. Automatic simplification is always initially turned on and traces are "simplified" after drawing. Simplification eliminates loops and other artifacts that can reduce the accuracy of measurements and three-dimensional representations. It is advisable to keep this option on at all times. For more details on the other simplification options see Chapter 8: Traces. Wildfires are a form of autotracing based on a region growing algorithm. The "Stop..." parameters control when region growing will terminate. These parameters may need to be set carefully to get the Wildfire Tool to work at all. For more details see Chapter 8.

TABLE 5.1. Series Options Stored in the Series File

Option	Type	Default(s)	Description
index	integer	0	current section number
units	string	microns	units of measurement
viewport	floats	0, 0, 0.00254	position and magnification of view
defaultThickness	float	0.05	default section thickness
zMidSection	boolean	false	select top or middle for z-distances
autoSaveSeries	boolean	true	automatic series saves
autoSaveSection	boolean	true	automatic section saves
warnSaveSection	boolean	true	warn if section save needed
beepDeleting	boolean	true	beep when item is deleted
beepPaging	boolean	true	beep when paging fails
hideTraces	boolean	false	hide all traces when paging
unhideTraces	boolean	false	unhide all traces when paging
hideDomains	boolean	false	hide all domains when paging
unhideDomains	boolean	false	unhide all domains when paging
useAbsolutePath	boolean	false	absolute vs. relative path to image

thumbWidth	integer	128	width of thumbnail buttons
thumbHeight	integer	96	height of thumbnail buttons
displayThumbContours	boolean	true	display traces on thumbnail images
useFlipbookStyle	boolean	false	thumbnail overlay flag
fitThumbSections	boolean	false	center sections on thumbnails
firstThumbSection	integer	1	first section for thumbnails
lastThumbSection	integer	9999	last section for thumbnails
skipSections	integer	1	distance between thumbnails
flipRate	integer	5	frequency of shuttle play in Hz
useProxies	boolean	true	proxy images will be used
widthUseProxies	integer	2048	threshold width for creating proxy
heightUseProxies	integer	1536	threshold height for creating proxy
scaleProxies	float	0.25	scale factor for proxy dimensions
defaultBorder	floats	1, 0, 1	default trace border color
defaultFill	floats	1, 0, 1	default trace fill color
defaultMode	integer	9	default trace fill mode
defaultName	string	domain\$+	default trace name
defaultComment	string		default trace comment string
listSectionThickness	boolean	true	include thickness in section list
listDomainSource	boolean	true	include source path in domain list
listDomainPixelsize	boolean	true	include pixel size in domain list
listDomainLength	boolean	false	include length in domain list
listDomainArea	boolean	false	include area in domain list
listDomainMidpoint	boolean	false	include midpoint in domain list
listTraceComment	boolean	true	include comments in trace list
listTraceLength	boolean	false	include lengths in trace list
listTraceArea	boolean	true	include areas in trace list
listTraceCentroid	boolean	false	include x centroids in trace list
listTraceExtent	boolean	false	include y centroids in trace list
listTraceZ	boolean	false	include z distance in trace list
listTraceThickness	boolean	false	include sect thickness in trace list
listObjectRange	boolean	true	include section range in object list
listObjectCount	boolean	true	include trace count in object list
listObjectSurfarea	boolean	false	include surface area in object list
listObjectFlatarea	boolean	false	include flat area in object list
listObjectVolume	boolean	false	include volume in object list
listZTraceNote	boolean	true	include z-trace comment in list
listZTraceRange	boolean	true	include section range in z-trace

			list
listZTraceLength	boolean	true	include length in z-trace list
borderColors	floats	0, 0, 0	custom colors for trace borders
fillColors	floats	0, 0, 0	custom colors for trace fills
offset3D	floats	0, 0, 0	positional shift applied to object
max3Dconnection	float	-1	Boissonnat surfacing parameter
first3Dsection	integer	1	first section for 3D generation
last3Dsection	integer	9999	last section for 3D generation
type3Dobject	integer	0	type of 3D representation to use
upper3Dfaces	boolean	true	Boissonnat surfacing parameter
lower3Dfaces	boolean	true	Boissonnat surfacing parameter
faceNormals	boolean	false	whether to create face normals
vertexNormals	boolean	true	whether to create vertex normals
facets3D	integer	8	density of faces in 3D output
dim3D	floats	-1, -1, -1	dimensions of 3D object
gridType	integer	0	default type of grid trace to draw
gridNumber	integers	1, 1	number of elements in x and y
gridSize	floats	1, 1	size of elements in x and y
gridDistance	floats	1, 1	separation between elements in x,y
hueStopWhen	integer	3	wildfire hue stop criterion
hueStopValue	integer	50	wildfire hue stop value
satStopWhen	integer	3	wildfire saturation stop criterion
satStopValue	integer	50	wildfire saturation stop value
brightStopWhen	integer	0	wildfire intensity stop criterion
areaStopPercent	float	999.0	wildfire autotrace area change
areaStopSize	float	0.0	wildfire minimum region size
brightStopValue	integer	100	wildfire intensity stop value
tracesStopWhen	boolean	false	wildfires stop at traces if true
mvmtIncrement	floats	0.022 1 1 1.01 1.01 0.02 0.02 0.001 0.001	sizes of keyboard movements
ctrlIncrement	floats	0.0044 0.01 0.01 1.002 1.002 0.004 0.004 0.0002 0.0002	sizes of keyboard movements when Ctrl key is held down
shiftIncrement	floats	0.11 10 10 1.05 1.05 0.1 0.1 0.005 0.005	sizes of keyboard movements when the Shift key is held down
contours	various	various	palette contour attributes and stamp shapes

Exporting and Importing Images

Sections contain a set of transformed images (domains) and traces. Importing multiple images into sections can be carried out by opening the *Images...* dialog from Series Import menu. This dialog is documented in *Creating a Domain* in Chapter 7: Domain Images. Entire sets of section data can be exported to image files using the *Images...* dialog from the Series Export menu. This dialog is described in *Exporting Section Images* in Chapter 6: Sections.

Exporting Trace Lists

The Trace List window provides a listing of traces and trace measurements from the current section (see Chapter 4: User Interface and Chapter 8: Traces). These individual Trace Lists can be saved to a file (and then imported into Excel) by using the *Save...* menu item at the top of the list window. The Series Export of all Trace Lists provides a method for collating all the individual section lists in one output file. Choose the filename in the *Save As...* dialog and the export operation will automatically generate and save the Trace Lists for all the sections in the series. The output is constrained by the options selected in the Lists tab of the Series Options, so the export can be limited to a subset of trace names and a subset of measurements. The output is a Comma Separated Values file (.csv) identical to the format for an individual Trace List with the addition of a section number field at the beginning of each row.

Exporting Traces as Line Drawings

The *Lines...* command in the Series Export menu allows the traces in a range of sections to be exported as .dxf files, one file for each section. All the non-hidden traces within a section are exported as POLYLINES with x, y, and z values in series coordinates. See the description of the DXF Export format in Software Overview for more details.

Importing Line Drawings

Line drawings in DXF format can be imported into a series using the *Lines...* item from the Series Import menu. The format of the line importing dialog takes the same form as the Import Images dialog. DXF files are selected into the dialog and then section numbers are assigned to these files using the dialog parameters. When section numbers are correctly specified, the Import button stores the DXF data as traces on the corresponding sections. The DXF data is assumed to be in the same units as the series.

Importing Series Data

Elements from an existing series can be copied into the current series using the *Series...* dialog from the Series Import menu (Fig. 5.3). Use the Browse button to navigate to a folder and select an existing .ser file to copy from. Then select what elements are to be copied by checking the appropriate boxes. Importing the Series Options or Trace Palette will overwrite the values in the current series. Importing z-traces will add to the current series only the z-traces with names that match the string. String matching of names follows the same rules as for limit strings in the Lists tab of the Series Options (above).

WARNING: Importing series elements cannot be undone!

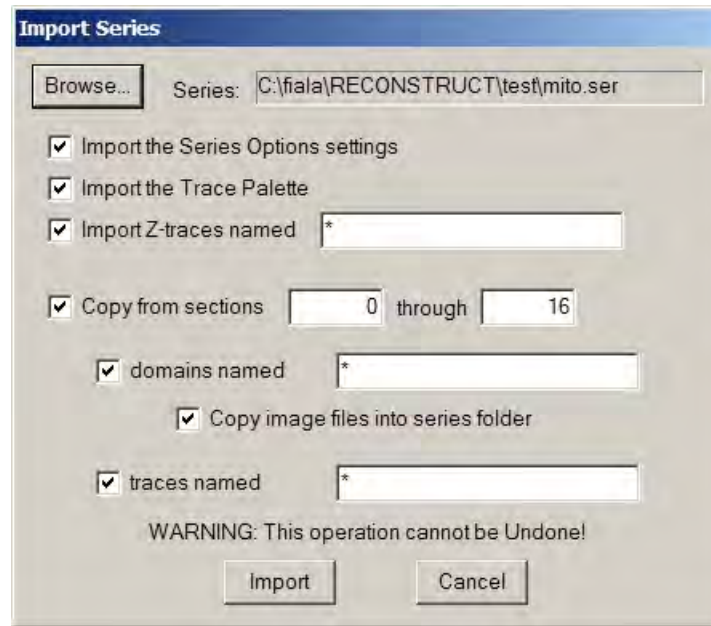


Figure 5.3. The Import Series dialog. The Series name is filled by selecting a series file with the Browse button. When copying section data, be sure to set the section range. Use the limit strings to select a subset of the data for import. Check the “Copy image files” box to copy the image data into the same folder as the current series. Without this, the current series will create a relative path to the images in the original folder.

Checking the "Copy from Sections" box enables importing of section data from a series. Data will be imported only from existing sections within the range specified. Importing domains or traces will add these elements to the corresponding sections in the current series. If a corresponding section does not exist in the current series, then the section will be created. The matching strings in the Import Series dialog can be used to select which domains or objects to import. This trace matching string works just like the limit string in the Object List. So if you can limit the Object List to just the objects desired to be exported from one series, using that same limit string in the Import Series dialog will ensure that only those objects are imported. Be aware that importing an existing object will create multiple copies of the object traces in the current series. To merge an object from two series, first rename the object in the current series, then import it from the other series and perform a section-by-section comparison of traces to decide which ones to keep.

The Import Series dialog can be used to copy and rename an entire series. This includes copying the actual image source files, but it does not facilitate renaming these files.



Chapter 6: Sections

Sections are the main data containers of a series. A section contains a set of images (domains) and a set of polylines (traces). The domains and traces for each section are stored in a separate file, the section file (Chapter 3).

Section Indexing

Since serial sections have an intrinsic order, the position of a section in this order is used as the file extension of the section filename. Sections must be numbered with non-negative integers in the range of 0 through 2147483646. The section with the smallest section number is referred to as the *first* section, while the highest numbered section is the *last* section.

The section index or number is also stored internally in the section file. The internal index and the file extension must be the same. Leading zeros and spaces should be avoided in the section extension since these will not be maintained when saving the section file.

Sections do not have to be numbered sequentially. Gaps in section numbering are generally ignored. But note that gaps may require careful interpretation for 3D reconstructions and measurements. The position of a section in the volume is determined by the combined section thicknesses of lower numbered sections (see *Z Distance* in Chapter 9: Calibration and Measurement), so the section thicknesses must account for any gaps in the series to get correct 3D distances.

To reorder some sections in a series, use the Numbering item from the Modify menu of the Section List. First highlight the sections you want to move in the list, then open the Numbering dialog and enter the amount to shift the section numbers. You may need to shift parts of your series to create a number gap into which you can move the sections.

Creating a Section

A new section can be created by selecting the *New...* item from the Section menu. Sections can also be created by importing sections, images, or lines into multiple sections using the dialogs from the Series Import menu.

The New Section dialog asks for the number of the new section. The default value is the new sequential number after the index of the current section. After a number is entered and the OK button pressed, a section file with default attributes is created with that index. If a section with the index already exists, no new section is created and an error message is displayed.

After creating the new section, the main window switches to display it. Initially the new section contains no data. A section without images or traces displays as black. Data can be added to the section by importing domain images (Chapter 7: Domain Images) and by drawing traces (Chapter 8: Traces). The new section is given the default section thickness defined in the Series Options General tab.

Saving a Section

By default, section data is automatically saved to the section file before it is deleted from memory. This can be changed in the General tab of the Series Options. When sections are not automatically saved, the user will be warned before any changes are deleted provided the "Warn before deleting unsaved sections" checkbox is set. A section can be manually saved by selecting *Save* from the Section menu. Since the filename is already specified there is no dialog, the save operation will be executed immediately.

Displaying a Section

Images and traces within a section are defined in a local coordinate system and transformed into position on the section by a general nonlinear transformation. To display a section, each image data file is read from disk and transformed pixel-by-pixel into position on the section's display bitmap. To support color images with different color formats, the display bitmap uses 24-bit color. After all images are rendered, the line traces are transformed into position and drawn using the windows graphics device interface. Traces are drawn with or without a colored fill to highlight the interior of the object. Finally, the rendered bitmap is copied to the client area of the window. The process of rendering a section can be slow depending on how much data is to be displayed and the media on which the data is stored. When a section is being rendered for display the cursor will change to the wait/busy cursor, typically an hourglass. When rendering is complete the section will be displayed and the cursor will be restored.

Centering

The set of images that comprise the section can be centered in the main window by selecting the *Center* item from the Zoom submenu of the Section menu. Pressing the Home key will accomplish this quickly. To center a section that contains only trace data, use *Select All* from the Trace menu to select all the traces then use *Zoom To* from the Traces menu (Ctrl-Home),

Zooming and Panning

The magnification of the section display can be changed by using the Zoom submenu of the Section menu. Using *Center* or *Actual Pixels* will change the magnification to either center the section image data or display the image data at the screen resolution, respectively. The previous magnification and position can be restored by using the *Last Zoom* item, or equivalently hitting the End key. Using the *Magnification...* dialog will allow you to set the magnification to a specific number of screen pixels per series unit.



The magnification can also be changed with the Select Zoom Region Tool from the Tools window. Select the tool that looks like a magnifying glass as shown to the left. Depress and hold the left mouse button when the cursor is at the corner of the desired region in the main window. Drag the rectangle to cover the desired region and release the button. The section will be re-rendered to the largest magnification at which this region still fits entirely within the main window. The previous magnification can be regained by clicking the right mouse button.



The Pan and Zoom Tool can be used to dynamically adjust the section magnification and position. When this tool has been selected the cursor will be crossed arrows. Drag with the left mouse button depressed to pan the section. After panning, the hourglass may again appear while the section is re-rendered with the new settings. Holding the right mouse button down changes the tool to the zoom cursor. Dragging with the right mouse button zooms the section. Drag upward to decrease the

magnification, and downward to increase the magnification. A particular point in the section can be zoomed by putting the cursor over the point before pressing the right mouse button.

Zooming in will ultimately lead to a pixelized image where each pixel of the original data appears as a square on the screen. This is normal. But zooming in while proxy images are used for low magnification (Chapter 7: Domain Images) may result in premature pixelization. To make sure the display is at the maximum resolution of the data, page through the sections to force a load of the full image instead of the proxy. A section reload can also be accomplished by selecting the section from the section list.

Note that panning and zooming is entirely different from moving a section (Chapter 10: Movements and Alignments). Panning and zooming temporarily changes the view point for the whole series. All sections are affected in the same way. The relative position of a section is not changed by panning or zooming. Moving or aligning a section permanently changes the relationship of the section to the rest of series.

Paging Sections

Sections in the series can be accessed sequentially by using the Page Up or Page Down keys. Page Up changes the displayed section to the next higher section number, while Page Down changes to the next lower section number. If no higher or lower section is found for Page Up or Page Down, respectively, a warning beep may be issued (depending on the Series Options settings). Paging can also be performed by the mouse wheel when the mouse is in the main window.

Blending and Flickering Sections

RECONSTRUCT™ maintains two sections in memory, the currently displayed section and the section previously displayed. It is therefore possible to rapidly switch the main window display between these two sections. This can be done with the / key. Pressing and releasing the / key will switch the display back to the previous section. If no previous section is available a warning beep will be issued provided the "Beep when paging fails" option is turned on.

Holding the / key down will flicker between to the two memory sections. Flickering sections is sometimes useful to detecting misalignments by observing apparent motion of objects in the images (Chapter 10: Movements and Alignments). The first / after paging may require re-rendering of the previous section. Since this operation may be slow the hourglass may appear briefly during the first toggle of sections. Subsequent switching will be very fast since both sections will be fully rendered in memory.

The two memory sections, sometimes referred to as the current and previous sections, can also be blended together by selecting *Blend* from the Section menu. Once in blend mode, all displays will be the blending of the last two sections. Thus, Page Up or Page Down can be used to sequentially compare pairs of sections in a series by blending. Blending can be turned off by selecting the *Blend* item again, or by hitting the spacebar or the / key.

Listing all Sections

All the sections in a series can be listed using the Section List window described in Chapter 4: User Interface. The Section List window is opened by selecting List from the

Section menu. The columns and rows displayed in the Section List are determined by the settings in the Lists tab of the Series Options. The entire list can be saved to file using the Save... command from the List menu at the top of the Section List window.

To speed up the display of the Section List, it is created from the section files in the series folder when the series is opened. The list is updated as changes are made within RECONSTRUCT™, but changes made outside the program (such as deleting section files) will not be incorporated in the list until the series is reopened from the Series menu.

Editing Section Attributes

Each section has three modifiable attributes: the section thickness, the movement lock flag, and the section index. Changes to these attributes can be made from the Section menu or from the Modify menu at the top of the list.

Movement Lock

The lock/unlock attribute can be toggled by selecting the *Lock* or *Unlock* item from the Movement submenu of the Section Menu. When a section is unlocked, movements of the entire section are permitted. Accidentally hitting an arrow or function key on an unlocked section will move the section relative to the rest of the series (Chapter 10: Movements and Alignments). Therefore, it is a good idea to maintain the movement lock at all times. The lock/unlock attribute is reflected in the icon on the left end of the status bar.

The lock/unlock attribute of multiple sections can be changed by one command from the Section List. Open the Section List and highlight the sections to be changed. Select *Lock* or *Unlock* from the Modify menu at the top the Section List. The changes will be immediately reflected in the Section List icons. A red, closed padlock appears by sections that are locked, while unlocked sections display a clear, open padlock.

Thickness

The section thickness is changed by selecting the *Thickness...* item from the Section menu and typing in a new value. The section thicknesses of many sections can be changed to the same value using the *Thickness...* item from the Modify menu of the Section List (Fig. 6.1).

Warning: There is no Undo for Modify > Thickness... from the Section List.

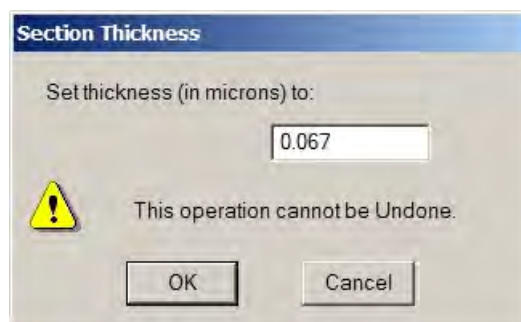


Figure 6.1. Dialog for modifying section thickness from the Modify menu of the Section List. All the sections which are highlighted in the list at the time the dialog is invoked will be changed and this change cannot be globally recovered with an Undo command.

Section Numbering

Sections can be renumbered using the *Shift Numbering...* command from the Modify menu of the Section List. This dialog allows a shift to be applied to the section numbers of all sections highlighted in the list. The example illustrated in Figure 6.2 is an increase in section number. The integer value entered (e.g. 1) is added to each section number. Thus, a positive value shifts the section numbers up, and a negative value shifts the sections numbers down. Only constant shifts of section numbers are performed and existing sections cannot be overwritten. When a section would overwrite another section, renumbering is terminated.

The order of a group of sections can be reversed using the *Reverse Order...* command from the Modify menu of the Section List. Highlight the sections to be reversed in the Section List and then select the command. The data in the highlighted sections will be swapped, the lowest numbered section with the highest section, then the second lowest with the second highest, and so on.

Warning: There is no Undo for any of the section renumbering commands.

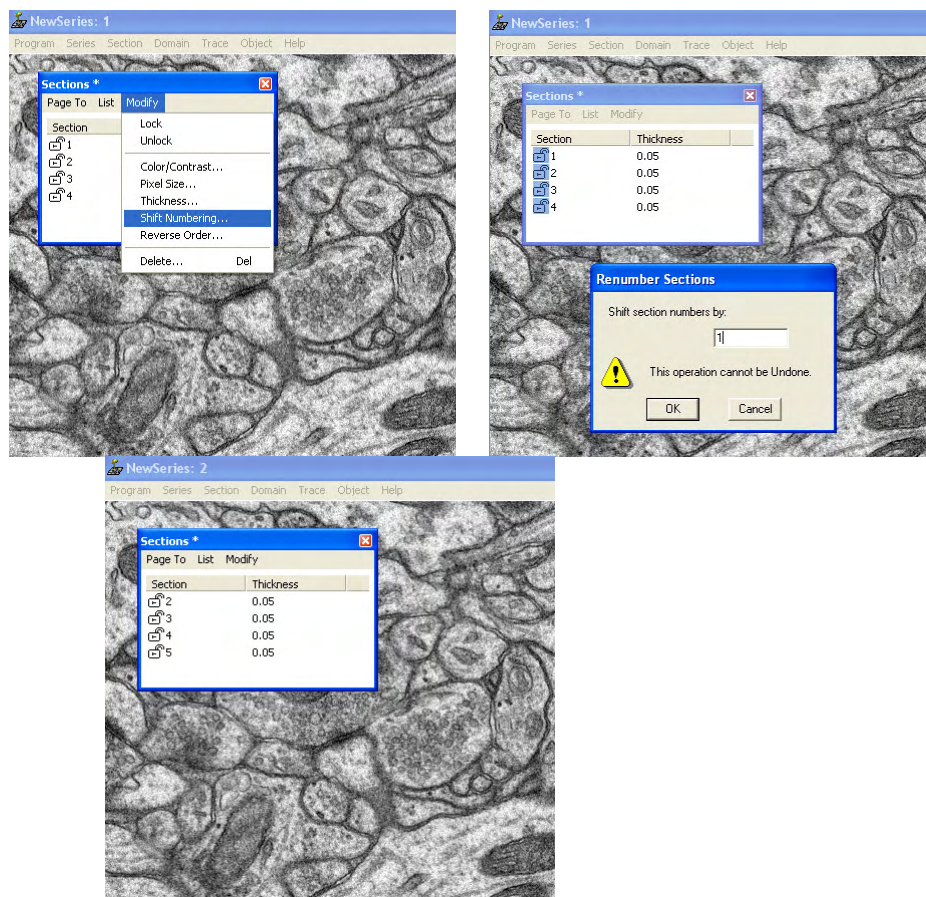


Figure 6.2. Dialogs for modifying section numbering from the Modify menu of the Section List (left to right). All the sections which are highlighted in the list at the time the dialog is invoked will be changed and this change cannot be altered with an Undo command.

Deleting Sections

Sections can be deleted from the series using the Section List. Highlight the sections in the list to be deleted and select the *Delete...* item from the Modify menu or use the Delete key. This will delete the section files permanently from the series (Fig. 6.3).

Warning: There is no Undo for Delete Sections.

Section files refer to image data that is stored in separate image files. To delete these image files along with the section files, set the checkbox in the Delete Sections dialog.

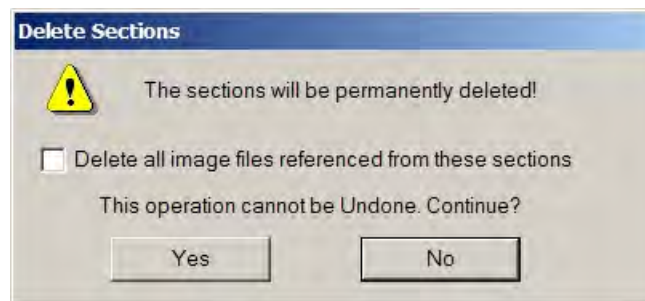


Figure 6.3. The Delete Sections dialog invoked whenever the user requests to delete sections from the Section List.

The Undo Stack

Each section has an undo stack associated with it. The undo stack keeps copies of the section data as changes are made, allowing the user to return to previous versions of the section. Changes saved in the undo stack include movements of sections, domains, and traces, adding or deleting traces and domains, etc. The undo stack can hold up to 16 changes. The undo stack for a section is cleared each time the section is loaded into memory. A section is loaded into memory whenever it is displayed in the main window as the current section. The exception is when it was the previous section displayed, in which case the undo stack is retained in memory. Only when two new sections have been displayed is the undo stack for a section finally deleted from memory.

To restore the section to the state prior to the last change, select the *Undo* item from the Section menu. This can be accomplished more quickly by using Ctrl-Z on the keyboard. If the *Undo* item is disabled then the undo stack is empty. To return to the state prior to the first change in the undo stack, select the *Reset* item from the Section menu.

Immediately after an undo is made, the state of the section prior to the undo can be restored using the *Redo* command from the Section menu. Once additional changes are made the *Redo* option is no longer available and the *Redo* item will be disabled. *Redo* can also be used to restore the previous state after a *Reset* operation.

Thumbnails

Thumbnail representations of the sections are displayed in the Thumbnails window as described briefly in Chapter 4: User Interface. Two modes of display are possible. Thumbnail buttons can be displayed in a sequential array in which many sections are

visible simultaneously. Alternatively the thumbnails can be overlaid on top of each other in a flipbook arrangement. This mode of display and other parameters of the thumbnails are set in the Thumbnails tab of the Series Options dialog (Fig. 6.4).

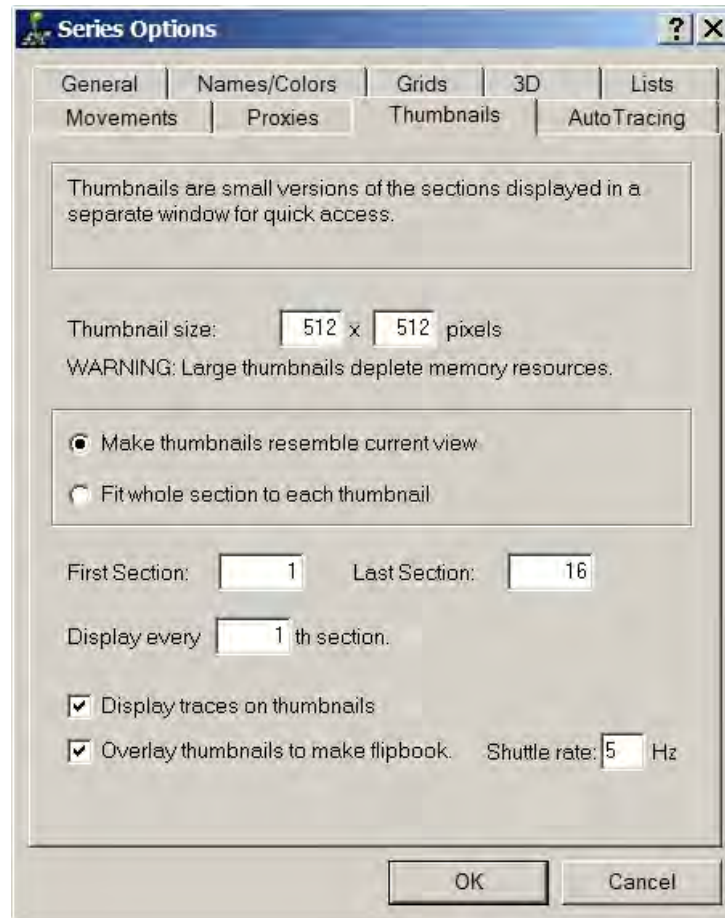


Figure 6.4. The Thumbnails tab of the Series Options dialog. These options allow the thumbnail size and style to be changed. Changes are accepted by selecting the OK button, and rejected by the Cancel button. Changes will appear the next time the thumbnails window is opened.

The thumbnail size parameters specify the width and height of the images in screen pixels. Any size and aspect ratio can be specified but consider that the thumbnail images will all be maintained in memory. Thus 100 sections rendered at 128x96 would use about 4 Mb of memory, while the same number of thumbnails at 512x384 would require an additional 52 Mb. Also rendering of the buttons will take longer for larger image sizes. Rendering of the buttons may require nearly as much time as paging the sections, depending on how slow the retrieval of image data is. However once the thumbnails are rendered access will be much faster than paging.

There is an absolute limit on the amount memory Windows will allocate for button images. This value is independent of the amount of physical memory installed in the computer. If Windows runs out of memory for button images, the first part of the thumbnail series will have sections images but the remaining buttons will all be gray. This is more likely to occur when you ask for large thumbnail images. This problem appears to be absent from the Windows Vista operating system.

Thumbnail buttons are created initially as gray buttons without images. The images are added sequentially by a background thread that runs at a lower priority to the main window. Thus, you can continue working while the thumbnail images are being rendered. When the Thumbnails window is closed, all of the thumbnail images are deleted from memory and they will need to be recreated the next time the Thumbnails window is reopened. Changes to the thumbnail options will only be realized when the Thumbnails window is opened and the thumbnails are recreated.

Two ways of rendering the thumbnails are selectable. Thumbnails can be made to resemble the current view or the section can be centered on the button. In the former case the thumbnail images will maintain the existing section-to-section alignment and can be zoomed into particular parts of the section. While in the later case the entire section will be visible in each thumbnail. Making the thumbnails resemble the current view is most useful for flipbooks, while the other mode can be used for a quick visual reference in the sequential mode.

Limiting the range of sections for thumbnails or skipping sections can speed up thumbnail generation. When traces are displayed on the thumbnails they are drawn using the normal trace attributes for color and filling.

By choosing the flipbook style, thumbnails will be overlaid rather than arrayed sequentially. Thus the thumbnail window will be sized to display only one thumbnail button. Resizing the window is not possible. Selecting which section is visible in the flipbook can be done by using the scrollbar at the bottom of the window or by using the mouse scroll button. Clicking the scrollbar with the left mouse button will begin shuttle play of the thumbnails in the flipbook. In shuttle play the thumbnails are displayed first in one direction and then in the other, continuously. The rate of switching images is determined by the Shuttle parameter in the Thumbnails tab of the Series Options. Shuttle play can also be started and stopped by the Tab key when the thumbnails window has the keyboard focus.

Exporting Section Images

Sections can be exported to image files in much the same way they are rendered on screen or on thumbnail buttons. This provides a way to export aligned montages created with RECONSTRUCT™ to other programs for further analysis. To create new image files from the sections, select the *Images...* item from the Series Export menu. The Export Images dialog (Fig 6.5) allows the range of sections to be selected. By default, all sections in the series will be rendered.

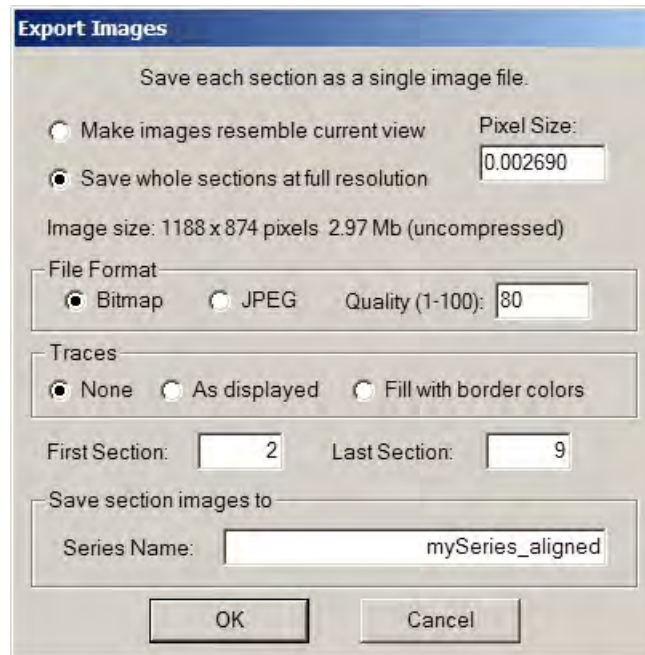


Figure 6.5. The Export Images dialog. The Pixel Size parameter determines the output file size. Output will be a 24-bit Windows Bitmap or JPEG image file. The output image can include a representation of the traces if desired. Each section in the specified section range will be rendered to a file with the Series Name and the section number as the file extension.

When the dialog is opened, the series is scanned for all sections and the image dimensions necessary to contain all the series data is determined. The output images will be sized accordingly so that every section is completely visible. The images will be given the same orientation and position as they appear when displayed in the main window, so aligned sections will be aligned in the output images.

Note when images are exported, they are shown with the identity transformation and a black border may appear on the image (Fig. 6.6; “an irregular space around the image”). The position and shape of the image relative to the black border is determined by the original transformation used to align the sections. This transformation represents a mapping from the transformed image back into the original image, telling the program where to find the data corresponding to a particular location on the section. For details about the transformation itself see Fiala and Harris, 2001; (<http://www.jamia.org/cgi/content/abstract/8/1/1>). If one zooms into a section such that there is no black border on screen and selects “Make Images Resemble Current View” from the Export Images dialog (Fig. 6.5) there will not be a black border on the exported images.

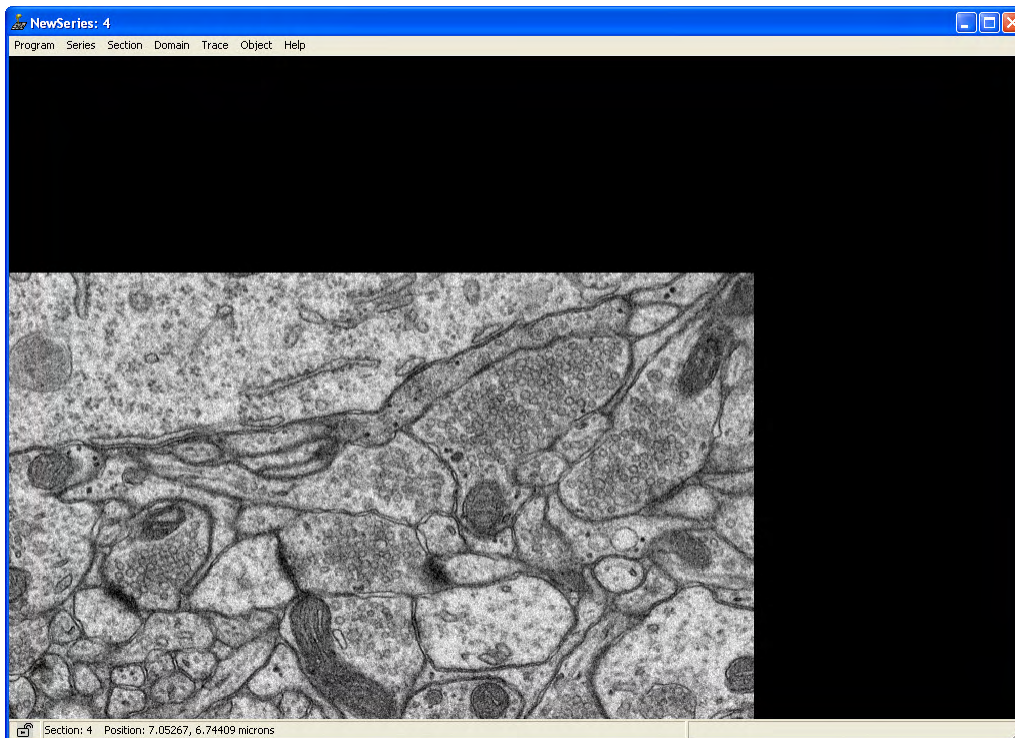


Figure 6.6. A black border may appear around an image. The position of the image relative to the black border is determined by the original transformation used to align the images.

The sections can be rendered at any resolution by adjusting the pixel size. The pixel size is the number of section units per pixel (Chapter 9: Calibration and Measurement). By decreasing the pixel size, more pixels will be used to render the data. A larger pixel size will result in fewer pixels being used to represent the same image data. An arbitrary positive number can be typed in the Pixel Size box to specify the number of pixels used in the output.

For convenience two different pixel sizes are automatically provided for full and current view resolutions. Full resolution uses the smallest pixel size detected among all domains in all sections. Essentially every pixel in the original data will have at least one full pixel in the rendered output, (provided there are no scaling transformations applied to the data.). For any pixel size, the resulting width and height (in pixels) of the output files is given in the Export Images dialog (Fig. 6.5).

When the option to *"Make images resemble current view"* is selected the pixel size width and height is set to that of the main window, and the offset is set to the current view offset. Thus each section will be made to appear as it does in the main window. For all other pixel sizes, the width, height, and offset is set to fully contain all domains in all sections at the selected resolution.

The file format of the images can be selected to be either an uncompressed 24-bit Windows Bitmap or a compressed JPEG image file. The Quality factor determines how much compression is applied to the JPEG output. A value of 100 has the least compression and the best preservation of the image data, while a value 1 has maximum compression at the expense of image quality. A value of 80 is usually appropriate for compression while maintaining most of the image quality.

Each image filename is given a leading underscore, followed by the string entered in the Series Name box, followed by a numerical extension indicating the section number, followed by the file type. To allow verification of the rendered images using RECONSTRUCT™, new section files are created along with the image files. A new series file is also created with the name entered in the Series Name box. All the files are created in the same folder as the original series. Thus, an attempt to overwrite the series by giving the same name in the Series Name box will result in an error message.

To display only the image data in the section, select None in the Traces option box. To display traces in the same way they appear in the main window, select the "As displayed" option. This will produce a similar effect that seen in the Thumbnails window when the "Display traces on thumbnails" option is checked. Hidden traces will not appear in the exported image, and the Series Options to hide/unhide all traces will effect whether traces appear. Any changes to the current sections must be saved to the section file to be reflected in the exported images. This includes traces which have just been hidden or unhidden.

An alternative method of displaying trace information on rendered sections is the "Fill with border colors" option. With this option, traced regions will be colored with the border color using the Colorize algorithm. Colorize replaces only the most intense (most white) colors in the region with the border color. Using very light, pastel colors for the border colors will result in colored regions in which much of the underlying contrast is still visible, as shown in Fig. 6.7. This algorithm also allows coloring with subtle shades that would not give nice fills in the main window. Like all fills, the Colorize fills may give different results for different trace orders (Fig. 6.7). A trace can be moved to the top of the implicit ordering by first selecting it, then deselecting it, and then saving the section.




Figure 6.7. The effect of Fill with border colors. Right: Original section as it appeared in the main window. Middle: After Render Section with border color fill. Left: The effect is altered by a different order of the traces. In this case the yellow region was on top, so the interior regions were colored first and then yellow colored over everything, so the interior appear only lightly darkened. In the middle image, the interior regions were on top so they were colored last.

After the OK button is selected in the Export Sections dialog, the sections begin being rendered and the new files created. A small dialog box appears showing the progress of the rendering. The operation can be aborted by using the Abort button in the dialog or the Ctrl-Pause key sequence. Aborting may be slow, but additional sections will not be processed after the Aborting message appears in the dialog.

When all sections have been rendered, a message is displayed indicating that the operation was completed.

The new exported images may be reviewed by opening the new series from the Series menu.

Copyright © 2006 by John C. Fiala
Revised by K. Harris and K. Sorra; 2009



Chapter 7: Domain Images

Domains are image regions defined by an image file (.bmp, .jpg, etc.), and the boundary outlining the region to be displayed on the section.

Creating a Domain

Domains can be created within a section by selecting *Import image...* from the Domain menu. This will open an Import Image dialog for selecting the image file for the domain (Fig. 7.1). Specify the estimated pixel size of the image (Chapter 9: Calibration and Measurement) at the bottom of the dialog before importing. Backing up or transferring series can be facilitated by keeping all component image and series files in one folder. For this reason it is recommended that the "Copy file to series folder" option is used.




Figure 7.1. The Import Images dialog from the Domain menu. The list shows three image files that may be selected as domain images to import.

When a domain image is imported, it is placed at the origin, i.e. $x=0$ and $y=0$, of the section. If the current view does not include the origin you may not be able to see the new domain image. Use the Home key to center the view of the entire section including the new domain.

Importing Multiple Images

Multiple images are added to the sections using the *Images...* dialog from the Series Import menu (Fig. 7.2). This dialog allows a set of images to be added to sequential sections in the series, one domain image per section. When a candidate list of images is ready for import, use the Import button in the upper right of the dialog to carry out the operation.

The Pixel Size box specifies the width of an image pixel in series units (Chapter 9: Calibration and Measurement). All domains created by Import Images will receive this pixel size attribute. The pixel size of a domain can be changed later by editing the domain's attributes or calibrating the pixel size.

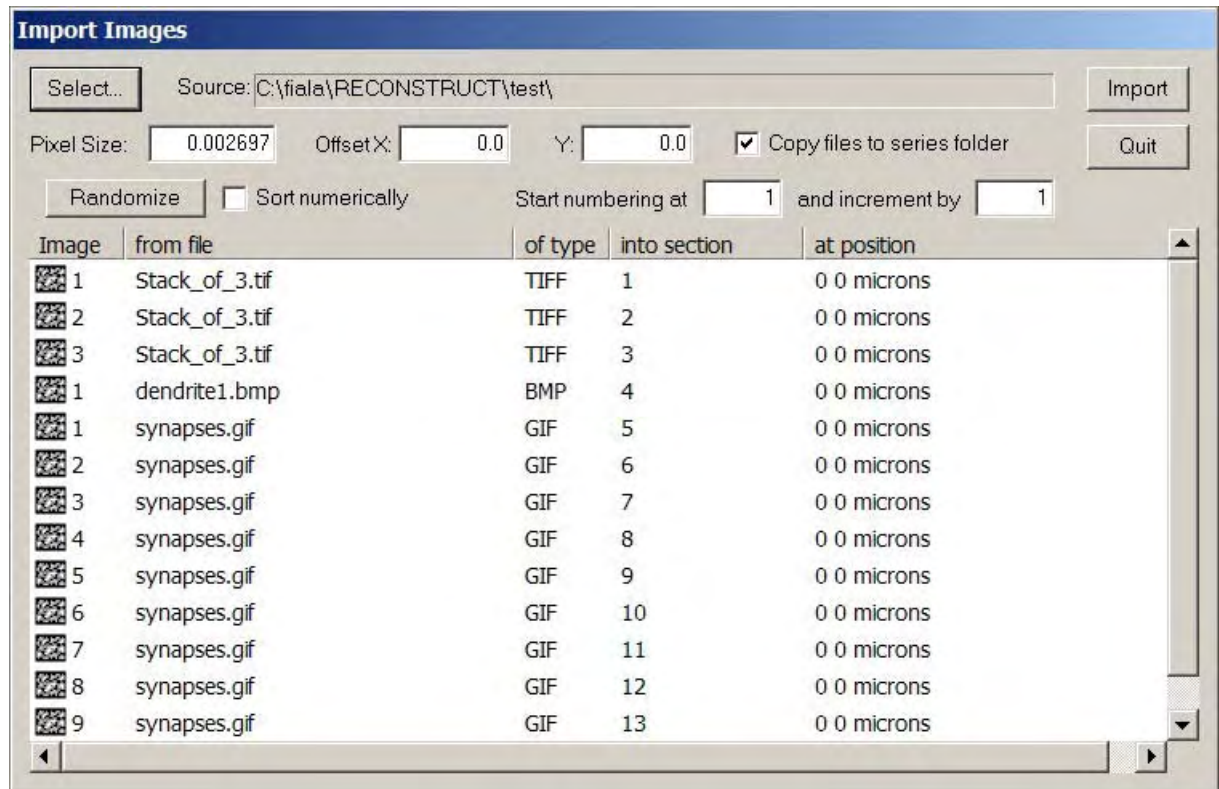


Figure 7.2. The Import Images dialog. The list shows three types of image files (*Stack_of_3.tif*, *dendrite1.bmp*, and *synapses.gif*) that were selected into the dialog. *Stack_of_3.tif* is a TIFF stack of 3 images numbered 1, 2, 3 in the leftmost column. *Dendrite1.bmp* is a single image. *Synapses.gif* is a GIF stack of 9 images. The images will be imported in the order shown into sections 1 through 13.

To fill the Import Images list with candidate images, use the Select button in the upper left corner of the dialog. Highlight one or more images in the Select Files dialog (Fig. 7.3). Use the Open button to close the dialog and add the candidate images to the list. Only files that the program can recognize as readable image files will be added to the list. Acceptable image file formats are Windows and OS/2 bitmaps (.BMP .VGA .BGA .RLE .DIB .RL4 .RL8) except for CCITT G3 MH Fax compressed and 24 bit RLE files, CompuServe GIF files (.GIF), JPEG File Interchange Format files (.JPG .JPEG .JPE), TIFF files (.TIF .TIFF) with 1, 4, 8 and 24-bits per pixel, Truevision Targa/Vista bitmaps (.TGA .VST .AFI), ZSoft PC Paintbrush Image format (.PCX .PCC), GEM Raster (.IMG), Portable Greyscale-map (binary P5 type) (.PGM), and Portable Pixel-map (binary P6 type) (.PPM).




Figure 7.3. The Select Files dialog. This dialog is invoked by the Select button in the Import Images dialog. Multiple image files have been selected from the folder “test”. These files will be placed in the Import Images list by clicking the Open button.

Once the Import Images list has been filled with one or more images, specify the first section into which the images will be imported by typing a number in the First Section box. The target sections for the images in the list will be numbered sequentially from this first section. If the first section is negative, the parameter is ignored and an attempt is made to extract the desired section numbers from filenames. This works when the filenames already have the desired section number as the file extension, for instance.

The order of image files in the list can be modified by using the Sort numerically checkbox. This option treats the filename as having an embedded numerical value and resorts the list accordingly in ascending numerical order. The Randomize button puts the image files in random order. It is not possible to reorder individual list items. To get finer grain control over the ordering of the images, just select one image at a time.

GIF and TIFF stacks containing multiple images per file can also be imported through Import Images (Fig 7.2). During the import process, individual subimages of a stack are copied to BMP files in the series folder. This occurs regardless of the “Copy files...” checkbox since image files need to be individually manipulated. Each subimage of the stack can be imported into a different section as specified in the “into section” column of the list.

When all domains have been imported, exit the Import Images dialog by selecting the Quit button. If not already displaying a section in the main window, use Page Up and Home to display the first section.

Montages

Images can be tiled within a section to create a montage. A montage is created by importing multiple domains and placing them at the correct positions within the section. A section can contain no more than 254 image domains. The initial position of each domain can be set in the Import Images dialog by using the Offset X and Offset Y values. Multiple domains can be added to a section positioned at different offsets by filling the list and clicking the Import button repeatedly, each time with a different offset specified. To learn how to montage domains using the point correspondence method, see Chapter 10: Movements and Alignments.

Listing Domains

The Domain List window is opened by selecting *List image domains...* from the Domain menu. Every domain in the section is included provided the domain name matches the limit string in Lists tab of the Series Options. The limit string appears at the end of the text in the title bar (Fig. 7.4). When changes are made to domains in the section (e.g. by importing images), or a change is made to the limit string in the Lists tab, the list may need to be regenerated.

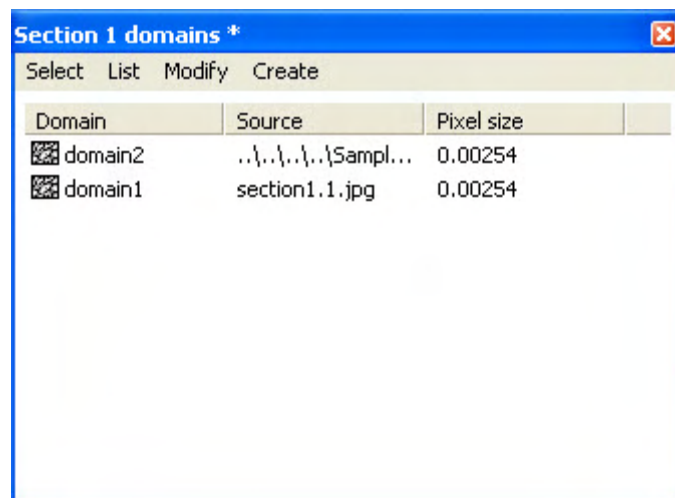


Figure 7.4. The Domain List window. Section 1 has two domains, named domain1 and domain2. Domain1 has an absolute source path, while the source path for domain2 is relative to the series folder. The limit string "*" appearing in the title bar ensures all domains are listed.

Each domain is listed by name in the order that they are painted onto the section (Fig. 7.4). The domain name is the name of the domain boundary trace, so it will be created from the default trace attributes at the time the domain is added to the section. In addition to the domain name, the list can contain columns denoting the path to the source image file, the pixel size of the image, the domain boundary length, the domain boundary area, and the midpoint of the domain (all in series units). The columns displayed in the Domain List are determined by the checkboxes in the Lists tab of the Series Options. The entire list can be saved to a file using the command from the List menu at the top of the list window.

Only a single domain can be highlighted in the list. Pressing Enter or clicking the *Select* menu item at the top of the list will select the domain from the section and display it in isolation. Whenever a domain is selected, the Domain List will be automatically closed.

The *Modify* menu at the top of the Domain List window allows the highlighted domain to be modified in several ways. The domain can be deleted, hidden from view, or unhidden. The *Attributes...* choice on the *Modify* menu allows the attributes (including the domain name, the source path, and the pixel size) of the highlighted domain to be edited without selecting the domain from the section.

A copy of a domain is made by highlighting it in the list and selecting *Create*, then *Copy*.

Selecting a Domain



A domain in a section can be selected using the Select Domain Image Tool from the Tools window. Click on an image in the section with the left mouse and the domain will be selected. The boundary of a domain is visible only when it is selected. In addition, this boundary can be edited only when it is selected (see below).

A selected domain is removed from the section and displayed alone in the front display. The back display contains the remaining components of the section. So the same techniques for comparing sections can be used to compare a domain to its section or to an adjacent section. Use the *Blend* command from the Section menu (spacebar) to overlay the domain and the section. To flicker between the domain and the rest of the section use the / key.

A domain can also be selected from the Domain List, as mentioned above. Only one domain can be selected from the section at a time. In order to select a different domain, the currently selected one must be merged back into the section. A selected domain can be permanently deleted from a section by using the Delete key when the domain is displayed, or by using the *Delete* item from the Domain menu when the domain is selected.

Domain Boundary

The domain boundary is a trace that defines what part of the image will be displayed. The initial default boundary is a rectangle surrounding the entire image. To change the domain boundary, first select the domain from the section. Then select a drawing tool that creates a closed boundary such as the Freehand Pencil or the Point-by-Point Tool. Define the default trace attributes to use for the boundary trace including the name that will be given to the domain. Draw the new boundary on the domain image. When the boundary is completed, the domain image will be repainted with only the portion interior to the boundary displayed. To return to the original domain boundary, select *Boundary* from the Reinitialize submenu of the Domain menu.

A domain boundary can also be created from an existing trace by pasting the trace from the trace clipboard onto the selected domain. Place the trace on the trace clipboard before selecting the domain by selecting the trace and copying it (Ctrl-C).

Domain Attributes

When a domain is selected, edit its attributes by selecting *Attributes...* from the Domain menu at the top of the main window (Fig. 7.5). These attributes can also be changed without selecting the domain by using the *Attributes...* item from the Modify menu at the top of the Domain List window.

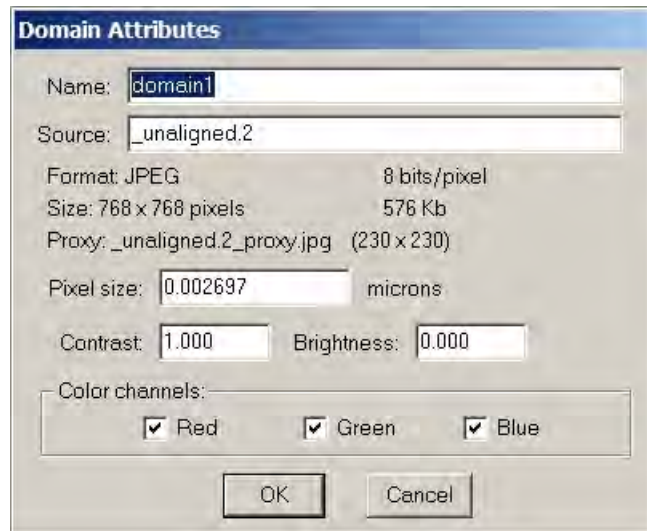


Figure 7.5. The Domain Attributes dialog. This dialog allows the attributes of a domain to be changed. By selecting Cancel no changes are made.

Name

The current domain name is shown in the Name box. A new name for the domain may be typed in the box. The rules for domain names are the same as for traces. Domain names are limited to a maximum of 63 characters. Illegal characters {, " = \ < >} are replaced with underscores. The special tokens {+ \$ \$+} are interpreted for domain names entered in the Attributes dialog. These tokens allow domain names to be automatically numbered just as for trace names, but the \$+ token is interpreted independently for domains and traces. See Chapter 8: Traces for details.

Source

The source path is a string of up to 255 characters that defines the location of the image file. This path can be an absolute path or a path relative to the series folder (Fig 7.4). When a new domain is created, by default the program stores the relative path from the series folder to the image in the section file. Thus, if the image file is copied into the series folder, the domain path to the image simply the file name. By checking the "Use absolute paths for new domains" option in the General tab of the Series Options, the absolute path including both the folder name and the file name will be stored as the domain path. This option would allow the series folder to be moved around while the image files remain in a fixed location.

Information about the source image file is also included in the attributes dialog. This information included the image file format, the number of bits per pixel, and the file size. The size is given both as width and height dimensions in pixels, and as the number of kilobytes of memory the image requires when decompressed. The Attributes dialog also reports the name of the proxy image file if one exists, and the dimension reduction factor of the proxy image. These values can only be changed in the Proxies tab of the Series Options (see below).

Pixel Size

The pixel size specifies the size of the image picture elements (pixels) in terms of the units of measure of the series. This can be changed for a domain by typing a new value in the Attributes dialog. The pixel size can also be changed by using the Pixel Size dialog from the Section List, or by calibrating traces. It is best to calibrate the pixel size at the time images are imported. For more information on the pixel size parameter and its calibration see Chapter 9: Calibration and Measurement.

Note: Changing the pixel size of a domain will effect the alignment of the domain image with traces and with other images.

Contrast and Brightness

The contrast and brightness of the image can be adjusted by typing in new values in the Domain Attributes dialog. The intensity of each pixel is adjusted by first multiplying it by the contrast value. A contrast value of 1 indicates no change. Then the brightness value is multiplied by 255 and added to the contrast-adjusted intensity.

The brightness and contrast of a domain can also be changed in the main window when a domain is selected. This is done with the hyphen, equal sign, and bracket keys as described in [Table 4.2](#). Pressing the hyphen key decreases the brightness, while the equal key increases the brightness. Contrast is increased by the right bracket (]) and decreased by the left bracket ([). To return to the original brightness and contrast values, select Contrast from the Domain Reinitialize menu.

The brightness and contrast can also be changed for all the domains contained in a range of sections by using the Brightness/Contrast dialog from the Section List. First, highlight the range of sections to be modified, then open the dialog. Enter the desired contrast and brightness values and select OK. It may be helpful to first determine appropriate values by use the key commands on a selected domain and then checking the final values in the Domain Attributes dialog.

Color Channels

The red, green, and blue color channels can be turned on or off for the display of individual domains by using the checkboxes at the bottom of the dialog. For example, a grayscale image can be displayed as red by unchecking the Green and Blue checkboxes. Make sure the boxes are completely empty to turn off the channels. A grayed check mark will keep the channel state unchanged.

Moving and Montaging Domain Images

Each domain in a section is associated with a transformation that places the image on the section at a particular location and orientation (Chapter 10: Movements and

Alignments). To change the position of an image on the section, pick up the Select Domain Image Tool and select the domain image using the left mouse button. The domain image will be displayed by itself inside its domain boundary. The position of the domain relative to the rest of the section can be seen by blending the domain with the section (spacebar). Translate the domain image by using the arrow keys or by dragging the domain with the left mouse button. When done moving the domain, merge the domain with the rest of the section by clicking the right mouse button.

Multiple domains can be created for an image by importing the image into the section multiple times. Up to 254 images may be arranged as separate domains within a single section to create a montage. Each image is imported as a separate domain (see above) and then moved into position with the other images by dragging it with the Select Domain Image Tool, by using keyboard commands or by using one of the other methods discussed in Chapter 10: Movements and Alignments. If the precise arrangement of sections is known in advance, e.g. each image is offset by 10 microns along the x-axis, use the Import Images dialog from the Series menu with particular X and Y offsets to place the domains in the correct locations on the section.

Merging a Domain

To return a selected domain to the section (i.e. to deselect it), use either the *Merge front* or *Merge rear* commands. Domains have an implicit order that determines which image will be displayed on top of (and occluding) other images. This order is reflected in the Domain List. The order of domains can be changed by selecting and merging the domains to either the front or rear of the domain order. *Merge front* puts the domain at the end of the domain list where it will be painted on top of all other domains. *Merge rear* puts the domain at the start of the domain list where it will be painted over by all other domains.

When the Select Domain Image Tool is being used, the *Merge front* operation can be carried out by clicking the right mouse button.

When to Merge a Domain

The appearance of either a folded section (Fig 7.6) or cracked section (Fig. 7.7) can be "fixed" in RECONSTRUCT™ by merging domains. This is done, as needed, to preserve a region of interest through serial sections. Start by creating domains for each piece where the fault exists and shifting the pieces to align with adjacent sections in the series. This is illustrated in more detail below.




Figure 7.6 An image of a folded serial section.




Figure 7.7 An image of a cracked serial section

Figures 7.7a to 7.7d illustrate key steps for fixing a cracked section. The objective is to “fix” the section by drawing domains on one or both sides of the crack (depending on the region of interest), and shifting the pieces back together in alignment with adjacent section(s).

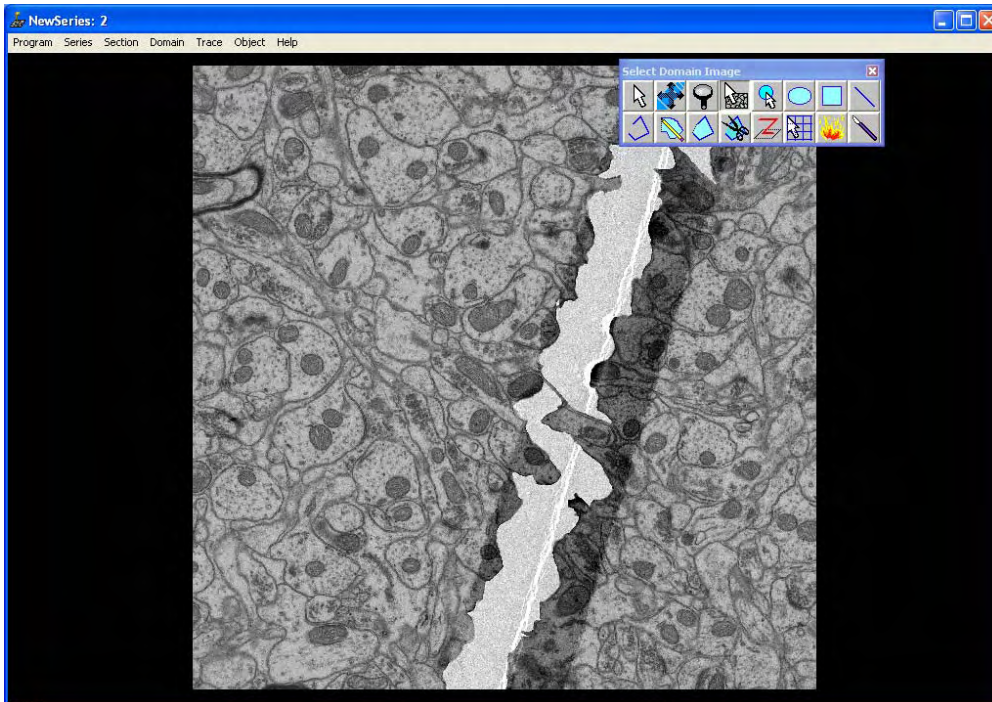


Figure 7.7a. Image of a cracked section.




Figure 7.7b. A separate domain is created from the region of interest to the left of the crack. The perimeter of the domain is illustrated in pink.




Figure 7.7c. An adjacent section is imported as a separate domain to facilitate image alignment. Alignment is realized by moving the first domain in relation to the second, adjacent section. The domain of the cracked image appears as an overlay, and movement is accomplished using the Up, Down, Left, Right, Arrows on the keyboard. The size of the movement increment is modified by holding the Ctrl or Shift key down while pressing a movement key. The Blend function is used to facilitate alignment (Chapter 10). Misalignment produces a blurry image as illustrated above.




Figure 7.7d. Proper alignment results in sharp, clear features where regions of interest are in cross-section.

To “fix” the appearance of a cracked section, import the section image (Fig. 7.7a) to create a new domain. Redraw the domain boundary around the section on one side of the crack (Fig. 7.7b). Align the domain with the section above or below it (Fig. 7.7c) in the series. Alignment can be accomplished by translating the domain by using keyboard commands (e.g. up and down arrows) or left mouse button. Blend the new domain with the section and shift it until the edges of the two pieces match (Fig. 7.7d). Merge the domain into the section. If desired, import the cracked image a second time and redraw the domain around the other side of the crack. Merge the second domain into the section to complete the process and align. Use similar methods to fix undesirable folds in a section.

Reinitializing a Domain

When a domain is selected, the Reinitialize submenu becomes accessible on the Domain menu. This submenu provides a way to reset all of the changes in contrast and brightness, the domain boundary, and in the positioning of an image.

Contrast

The *Contrast* item of the Reinitialize menu clears the contrast setting to 1 and the brightness setting to 0. These are the values when the domain image is first imported.

Boundary

Reinitializing the boundary sets it to a rectangle that borders the entire image. The last boundary can be recovered by *Undo* (Ctrl-Z) until the section is saved and removed from memory.

Transform

By reinitializing the transform, all positioning and alignment information will be cleared from the domain and it will be set to the origin of the section. The pixel size information will not be changed, but any magnification information stored in the transform will be removed as well. Ctrl-Z (*Undo*) can be used to recover the transform while the section is still in memory.

Copying Domains

When an image has severe distortion due to cracks or folds in the section, it may be useful to define two or more domains for the image that can be moved or aligned independently. To make this task easier, a command is provided in the Domain List to copy a domain image within a section. To copy a domain, open the Domain List and highlight the name of the domain to be copied, then select Copy from the Create menu. A domain identical to the highlighted domain will be added to the section. Since the domain has the same image and transform as the original domain, it will exactly overlay the original image so that the section does not appear different. You can select the domain and move it to verify that two domains are now overlaid in the image, but both domains should also appear in the domain list. By drawing a different domain boundary for each

copy of the domain image, you can define sub-regions of the image that can be moved independently.

Hiding Domains

A domain can be hidden from view even though it still exists in the section. This can be used to reveal images underneath a domain or to speed up section display by reducing the amount of image data displayed. Hidden domains appear in the Domain List with a red "H" icon. Thus, the Domain List can be used to hide and unhide domains.

All domains can be hidden at the time the section is first displayed. This allows all images to be removed from display, leaving only traces that can be displayed very quickly. Hide all domains with the checkbox in the General tab of the Series Options. To prevent a permanent change to the section, make sure the Autosave checkboxes for sections are unchecked. All domains can be unhidden with an option as well. If both checkboxes are checked simultaneously, all hidden domains will be revealed while all visible domains will be hidden. This feature can be used to quickly switch between two sets of domains for the series.

Proxies

Proxy images are smaller versions of domain images that can be quickly displayed for low magnification views of the section. When proxies are in use, the paging of sections at low magnification and the rendering of thumbnails can be faster because less image data needs to be read from the hard disk. Initially, no proxy images exist. You must explicitly create them.

Proxy images are created and deleted for all sections from the Proxies tab of the Series Options (Fig. 7.8). When the Proxies tab is accessed, there may be a pause with the hourglass displayed at the cursor. During this pause the program is examining all the domains in the series and determining how many domains have proxy images. When the examination is complete the information will be reported in the dialog. The use of proxies can be turned off by unchecking the "Use proxy images when displaying sections".

Warning - (Re)Create Proxies may be slow and it cannot be aborted.

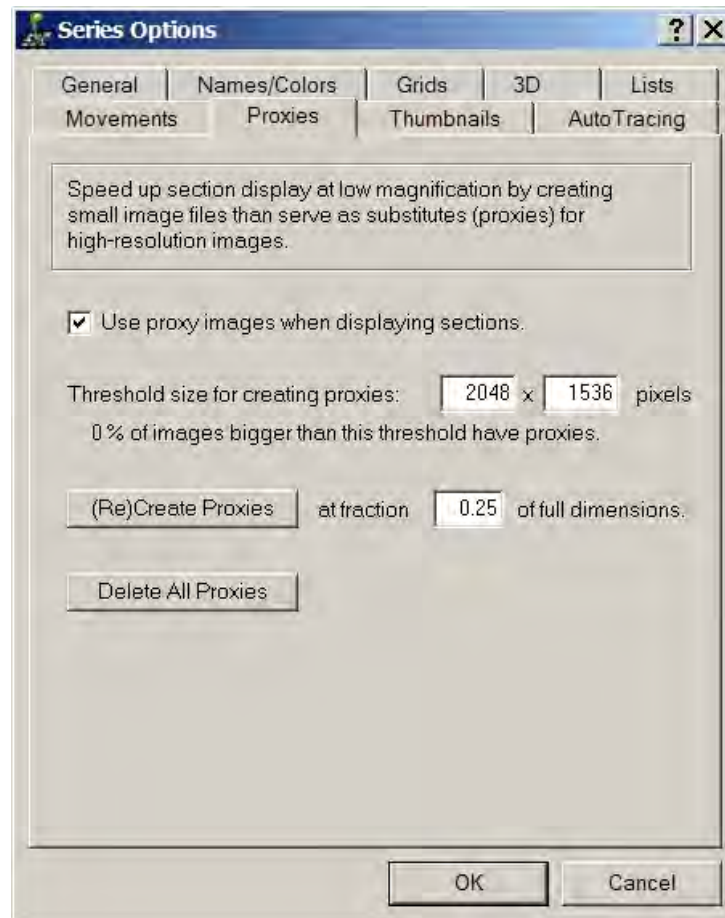


Figure 7.8. The Proxies tab of the Series Options dialog. The use of proxy images can be turned on or off by the checkbox. The proxy images can also be created or deleted by the buttons. Note that creating proxies may be very slow, and the program will be disabled until the operation is complete.

Proxy images can be created by pressing the "(Re)Create Proxies" button. A proxy image is then created for every domain image in the series for which the domain image dimensions equals or exceeds the threshold dimensions. A dimension of the proxy image is the fraction of the domain image dimension, as specified in the dialog. For a domain image 4000 pixels wide by 2000 pixels high, a fraction value of 0.25 would produce a proxy image of 1000x500 pixels, so that whenever the domain is displayed at a screen resolution of 1000x500 or smaller, the proxy image would be accessed and displayed instead of the full domain image.

For proxies to speed up section display you must carefully specify the fraction and threshold prior to creating the proxies. For example, suppose you want to work with the main window sized to full screen (1600x1200), and you want to display quickly the entire section (73220x4840) in the full screen window. First, calculate the fraction dimensions of the section relative to the full screen:

$$\text{width fraction} = \frac{\text{screen width}}{\text{section image(s) width}} = \frac{1600}{7322} = 0.2185$$

$$\text{height fraction} = \frac{\text{screen height}}{\text{section image(s) height}} = \frac{1200}{4840} = 0.2479$$

The fraction value in the Proxies tab must be set to a value higher than both the width and height fractions in order for proxies to be used at this magnification. For the example numbers given, a fraction of at least 0.28 is required for proxies images to be loaded during full section display. Without proxies, each 7322x4840 section would require a long time to load and display. After creating proxies at fraction 0.28, the image loading time would decrease by a factor of 4 or 5. But if you zoom in to the section, the display would slow back down because the whole image must loaded to display the data at the higher magnification.

In addition to section display in the main window, proxies are also used when sections are displayed on thumbnails or exported as images. If desired, you could set the fraction parameter to a much smaller value to increase the speed of thumbnail generation. For the 7322x4840 sections, suppose the full image load (without proxy) requires 800ms. Generating 100 thumbnails at a size of 128x96 might therefore require more than a minute without proxies. Creating proxies at 0.28 of full dimensions would allow the 100 thumbnails to be generated in 15s, and speed up the full screen display of sections as well. By making the proxy fraction even smaller, say 0.028, the thumbnails could be generated in just a second. However, at 0.028 our full screen display of sections would slow down again because these proxies are too small to be used.

Proxy Rule of Thumb for single domains

If your domain image is at least twice the screen resolution, use this formula:

$$\text{proxy fraction} = 1.2 \times \frac{\text{screen height}}{\text{domain height}}$$

If your domain image is not twice the screen size, set the threshold to less than the domain dimensions and set the proxy fraction for thumbnail display

$$\text{proxy fraction} = 1.2 \times \frac{\text{thumbnail height}}{\text{domain height}}$$



Chapter 8: Traces

Traces are the basis for quantitative measurements, calibrations, and alignments. A series must be calibrated in order to obtain accurate measurements. The Line Tool, described below, is useful for creating calibration lines and is described in detail in Chapter 9: Calibration and Measurement.

When traces outline the profiles of sectioned objects, these data elements are used for three-dimensional visualizations of object structure. Each trace is a polyline or polygon, represented as an ordered list of points in series units connected by straight line segments. Point order is clockwise for traces with a positive orientation. Clockwise traces have positive area and enclose the interior of the object. Counterclockwise traces have negative area and enclose holes.

Creating a Trace

Traces are created on sections by drawing with the mouse (or pen). Drawing when zoomed at high magnification allows for subpixel traces with respect to the image data. The resolution and density of trace segments is generally greater when drawing at higher magnification. When a trace is created, it is added to the set of *selected* traces on the current section. Selected traces are described in detail elsewhere in this chapter.

The *Tools* window (Fig 8.1) opens automatically when starting RECONSTRUCT™ or can be accessed by selecting Program > Windows > Tools window. Several drawing tools are provided in the *Tools* window. Each *Tools* window button indicates the type of trace to be drawn: Point-by-Point Drawing Tool, Freehand Drawing Tool, Closed Point-by-Point Drawing Tool, Ellipse Tool, Rectangle Tool, and Line Tool. Pressing a tool button places the description of the tool in the title bar of the Tools window. Moving into the main window changes the mouse pointer to the tool cursor. Most drawing tools also support a precision crosshair cursor (Ctrl-P).

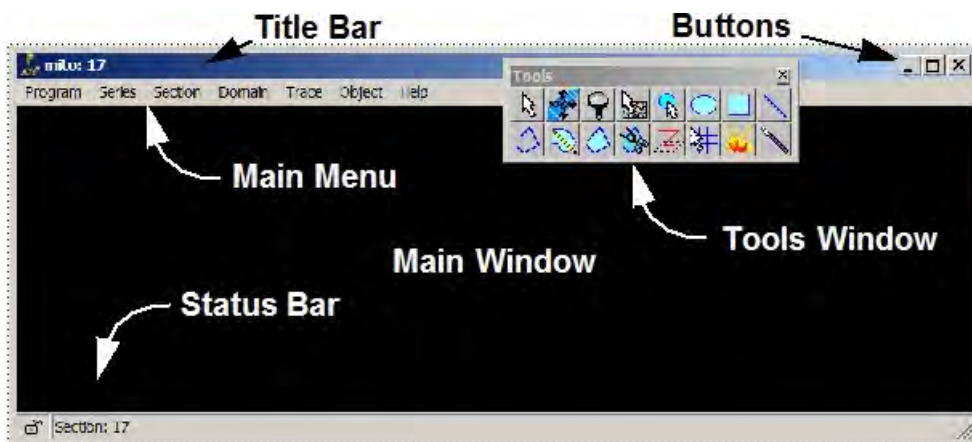


Figure 8.1. Initial window layout in RECONSTRUCT™ showing the floating Tools window.

Open Point-by-Point Drawing



This tool is used for drawing an open trace as a sequence of line segments. Select the tool and begin drawing in the main window by pressing and releasing the left mouse button. A candidate line segment is shown as the mouse is moved. When the line segment is as desired, click the left mouse button again to add the segment to the trace. Segments are sequentially added in this manner until the right mouse button is pressed to complete the trace.

During the entering of segments, preceding segments can be removed by using the Backspace key. The starting point cannot be removed using Backspace. Use the Esc key to abandon the trace and start over.

It is possible to change sections while drawing point-by-point. The trace will be created in the section in which the trace began, even if points were added (or the trace completed) while other sections were displayed. This is true of both open and closed point-by-point drawing and the Line Tool. To create a trace that actually spans multiple sections, use the Z-Trace Tool (see below).

Closed Point-by-Point Drawing



The closed Point-by-Point Tracing Tool works in the same way as the open one, but the candidate completed trace is shown along with any applicable fill color. To avoid the fill color while tracing, select None for the fill mode in the default attributes. This attribute can be easily changed after tracing.

To accept the candidate trace press the right mouse button. Use Backspace to remove segments and Esc to abandon the trace altogether. If automatic simplification is enabled, the final trace is given a positive orientation and all internal loops are removed. Thus, a trace with a single closed path is always generated.

Freehand Drawing



The Freehand Drawing Tool has a cursor that looks like a pencil. The tool is used much like a pencil in that a single drag operation by the mouse generates the closed trace. Pressing and holding the left mouse button starts the trace. The path of the trace follows the mouse movement. The drawing operation can be aborted by using the Esc key.

When the mouse button is released the trace is closed and simplified (if enabled) to produce a single, clockwise loop. It is possible to change sections while drawing with the pencil. The trace will be created only on the section in which the trace is terminated, even if points were added while other sections were displayed. This is similar to the behavior of other tools that involve mouse dragging such as the Ellipse and Rectangle Tools.

Like the Stamp and Grid Tools, the section can be panned by dragging the right mouse button near the edge of the main window when the Pencil Tool is used.

Ellipse Tool



The Ellipse Tool is used to create circles and ellipses quickly. Press and hold the left mouse button while moving the mouse to define the shape. Release the

mouse button to create the trace. To abort the ellipse without creating a trace hit the Esc key before releasing the mouse button.

Rectangle Tool



The Rectangle Tool is used to generate a closed rectangular trace. Drag (press and hold) the left mouse button to define the rectangle. Release the mouse button to create the trace. To abort without creating a trace, use the Esc key.

Line Tool



The Line Tool is used to create a trace consisting of a single line segment. This tool is useful for creating calibration lines as described in Chapter 9: Calibration and Measurement. Press and release the left mouse button to start the segment. A candidate line will then follow the cursor until the left mouse button is clicked again to create the line trace. To abort without creating a trace, use the Esc key.

Stamping Predefined Shapes



The Stamp Tool allows creation of traces with a single mouse click. The predefined shape is centered at the position of the Stamp Tool cursor. The shape is drawn at a fixed pixel size such that the shape is a constant size with respect to the screen. This means that stamping at high magnification will produce a much smaller trace in series units, than stamping at lower magnification. Shapes with a predefined fixed size in series units can be created using the Grid Tool.

The right mouse button allows panning the section while using the Stamp Tool. Press and hold the right button and drag the cursor to the edge of the main window. The section will scroll as long as the cursor is near the edge and the right button is held down.

Stamp shapes can be selected from the Trace Palette by selecting one of the palette buttons or equivalently by using the number keys. A stamp shape can also be obtained from the traces in the main window by clicking the left mouse button while holding down the Ctrl key on the keyboard. The trace closest to the cursor point replaces the current stamp shape. This shape can replace the shape in a palette entry by using the Ctrl-number key sequence to update the palette. Stamp shapes can be either open or closed depending on how the shape was originally drawn or how its attributes were set prior to selecting it. To stamp shapes consisting of more than one trace use the Grid Tool.

Creating a Grid of Traces



The Grid Tool is used to create a repeating pattern of traces by placing multiple copies of a set of traces a certain distance apart. The Grid Tool is designed to generate grids primarily for facilitating stereological measurements on sections, but it can also be used to create lines, rectangles, and circles of a specific size. The left mouse button creates the grid on the section. The right mouse button can be used to scroll the section by dragging at the edge of the main window.

The type of grid generated by the Grid Tool is determined by the settings in the Grids tab of the Series Options (Fig. 8.2). Five predefined grid types consist of rectangles, crossed lines, ellipses, sampling frames, and cycloids (Fig. 8.2). The number of elements in the grid is specified in the x- and y-directions, along with the size of each element and the distance between elements. For rectangular and sampling frame grids the left mouse

button places the upper left corner of the first grid element. For cycloidal grids the left mouse button places the left-most point of the first grid cycle. For the remaining grid types a left mouse button press places the center of the first grid element. Additional grid elements are added to the left (x-direction) and below (y-direction) the first element.

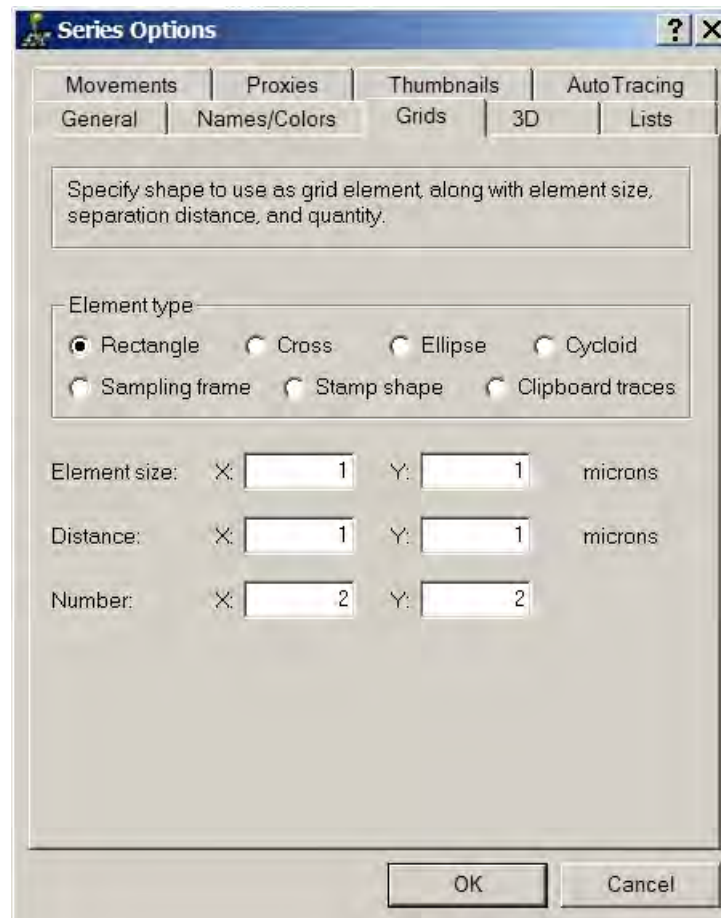


Figure 8.2. The Grids tab of the Series Options dialog. The Element type determines the shape to be drawn (as shown in Fig. 8.3). The shape is sized in the x- and y-directions according to the Element size boxes. Shapes are repeated Number times and at Distance intervals. The settings are used only when OK is selected at the bottom of the tab.

Arbitrary grid elements can be defined using one of the two remaining element types (Stamp shape or Clipboard trace). When the stamp shape type is selected, the Grid Tool centers the current stamp shape at the grid positions. The stamp shape can be selected from the Trace Palette or by using the Stamp Tool to pickup a trace from the section. The stamp shape is scaled to the Element size parameters in the Grids tab. This allows the Grid Tool to be used to stamp shapes with a fixed size, unlike the Stamp Tool which draws in pixels. When the Element size parameters are zero or negative, the stamp shape is drawn in pixels just as with the Stamp Tool.

A group of traces can be used as the grid element by using the Clipboard trace(s) option. When this option is selected in the Grids tab, the Grid Tool copies the contents of the Traces Clipboard to the grid. The group of traces is centered at the cursor using the midpoint of the x- and y-extent of the set of clipboard traces. The group is then repeated

at each point specified by the Distance and Number parameters. By selecting the grid Number to be 1, only 1 copy of the clipboard traces will be produced and the Grid Tool acts like a Stamp Tool for multiple traces.

Grid elements created with the Grid Tool are given the name and color attributes of the default tracing parameters. Each grid element is given the default trace name, allowing for interpretation of any special characters. Thus, using the default name "grid\$+" will result in elements being named "grid1", "grid2", "grid3" and so on. The sampling frame element is given a fixed color pattern (Fig. 8.3) where exclusion lines are colored red and inclusion lines are colored green.

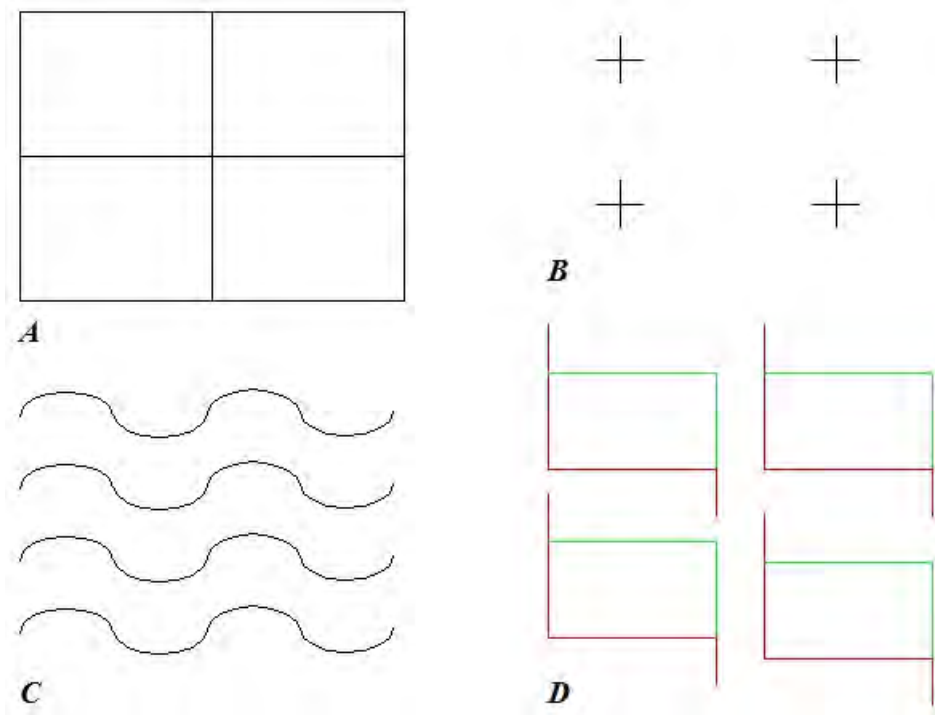


Figure 8.3. Some types of predefined grids. **A.** A 2x2 grid of rectangles. **B.** A 2x2 grid of cross points. **C.** A 4x4 grid of cycloids. **D.** A 2x2 grid of sampling frames.

The Clipboard trace(s) grid element does not use the default tracing name or colors. Rather the grid elements replicate the clipboard attributes for each grid element. Also unlike the other grid types, the Clipboard traces are not resized using the Element size parameters. The Element size parameters are simply ignored and Clipboard traces are reproduced at their actual sizes.

Semi-Automatic Tracing

RECONSTRUCT™ has a semi-automatic tracing tool called Wild Fire.



The Wildfire Tool has a cursor that looks like a match. It is not recommended to use the Wildfire Tool for most transmission electron microscopy images. Traces often go astray due to the complexity of electron microscopy images. Wildfire may be helpful if you have other types of images (e.g. fluorescence, light microscopy or other thresholded images). Please see a full description of this tool at the end of this chapter.

Z-trace as a 3-D Ruler



The Z-Trace Tool resembles the letter Z on the Tools window. Z-trace is a 3-D ruler and can be used to measure lengths across serial sections. The tool is used just like any of the other Point-by-Point Tracing Tools, but each point can be entered on a different section. These sections are stored along with the x and y positions of the trace points so that 3D measurements of shapes can be generated from the Z-trace data.

Z-traces are not stored or displayed on sections. They are stored in the series file and are listed in the Z-Trace List opened from the Object menu. Pressing the left mouse button starts the trace. A drag line follows the cursor to the next point. Page up or down to change sections and add trace points with left mouse clicks on the desired sections. Clicking the right mouse button ends the Z-trace at the last point entered with the left mouse button. The creation of a Z-trace can be aborted by using the Esc key.

- Z-traces are properly updated when section numbers are shifted but not when the order of section numbering is reversed.
- Z-traces remain valid when section thicknesses are modified.
- Z-traces are not automatically adjusted when the section image domains are transformed, calibrated, or re-scaled. Therefore, Z-traces should be created only after the order of sections, and the alignment and calibration have been finalized.

The position and structure of a Z-trace can be visualized by adding it to the 3D Scene. Note that if the 3D Scene Window is open when a Z-trace is added, it may be out of the field of view and so will not appear in the window. The 3D Scene may need to be rotated or reset in order to see the new Z-trace. Another important feature is that only one object in the same series can have the same name. Creating a Z-trace with the same name as an existing object will preclude them both being visualized together in the same scene. Consequently, Z-traces should be uniquely named. The special character + can be used in the default name when repeatedly drawing Z-traces to create unique names.

Note: Z-traces should be uniquely named.

Z-traces and Objects can be displayed together in a 3D scene but to see where Z-trace points are on the sections, one must convert Z-traces to Objects by using "Create Grid at Z-trace points" from the Z-trace list. First, setup the desired Grid shape in the Grids tab from the Series Option menu, and the desired object name in the Names/Colors tab also on the Series Option menu. It is also possible to create the Z-trace from a traced object, rather than drawing the Z-trace directly, but this requires that your object is unbranched and with convex traces.

Setting Trace Attributes for Drawing

When a trace is drawn, it is given the current default trace attributes. These are set in the Names/Colors tab of the Series Options (Fig. 8.4). The default tracing attributes are saved with the series so that tracing can easily be continued when a series is reopened.

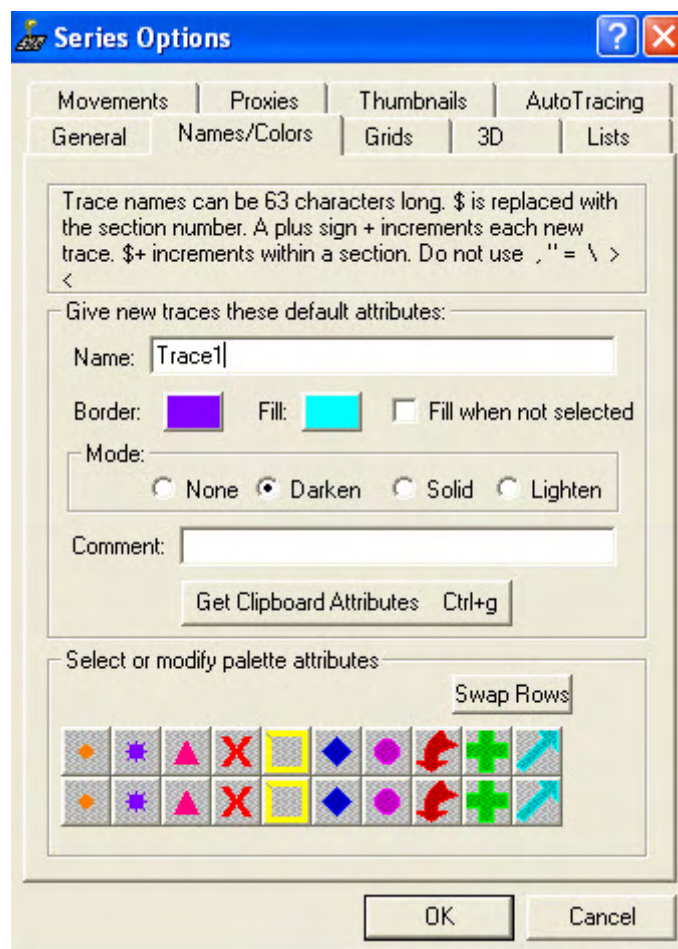


Figure 8.4. The Names/Colors tab of the Series Options dialog. The settings in the "Give new traces these default attributes" box determine the name, color, and comment that will be used to label traces when they are drawn. The attributes of a palette entry can also be set by pressing one of the palette buttons and then modifying the "Attributes of new traces" settings.

The attributes in the Names/Colors tab are essentially the same as those in the Attributes dialog from the Trace menu described in Chapter 3. Absent from the default attributes are Hidden, Simplified, and Closed checkboxes. Drawn traces are always initially unhidden. The closed and simplified attributes of new traces are determined by the drawing tool used and the status of automatic simplification.

The default trace attributes can be filled from the first entry on the trace clipboard by selecting the "Get Clipboard" button. This provides a way to copy the attributes of an existing trace and use them for subsequent drawing of new traces. This also provides a mechanism for copying attributes of existing traces to the Trace Palette (see below).

The Trace Palette

A set of predefined default trace attributes can be quickly accessed using *Palette...* from the Trace menu. The Trace Palette window has 20 user-definable palette entries. A new series begins with user entries set to predefined values (Fig. 8.5).



Figure 8.5. The Trace Palette window. The top row of buttons displays the user defined palette entries, while the bottom row displays the predefined palette entries. Each button shows the default border and fill colors, fill style, and Stamp Tool trace shape.

To use the trace palette, simply select a button in the palette window using the mouse. The default trace attributes will be immediately set to the name, colors, and fill mode defined for that palette button. Holding the mouse over a button when the palette window is active shows the name attribute in the title bar. The colors and fill styles are rendered on the buttons. The shape on the palette button is the shape of the stamp trace that will be adopted whenever the Stamp Tool is used (see above).

Define the trace palette entries with the Names/Colors tab of the Series Options (Fig. 8.4). Each user-definable palette button is shown in this tab. Select one and edit the attributes in the Attributes box. The palette entry attributes are changed immediately; an OK in the Series Options dialog is not required. To copy the attributes from an existing trace to the palette, copy the trace to the trace clipboard and then select the palette entry in the Names/Colors tab. Use the "Get Clipboard" button to copy the clipboard attributes immediately to the palette entry.

The palette entries can also be selected and set using the keyboard (Table 4.2), even when the Trace Palette window is not open. Pressing the number keys 0 through 9 when the main window has the keyboard focus, sets the default trace attributes for drawing from the corresponding palette entry of the first ten entries. The Ctrl key in combination with the number key sets the current default drawing attributes onto the corresponding palette entry. Although only the first ten palette entries can be accessed using the keyboard, the last ten entries can be switched with the first ten by using the "Swap Rows" button in the Names/Colors tab of the Series Options. This also allows stamp shapes to be set (as described above) on the last ten palette entries.

When the palette name contains automatic numbering it does not get updated like the default trace name does during tracing. Use the Ctrl-number key command to update the palette after tracing. For example, if the 0th palette entry name is set to d00+, then selected and used for drawing 4 traces, the default trace name would be incremented (d01+, d02+, d03+, d04+) but the palette entry would remain d00+. Pressing Ctrl-0 would update the palette entry to d04+.

The palette is mainly for quickly changing the default attributes when drawing new traces. However it can be used to modify existing traces as well. To give the attributes of a palette entry to an existing trace, select the trace in the main window and select the palette entry from the Trace Palette. Then open the Trace Attributes dialog (Ctrl-A) and use the Get Defaults button to fill the dialog with the attributes from the palette entry.

About Trace Names

The name of a trace identifies the object to which the trace belongs. Traces with the same name are grouped together to make 3D objects (Chapter 11: Objects). Trace names are also used to make correspondences between sections for alignments (Chapter 10: Movements and Alignments).

Trace names are limited to a maximum of 63 characters. Certain characters are not allowed in trace names. Illegal characters include the quotation marks ("), the equal sign (=), the comma (,), the backslash (\), and the less than (<) and greater than (>) signs. These characters entered in a trace name will be replaced with underscores (_). A few other characters have special interpretations. These special characters include the asterisk (*), the question mark (?), the pound sign (#), the plus sign (+), and the dollar sign (\$). The dollar and plus signs can be used anywhere in the name string to automatically number trace names.

The dollar sign in a trace name string is replaced with the current section number when a trace is created. For example, the name string "s\$" would result in the string "s101" for all traces created on section 101 and "s45" for all traces created on section 45.

A plus sign in a trace name results in the immediately preceding integer being incremented. The number is incremented every time a new trace is drawn. For example, the default name string "vesicle+" would result in "vesicle1" for the first trace, "vesicle2" for the second trace, and so on. Multiple plus signs can be used, but plus signs occurring together will be reduced to a single one. Initialize the increment at a particular value by preceding the plus sign with an integer. For example, the default name string "my24+th_vesicle" would result in "my25th_vesicle" for the next trace created, "my26th_vesicle" for the next, and so on.

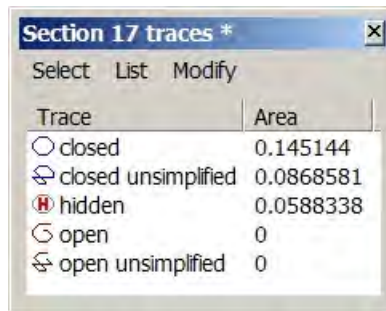
When the dollar and plus signs are used together (\$+) incrementing is restricted to individual sections. Each new trace created is given a number higher than all the other traces in the section that have the same leading characters. Unlike the previous special characters, the section is searched for similar trace names every time a trace is created. So while a plus sign can generate duplicate trace names within a section, dollar-plus names tend to be unique. For example, a default name string "vesicle\$+" will generate "vesicle1", "vesicle2", "vesicle3" on section 45, and then on section 101 will generate "vesicle1", "vesicle2". Returning to section 45 would then generate "vesicle4", "vesicle5", and so on.

The Trace List

Traces can also be selected by name using the Trace List. The Trace List window is opened by selecting *List traces...* from the Trace menu. Every trace in the section that matches the limit string specified in the Lists tab of the Series Options ([Fig. 5.2](#)). When new traces are drawn, the list will be updated only when it is regenerated, either by refreshing it, opening it again, or by paging sections.

Each trace is listed by name in alphabetical order along with an icon that indicates some attributes of the trace ([Fig. 8.6](#)). A closed blue figure indicates that the trace is closed, while an unclosed red line indicates that the trace is open. Traces that are not simplified may contain loops (see below), so unsimplified traces are symbolized by a figure-eight shape. Traces that are hidden from view are also listed with a special "H" icon.

In addition to the trace name and icon, the list may contain several columns of additional data about each trace. These Trace List data columns are selected in the Lists tab of the Series Options. In addition to the area ([Fig. 8.6](#)), the Trace List can include comment, length, centroid, extent, section thickness, and z distance. For more on these values see [Chapter 9](#). The entire list can be saved to file using the *Save...* command from the List menu at the top of the list window.



Trace	Area
<input type="radio"/> closed	0.145144
<input type="radio"/> closed unsimplified	0.0868581
<input checked="" type="radio"/> hidden	0.0588338
<input type="radio"/> open	0
<input type="radio"/> open unsimplified	0

Figure 8.6. The Trace List window. This example list includes each of the possible list icons and trace names that correspond to the attributes for these icons. In this list only the Area data column was included. Note that open traces have zero enclosed area. The "*" in the title bar is the default limit string which effectively displays every trace in the section.

Multiple traces can be selected from the Trace List by highlighting trace names with the mouse and then pressing Enter or the *Select* menu item at the top of the list. When traces are selected from the list, any traces previously selected in the section are deselected, the list traces are selected (and unhidden if necessary), and control is returned to the main window.

Editing Traces



Existing traces can be edited using the tool that looks like scissors. The Scissors Tool is used to break a trace at the cursor point and start redrawing it using the point-by-point method. Backspace can be used to undo a portion of the trace after breaking it with the scissors. If the trace is clockwise (as most are) then be sure to select a point further clockwise than the part of the trace to be redrawn. This will ensure that backspace (which is usually counterclockwise) gets rid of the unwanted part of the trace.

When the left mouse button is pressed the segment of the trace nearest the cursor is removed and the tool is immediately switched to a Point-by-Point Drawing Tool. If an open trace is edited the portion of the trace after the break point is removed before redrawing. For a closed trace, only one trace point is removed by the Scissors Tool. The subsequent drawing operation can be aborted by using the Esc key, but this will completely delete the trace being edited. Use *Undo* (Ctrl-Z) to recover the original trace.



Existing traces can also be quickly sliced into separate pieces using the Scalpel Tool (Fig. 8.7). The tool is used to divide a closed trace into two or more pieces based on a separation line drawn with the cursor. To divide a closed trace, place the cursor outside of the trace near the point where the trace is to be divided. Press and hold the left mouse button and drag the cursor through the interior of the trace. A temporary line will show the intended path of the division. When the cursor reaches the outside of the trace again, release the left mouse button and the trace will be divided along the path. If the division path is wrong, press the Esc key before releasing the mouse button to abort the operation. After a successful division, the middle mouse button can be used to select one of the pieces for further processing such as deletion or attribute changes.




Figure 8.7. The process of dividing a trace using the Scalpel Tool. A. A left mouse drag is executed along the line shown. B. When the left mouse button is released the trace is divided into three pieces along the line. C. The middle piece is selected by using the middle mouse button of the Scalpel Tool. D. The final product after deleting the unwanted piece with the Delete key.

Note that the division of a trace may not always succeed. Carefully observe the result to make sure it is as intended. The Scalpel Tool only works as described when the trace has a clockwise (positive) orientation. On a counterclockwise trace the Scalpel Tool will add a piece to the trace rather than divide it. The left mouse button press must occur near the trace to be divided and not closer to another trace. The trace closest to the starting point is the one that will be divided. The division path should be perpendicular to the trace at the points where it crosses it. Under some circumstances the Scalpel Tool will fail to find a suitable division and will return fewer pieces than expected. This is usually due to the fact that the division line is parallel to the trace at the point where it crosses.

The Scalpel Tool is often useful for quickly removing an external evagination of a trace, but it does not allow an internal invagination to be sliced off. To remove an internal invagination, select the trace and use the Line Tool to make a new trace (with the same trace name) that would produce the desired cut. Then use *Merge* from the trace menu to merge the line trace and original trace, effectively removing the invagination.

Merge

The *Merge* operation attempts to combine all selected traces into one closed trace. Traces to be merged should be closed and should overlap. If some traces do not overlap the effect of merge will depend on the ordering of the selected traces. Some non-overlapping traces may be deleted while overlapping ones are merged. Another possibility is that none of the traces will be changed.

Reverse

One can verify that a trace has a clockwise or counterclockwise orientation by checking the area measurement in the Trace List. Clockwise traces have positive areas whereas counterclockwise traces have negative areas. Typically traces are generated as clockwise boundaries of the exterior surface of an object. Note a closed trace with a negative value assigned to it (i.e. was originally traced counterclockwise) as an interior boundary will surface as a hole. *Reverse* changes the orientation of selected traces. Apply the *Reverse* command from the Trace menu to reverse the order of a trace and change its characteristics. Multiple traces can be changed this way from the Objects List to avoid editing section by section.

Simplify

The *Simplify* command is used to reduce the complexity of traces. *Simplify* provides on-demand simplification of selected traces provided the Simplified attribute of the traces is false (unchecked). If a trace is already simplified, then the *Simplify* command has no effect. If no traces are selected that have the Simplified attribute unchecked, then the *Simplify* operation will generally report "Nothing to simplify". To enable further simplification on an already simplified trace, use the Trace Attributes dialog to first uncheck the Simplified attribute.

Simplify applies a simplification process similar to that normally used when completing a drawing operation, (see the Simplified attribute section below). The trace is re-pixelized and then reduced to a set of line segments that match this path. If the trace is closed, all loops are removed and the trace is reoriented to a clockwise direction.

The resolution of the re-pixelized line segments is determined by the resolution parameter for "Simplify on demand" in the AutoTracing tab of the Series Options window (Fig. 8.8). *Simplify* uses the resolution set in the Series Options menu and Autotracing tab (Fig.8.8), while simplification during drawing is done relative to the current screen pixel size. If the resolution parameter is negative or zero, then a resolution equal to 1% of the largest extent of the trace in the x- or y-direction is used.

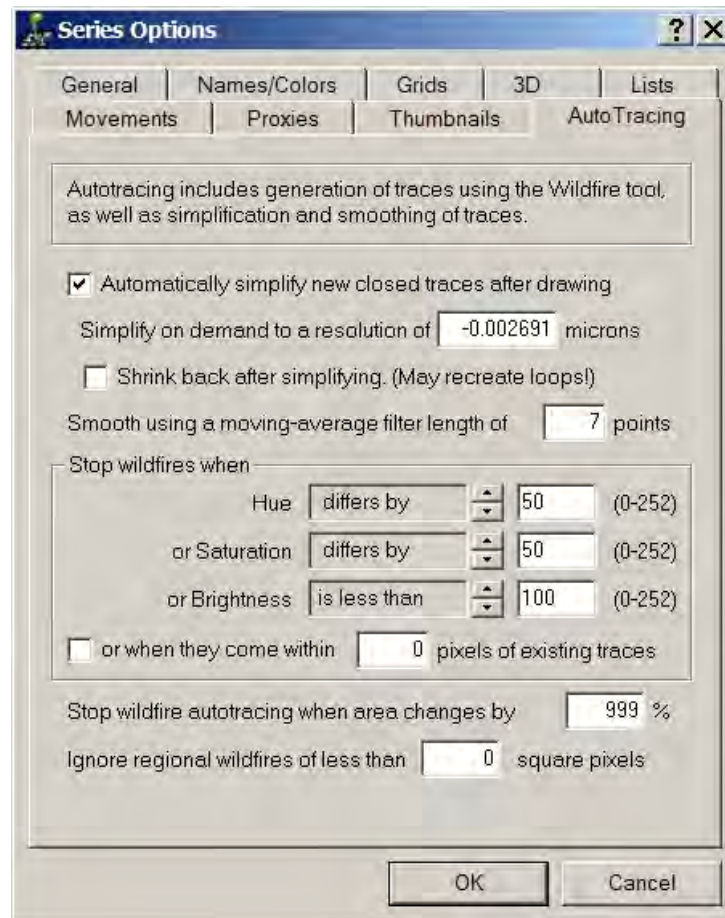


Figure 8.8. The AutoTracing tab of the Series Options dialog. Automatic simplification can be disabled by unchecking the option, but this effect is only temporary. Whenever a series is reopened Automatic simplification will be turned back on. This is because simplification is important for eliminating loops that reduce the accuracy of measurements and prevent proper 3D surfacing of the traces. The Stop Wildfires box specifies the Hue, Saturation, and Brightness (HSB) criteria for halting region growing. Valid HSB are between 0 and 252. By setting a criterion to "differs by 255", that criterion is disabled and will not affect region growing. Default settings are shown.

Simplify will normally expand the entire perimeter of a trace by a distance equal to the resolution (e.g. 1 screen pixel size). Expansion of the trace is caused by the algorithm that is applied to the trace to eliminate loops. This is because the simplification process ensures a clockwise orientation without loops by tracing around the outside of the existing trace.

Tips when using *Simplify*:

- Trace the inner surface of your region of interest. For example, trace along the inner sheet of a lipid bilayer (cell membrane) so that the expansion of 1 screen pixel will have a minimal effect on the trace if you choose to *Simplify*.
- Avoid using *Simplify* repeatedly on a single trace as this may cause an undesirable trace expansion for your region of interest

- Draw or edit a trace when fully zoomed on it so that the expansion of 1 screen pixel will have minimal effect on the trace.

Trace expansion can be reduced by checking the box next to "Shrink back after simplifying" in the AutoTracing tab. When this box is checked the clockwise outer expansion is followed by an inner clockwise traversal around the interior of the trace. Shrink back should produce a reasonable result for convex traces, but traces that are highly concave with narrow interior regions exhibit loops after shrink back. For this reason, the shrink back command should be used sparingly.

Simplifying traces using a large, positive resolution value results in traces with fewer trace points. A smaller resolution value will produce many more points. If the trace is large and the resolution is very small, so many points may be generated that the response to a *Simplify* operation will become painfully slow. To avoid slow simplifications, make sure the resolution parameter is large enough. If simplification is allowed to be performed automatically (by checking this box in the Series Options), the resolution will be chosen appropriately for the zoom factor, allowing traces to be quickly simplified to a reasonable fidelity.

For simple shapes such as lines and rectangles, the simplest possible representation is automatically generated. For this reason these shapes are marked as Simplified as soon as they are created. Unchecking the Simplified state of these traces and using the *Simplify* menu item will actually add trace points to these traces and may significantly change their shape. For most other trace shapes some reduction in trace points without severe shape change may be possible provided the simplify on-demand resolution is appropriately chosen.

Traces created by open point-by-point drawing are not automatically simplified on completion. The trace may consequently contain loops that will adversely affect the interpretation of the trace. These loops can be eliminated by closing the trace and using the *Simplify* menu item. Drawing an open point-by-point trace, closing it by changing its attributes, and then using *Simplify*, is essentially the same as using the closed Point-by-Point Tracing Tool in the first place. The main difference is that *Simplify* will use a resolution set by the user in the Series Options, while simplification during drawing (see below) is done relative to the current screen pixel size. Note that the simplification process applied to an open point-by-point trace will very likely add points so that the trace more closely matches the option-specified resolution.

Smooth

The smoothing operation applies a moving-average filter to the locations of the trace points. This effectively smoothes the transition between points, and converts a polygonal shape with corners into a more rounded profile. The magnitude of the effect depends on the density of trace points and the length of the moving-average filter, i.e. how many adjacent points are used to adjust the position of each point on the trace. When the filter is long the trace will be smoothed more than when the filter is short. The effect also depends on the density of trace points. For a given filter length, sparsely distributed trace points, as generated through point-by-point drawing, will be more drastically modified than dense numbers of trace points generated by region growing.

The length of the moving average filter is specified in the Autotracing tab of the Series Options. The length is the total number of points around a given trace point (including that point) that are averaged together to get the new trace point location. For example, a filter length of 9 will average the 4 preceding points, the current point, and the 4

succeeding points along the trace to calculate a new position for the current point. For simple traces, such as squares the effect will generally be undesirable no matter what filter length is chosen. This is because the smoothing operation does not add or subtract points to make the trace smoother. It simply shifts the existing trace points to new locations. Note that this fact means that the smoothing operation also has the potential for creating loops in the trace. It is recommended that the Simplify command be (re)applied immediately after smoothing.

Modifying Trace Attributes

Every trace has a set of attributes that specify its name, color, fill style, and whether it is closed, hidden, or simplified. Newly created traces are given the current default attributes (see below). Trace attributes can be modified at any time by selecting the trace(s) and then opening the Attributes dialog from the Trace menu (Fig. 8.9). The mouse or keyboard can be used to access the fields of the dialog. Use the Tab key to move through the fields. For the Fill mode box use the arrows to highlight a mode and the space bar to select. If any fields are grayed (as shown for the Fill mode items in Fig. 8.9) then the current state of these items has mixed values among the selected traces. If grayed fields are not changed, there will be no change to these attributes of the selected traces.



Figure 8.9. The Trace Attributes dialog. The title bar indicates how many traces were selected when the dialog was invoked. All of these traces will receive the changes made in the dialog. Grayed checkboxes or buttons (e.g. the Hidden checkbox) indicate a don't care condition. As long as these boxes are not modified these parameters will not be changed on any of the selected traces.

The "Get Default" and "Get Clipboard" buttons can be used to fill the attribute fields with the values of the default attributes or the attributes of traces placed on the trace clipboard. Alternatively, the individual fields can be filled as described in the following.

Trace Name

The current trace name is shown in the Name box. A new name for the trace may be typed in the box. If more than one trace was selected then all the traces will be given the new name. The name in the Trace Attributes dialog only modifies the selected traces. When new traces are drawn they use the default trace attributes for the series which are set in the Names/Colors tab of the Series Options or with the Trace Palette (see above).

If the selected traces had different names when the Attributes dialog was opened, then an asterisk will be displayed in the Name box. Only when the asterisk is replaced with a new string will the names of the traces be changed. This allows the colors or other attributes of a set of traces to be modified at once even though the traces have different names. The asterisk notation can also be used to modify a set of trace names. Any text after the asterisk is appended to all the existing trace names.

For example, suppose traces named "d1", "d2", and "d3" were selected prior to invoking the Attributes dialog. The Name box would have an asterisk inside because the names are different. Changing the Name box string to "*spine" would change the trace names to "d1spine", "d2spine", and "d3spine". Leaving the Name box string as "*", would leave the trace names as "d1", "d2", and "d3". Replacing the asterisk with "spine" would result in all traces being named "spine".

Border Color

The border color is the color of the line segments that are drawn to represent the trace on-screen. Every trace has a visible border with a color. The Border Color button indicates the current color of the selected trace. If this button is gray with a diagonal line through it, then there is more than one color for the borders of the selected traces.

The border color can be changed by pressing the Border Color button. This opens the Color dialog shown in Fig. 8.10. The Color dialog allows a new color to be chosen from the Basic colors, Custom colors, or created from the color palette. Initially the Custom colors are all black. These can be modified by using the Add to Custom Colors button. The set of custom trace border colors is saved with the series. These values can also be modified in the Tracing tab of the Series Options dialog.




Figure 8.10. The Color dialog. This dialog is used for setting trace border and fill colors. Colors added to the Custom colors will be saved with the series.

Fill Color

The second color button in the Attributes dialog shows the current fill color if one consistent color is defined among all traces. Clicking on the button will open the Color dialog with a separate set of Custom colors for fill colors. These values are also saved with the series and can be set as well in the Tracing tab of the Series Options.

Fill color defines the shading for the interior of a closed trace. An open trace is not filled. The fill color can be any red-green-blue combination but the appearance of the fill is affected by both the choice of color and the fill mode.

Fill Style

The style of filling can be modified with the checkboxes in the dialog. A closed trace can be either filled when it is selected, when it is not selected, or not filled at all. An open trace is never filled.

Normally, the interior of traces are filled when the trace is selected, and unfilled when the traces are not selected. To turn off the filling of traces completely, change the Fill mode to be None.

The selected traces are filled when they are not selected by checking that box in the Attributes dialog. The trace will be given a transparent interior when they are selected. This is useful for showing particular regions as highlighted on the sections or in the thumbnails.

The Fill mode determines how the fill color will be combined with the underlying image or other trace data to produce the final display. A Fill mode of Solid will color the interior with a solid color, no image data will be visible. A Fill mode of None means that no filling will be applied to trace regions.

Darken fill mode sets the display to a combination of the colors common to both the fill color and the underlying image, such that a fill color of white would result in a completely transparent fill while a black color would result in a solid black fill. A pure blue (0,0,255) fill on a grayscale image would result in shades of blue replacing the grayscale. Lighten uses a combination of the fill color and the image color such that the opposite effect is produced. Black fill results in transparency, while white fill color produces a solid white interior. The basic colors that have either 0 or 255 for the red, green, and blue channels work best for fill modes using Lighten or Darken.

Hidden

A trace can be hidden from view in a section by checking the Hidden checkbox in the Attributes dialog. A hidden trace will not be drawn in the main window but will appear in the trace list with a special "H" icon (Fig. 8.6). To unhide a trace simply select it from the Trace List, or use *Unhide* from the list Modify menu.

All traces can be hidden at once whenever a section is first displayed by using the checkbox in the General tab of the Series Options. To prevent a permanent change to the section, make sure the Autosave checkboxes for sections are unchecked as well. Subsequently paging through the series will result in sections being displayed without any visible traces. Go back to the General tab and uncheck the "Hide all traces when loading section" box. Subsequent paging will display the traces again. If the traces remain hidden,

then the sections were saved with Hidden attribute set on all traces. To unhide them check the "Unhide all traces when loading section" box in the Series Options. If both hide and unhide checkboxes are checked simultaneously, all hidden traces will be revealed while all visible traces will be hidden. This feature can be used to quickly switch between two sets of traces for the series.

Simplified

Trace simplification occurs automatically after traces are drawn, except with Open Point-by-Point drawing. Trace simplification means that a trace is re-pixelized to the resolution of the screen and then reduced to a set of line segments that match this path. If the trace is closed all loops are removed and the trace is reoriented to be clockwise. Traces created as predefined shapes (e.g. ellipses, rectangles, and stamp shapes) do not undergo the simplification process, as they are assumed to be already simplified.

Once a trace has been simplified, the Simplified attribute is checked to prevent further, unnecessary simplification (Figure 8.11). The possibility of further simplification can be reestablished by unchecking the Simplified attribute.

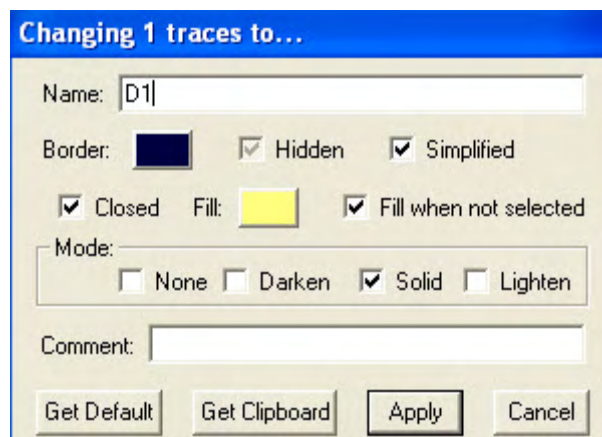


Figure 8.11. The Simplified attribute is checked in the Trace Attributes dialog.

Automatic simplification can be disabled by unchecking the option in the AutoTracing tab of the Series Options (Figure 8.11). Disabling simplification will allow the creation of traces with internal loops (Fig. 8.12). **Be careful not to create traces with loops!** The behavior of many subsequent measurements and operations are adversely affected by traces with loops, e.g. self-intersections.

For best results keep automatic simplification on and set the associated Simplify on Demand value to a negative number; which forces simplification to be at screen resolution.

NOTE: Keep automatic simplification on at all times for best results.




Figure 8.12. The closed trace on the left has loops, while the closed trace on the right does not.

Certain trace operations are not automatically followed by the simplification procedure. Therefore traces with loops can be created even when automatic simplification is enabled. A closed trace with loops can be created by drawing an open loop trace and then converting it to closed by changing its attributes. Doing this does not eliminate the loops even when automatic simplification is turned on. To eliminate the loops in a closed trace originally drawn as open, apply the *Simplify* command from the Trace menu. More information on the *Simplify* command is provided above.

Closed

A trace can be either closed or open. An open trace does not have the end of the trace connecting back to the start of the trace. Generally the open vs. closed characteristic is determined by the choice of drawing tool used to create the trace. For example, ellipses and rectangles are created as closed traces, while lines are open. The open vs. closed characteristic of traces can be changed by using the Closed checkbox in the Attributes dialog.

Adding a Comment

A comment string of up to 127 characters can be associated with each trace. This comment can be displayed in the Trace List window. The rules for comment strings are similar to those for trace names. Illegal characters {, = " \ < >} are replaced with underscores. The special characters {+ \$} are not interpreted for comment strings.

Trace Clipboard

The trace clipboard is an internal memory of the most recent set of traces copied or cut from the series. This clipboard is distinct from the Windows clipboard and cannot be used to paste data into another application. It can be used to paste trace data between sections within a series, or even between series.

Put the selected trace on the trace clipboard by selecting *Copy* from the Trace menu (Ctrl-C). This item is only enabled when one or more traces have been selected in the section. To also remove the selected traces from the section, use *Cut* (Ctrl-X). Copying or cutting to the trace clipboard deletes any traces currently in the clipboard before storing the new ones. Traces deleted with the Delete key or menu item are not put on the clipboard.

When the trace clipboard has data that can be pasted into the section, the *Paste* item in the Trace menu is enabled. Select *Paste* from the menu or use Ctrl-V to add the traces to the section. Any traces that are selected on the section are deselected and the pasted traces are selected. Pasting data does not empty the trace clipboard and the same traces can be repeatedly pasted into sections. Use *Paste attributes* (Ctrl-B) to copy the attributes of the first clipboard trace onto the selected traces.

The "Get Clipboard" button in the Trace Attributes dialog also allows the attributes of a clipboard trace to be copied to other traces or palette entries. To see what is on the trace clipboard, open the Trace Attributes dialog and use the Get Clipboard button. This will fill the dialog fields with the values of the first trace on the clipboard, revealing the trace name, color, etc.

Selecting Traces

When a section is first displayed, none of the traces are selected. When no traces are selected most of the operations in the Trace menu are disabled, since the operations are defined on selected traces only. To modify or use traces, it is necessary to first select them.

Traces can be selected in several ways. One way is to list the traces by selecting *List traces...* from the Trace menu. Select one or more traces from the list and then click on the *Select* item at the top of the Trace List window. When traces are selected they are drawn with a dashed border and possibly a colored interior. Selecting other traces from the Trace List will first deselect any existing selections.

To quickly select all the traces on the section, use the *Select All* item from the Trace menu. This operation is accelerated by the Ctrl-S key sequence.



The Selection Tool signified by the arrow pointer can also be used to select traces in the section. Place the tip of the arrow near a trace and press the left mouse button to select it. Use the right mouse button to deselect it. Pressing and holding the left mouse button allows a rectangular region to be delineated. When the button is released, all traces that have a trace point inside the region will be selected.

When using any other tool, depress the central mouse button or wheel to select the nearest trace. This allows traces to be selected for deletion or moving without changing back to the Selection Tool. For example, when cutting traces with the Scalpel Tool, one piece can be selected with the middle mouse button and deleted. It is not possible to deselect individual traces without using the Selection Tool.

Deselecting Traces

Traces can be deselected one at a time using the Selection Tool right mouse button. The selected trace nearest the tip of the arrow cursor is deselected when the right mouse button is pressed. To deselect all the selected traces simultaneously, use the *Deselect All* item from the Trace menu (Ctrl-D). *Deselect All* can be used with any tool.

Zooming to Selected Traces

The set of selected traces can be brought into maximum magnification by using the *Zoom To* item from the Trace menu. This operation is accelerated by the Ctrl-Home key sequence.

Finding Traces

A trace in the series can be located by using the *Find...* dialog from the Trace menu. When you type a trace name into the dialog and select OK, Find will search all the sections from the beginning of the series to find the first occurrence of that trace name, and then display the section with the trace selected. Selecting the *Find...* command again with the trace still selected, will find the first occurrence of the trace on the next higher section in the series. So, by repeatedly pressing Ctrl-F and Enter, the first instances of a trace can be tracked through the series.

The trace name entered in the *Find...* dialog can be any valid limit string as defined for the List tab of the Series Options. For example, searches for *nuc** will find the first occurrence of any trace with a name beginning with the letters n-u-c. Only the first instance of a trace within a section will be found, however. To locate and select all the instances of a particular trace name within a section, use the Trace List.

The first traces of an Object in the series can be located by selecting the *Find 1st!* command from the Object List menu. Selecting this command with an object name highlighted in the list will display the first section and first trace of the object in the main window. Repeating this command for an object will cycle the display through all the traces of the object that are found on the first section. This command cannot be used to find traces in the other sections of the object.

Tracing with Wildfire

The Wildfire Tool looks like a match on the Tools window. Pressing and releasing the left mouse button can be thought of as "igniting a fire" that spreads in all directions from the cursor point. When the spreading stops at the boundary of a region, a trace is created that surrounds this boundary. After the trace is created it is added to the selected traces and drawn on the section. If the trace is unacceptable it may be deleted by using the Delete key, but be careful that no other traces are also selected.

With the default settings, the Wildfire Tool ignites on bright regions and stops spreading at black regions. If a black region is completely surrounded by the bright region, the wildfire will grow completely around it and the black region will be ignored, i.e. it won't be part of the final boundary. Wildfires will not spread beyond the display of the window (except for the part under the status bar) and will not spread into a region in which there is no domain image. If an attempt is made to ignite a wildfire trace on a black region, no trace will be created.

The stop criteria for the Wildfire Tool can be changed in the AutoTracing tab of the Series Options (Fig. 8.8). Existing traces can be used as stop criteria by checking the appropriate box in the AutoTracing tab, but the main criteria that stop region growing are the relations set on the hue, saturation, and brightness values.

The default stop value of "Brightness is less than 100" means that Wildfires will not pass through pixels that have an intensity less than 100, where 0 is black and 255 is white. Changing this option to "Brightness differs by 50" will allow either white or black regions to be automatically traced with the Wildfire Tool. Saturation refers to the amount of color. A saturation value of 0 has no color, i.e. it is grayscale. Hue is actual color (red, green, blue, etc.). To get a visual picture of the colors defined by different hue, saturation, and brightness values, open a color dialog by clicking on the Fill color button in Names/Colors tab. (Note that brightness is sometimes referred to as luminance.) A stop criterion can be disabled by selecting the "differs by" relation with a large (255) value. For example, to

ignore color during region growing set "Hue differs by 255". Since this can never be true, a change in hue will never stop a region from growing.

In addition to the stop criteria, Wildfires can be contained by creating firebreaks with the right mouse button. This is typically used to separate two regions that are connected by a small gap. Such gaps are often discovered after having created a Wildfire that spread more than expected. To correct the spread of the Wildfire, delete the bad trace and place the cursor at one edge of the gap. Press and hold the right mouse button. The cursor will change to a burnt-out match. Drag the cursor across the gap leaving a small black line. Release the mouse button to create the firebreak. Use the left mouse button to repeat the original Wildfire. The region growth will not pass through the firebreak, although it can still grow around it if additional gaps remain.

During Wildfire region growing the cursor will change to an hourglass. This should last for only a second, but during this time region growing cannot be aborted by the Esc key. The drawing of a firebreak with the right mouse button can be aborted by hitting Esc before releasing the button.

Wildfire Editing

After a Wildfire trace is generated, examine the border for indentations with wide mouths (greater than about 1% of the total width of the object). Pick up the Pencil Tool and draw across these gaps, then hit Ctrl-M to merge the closed traces with the Wildfire trace. This will close the big gaps. To get rid of any remaining small gaps, use Ctrl-A to open the Attributes dialog and uncheck the Simplified flag. Click the Apply button and then select Simplify from the Trace menu. This will trace around the entire boundary of the Wildfire trace, close the small gaps and produce a simplified final trace.

Propagating Wildfires to Adjacent Sections

Wildfires can be automatically propagated to adjacent sections using the F or B keys on the keyboard. The F key propagates any recent wildfires forward to higher sections, while the B key propagates any recent wildfire backward to sections with a lower index. Once initiated, automatic tracing of wildfires will continue section-by-section to the end of the series, unless the user terminates the process (Ctrl-T) or the wildfire algorithm fails to generate a new region.

Note: When the Wildfire Tool is in hand, any selected traces will be propagated to adjacent sections when the F or B key is touched, so care should be used when using other Ctrl key sequences while holding the Wildfire Tool.

Automatic tracing of wildfires proceeds by using the centroid of the selected traces on the current section as ignition points of new wildfires on the next section. If ignition is successful, each wildfire grows to fill an area defined by the stop criteria. The area of the resulting new trace is compared to the area of the trace on the previous section. If this area differs by more than a certain percentage, the wildfire propagation is aborted. The percentage value can be specified in the AutoTracing tab of the Series Options. By setting this value high enough, say 999%, the area test can be effectively turned off. Tracing of very small regions can be prevented by setting the minimum number of square pixels a region should contain for tracing by wildfire in the Autotracing tab as well.

Generating Wildfires throughout a Region

By pressing and holding the left mouse button, the Wildfire Tool can be used to drag out a rectangular region (Fig. 8.13). When the mouse button is released all points inside the rectangle are tested for wildfire ignitions. Any regions satisfying the stop criteria are automatically traced using the region growing algorithm. When these regional wildfires are generated, firebreaks are respected but existing traces are not. Wildfires generated by the regional method, like any other selected traces, can be automatically propagated to adjacent sections using the F or B key.

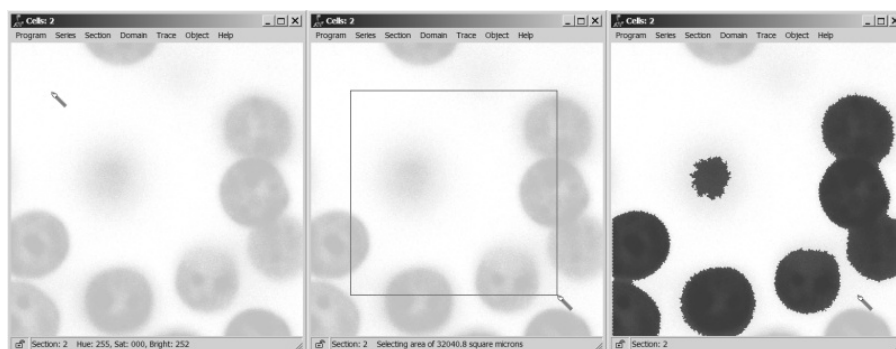


Figure 8.13. An example of using the Wildfire Tool to trace multiple areas. To begin, the cursor is positioned and the left mouse held down (left). A rectangle is created by dragging the mouse (middle). When the left mouse button is released (right), multiple areas are grown up to the stop criteria of "Hue differs by 40", "Saturation differs by 50", and "Brightness is greater than 200". To prevent creation of small areas, a minimum size of 100 square pixels was required for trace creation.

Regional wildfires can generate a very large number of traces and this can take a long time. Make sure the stop criteria are set appropriately by testing them with a single mouse click before attempting to trace multiple wildfires using the mouse drag. Use "less than" or "greater than" stop relations rather than "differs by" to generate fewer regions (by tracing only white areas rather both white and black areas). To further reduce the time required for regional wildfires, work at lower magnification.

The large number of small areas sometimes generated by a regional wildfire can be controlled by setting the "Don't trace areas less than..." parameter in the AutoTracing tab to a larger value. A value of zero (the default) allows all areas to be traced. A value of 100 square pixels, on the other hand, will limit tracing to just the areas containing at least 100 pixels at the current magnification (Fig. 8.13). This parameter also applies wildfires initiated by single mouse clicks or propagated from adjacent sections.



Chapter 9: Calibration and Measurement

This chapter deals with the quantitative aspects of RECONSTRUCT™ such as calibrating the dimensions of a series and getting accurate measurement data from the traces.

Units of Measurement

RECONSTRUCT™ works in an arbitrary system of units. The user can specify the name of the units in the General tab of Series Options. Each series has its own units. Note that RECONSTRUCT™ does not interpret or "understand" the units string. It is only a comment entered by the user to signify what units are used when data is entered. Nonetheless, the choice of units should be made carefully. Some computations are dependent on discretizing the data and this is sometimes done at relative fraction of 1 unit. Thus, if the objects to be reconstructed are about 1 micron in size, use microns as the units rather than millimeters or angstroms. A good rule of thumb is to choose units so that the image pixel size is a small fractional value (e.g. 0.01).

All input and output values are consistently in series units. For example, if the units are declared to be "microns" then all values entered for section thickness, movement values, image pixel size, etc. should be in microns. A section thickness value of 50 nm would be entered as "0.05" microns. A deformation would be entered as microns⁻¹. Trace area would be reported in the Trace List window as square microns. Similarly, the cursor position would be given on the status bar as the x- and y-position on the section in microns.

Data Precision

Section data is displayed and stored in series units. For example, trace points are entered as screen pixels and these values are immediately transformed into the corresponding series units on the section. The traces are stored in the section file as a sequence of (x,y) points in series units. These points are stored as ASCII strings with a fixed number of significant digits. The default number of significant digits (6) would allow values between -999 and 999 to be represented at a resolution of 0.001 series units. A trace point of more than 6 significant digits would be rounded when it is stored (e.g. 10000.059 to 10000.1), so the significant digits parameter must be large enough to capture the desired precision of the data over the whole range of the series volume.

The number of significant digits that are displayed and stored can be modified in the General tab of the Series Options. Increasing the significant digits will allow a larger range and/or increased resolution, but will also increase the storage requirements for the section files. The significant digits parameter also determines how many digits are displayed in lists and how many digits are output to .csv, .dxf, and .wrl files. Decreasing the parameter is generally not a good idea as this can result in making traces inaccurate, but saving fewer significant digits has the benefit of reducing the size of the files.

The actual position of a trace on the section is additionally determined by the associated transformation also stored in series units. Because the transformation consists of just a

few parameters it is always stored at the full resolution. Since all real values are represented by 64-bit floating point numbers, the full resolution is 15 significant digits (plus an exponent).

Image Pixel Size

Images are stored as raster bitmaps in the original image files. When this data is imported into a section, the pixels are mapped into series units using a magnification factor, the pixel size. The pixel size parameter tells how big a (square) pixel of the image is in series units.

Each image has its own pixel size parameter. This allows images obtained at different magnifications or scanned with different dpi settings to be accurately incorporated into the same series. A domain is stored as a pointer to the pixel data file and a pixel size parameter. The actual position of the domain on the section is additionally determined by the associated transformation (also stored in series units).

The pixel size of an image can be determined from the magnification factor of the microscope and the dpi used to digitize the optical image into pixels. Many computer-controlled imaging systems report the pixel resolution as part of the data, so that it doesn't have to be calculated. It can just be entered directly into the pixel size input field when importing images or domains.

The basic formula for calculating pixel size is

$$\text{pixel size} = \left(\frac{\text{units}}{\text{inch}} \right) \div \text{dpi} \div \text{magnification}$$

where "units" are the chosen series units, "dpi" is dots or pixels per inch, and magnification is the optical magnification from the specimen to the digitizer. A reminder of this formula is provided in the Domain Import dialog (Chapter 7: Domain Images).

For example, suppose an electron microscope is used with a known magnification factor from specimen to photographic film of 10,000. The film is then developed and scanned at 1000 dpi. If the series units were chosen to be microns then the pixel size would be

$$0.00254 = \left(\frac{25400 \text{ microns}}{\text{inch}} \right) \div 1000 \div 10000$$

While this estimate is a useful first approximation, magnification factors of instruments are rarely known precisely. Accurately determining the pixel size from the data is the main goal of calibration. If no image data is entered or calibrated, the initial screen resolution is assumed to represent a pixel size of 0.00254 in the default series units (microns).

Note that the final size of an image pixel in a section is determined by a combination of the pixel size and the transformation that maps the image to the section (Chapter 10). If this transformation contains scaling, the pixel size parameter will be different than the actual pixel size in series units. For this reason it is possible to calibrate the image by a combination of transformation scaling and pixel size but the preferred mode will be a fully calibrated pixel size.

Calibration Grids and Scale Bars

The pixel size can be accurately determined from calibration images. These can be scales or grids of a known physical size that were imaged at the same magnification and digitized at the same settings as the serial sections.

With some imaging systems a scale bar is imprinted on the section data (Fig. 9.1). RECONSTRUCT™ is flexible enough to allow the use of this scale bar data, but in most instances a calibration grid image is preferred.




Figure 9.1. Image of calibration grid. An electron micrograph of a diffraction grating replica with a known line spacing of 0.463 microns. The scale bar at the left (below the arrow) was generated by the instrument at a nominal length of 1 micron.

Calibration images can be stored anywhere within the series, in a few sections at the end of the data section, or in some empty corner of the data sections themselves. By tradition, a single calibration grid for the series is placed in an empty section at the start of the series in section 0. Section 0 is a special section that is not used to define objects or compute the z-distance. The thickness of section 0 is therefore arbitrary.

A section 0 can be created using *New...* from the Section menu and entering 0 for the section number. In section 0 use *New...* from the Domain menu to import the calibration grid image. The calibration grid can be imported with any pixel size as this will later be changed to the correct value.

Calibrating Traces

The process of calibration in RECONSTRUCT™ involves drawing traces of known length and then entering that length value using the *Calibrate...* operation from the Trace menu. If multiple traces are selected, *Calibrate...* will ask for the length of each trace by name and compute the average scale factor or pixel size for these values.

Drawing Calibration Traces

The easiest traces to create with a known length are usually line traces but other types of traces can be used. First, set the default drawing parameters using the Names/Colors tab of the Series Options. When drawing multiple calibration traces it may be convenient to use a "\$+" default trace name to allow the names to be incremented in the section

(Chapter 8: Traces). When all calibration traces will be of the same length, they can all have the same name if desired.



Select the Line Tool from the Tools window and draw a line trace on the calibration grid image. Place the ends of the line at the edges of the grid elements. Make the trace as long as possible to minimize the effect of inaccurate point placement at the ends. Draw more lines to also reduce inaccuracies of a single trace. When possible, draw lines in both directions on the calibration grid.

Calibrate Selected Traces

Once calibration lines are drawn, select the traces on the section. With only the calibration traces selected, select the *Calibrate...* item from the Traces menu. A dialog box will appear for each selected trace displaying the current length of the trace. Enter the actual length of the trace in place of the current length. For an open trace the length is the sum of the distances between trace points. The length of closed traces is the entire perimeter including the closing segment.

Resize Traces, Set Pixel Size, or Scale Section

The end result of calibration can be one of three things: resizing the selected traces to the length specified, changing the pixel size of the domain in which the traces were drawn, or scaling all the transformations in the section (domains and traces). Which of these three operations will be performed is determined in the Calibrate Traces dialog (Fig. 9.2), presented after the lengths of all selected traces have been entered. Complete the calibration operation by choosing one of the options and selecting OK. No changes will be made if the Cancel button is selected.

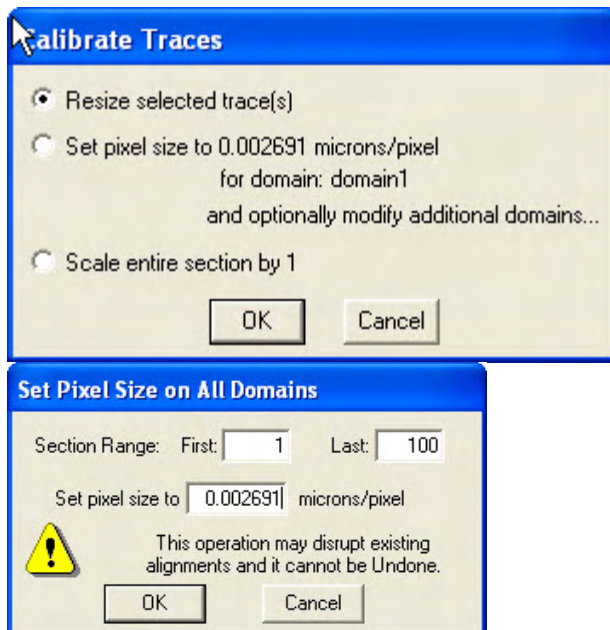


Figure 9.2. The Calibrate Traces dialog (left). Three options are provided for adjusting the section data based on the entered trace lengths. When the second option is chosen a subsequent dialog (right) allows for applying the pixel size adjustment to all domains in the series.

The "Resize selected trace(s)" option will only modify the selected traces. The transformations of the traces are scaled so that the traces will be redrawn with the specified new lengths. Each trace is scaled independently rather than to an average rescaling value as in the other options.

The domain on which the selected traces are centered is listed in the second option. If the calibration traces are not centered on a domain, then this option will be disabled. The correct pixel size of the domain will be set by selecting the second option. This pixel size value is computed by transforming the selected traces into the domain and computing their lengths on the image in pixels. The sum of the entered lengths (in units) is divided by the sum of lengths in pixels to get the correct pixel size for the calibration grid. This pixel size can be subsequently applied to all domains as discussed in the next section.

Selecting the second option will only change the pixel size of domains, not of traces in the section. If the section has lots of traces and domains it may be advantageous to choose the final option instead. This option computes a scaling factor based on the ratio of the sum of the entered lengths divided by the sum of the current lengths. The scaling factor is applied to the transformations of all domains and traces in the section, effecting the calibration by scaling rather than setting the pixel size.

It is desirable to have the pixel size parameter accurate for later reference, since the scaling factor will not be visible to the user. So the second option should be used in the Calibrate Traces dialog for the initial calibration. With this second option existing traces will no longer be correct and existing alignments between domains will be disrupted. This is not a problem if calibration is carried out before any domain movements, section alignments, or tracing. However, the calibration traces will be incorrectly sized after a pixel size adjustment. These calibration traces may be simply deleted. To verify the calibration, redraw the calibration traces and check their length to see that they match the known lengths in the calibration image.

Apply the Calibration to Other Sections

The pixel size set by the "Set pixel size..." option can be subsequently extended to other domains in the series. When this option is chosen the calibration automatically opens the "Set Pixel Size on All Domains" dialog box (Fig. 9.2). This dialog displays the same pixel size value as in the Calibrate Traces dialog and allows a range of sections to be selected. All domains in the included sections will be set to the pixel size value.

Note: This operation cannot be undone when executed over many sections.

The rescaling performed by the third option of the Calibrate Traces dialog can be applied to other sections in the series by using the *Repeat* or *Propagate...* commands from the Movement submenu. To rescale one other section, unlock the section and select *Repeat* (or Ctrl-R). To rescale a range of sections, use *Propagate...* and specify the range of sections to modify. Be careful not to rescale the current section again, because this section was already rescaled by the Calibrate Traces dialog.

If the first calibration option "Resize selected trace(s)" was chosen in the Calibrate Traces dialog, then no multi-section application of the calibration is possible.

Section Thickness

Section thickness is critical for determining the 3D representation and for getting accurate 3D measurements. If section thickness is known in advance for all sections it can be set in the General tab of the Series Options. As mentioned above, all input and output values are consistently in series units. For example, if the units are in "microns" then all values entered for section thickness, movement values, image pixel size, and so on should be in microns. RECONSTRUCT™ will prompt you to enter units when you modify section thickness using the *Thickness...* menu item from the Modify menu of the Section List. This value is applied to new sections. If a default section thickness is not specified a value of 1/20 of a section unit will be used for a new series.

The section thickness of individual sections can be set at any time using the *Thickness...* item from the Section menu. For many sections, the thickness can be set by selecting the sections from the Section List, then using the *Thickness...* item from the Modify menu of the Section List. One cannot assign different thicknesses within an individual section.

Calibrating Section Thickness

Thickness is a property of a section that needs to be calibrated or estimated. Microtome settings can be used as an estimate but accuracy may be an issue. The thickness of individual sections can be measured from minimal folds in the section when these are present. An alternative method for estimating average section thickness takes advantage of longitudinally sectioned cylindrical objects such as mitochondria (Fiala & Harris, 2001). This method is applicable to series in which the section dimensions (e.g. pixel sizes) have already been calibrated and the sections have been aligned.

A cylinder sectioned longitudinally has its diameter appear in the series in two directions (Fig. 9.3). First, in the section passing through the middle of the cylinder, the diameter (d) can be measured in the plane of the section. The diameter also extends orthogonal to the plane of sectioning, causing the cylinder to appear across several sections. By counting the number of sections (s) through which it appears, section thickness can be estimated as the ratio d/s . Taking the average of many such measurements distributed throughout the series results in an estimate of the average section thickness. This value can be applied to every section in the series using the *Thickness...* dialog from the Modify menu of the Section List.




Figure 9.3. The method of cylindrical diameters involves computing the ratio of the diameter of a cylindrical object at its widest point on a section (arrows) relative to the total number of sections the object spans (e.g. section thickness = $d/\#$ sections occupied)

Z-Distance

The z-distance of a section is the distance along an axis perpendicular to the plane of sectioning from the start of the series to the section (Fig. 9.4). This z-distance is used to compute 3D values from traces. The z-distance for any section is computed by summing the thicknesses of preceding sections. The section indices are not used in the z-distance calculation. Section indices only indicate the order of sections, not the relative locations of sections in space. Section locations and their relative distances are entirely determined by the sum of thicknesses. The z-distance is computed relative to the section with the lowest index greater than zero. Changing section indices without altering the order does not affect z-distance. However, adding sections with lower indices will change the z-distances.





Figure 9.4. Two methods are available for computing the z-distance of section n. One computed the distance to the middle of the section: $z = t_1 + t_2 + t_3 + \dots + t_{n-1} + t_n/2$. The other method computes the distance to the top of the section: $z = t_1 + t_2 + t_3 + \dots + t_{n-1} + t_n$.

Two methods are provided for the computation of z-distances (Fig. 9.4). One computes the distance to the middle of the section, while the other computes the distance to the top of the section. The former method gives more accurate z-distances and 3D representations when sections are not uniform in thickness. Setting z-distances to the tops of sections provides an easy way to represent gaps of missing sections by giving a larger thickness to the sections at the top of the gaps. When section thickness is uniform, the only difference between the methods is a half thickness offset in z values. The method of z-distances calculation can be selected in the General Tab of the Series Options.

Physical gaps in the series can be represented by empty sections of an appropriate thickness. For example, suppose a series was cut at 50nm section thickness but only every fifth section was imported. The sections were numbered 1, 6, 11, 16,... to reflect that every fifth section was used. Assuming uniform section thickness, these sections could all be given a thickness of 50nm. New sections could be created at 2, 7, 12,... with thickness 200nm to represent the gaps. To get accurate object measurements with this method, traces would need to be created in the gap sections so that the missing thickness is accounted for in the surface area and volume calculations (see below). Alternatively, only sections 1, 6, 11, 16,... could be created but with thicknesses 50nm, 250nm, 250nm, 250nm,..., reflecting the gaps between these sections. The z-distances to the tops of these sections would then be 50nm, 300nm, 550nm, 800nm, etc. No additional gap traces would be needed in this case to get accurate surface areas and volumes.

Status Bar Measurements

Traces are quantified in various ways and these values are reported in the lists, as discussed below. In addition, the status bar reports incremental measurements related to the position of the mouse pointer on the screen. When no tool is active, the status bar reports the name of the trace nearest the pointer and the position of the pointer on the section in series units. When a selected domain is displayed, instead of a trace name the domain is reported along with the position of the pointer in image pixels. When a drawing tool is active, the reported value depends of the particular tool in use.




These tools create a rectangular region when dragging while the left mouse button is depressed. The status bar reports the area of this rectangle in series units (squared), except in the case of the ellipse. For the ellipse, the rectangular area is multiplied by $\pi/4$, giving the area of the ellipse that would be created when the mouse button is released.



For the Panning and Domain Selection Tools, the status bar reports the total distance moved by a left mouse drag.



When a line is being drawn with the Line Tool, the status bar reports the length of the candidate line on the section. If the section is changed during drawing, the 3D length is not reported. For a dynamic length measurement that takes into account the section thicknesses, use the Z-Trace Tool.



For both the open and closed point-by-point drawing tools, the status bar reports the position of the pointer on the section. This allows points to be entered at particular locations.



When a z-trace is being drawn the status bar reports the length of the next segment including the difference in the z-distances of the start and end sections:

$$\text{length} = \sum_{i+} \sqrt{(x_{i+1} - x_i)^2 + (y_{i+1} - y_i)^2 + (z_{i+1} - z_i)^2}$$

Domain List Measurements

Two dimensional measurements on the domain boundary, including length, area, and the x- and y-position of the midpoint, are reported as columns in the Domain List window. The columns are selected by checkboxes in the Lists tab of the Series Options. All values are computed at the time the Domain List window is created. To make sure that the values are up-to-date, refresh or reopen the Domain List window after any changes to the domains are made. All values are in series units.

Length

The length of the domain boundary is the sum of the distances between boundary points.

Area

The area is the size in square units of the region enclosed by the boundary. This area is computed using a formula based on Green's Theorem.

$$\text{area} = \frac{1}{2} \sum_i (x_i y_{i+1} - x_{i+1} y_i)$$

for each segment from (x_i, y_i) to (x_{i+1}, y_{i+1}) in the boundary. The accuracy of this computation is dependent on the boundary being free of internal loops and the values (x, y) not being too large. Since 64-bit floating-point values are used for all such calculations, the second factor should not be an issue for most series.

Midpoint X

This column reports the x position of the midpoint of the extent of the boundary in series units. The calculation is

$$x = \frac{\max(x_i) + \min(x_i)}{2}$$

Midpoint Y

This column reports the y position of the midpoint of the extent of the boundary, calculated by

$$y = \frac{\max(y_i) + \min(y_i)}{2}$$

Trace List Measurements

Two dimensional measurements are reported for traces within a section in the Trace List window. These include trace length, area, and the x- and y-position of a trace. These values are reported as columns in the Trace List window. The columns displayed are determined by checkboxes in the Lists tab of the Series Options. All values are computed from the section data at the time the Trace List is filled. Refresh, reopen, or page sections to bring the Trace List window values up-to-date after any changes to the traces. All values are in series units. Use the Series>Export>Trace Lists... function to put all trace length data into a format for use in a spreadsheet such as Excel.

Length

The length of a trace is the sum of the distances between trace points.

$$\text{length} = \sum_{i+} \sqrt{(x_{i+1} - x_i)^2 + (y_{i+1} - y_i)^2}$$

Area

The area is the size in square units of the region enclosed by a trace. A negative area indicates that the trace has a counterclockwise orientation. A positive area indicates a clockwise orientation. The area of an open trace is zero.

Area and centroids are calculated from a formula based on Green's Theorem.

$$\text{area} = \frac{1}{2} \sum_i (x_i y_{i+1} - x_{i+1} y_i)$$

for each segment from (x_i, y_i) to (x_{i+1}, y_{i+1}) in the trace. The accuracy of this computation is dependent on the trace being free of internal loops and the values (x, y) not being too large. Since 64-bit floating-point values are used for all such calculations, the second factor should not be an issue for most series.

Centroid X

This column reports the x-position of the centroid of the trace in series units. The calculation is

$$x = \frac{\sum_i (x_i + x_{i+1})(x_i y_{i+1} - x_{i+1} y_i)}{3 \sum_i (x_i y_{i+1} - x_{i+1} y_i)}$$

Centroid Y

This column reports the y-position of the trace centroid in series units, calculated by

$$y = \frac{\sum_i (y_i + y_{i+1})(x_i y_{i+1} - x_{i+1} y_i)}{3 \sum_i (x_i y_{i+1} - x_{i+1} y_i)}$$

Min X

This column reports the minimum x-position of the trace in series units.

Min Y

This column reports the minimum y-position of the trace in series units.

Max X

This column reports the maximum x-position of the trace in series units.

Max Y

This column reports the maximum y-position of the trace in series units.

Z Position

The z-distance of the section, as described in the previous section.

Thickness

The section's thickness in series units.

Object List Measurements

The Trace List column data can be used to compute 3D quantities by utilizing the length and area columns in combination with the section thickness column. Since these quantities depend only on the section thickness and not the z-distance, they are not affected by changing the method of z-distance calculation. A few of these common computations are done automatically in the Object List.

Whenever the Object List window is cleared and refilled, all section files are scanned for trace data and these are grouped into objects. The measurements for these objects are then generated from the traces and section thicknesses, and this data is added to the list. To make sure that the measurement values are up-to-date after any changes to the traces are made, refresh the Object List window.

The time it takes to fill the Object List depends largely on how fast the section files are scanned, but it also depends on how much data will be displayed in the list. The number of list entries can be reduced by first setting a limit string in the Series Options Lists tab.

A small speed-up can also be obtained by reducing the number of columns displayed for each entry. The area and volume measurements are computed only if these columns are displayed in the list, so unselecting these columns will decrease the time it takes to fill the list.

Count

The count is the total number of traces in all sections of the object.

Surface Area.

Each trace of an object contributes to the surface area an amount equal to the product of the trace length (as computed in the Trace List) and the section thickness.

$$\text{Surface Area} = \sum_{\text{all sections}} \text{length} \times \text{thickness}$$

Flat Area

The enclosed areas of closed traces in the object are summed with the surface areas of the open traces to get the total "flat" area.

$$\text{Flat Area} = \sum_{\text{all sections}} \left[\sum_{\text{open traces}} \text{length} \times \text{thickness} + \sum_{\text{closed traces}} \text{area} \right]$$

The area is computed using the Green's theorem formula as in the Trace List. Negative areas detract from the total flat area. This measurement is useful for objects that are known to be flat but may sometimes appear in the plane of section as well as perpendicular to it (Fig. 9.5). The flat area can also be used to include end caps on surfaces.

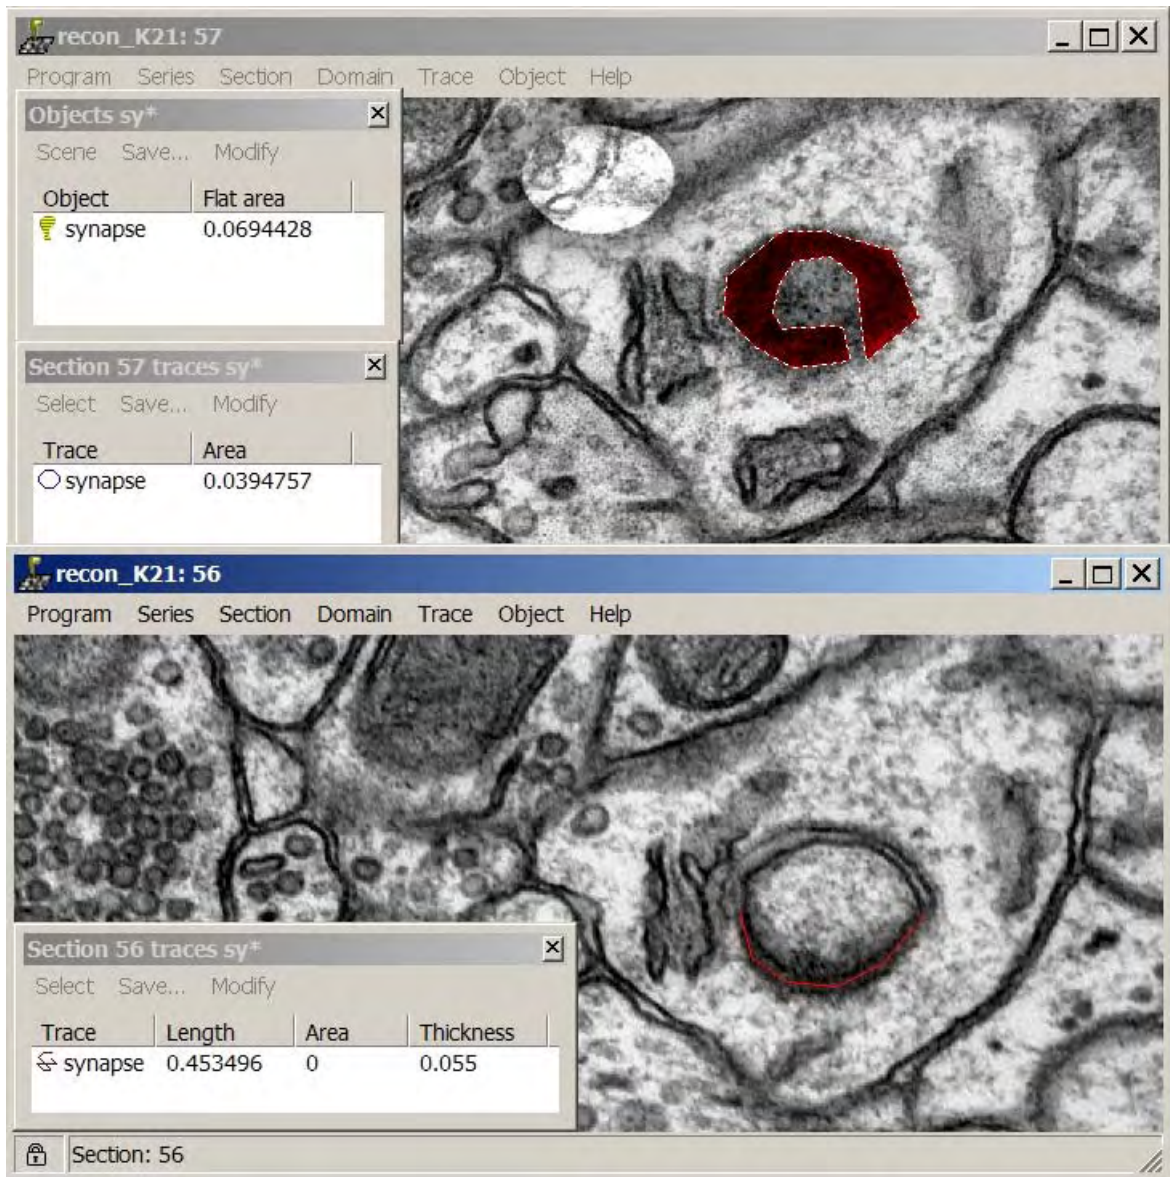


Figure 9.5. Computation of the flat areas of a cell-to-cell junction imaged by electron microscopy. The junction lies on the surface of two cell membranes. In section 57 (top) the junction lies mostly within the section and is represented as a closed trace (red). In section 56 (bottom) the junction lies mostly perpendicular to the section and is represented as an open profile (red). The flat area is the enclosed areas of the closed traces plus the lengths of the open traces times the section thicknesses. (Some sections of the object not shown.)

See Chapter 8 for detailed information on Open and Closed traces.

Volume

Volume is computed using the Cavalieri formula:

$$\text{Volume} = \sum_{\text{all sections}} \text{area} \times \text{thickness}$$

Each trace contributes the product of its enclosed area (as computed in the Trace List) times the section thickness. Traces with negative areas diminish the volume and open traces contribute nothing.

Z-Trace List Measurements

Lengths of z-traces are reported in the Z-Trace List window. All values are computed at the time the Z-Trace List window is filled. To make sure that the values are up-to-date after adjusting section thicknesses, refresh the Z-Trace List. All values are in series units.

Length

The length of a z-trace is the sum of the distances between points.

$$\text{length} = \sum_{\#} \sqrt{(x_{i+1} - x_i)^2 + (y_{i+1} - y_i)^2 + (z_{i+1} - z_i)^2}$$

The x- and y-positions of the points were determined when the data was entered, as was the section index of the point. When the Z-Trace List is filled these section values are converted to z-distances (z_i). Note that z-trace lengths may be affected by the method of z-distance calculation when section thickness is not uniform (see above.)

Distances List Computation

Three-dimensional distances between objects are listed in the Distances List window. Be sure to use the limit strings in the Lists Tab of the Series Options to limit the number of distance computations because the operation will be slow when the distances from every object to every other object are included. All distances are computed when the Distances List window is filled. You can monitor the progress of the computations by looking at the first line of the list. This line shows the objects between which distances are currently being computed. To abort the computations, close the list window. When all distances have been computed, the list will be filled with the results. All distances are in series units.

Algorithm

The distance between two objects, O_1 and O_2 , is the minimum 3D distance between any two traces of the object. If any two traces intersect on a section, then the distance between the objects is zero. If traces never intersect, then the distance is a positive value. Thus, if one object is located inside another there will still be a distance between them that reflects the minimum distance between their boundaries.

$$\text{Distance}(O_1, O_2) = \text{Min}(\text{Distance}(T_i, T_j) \forall (T_i \in O_1) \text{ and } (T_j \in O_2))$$

where the distance between any two traces, T_1 and T_2 , is the minimum 3D distance between any two line segments (S_i, S_j) of the traces:

$$\text{Distance}(T_1, T_2) = \text{Min}(\text{Distance}(S_i, S_j) \forall (S_i \in T_1) \text{ and } (S_j \in T_2))$$

Since all the segments of T_1 and T_2 lie on two parallel planes, we can defer dealing with the z-component of the distance and first compute the minimum squared x-y distance. The squared x-y distance between two line segments, S_1 from point P to point Q and S_0 between points R and T, is computed by first computing the vectors

$$\begin{aligned} d0 &= T - R \\ d1 &= Q - P \\ d &= R - P \end{aligned}$$

then, using the x- and y-components of these vectors, computing

$$\begin{aligned} a00 &= d0_x d0_x + d0_y d0_y \\ a11 &= d1_x d1_x + d1_y d1_y \\ a01 &= d0_x d1_x + d0_y d1_y \\ b0 &= d0_x d_x + d0_y d_y \\ b1 &= d1_x d_x + d1_y d_y \\ c &= d_x d_x + d_y d_y \\ e &= (d0_x d1_y - d1_x d0_y)^2 \\ f &= a01 \cdot b1 - a11 \cdot b0 \\ g &= a00 \cdot b1 - a01 \cdot b0 \\ h &= a00 - 2 \cdot a01 + a11 + 2 \cdot b0 - 2 \cdot b1 + c \end{aligned}$$

The segments are parts of two lines. The squared x-y distance between the lines is given by (Schneider & Eberly, 2003)

$$a00 \cdot t_0^2 - 2 \cdot a01 \cdot t_0 \cdot t_1 + a11 \cdot t_1^2 + 2 \cdot b0 \cdot t_0 - 2 \cdot b1 \cdot t_1 + c$$

the minimum of which occurs at the values t_0 and t_1 where

$$\begin{aligned} a00 \cdot t_0 - a01 \cdot t_1 + b0 &= 0 \\ a11 \cdot t_1 - a01 \cdot t_0 - b1 &= 0 \end{aligned}$$

The correct values of t_0 and t_1 must take into account the end points of the line segments. This constrains the solutions of t_0 and t_1 to be between 0 and 1, inclusive.

If e is zero, then the lines are parallel, and the squared x-y distance D is

$$\begin{aligned}
 &\text{if } (a_{00} = 0) \text{ and } (a_{11} = 0) \text{ then } D = c \\
 &\text{else if } (a_{00} = 0) \text{ then } D = \begin{cases} c & \text{if } (b_1 \leq 0) \\ a_{11} - 2 \cdot b_1 + c & \text{if } (b_1 \geq a_{11}) \\ c - (b_1)^2/a_{11} & \text{otherwise} \end{cases} \\
 &\text{else if } (a_{11} = 0) \text{ then } D = \begin{cases} c & \text{if } (-b_0 \leq 0) \\ a_{00} + 2 \cdot b_0 + c & \text{if } (-b_0 \geq a_{00}) \\ c - (b_0)^2/a_{00} & \text{otherwise} \end{cases} \\
 &\text{else } D = \begin{cases} c & \text{if } (a_{01} < 0) \text{ and } (-b_0 \leq 0) \\ a_{00} + 2 \cdot b_0 + c & \text{if } (a_{01} > 0) \text{ and } (-b_0 \geq a_{00}) \\ a_{11} - 2 \cdot b_1 + c & \text{if } (a_{01} > 0) \text{ and } (b_0 \geq a_{01}) \\ h & \text{if } (a_{01} < 0) \text{ and } (a_{01} - b_0 \geq a_{00}) \\ c - (b_0)^2/a_{00} & \text{otherwise} \end{cases}
 \end{aligned}$$

If e is not zero, then the lines are not parallel. If f is less than or equal to zero, then D is computed by

$$\begin{aligned}
 &\text{if } (g < 0) \text{ and } (b_1 < 0) \text{ then } D = \begin{cases} c & \text{if } (-b_0 < 0) \\ a_{00} + 2 \cdot b_0 + c & \text{if } (-b_0 > a_{00}) \\ c - (b_0)^2/a_{00} & \text{otherwise} \end{cases} \\
 &\text{else if } (g > e) \text{ and } (b_1 > a_{11}) \text{ then } D = \begin{cases} a_{11} - 2 \cdot b_1 + c & \text{if } (a_{01} < b_0) \\ h & \text{if } (a_{01} - b_0 > a_{00}) \\ a_{11} - 2 \cdot b_1 + c - (a_{01} - b_0)^2/a_{00} & \text{otherwise} \end{cases} \\
 &\text{else } D = \begin{cases} c & \text{if } (b_1 < 0) \\ a_{11} - 2 \cdot b_1 + c & \text{if } (b_1 > a_{11}) \\ c - (b_1)^2/a_{11} & \text{otherwise} \end{cases}
 \end{aligned}$$

If f is greater than or equal to e , D is instead computed by

$$\begin{aligned}
 &\text{if } (g < 0) \text{ and } (a_{01} + b_1 < 0) \text{ then } D = \begin{cases} c & \text{if } (-b_0 < 0) \\ a_{00} + 2 \cdot b_0 + c & \text{if } (-b_0 > a_{00}) \\ c - (b_0)^2 / a_{00} & \text{otherwise} \end{cases} \\
 &\text{else if } (g > e) \text{ and } (a_{01} + b_1 > a_{11}) \text{ then } D = \begin{cases} a_{11} - 2 \cdot b_1 + c & \text{if } (a_{01} < b_0) \\ h & \text{if } (a_{01} - b_0 > a_{00}) \\ a_{00} + 2 \cdot b_0 + c - (a_{01} + b_1)^2 / a_{11} & \text{otherwise} \end{cases} \\
 &\text{else } D = \begin{cases} c & \text{if } (b_1 < 0) \\ a_{11} - 2 \cdot b_1 + c & \text{if } (b_1 > a_{11}) \\ c - (b_1)^2 / a_{11} & \text{otherwise} \end{cases}
 \end{aligned}$$

Finally, if f is between zero and e , D may be obtained from

$$\begin{aligned}
 &\text{if } (g \leq 0) \text{ then } D = \begin{cases} c & \text{if } (-b_0 < 0) \\ a_{00} + 2 \cdot b_0 + c & \text{if } (-b_0 > a_{00}) \\ c - (b_0)^2 / a_{00} & \text{otherwise} \end{cases} \\
 &\text{else if } (g \geq e) \text{ then } D = \begin{cases} a_{11} - 2 \cdot b_1 + c & \text{if } (a_{01} < b_0) \\ h & \text{if } (a_{01} - b_0 > a_{00}) \\ a_{11} - 2 \cdot b_1 + c - (a_{01} - b_0)^2 / a_{00} & \text{otherwise} \end{cases}
 \end{aligned}$$

In the final case, where g is also between 0 and e , the segments probably intersect, but to make sure, the full distance is calculated and the magnitude checked.

$$\begin{aligned}
 x &= R_x - P_x + (f \cdot d_{0x} - g \cdot d_{1x}) / e \\
 y &= R_y - P_y + (f \cdot d_{0y} - g \cdot d_{1y}) / e \\
 D &= x^2 + y^2
 \end{aligned}$$

If this value is very small ($< 10^{-14}$) then D is explicitly set to zero. So some objects that touch may have small, non-zero distances, but objects that don't touch will never accidentally be given a distance of zero.

The minimum squared x-y distance $Min(D)$ between two traces is determined by repeating the above calculation for every pair of segments in the two traces. The actual 3D Distance between the traces is then computed by

$$\begin{aligned}
 z &= ZDistance(T_1) - ZDistance(T_2) \\
 Distance(T_1, T_2) &= \sqrt{Min(D) + z^2}
 \end{aligned}$$

where the z-distance of the section of each trace is computed as described above. Thus, inter-object distances will be dependent on the method of z-distance calculation only when the section thickness is not uniform throughout the series.

Copyright © 2006 by John C. Fiala
Revised by K. Harris and K. Sorra; 2009



Chapter 10: Alignments and Movements

One of the principal features of RECONSTRUCT™ is the ability to align serial section images. Bringing data on adjacent sections into alignment facilitates measurement, analysis and 3D reconstruction. Data within a section can also be arranged by moving image domains and traces in the main window. This process of moving data elements in a section is different than panning or zooming the view. Changing the view point affects the display of all sections in the same way, whereas moving a section or an element within a section changes the structure of the representation.

Alignments

Alignment of adjacent serial sections is accomplished by moving one section in relation to another with the keyboard. This process can be time-consuming, so alternative methods are provided. Sections may be aligned *By correlation* from the Movement submenu of the Section menu, or by trace correspondence using the Align submenu of the Trace menu.

The Align submenu provides four alignment operations (Table 10.1). The alignment operations adjust different numbers of basis functions in the transformation. Using more basis functions in a transformation allows more degrees of freedom but requires more data points for the Align computation.

TABLE 10.1. The Align Submenu Commands

Command	Basis functions	Number of data points required	Movement potential
Rigid	1,x,y	2	Rotations and translations
Linear	1,x,y	3	+ scaling and skew (slant)
Deformal	1,x,y,xy	4	+ deformations
Quadratic	1,x,y,xy,x ² ,y ²	6	+ bending

Aligning Selected Traces

The Align submenu commands can be used to put the centroids of two or more traces into alignment with the centroids of corresponding traces in another section. When multiple traces are to be aligned they must have unique names, such that traces can be paired by name. This can often be achieved by using the "\$+" naming convention when creating traces (Chapter 8: Traces).

To align traces between sections, select in one section the traces that will be the stationary reference. Jump to the other section and select the traces which will be moved. With the section locked, open the "Align traces" submenu and select the desired alignment operation, e.g. *Rigid*. The selected traces will be immediately repositioned.

Aligning Entire Sections

When the section is unlocked the Align submenu is changed from "Align traces" to "Align section". This indicates that the entire section will be moved by the alignment command rather than just the traces.

The same procedure is used for aligning sections as with aligning traces. Select in one section the traces that will be the stationary reference. Page to the other section and select the traces which will be aligned to this reference. Unlock the section (Ctrl-L). From the "Align section" submenu select the desired alignment operation. The selected traces and the rest of the section will be immediately repositioned. If the message "Alignment needs more traces" is displayed, you may need to enter more traces on each section to have a sufficient basis for alignment.

Evaluating Alignments

As discussed in Chapter 6: Sections, the alignment between two sections can be compared using either blending or flickering. Blending overlays the two sections (Fig. 10.2), while flickering rapidly switches between the two sections. Blend the sections just aligned using the spacebar. To flicker, hold the / key down.




Figure 10.1. Blending of unaligned (left) and aligned (right) sections. Misalignments produce a blurry image on the left, while the blending of aligned images results in sharp, clear features where they are in cross-section (right).

While either blending or flickering alone can be used to evaluate the alignment, a combination of the two methods is advantageous. With blending it is relatively easy to tell when there is good alignment. Many of the cross-sectioned features (edges) come into sharp focus, while misaligned features are blurry. However, with blending it is often difficult to tell the direction of misalignment. Flickering between sections generates apparent motion of misaligned features that quickly indicates the direction and magnitude of the offsets. Some misalignment between features is to be expected even when the sections are correctly aligned since not every object in the series will be perfectly cross-sectioned. If the apparent motion generated during flickering is not uniform in direction and magnitude throughout the section, but is instead locally random, this indicates that the alignment is correct even though there is some blurriness with blending.

If an alignment command produces a bizarre distortion of the section, this probably means that there is a serious problem with the traces used for correspondence. Undo the alignment and carefully check the trace names to make sure that the desired correspondences were correctly specified.

Undoing and Adjusting Alignments

If the alignment produced by one of the alignment commands is not adequate it can be modified or undone. To undo it, first make sure the section number in the title or status bar is the one on which the alignment was done, i.e. the section that was moved. Then use Ctrl-Z (*Undo*) to return the section to its previous state.

If you plan to try the same alignment operation again after modifying or adding traces for correspondence points, it is not necessary to first undo the alignment. Points can just be added to each section and the alignment operation repeated. If a *Quadratic* or *Deformal* alignment is followed by a *Rigid* or *Linear* alignment it will be necessary to clear the nonlinear components from the transformation (by Ctrl-Z) since they are not cleared by executing a subsequent *Linear* or *Rigid* alignment. After a *Linear* alignment, an *Undo* is not necessary prior to adding points and trying again with a *Linear*, *Deformal*, or *Quadratic* alignment.

Repeating an Alignment

A command from the Align submenu can be repeated and propagated just like any other movement. Use the *Repeat* item from the Movement submenu to apply the same adjustment to subsequent sections. Use *Propagate...* to apply the same adjustment to many sections. Remember, *Repeat* can be undone but *Propagate...* cannot.




Figure 10.2. A series with a misalignment between sections 103 and 104. By comparing the view of the faces, all the lower sections 101 through 103 are aligned to each other. All the upper sections 104-200 are also aligned, but a misalignment between sections 103 and 104 spoils the complete series alignment.

Sometimes a misalignment is found in the middle of an aligned series, creating the scenario depicted in Figure 10.2. To fix such an alignment mistake, align the section on the upper side of the mistake (section 104 in the figure) to the section on the lower side (section 103). Keep undoing the alignment command, fixing the points, and repeating the alignment command until completely satisfied with the result. (Only the last alignment command can be propagated, so using multiple commands from the Align submenu without *Undos* between them will not work in this case.) Once 104 is aligned to 103, select *Propagate...* from the Movement submenu and enter the range for the remaining sections on the upper side (sections 105-200). The result will put the entire series in alignment.

Montaging Domains by Point Correspondences

Quick Start gave a protocol for aligning sections using point correspondences. The point correspondence technique can also be applied to the problem of montaging multiple images within a section. This is done by repeating on the selected domain the alignment movement generated from "Align Traces" on corresponding points.




Figure 10.3. Montaging by point correspondence. A. Adjacent section (low resolution field or completed montage). B. New domain on section with yellow triangles representing correspondences to points in A. C. Align Traces. D. Repeat the movement on the selected domain to place image at the correct location on the section.

Suppose an adjacent section is already montaged (Fig. 10.3A). This can be accomplished by manual movements, or by using an image that contains the entire field at low magnification. To align domains in the next section to this adjacent section, first enter corresponding point traces at places on the domain (Fig. 10.3B) and the adjacent section (Fig. 10.3A) that should be put into alignment. This is done on the sections without domains selected, rather than on selected domains. With the point traces selected in both sections, use *Align Traces* from the *Trace* menu to generate the correct alignment movement for the domain. This will move the traces from the domain to the correct positions on the section (Fig. 10.3C). To move the domain to match the trace movement, select the domain and use *Repeat* from the *Movement* menu. This will place the domain at the location of the moved traces (Fig. 10.3D).

The technique can be repeated with each domain to align each element of a montage separately. Just make sure that a unique trace name is selected for entering trace points on each domain, so that the traces for each domain can be distinguished.

A variation of the above method can be used to add and align a domain to an existing section. Assume that goal is to import a new domain image into section I and align it to the existing part of section I, instead of to an adjacent section.

Step 1.

Import the domain to a different section (II). This section is arbitrary, but it may be convenient to use an adjacent section for paging. Select the domain and move it to the approximate desired location on the section. Merge the domain into the section.

Step 2.

Enter correspondence points on sections I and II that would align the domain on section II with desired features on section I. Then use *Align Traces* from the *Trace* menu on section II. This will move the correspondence points into alignment which can be verified by flickering.

Step 3.

Select the domain from section II. Use *Ctrl-R* to repeat the alignment that was applied to the traces. The domain will now have the correct transformation on section II that it should have on section I. If desired, the alignment can be verified by merging and flickering, however at the end of this step the domain should be selected.

Step 4.

From the *Movement* submenu, use *Record from selected* to capture the domain transformation from section II. Merge the domain back into section II.

Step 5.

Import the domain image into section I using the same pixel size as used when importing it into section II. Select the new domain on section I and repeat the transformation recorded from section II by using *Ctrl-R*. The domain will now be correctly positioned on section I. Go back and delete the domain from section II. If desired, the deletion can occur in Step 4, immediately after recording the transformation.

Aligning By Correlation

A simple alternative to the point correspondence technique can sometimes be used to align domains or sections. This technique uses the *By correlation* movement command (Ctrl-). *By correlation* computes the peak of the cross-correlation function between the front section and the back section, so two sections (or a domain and a section) must be available for the operation to do anything. (You can be sure that two are available if you flicker between them.) The algorithm determines the translation required to move the front (current) section or domain to the peak of the cross-correlation and immediately applies this shift to the section. The entire computation will typically require several seconds, but the end result will be a movement of the current section into alignment with the previous section, provided the current section is unlocked.

By correlation only works for pure translational offsets between two sections, but it can sometimes be used iteratively with keyboard commands to get more general alignments. A protocol for aligning *By correlation* could be as follows:

Step 1.

Go to a section which will serve as the stationary reference, and lock it (Ctrl-L). Then page to the adjacent section that will be aligned to the stationary one. Unlock the section if necessary. Blend the two sections (spacebar).

Step 2.

Align the sections *By correlation* (Ctrl-).

Step 3.

Inspect the sections to determine if the alignment is acceptable. Use flickering (/) as needed. If alignment is acceptable, go to step 1 with the current section as the new stationary reference.

Step 4.

Estimate the rotational adjustment that would improve the alignment. Blend (spacebar) and then rotate the section (F3 or F4). Go to step 2.

By correlation uses only image intensity information visible in the Main Window to compute the correlation. Color and trace information is not used. Color images will be converted to grayscale intensities for the correlation computation. You can restrict the region used for the alignment by sizing the window and/or zooming into a subregion of the section. For example, to align based on just a single object, zoom in until that object fills the Main Window, flicker to update the back view, and then use *By correlation*. Note that the object to be aligned must be fully visible in both sections. Another way to limit the alignment region is to edit the domain boundary. Only the part of the image within the new domain boundary will be displayed and therefore used in the *By correlation* operation.

Tips for Aligning Sections

Alignment and reconstruction of serial sections is an art. Obtaining accurate results depends on many factors, including the experience of the user and her familiarity with the material to be reconstructed. A number of good practices that can improve the outcome are presented here. For a simple protocol see Chapter 1.

Section the Material Appropriately

Experience has shown that sectioning is a very important consideration. Sections should be as thin as is reasonable, and at least thin enough that lots of common features are available in adjacent section. There should be no missing sections. Too much disparity between sections and alignment is very difficult if not impossible. In addition, sectioning should be done such that as many objects as possible are cross-sectioned, with the long axes perpendicular to the sections. This ensures that cross-sectioned features rather than oblique features will be available for alignment.

Beware of Ambiguous Alignments.

Ambiguous alignments arise when multiple interpretations are possible for the profiles of an object (Figure 10.4). Aligning whole fields of objects with a single smooth transformation can minimize the distorting consequences of these local ambiguities, since the ambiguities of different objects will likely be different and oriented randomly with respect the plane of sectioning.




Figure 10.4. Objects with ambiguous alignments and the resulting distorted reconstructions. (top) A circular cylinder on the left is sectioned obliquely so that on the sections it appears as a sequence of shifted elliptical profiles. By mistakenly aligning the elliptical profiles an elliptical cylinder is reconstructed instead of a circular one. (bottom) A conical object on the left appears as circles of increasing diameter on sections. If sections are scaled to align these circles, then the cylinder on the right would be erroneously reconstructed rather than the original cone. Many other types of locally ambiguous alignments are possible and these may combine (e.g. a cone sectioned obliquely) to produce complex distortions.

Avoid Aligning the Same Features Every Time

To minimize the effect of ambiguous alignments it is also a good idea to choose distinct objects for correspondence points on each pair of sections. Repeatedly aligning the same profiles is more likely to produce a pattern of shift or scaling that builds up to high levels over many sections. For example, two objects can be moving towards each other as the sections are ascended. By aligning only these objects an artifactual scaling may be introduced. Randomly distributing correspondence points as you align pairs of sections can help by using more of the whole field for alignments.

Don't Over-Align.

Ignore the urge to make every alignment perfect, some local misalignments are to be expected since objects will move relative to each other. Don't rely just on blend to evaluate alignment, use flicker to see if there is a pattern of motion consistent with a whole field misalignment.

Avoid Quadratic Aligns.

As much as possible use *Linear* or *Rigid* Align commands as these introduce the least distortion. Use *Quadratic* and *Deformal* alignments sparingly. It may be a good idea to skip sections that can only be aligned by *Quadratic* or *Deformal* alignments and come back to them once all the other sections have been aligned.

Choose the Starting Point Carefully.

Don't start the alignment from an obviously distorted section. Choose a stationary reference section without folds or cracks, and with a uniform, high contrast. It may be helpful to choose a section in the middle of the series to minimize the differences introduced in the series relative to the reference. For montaged sections it may be helpful to start from a low magnification section that covers the entire montage area. This will allow the montage to be arranged accurately as well.

One Strategy

One strategy that can facilitate aligning two sections by point correspondence is to begin by picking two points at low magnification that should obviously be near each other. (These points can be deleted later in favor of more accurate ones.) With the section unlocked, select *Rigid* (Alt-T-N-R) to bring the sections closer together. Now more precise correspondence points will be easier to determine. Use Blend to compare the sections. Some parts of the blended sections will be blurry if the alignment is inadequate. Use the apparent motion produced by flickering to decide whether there are scaling differences between sections that require a *Linear* alignment for correction.

Now zoom in to one of the blurry regions to add more correspondence points at high magnification. Select the Stamp Tool and add a pair of points to this region. Drag near the edge of the window with the right mouse button depressed. This allows rapid panning to the next area of the section that needs correspondence points without changing tools. After panning and adding points to several areas, use the End key to return to the low magnification display. Select *Rigid* (or *Linear* if required). Compare the sections and continue adding points until the alignment is satisfactory.

Movements

Transformations

To allow independent movement of data elements, the position, orientation and scale of each domain and trace in the section is defined by a transformation. Every domain and

every trace has a transformation associated with it. A single domain or trace can be moved within the section by changing its transformation. The entire section can be moved by changing all the transformations together. Transformations and coordinates are coded in the XML files of RECONSTRUCT™. Details on the representation are provided in Fiala and Harris 2001; J Amer Med Informatics Assoc 8:1-16.

(<http://www.jamia.org/cgi/content/abstract/8/1/1>).

Each transformation maps trace points or image pixels into the section using a combination of basis functions in two dimensions. Essentially, each basis function represents an elementary motion such as translation, orientation, slant, scaling, deformation, or bending. By combining these movement components in different proportions, a complex remapping of the underlying data is possible. This remapping can be used to move an image on the section or to increase the size of a trace, for example.

Keyboard Increments

As described in Chapter 4: User Interface the keyboard performs diverse functions in the main window, but an important additional function is making movements. Table 10.2 summarizes the movements that can be performed by keystrokes. This information is also available by selecting *Key commands...* from the Help menu. A keyboard overlay can also be used as a quick reference (see Appendix B).

Generally all keyboard movements can be applied to traces, domains, or entire sections. If a trace or domain is selected then keyboard movements apply only to the selection. Selected elements are always moveable; locking the section does not prevent movement of selected elements. If no elements are selected and the section is unlocked, then keyboard movements are applied to the entire section including all traces and domains.

TABLE 10.2. Movement keystrokes in the Main Windows

Keystroke	Action
F1	Flip horizontally.
F2	Rotate 90 degrees clockwise.
F3	Rotate counterclockwise increment.
F4	Rotate clockwise increment.
F5	Slant/deform horizontally in counterclockwise direction.
F6	Slant/deform vertically in counterclockwise direction.
F7	Slant/deform vertically in clockwise direction.
F8	Slant/deform horizontally in clockwise direction.
F9	Scale to shrink horizontally.
F10	Scale to shrink vertically.
F11	Scale to enlarge vertically.
F12	Scale to enlarge horizontally.
Up Arrow	Move upward.
Down Arrow	Move down.
Left Arrow	Move to left.
Right Arrow	Move to right.

Each keystroke produces a fixed increment in the transformation that maps the data into the section. The size of increments can be modified by holding the Ctrl or Shift key down while pressing one of the movement keys of Table 10.2. Typically, Ctrl reduces the size

of the movement while Shift makes the movement larger. The size of the movement increments, including the effect of Ctrl and Shift keys, can be adjusted by using the Movement tab in the Series Options. If the arrows keys do not produce a movement even when the section is unlocked, try increasing the movement parameters in the Series Options.

The Movement tab includes a checkbox which changes the operation of the F5-F8 keys. When the box is unchecked the F5-F8 keys produce skew. Unchecking the box allows deformations in which one side of the object is scaled relative to the other side. Figure 10.5 shows the difference between slants and deformations for the movement directions defined by the default movement parameters.




Figure 10.5. The effects of slant and deformation movements centered on a rectangle region. Shown are the default movement directions. When Deformations are selected, F5-F8 produce the deformations shown in the top row. In the default mode, the slants (skews) shown in the bottom row are produced.

All keyboard movements are performed with respect to x- and y-axes aligned with the main window. Movements are initially made relative to the origin of these axes at position (0,0) in series units. This movement center can be changed by using the Type-In dialog as described below. The exception to this is the horizontal flip (F1) and the quarter revolution (F2). These movements are always made with respect to the center of the image or section being flipped or rotated.

While there is no vertical flip key command, the F2-F2-F1 sequence of two 90 degree rotations followed by a horizontal flip yields a vertical flip. Likewise, F2-F2 will produce a 180 degree rotation and F2-F2-F2 will produce a counterclockwise 90 degree rotation. These movements can also be produced directly using the Movement submenu of the Section menu.

Movement Type-In Dialog

The dialog invoked by selecting *Type-In...* from the Section Movement submenu can be used to specify a movement of selected traces or domains, or of the entire section. As with keyboard movements, the entire section is only moved when it is unlocked AND nothing is selected in the section.

The Type-In Movement dialog allows input of an incremental rotation, translation, scale, skew, and/or deformation that will be executed when the OK button is selected. This movement is made in the section coordinates relative the specified movement center. The center of the movement can be automatically set to the center of the selected

elements by using the Set to Center button. This movement center point is used for all subsequent keyboard movements as well.

Undoing Movements

Movements of a selected trace or domain, or of an unlocked section, can be undone by using the *Undo* item from the Section menu, or equivalently Ctrl-Z, as described in Chapter 6: Sections. Each incremental movement is pushed onto the undo stack, so that multiple undos will back up through a movement sequence. The state of the section prior to any changes can be recovered by the *Reset* command.

To move a domain it must first be selected from the section by using the domain list or the domain selection tool (Chapter 7: Domain Images). An undo stack is created for a selected domain so that keyboard movements can be undone. Although the undo stack has a limited depth, the original position of the domain is retained at the bottom of the stack. Backing up through the entire stack will eventually lead to the original position, the one prior to selection of the domain. This starting position can also be restored by using the *Reset* item from the Section menu. Once the domain is deselected, the movements are no longer undoable. However the entire domain change can be undone by using *Undo* on the section.

In rare circumstances it may be necessary to completely reinitialize the domain transformation, returning the domain to the origin just as though it is newly imported. This can be accomplished by using the *Transform* item from the Reinitialize submenu of the Domain menu. This option is only available after the domain is selected from the section.

Repeating Movements

An individual movement can be repeated by using the *Repeat* item (Ctrl-R) from the Movement submenu. This allows a movement made on one section to be repeated on other sections in the series. Similarly, a movement made on one trace or domain can be repeated on other selected traces or domains. The movement to be repeated is the last one made by keyboard input or the Type-In Movement dialog. In addition, a movement sequence or an alignment command may also be repeated.

The entire transformation of a selected element can be applied to another object using *Repeat*. To copy a transformation use *Record from selected* in the Movement menu. This will set the transformation of the selected trace or domain into the buffer that stores the last movement. Selecting a new trace or domain and using *Repeat* will apply the buffered transformation to the selected item. Note that new transformation will be a combination of the old one and the buffered one, not just the buffered one alone.

Propagate

A repeated movement can also be applied to a range of sections within the series by using the *Propagate...* command from the Movement submenu. This is particularly useful for repeating a change made to one section on the rest of an aligned series. Such an operation allows a misalignment at one point in the series to be easily fixed. For example, suppose a small offset existed in the middle of the series but all the sections above the offset were aligned and all the section below the offset were also aligned. Fixing the offset by moving the upper section would now cause this section to be misaligned with the next higher section. By propagating the same movement to all higher sections, the entire series could be put into alignment.

NOTE: There is no undo for Propagate...

The *Repeat* menu item is only available when an unlocked section or a selected item is available for movement, whereas the *Propagate...* operation can be applied anytime after a movement has been made. This is because the *Propagate...* operation applies to entire sections and the sections do not have to be unlocked first.

Repeating Movement Sequences

When the keyboard is used to move a trace or domain, only the last keyboard increment is saved for repetition with *Repeat* or *Propagate...*, but it is sometimes desired to repeat a sequence of keyboard movements. A sequence can be captured using the Record submenu of the Movement menu.

Selecting *Start* from this submenu clears the recording buffer and prepares for accumulating movements. Any subsequent keyboard, *Type-In...*, or mouse drag movements are added to the recording until the *Stop* item is selected from the Record submenu. Note that alignments or repeated movements are not added to the recording. After *Stop* the recorded sequence is moved to the last movement buffer where it can be applied using *Repeat* or *Propagate...*



Chapter 11: Objects

An object is the set of all traces in sections with index greater than zero that have the same name. Thus, an object is defined by a trace name not by the color or other shared attributes of traces. Traces with different names belong to different objects.

The Object List

The list of all objects in the series is reported in the Object List window (Fig. 11.1). This list of objects is regenerated from the section data every time the Object List window is opened or refreshed. Since each section file in the series is opened and scanned for object traces, generating the Object List may be slow if there are large numbers of sections or access to the section files is slow. In addition, if a new object is created by tracing in the sections, it will not immediately appear in the Object List. Rather the Object List window will need to be closed and reopened (from the Object menu) in order for the new object to be recognized.

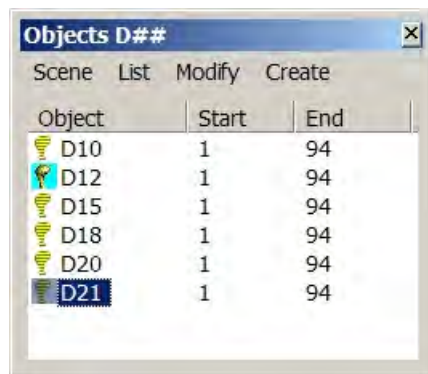


Figure 11.1. The Object List window. Objects are listed along with the range of sections in which traces of the object were found. The Scene and Modify menus appear above the list. The object "D21" has been highlighted in the list. The object "D12" has been added to the 3D Scene as indicated by the shaded icon with the cyan background. None of the other objects are in the scene, as denoted by the yellow icon with lines representing section profiles. The rows of the list were limited to objects with names that consist of the letter D followed by 2 digits, as indicated by "D##" in the title bar.

Multiple objects can be selected in the list by using the control or shift keys in combination with the left mouse button. Pressing the Enter key or double-clicking the mouse causes the selected objects to be added to the scene. Selected objects can also be added to the scene by using the *Add to scene* item from the Scene menu. When an object has been added to the scene, the object's icon in the Object List is changed to a shaded figure with a cyan-colored background (Fig. 11.1). By this mechanism, the Object List indicates which objects are in the 3D Scene. This allows particular scene objects to be selected and removed. Selected scene objects are removed from the scene using the *Remove from scene* menu item from the Scene menu.

The Object List window can be made to include 3D measurements of objects, such as the volume and surface area, or counts of the number of component traces of the objects. This is accomplished by checking the appropriate boxes in the Lists tab of the Series

Options before opening the Object List. When the Object List is filled, all sections are scanned for traces that are grouped into objects, so the list must be refreshed whenever there are changes to the trace data for the measurements to be accurate. (For details on the volume and surface area calculations see Chapter 9: Calibration and Measurement.)

Modifying Objects

The Modify menu at the top of the Object List window provides a convenient set of commands for making changes to ALL of the traces of an object. While convenient, there is no global undo for these multi-section commands.

Changing Object Attributes

Changing attributes of objects, changes the attributes of every trace of the object in all sections. This change is performed on the currently highlighted objects in the Object List by using the *Attributes...* dialog of the Modify menu. Note that color of an object can be temporarily changed in the 3D Scene without modifying the traces. This is done by the *Color/Transparency...* dialog of the Scene menu. Scene color changes are forgotten whenever the object is removed from or added to the scene. Scene changes are discussed in detail in Chapter 12: The 3D Scene.

The *Attributes...* command permanently modifies all the traces of the selected objects in all sections of the series. For example, to change the color of an object, select the object in the Object List and then open the *Attributes...* dialog. The Object Attributes dialog will appear with a title indicating that 1 object will be changed. Fields of the dialog will be grayed to indicate that no changes will be made to these fields unless a change is specifically entered in the dialog. Click on the Border button and select a color for the object. Close the color dialog and click the Apply button to change the border color of all the traces of the object in all sections.

NOTE: There is no undo for changing object attributes!

Changes made to the object traces will not be reflected in the Object List or 3D Scene immediately. If a change is made to the color of an object's traces and that object is already in the 3D Scene, the object will need to re-added to the scene to show the color change. Obviously, a simpler way to change the color in the scene is to use the *Color/Transparency...* command from the Object List Scene menu.

Renaming Objects

Objects are renamed by changing the name attribute of the object's traces. Unlike the *Attribute...* dialog, the *Rename...* dialog allows the operation to be applied to a subset of the object's traces. The operation is carried out over a range of sections specified in the dialog. The default section range will apply the name change to all traces of the object. To limit the name change to part of the object, entering a narrower section range in the dialog.

Note that when an object is renamed its entry in the Object List is no longer valid. Furthermore, such name changes are not immediately reflected in the list. The old object name is deleted from the list, but the list has to be regenerated to see the new object names. For this reason, after *Rename...* is complete a message will appear suggesting that any name changes will only be revealed by refilling the Object List.

By changing an object's name attribute, the object can be broken into its component traces. Suppose, for example, a set of vesicles were labeled with point traces all given the name "VESICLE". These traces would form a single object in RECONSTRUCT™. You can make each vesicle be its own object by using the *Rename...* dialog to change the Name string from "VESICLE" to "VESICLE+". This will rename the individual traces to "VESICLE1", "VESICLE2", "VESICLE3", and so on. *Refresh* the Object List, and you'll see a separate object for each vesicle.

To group the vesicles back into a single object, highlight all the vesicle names in the Object List and again use *Rename...* Change the Name string to "VESICLE" and press OK. Now all vesicles are back to being a single object. In general, anytime there are name changes to objects the Object List will need to be refreshed to reveal the changes.

Object renaming can also be used to divide an object into its component sections. In this case, rename the object using the special character \$ to add the section number (Chapter 8: Traces). For example, rename "myObject" to "myObject\$" to create a set of objects: "myObject1", "myObject2", "myObject3", etc. Each of these objects will correspond to one section of the original object. Now it is possible to rename a group of sections into one subpart. For example, "myObject1" through "myObject9" could be renamed to "myObject_part1" to create a slice of the original object from sections 1 through 9.

Hiding and Unhiding Objects

The *Hide* command in the Object List Modify menu allows all the traces of highlighted objects to be hidden from display when paging through the section. The *Unhide* command has the opposite effect. Be aware, however, that in either case the result can be overwritten by the Hide/Unhide flags in the General Tab of the Series Options.

Simplifying Objects

Selecting *Simplify...* from the Modify menu of the Object List will apply the operation to any highlighted objects in the list. Every trace of an object that is not already marked as "simplified" will have the simplification on-demand operation applied to it as described in Chapter 8. The "Simplify on demand" resolution specified in the Series Options will be used to determine the new traces for the object.

When traces are already marked as simplified, as is typically the case for most objects, the *Simplify...* command will have no effect. Therefore, to simplify an object you must first uncheck the simplified flag on all traces. This can be done by opening the Object List *Attributes...* dialog and clearing the simplified box.

NOTE: There is no undo for simplifying objects!

Deleting Objects

All the traces associated with a set of objects can be deleted from the series by using the Delete key on the Object List or selecting the *Delete...* item from the Modify menu. Deleting objects is permanent; the deleted traces cannot be recovered. For this reason any request to delete objects will be challenged by a warning dialog.

NOTE: There is no undo for deleting objects!

Creating Objects

The Create menu provides a set of commands for creating z-traces and objects from existing objects. In the later case, traces are created on multiple sections as needed to fulfill the command.

Copying Objects

Selecting *Copy* from the Create menu of the Object List will duplicate all the traces for every highlighted object in the list. The duplicate traces have the prefix "Copy of" added to the name. While there is no undo for this multi-section operation, the copied object can be easily deleted.

The *Copy* command will create a duplicate set of traces on the sections. These traces will exactly overlay the existing traces of the object, so be careful about selecting and editing traces of copied objects as changes can easily be made on the wrong object.

Smoothing Objects

The *Smoother version* command will copy every highlighted object in the Object List and apply the *Smooth* operation (Chapter 8: Traces) to every copied trace. The smoothed traces have the suffix "_smooth##" added to the name, where ## is replaced by the filter length used during smoothing. The filter length specified in the Autotracing tab of the Series Options will be used during smoothing. Since smoothing may modify the traces in such a way that loops are created, smoothed traces are marked as "unsimplified". Thus the *Smoother version* command can be followed immediately by the *Simplify...* command.

Note that the *Smoother version* command does not smooth the shifts between sections caused by misalignments of domain images. Smoothing of these section-to-section offsets requires the generation of z-traces along the object.

Creating Z-traces at Midpoints

A z-trace can be automatically generated along the midpoints of an object by using the *Z-trace at midpoints* command from the Create menu. The midpoints are the average of the maximum and minimum extents in the x- and y-directions of the traces of the object on each section. If an object has only one trace on a section, the z-trace point will be placed at the midpoint for that trace. If an object has multiple traces on a section, the midpoint will be placed at the center of the group.

Calculating the Distances Between Objects

A distance between two objects is defined to be the minimum 3D distance between all the line segments of all the traces of the objects (Chapter 9: Calibration and Measurement). If object traces intersect the distance is zero. This distance does not take into account whether one object is inside another. Rather, the distance between boundaries of objects is computed.

Three-dimensional distances between objects are listed in the Distances List window (Chapter 4: User Interface). Distances are calculated only when the Distances List is (re)filled. The list needs to be regenerated whenever there are changes to the data.

Object names and calculated distance (in series units) are displayed for each pair of objects that matches the list's limit strings. To limit the number of distances that are calculated, set the limit strings in the Lists tab of the Series Options BEFORE selecting *Distances...* from the Object menu. With the default settings, the distances between every object in the series and every other object in the series will each be computed twice: once with the object in the left column and once with the object in the right column. Use more specific limit strings to speed up the list calculations. For example, to generate only distances from numbered vesicles "vesicle1", "vesicle2",... to a synapse named "mySynapse", enter "vesicle*" for the first limit string and "mySynapse" for the second limit string.

3D Object Representations

An object is added to the 3D Scene (Chapter 12: The 3D Scene) using one of several possible 3D representations. Which representation is generated is determined by the selection in the Generate box of the 3D tab of the Series Options dialog (Fig. 11.2). The 3D representation is generated only for the traces of the object in the section range indicated.

WARNING: Closing the series during 3D generation will crash RECONSTRUCT™!

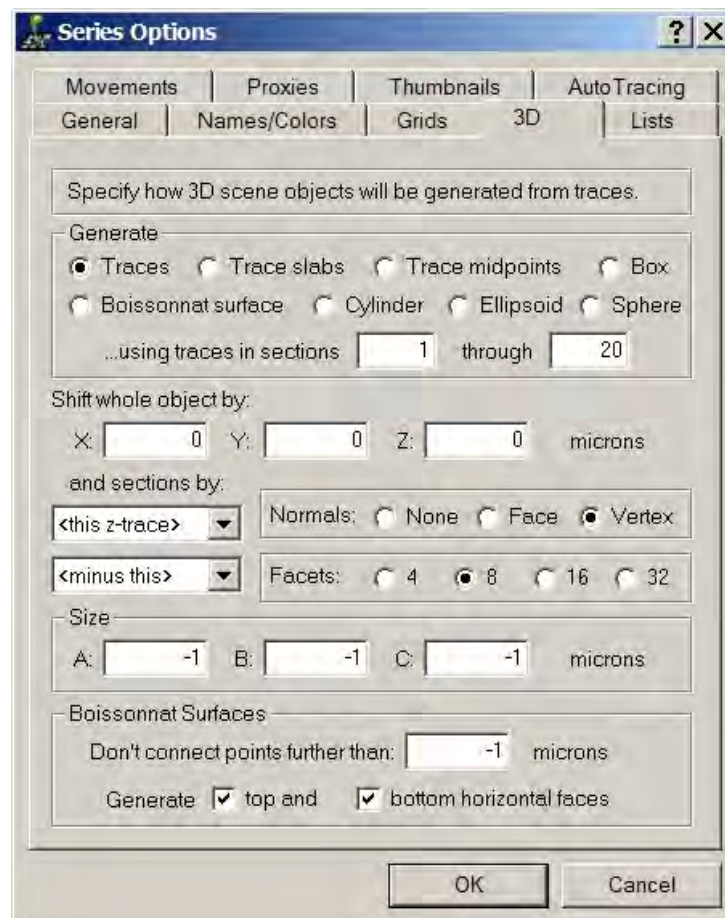


Figure 11.2. The 3D tab of the Series Options determines how an object will be represented in the 3D Scene.

Traces

The Traces option in the Generate box (Fig. 11.2) produces 3D lines corresponding to the traces defined in the sections. The lines are colored using the trace border color and spaced according to the Z-distance calculation for each section (Chapter 9: Calibration and Measurement). For this type of representation, the only additional 3D options used are the section range parameters and shift parameters.

Trace Slabs

The Trace Slabs option represents the object by extruding each trace by the thickness of the section. A slab is generated for each trace of the object, with no attempt to connect or smooth between traces. This representation may be advantageous for certain types of objects, particularly where interior holes need to be more faithfully represented. Reversed orientation (counterclockwise) traces will appear as holes inside the slabs. For example, an object drawn as two concentric circles with the inner circle reversed (using the Trace menu item), will appear as a ring in the 3D Scene with Traces Slabs.

The method used for generating Trace Slabs is Delaunay triangulation of the trace vertices followed by boundary completion, as in the first step in the Boissonnat surfacing algorithm. For Traces Slabs to generate a reasonable result, follow the tracing rules defined below for Boissonnat surfaces. In particular, traces must be properly oriented without intersections. This can be most easily accomplished by keeping the "Automatically simplify..." option turned on when drawing (Chapter 8: Traces). Traces Slabs can also be used to evaluate the quality of the input to the Boissonnat surfacing algorithm. Problems evident in Trace Slabs will be reflected as defects in the Boissonnat surface.

The thickness of slabs is the thickness of sections, unless the user specifies a different thickness using the Size A parameter in the Series Options 3D tab. When Size A is positive, all slabs will be as thick as the specified value. When Size A is zero, the object will appear as a series of filled profiles arranged at the correct z-distances. Entering a negative value for Size A (default) restores the generation of slabs with the same thicknesses as the sections.

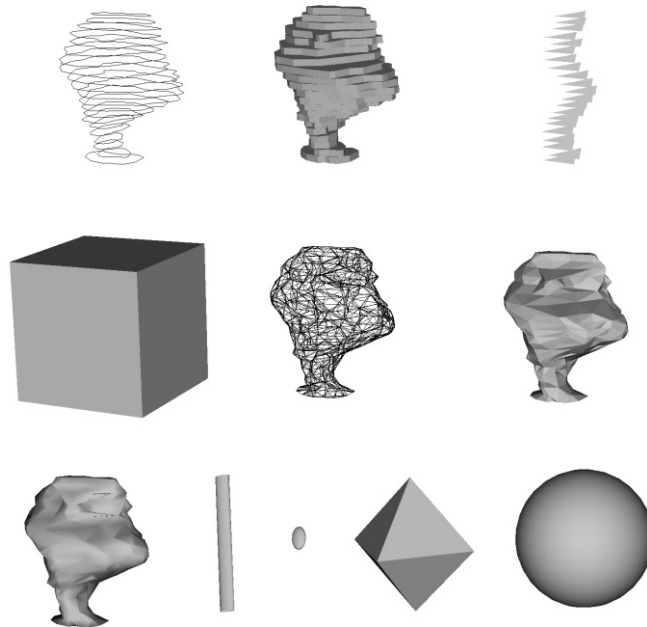


Figure 11.3 Examples of 3D representations of one object. The object is a dendritic spine about 0.5 microns tall, reconstructed from serial section electron microscopy. All images were saved from the 3D Scene at the same magnification and viewing angle. A. Traces. B. Trace Slabs. C. Trace Midpoints with a Size A and B set to 0.25 microns. D. Box (Size: $A=1, B=1, C=1$). E. Boissonnat surface drawn with front and back faces in wireframe mode instead of filled. F. Boissonnat surface with face normals. G. Boissonnat surface with vertex normals. H. Cylinder with the length (Size A) set to 0.5 microns. I. Ellipsoid (Size: $A=-1, B=-1, C=-1$). J. Sphere with a faceting density of 4 and face normals. K. Sphere with a faceting density of 32 and vertex normals.

Trace Midpoints

To represent Trace Midpoints a single triangular face is centered on each trace of the object. Each triangle lies at the z-distance of the section, and is centered at the midpoint of the trace extent in the x- and y-directions within the section. The width and height of each triangular face is equal to the width and height of the corresponding trace, provided the Size A and B parameters are zero or negative. If Size A is positive, the width of each triangular face is set to Size A. If Size B is positive, the height of each triangular face is set to Size B.

The representation is a single indexed face set object with each disjoint face centered on each trace. This can be used to represent trace positions in 3D. But more important, the object can be used as a distribution object when exported to 3D rendering software like 3D Studio MAX. 3D Studio MAX allows an arbitrary 3D shape to be centered on each face.

One application of Trace Midpoints would be to place a small sphere at each trace. If each trace was used to identify the location of a cell or vesicle, say, the sphere distribution would visualize the cells or vesicles in 3D. There are two reasons to use the Trace Midpoints/distribution method rather using Sphere representations directly. First, Trace Midpoints representations are much smaller, an important consideration when dealing with hundreds of spheres. Secondly, a single object name can be used for all

spheres, whereas if a Sphere representation were to be generated for each trace, each trace must have a unique name and be a unique object.

Box

The Box representation is centered at the midpoint of the object's extent in the x-, y-, and z-directions. The sides of the box are parallel to the axes. If Size A, B, or C parameters are positive, these values are used to dimension the box in the x-, y-, and z-directions, respectively. When the Size parameters are zero or negative, the box is sized to the extent of the object.

Hint: The Box object can be used to generate scale bars or cubes for the scene. Simply select the Box representation and set the desired Box dimensions with the Size parameters. Then select an object defined in the series and add it to the scene. The scale Box will be centered on the midpoint of the object extent in 3D. If necessary, the scale Box can be shifted to a different location in the scene by changing the "Shift whole object by" parameters (see below) and re-adding the object to the scene.

Boissonnat Surface

A Boissonnat Surface is generated from a 3D Delaunay triangulation of the traces using the algorithm invented by Jean-Daniel Boissonnat (Boissonnat, 1988). In general terms, the surface is constructed of triangular surface patches where an edge of each triangle lies on the traced profile. The Boissonnat Surface preserves the vertices created during tracing, although additional vertices may be added to the final reconstruction. The Boissonnat algorithm attempts to generate a maximum solid with no singularity. While this representation may not be ideal in some instances, the method can handle multiple traces per section in a general way that produces reasonable results for objects that are branched.

When the Boissonnat Surface is not reasonable, there are several things to try to improve the result. First and foremost, make sure that the traces used to generate the surface follow the rules outlined below. Commonly, counterclockwise traces or traces with internal loops cause the algorithm to generate invalid triangles outside of the objects intended surface. Since a 2D Delaunay triangulation of the traces is the first step in the surfacing algorithm, the Traces Slabs output can be used to see whether the problem lies with the traces or with the surfacing algorithm itself. If Trace Slabs does not produce a filled representation similar to Traces, locate and correct the traces that are invalid (Fig. 11.4). It may help to limit the section range in the 3D object to narrow down which traces are problematic.




Figure 11.4. Examples of invalid object traces for Boissonnat surfacing. A. A trace with an internal loop is invalid. Avoid this scenario by using the *Automatically Simplify* option when tracing. B. One trace intersecting another trace for the same object is invalid. Merge traces to eliminate this problem. C. A counterclockwise trace without a surrounding clockwise trace is invalid. The inner trace should be counterclockwise to mark the hole boundary and the outer trace should be clockwise to delineate the exterior of the object. Use *Reverse* to fix the orientations.

Draw closed traces

Boissonnat surfacing is designed to surface closed traces, however the algorithm will attempt to incorporate open traces as well. In the best case, open traces will be partially filled in the clockwise part of the traces. In general, however, the results with open traces will be difficult to control. Straight line traces will be ignored altogether since these are not capable of forming a solid connection with the rest of the object.

Use clockwise traces to define outer boundaries

Clockwise traces represent the exterior of the object. Clockwise traces must be wholly outside of all other clockwise traces of the same object in the section. The exception to this rule is when the clockwise trace is completely contained within a counterclockwise trace. A clockwise trace can occur inside another clockwise trace only if it is fully enclosed by a counterclockwise trace of the same object.

Use counterclockwise traces to define holes

Counterclockwise traces represent an interior (hole) boundary. Traces representing holes must be wholly contained within clockwise traces (Fig. 11.4.C). When traces are automatically simplified during tracing, they are also made clockwise. To change a simplified trace to counterclockwise use the *Reverse* command from the Trace menu.

Trace without loops or self-intersections.

As depicted in Fig. 11.4.A, traces with self-intersections do not produce a valid surface because of the ambiguity between interior and exterior regions. To avoid loops, all tracing should be performed with the "Automatically Simplify" option turned on (Chapter 8: Traces). In addition, traces should be drawn with one of the tools designed to produce closed profiles, rather than drawing open traces and then changing the trace attribute to closed.

Do not intersect traces of the object.

Traces within a section must not overlap as depicted in Fig. 11.4.B. Such traces also produce ambiguity about interior and exterior regions of the object. However, traces of one object may overlap traces of another object without problems because Boissonnat surfaces are generated independently for each object.

ACOS Domain Error.

If this error message appears when generating a Boissonnat surface, then probably the traces of the object contain intersections on one or more sections. **Check for overlapping and/or multiple copies of traces.** Use the section range parameters in the 3D tab of the Series Options to narrow down which section(s) contain the intersecting traces.

Cylinder

A Cylinder may be generated to represent an object. The midpoint/center of the Cylinder is given by the centroid of the trace points.

$$(x_c, y_c, z_c) = \frac{1}{n} \sum_{\text{all trace points}} (x, y, z)$$

where n is the total number of trace vertices, and the vertices of the traces are at the positions (x,y,z) in 3D space. The axes of the Cylinder are determined from the scatter matrix of the trace points:

$$\text{scatter matrix} = \frac{1}{n} \sum_{\text{all trace points}} (x, y, z)^t (x, y, z) - (x_c, y_c, z_c)^t (x_c, y_c, z_c)$$

The eigenvector corresponding to the largest eigenvalue of the scatter matrix determines the central axis of the Cylinder. The other two eigendirections specify the directions of the cylinder's diameters. These directions are scaled by the eigenvalues to get the final dimensions of the Cylinder.

The dimensions of the Cylinder can be explicitly defined using the Size parameters. Size A sets the length of the cylinder. Size B sets diameter along the secondary axis, while Size C sets the diameter along the minor axis.

Ellipsoid

The Ellipsoid representation is essentially the same as the Cylinder representation except the object is closed at the ends of the principal axis. The axes of the Ellipsoid are computed from the eigenvectors of the scatter matrix of the trace vertices, and the Ellipsoid is centered at the centroid of the trace vertices. The dimensions of the ellipsoid are determined by the scatter matrix eigenvalues, unless positive values are entered in the Size parameters.

Information about the calculated axes and dimensions of cylinders and ellipsoids can be obtained by looking at the comments in an exported VRML file for the object (Chapter 12: The 3D Scene).

Sphere

The Sphere is centered at the midpoint of the object extent in the x-, y-, and z-directions. The diameter of the Sphere is set to the maximum extent in x, y, or z. In effect, the Sphere is sized to the maximum dimension of the object. When the object is a single trace on a single section, this representation gives a reasonable point representation of the trace. Alternatively, a Sphere object can be given an explicit diameter by entering a positive value in the Size A box of the Options tab.

3D Parameters

In addition to the type of representation and the section range, several other parameters can be used to adjust the object's 3D representation (Fig. 11.2). Not all parameters are applicable to every type of reconstruction.

Shift whole object by

These parameters define an offset in the x-, y-, and z-directions to be applied to the 3D representation. These parameters let you position any type of 3D representation anywhere in the scene, but they should be used carefully. In particular, set the parameters to zero when done shifting objects, so that future reconstructions from the series will not be mistakenly shifted.

and sections by

This option allows different shifts to be applied at each section. The section-to-section shifts are represented by the distance between two z-traces. At each section for which both z-traces have a point defined, the differences in the x- and y-directions between the two points are computed to determine the shift within that section applied to the 3D representation. Z-trace shifts are only applied to Traces, Trace Slabs, Trace Midpoints, and Boissonnat Surface representations.

In combination with the smoothing of z-traces, z-trace shifts can be used to smooth the misalignments between sections using a three-step process. First, create a z-trace that follows the misaligned path of the object across sections. This z-trace should consist of one point entered on each and every section of the object. You may be able to generate the z-trace automatically using the *Z-trace at midpoints* command from the Object List menu, but the z-trace so created would only be appropriate for an object with a fairly regular, convex profile. A more general method for creating the z-trace is to manually enter it by following the center or edge of the main structure through the series. After creating a z-trace that follows the jagged object misalignment, smooth the z-trace using the *Smoother version* command from the Z-trace List menu. The difference between these two z-traces, the smoothed one and the original jagged one, defines the section-by-section offsets necessary to smooth the object reconstruction, so select the smoothed z-trace as <this z-trace> and the original jagged z-trace as <minus this>. Click OK and simply add the object to the scene to reveal the smoothed reconstruction. The effect is demonstrated in Fig. 11.5.




Figure 11.5. Reconstructions of the same data with and without smoothing using the difference of z-traces. The top object shows the original reconstruction, while the bottom one shows the dramatic effect of the smoothing operation. Note that the smoothing shift is applied to any objects reconstructed, in this case including the synapses on the surface of the dendrite.

Normals

All representations except Traces define surfaces. To properly reflect light and produce a shaded 3D shape, each part of the surface should have a surface normal vector. For Trace Midpoints no normals are generated because these representations use emissive color rather than diffuse color (see Chapter 12: The 3D Scene). For Trace Slabs face normals are generated. For all other representations, the Normals option allows normals to be generated for faces, vertices, or eliminated altogether. Eliminating normals will result in an unrecognizable 3D shape because there will be no shading. Face normals generally produce a faceted appearance to the 3D object. Vertex normals produce a smoother shading transition between faces. Consequently vertex normals are best except when the object really does have sharp corners (e.g. the Box object).

When exporting an object for use in another rendering program such as 3D Studio MAX, use face normals rather than vertex normals. This will enable the 3D rendering software to calculate the best normals for the final rendering.

Facets

Facets refer to the number of faces generated for the object. A low number of facets means fewer faces in the object representation. A high number of facets will produce a smooth surface representation. Facets only apply to Cylinders, Ellipsoids, and Spheres. For a Cylinder, it specifies the number of facets around the circumference. For instance, a value of 4 produces a 4-sided cylinder. Each facet is made up of 2 triangular faces. For Spheres and Ellipsoids, the total number of faces is $f_{\infty}(f-2)$, where f is the facets number.

Size

The size parameters allow you to specify the dimensions of the Trace Slabs, Trace Midpoints, Box, Cylinder, Ellipsoid, and Sphere objects. Size A specifies the width in the x-direction for Boxes and Trace Midpoint faces. For Trace Slabs, Size A can be used to override the section thickness parameter that normally specifies the thickness of slabs. For Cylinders, Ellipsoids and Spheres, Size A specifies the length or diameter along the principal axis. Size B specifies the height in the y-direction for Boxes and Trace Midpoint faces, and the diameter of secondary axis of Ellipsoids and Cylinders. Similarly, Size C specifies the diameter of the third axis of Ellipsoids and Cylinders. In all cases, the size parameters can be disabled by entering a negative number in these boxes. When sizes are negative, objects are sized according to the traces, as described above for each object type.

Boissonnat Surfaces

The parameters in the Boissonnat Surfaces box are only used when the Boissonnat Surface generation type is chosen. The surfacing algorithm attempts to connect vertices in one section to the nearest vertices in the adjacent section. You can limit these connections by setting a positive value in the "Don't connect points further than" box. When the value is negative or zero no limits are applied. A negative value is recommended, although unwanted connections can sometimes be eliminated by a positive value that is not so small that it interferes with reasonable connections.

Faces closing the ends of the object are generated by checking the two check boxes at the bottom of the dialog. Unchecking the boxes will leave open ends at the first and last sections of the object.

Z-Trace List

Open the Z-Trace List window by selecting *Z-Traces...* from the Object menu. If any z-traces have been defined for the series, a list will appear showing the z-traces arranged alphabetically by name (Fig. 11.6). Depending on the settings in the Lists tab of the Series Options, the range of sections and the length of each z-trace can be displayed in the list. In addition, each z-trace may be given a descriptive annotation of up to 127 characters. The entire list can be saved to a comma-separated values (.csv) file using the *Save...* command from the List menu. This file format can be readily imported into any spreadsheet application.

Z-Trace	Start	End	Length
Z D10length	1	94	4.66536
Z D18length	1	94	5.14913
Z D1length	1	94	4.83408
Z D20length	1	94	5.90529
D21length	1	94	5.43639
Z D3length	1	94	5.15151
Z D6changed	1	94	4.66478

Figure 11.6. The Z-Trace List. Every z-trace in the series is listed along with the range of sections which each trace spans. The Scene and Attributes menus appear above the list. The z-trace "D21length" is highlighted. The z-trace "D18length" has been added to the 3D Scene as indicated by the shaded icon with the cyan background. None of the other z-traces are in the scene, as denoted by the red Z icon.

Z-traces can be added or removed from the 3D Scene using the Scene menu of the Z-Trace List. Note that all z-traces should be given unique names so they can be independently displayed in the 3D Scene. A group of z-traces in the list can be renamed using the + special character to get a sequence of uniquely named traces. For example, renaming a list of z-traces to `length05+` would produce `length06`, `length07`, `length08`, etc.

Modifying Z-traces

The Modify menu at the top of the Z-trace List window allows the color, name, or annotation of highlighted z-traces to be changed. These changes permanently alter the attributes of the trace, and the changes are immediately reflected in the 3D Scene window if it is open.

Generating Z-traces

The Create menu at the top of the Z-trace List window provides commands for creating new objects or z-traces from existing z-traces. Z-traces may be copied and smoothed, and grid traces may be placed on the sections along the points of the z-trace.

Smoother Versions

Z-trace smoothing uses a moving average filter on the (x,y) positions without adjusting the positions in the direction perpendicular to the sections. That is, z-trace points remain on the sections in which they were originally drawn, but are shifted in the plane of the sections to bring them closer to adjacent z-trace points. The number of adjacent points that influence the position is determined by a filter length parameter requested when the command is executed. This filter length is independent of the filter length parameter in the Series Options Autotracing tab, but it works the same way. (See discussion of smoothing in Chapter 8: Traces).

Grid at Z-trace Points

The command for creating a *Grid at Z-trace points* places a copy of the current grid traces at each and every point of the highlighted z-traces. This command can be used to place a frame or circle around an object by first using the *Z-trace at midpoints* command in the object list to generate a z-trace. The Grid traces that will be used are defined in the Grids tab of the Series Options. Grid trace attributes are generated using the default settings in the Names/Colors tab.

Trace Substitutions

Because the grid trace options are so general, the *Grid at Z-trace points* command can be used to perform trace substitutions as follows. Suppose you have drawn a set of traces using circular shapes with the Stamp Tool, but you later decide that you want these traces to all be star-shaped instead. First, draw the desired star shape and put it on the trace clipboard (Ctrl-X). Then in the Grids tab of the Series Options select Clipboard traces for the element type. Make sure the element number is set to X=1 and Y=1 so only one copy will be generated. Next open the object list, select the names of all the traces that will be substituted (assuming each trace is uniquely named), and use the *Create Z-trace at midpoints* command to generate z-traces that define the location of each trace. Then open the z-trace list and select those newly created z-traces. Create a *Grid at Z-trace points* and the clipboard trace will be centered on every original trace. All that remains is to delete the original traces and the z-traces, leaving the new shapes in place of the old ones.



Chapter 12: The 3D Scene

The 3D Scene window displays the 3D representations of objects generated from the Object List. The window can be opened or closed by selecting *3D Scene...* from the Object menu. The 3D Scene window will be automatically opened whenever an object is added to the scene from the Object List.

The 3D Scene window can be resized and positioned arbitrarily on the desktop, but it cannot be minimized or moved behind the main window. To hide the scene temporarily, close the 3D Scene window. This can be done by selecting *Hide* from the 3D Scene window's Scene menu, or clicking on the X box on the right side of the window's title bar. Closing the 3D Scene window does not erase the scene. Reopening the window will show the same scene, although the viewing angle will be reset to the default position.

OpenGL Rendering

The 3D Scene is rendered using OpenGL (www.opengl.org). The OpenGL libraries provide a software interface to the computer's graphics hardware. These libraries must be properly installed on the computer to open and view the 3D Scene. OpenGL is shipped standard with every Windows (95/98/2000/NT/XP/Vista) computer and should work with RECONSTRUCT™ without any configuration. However, if opening the 3D Scene generates an error or crashes the program, a likely culprit is the OpenGL/hardware interface. An updated video driver might be available to fix the problem.

The speed and quality of the 3D Scene rendering is a result of the OpenGL implementation of the graphics hardware, a.k.a the video display adaptor. Most Windows computers are designed for 3D video games and have excellent OpenGL graphics displays appropriate for RECONSTRUCT™. For example, the 32Mb ATI Radeon VE graphics adaptor (www.ati.com) was used during code development and gave nice results. To render 3D scenes with even better resolution and quality, use third-party software specifically designed for that purpose such as 3D Studio MAX (www.discreet.com).

All OpenGL surfaces are generated as collections of triangular surface patches, as described in Chapter 11. This allows the surfaces to be quickly rendered and animated using common hardware graphics capabilities. It also allows the scene to be saved as a general indexed face set (in VRML format) for export to many other programs.

Composing a Scene

3D representations of objects are added and removed from the scene using the Scene menu in the Object List window (Chapter 11: Objects). The type of 3D representation generated is determined by the parameters set in the 3D tab of the Series Options. The positions of objects in the scene can also be modified by using the "Shift whole object by" parameters. Changes made in the 3D Options tab will only be realized after the object is regenerated for the scene.

Adding Objects

Objects are added to the scene from the Object List window. When an object is added to the scene, the display parameters of the scene are not modified. This allows objects to be added to and/or removed from a particular view or rotating scene. It is also possible that the added object is outside of the current view. To see objects newly added to a scene it may be necessary to zoom out or reset the scene view point using the *Reset* item in 3D Scene's View menu.

Note: Closing the series while objects are being added to the 3D Scene will cause RECONSTRUCT™ to crash!

When an object is added to the scene, it replaces any existing scene object with the same name. Any modifications to the object's in-scene attributes (color, transparency, etc.) will be lost when the object is re-added to the scene. The 3D representation or position of an object can be modified by specifying new parameters and then adding the object to the scene again. It is not possible to have two different representations for an object (e.g. traces and ellipsoid) in a scene, except by generating one representation and then renaming the object and generating the second representation without clearing the scene.

Removing Objects

Specific objects selected in the Object List are removed from the scene by the *Remove from scene* menu item in the Object List window's Scene menu. All objects can be removed the scene by selecting *Clear* from the 3D Scene window's Scene menu.

The objects in the scene are not removed when a series is closed and a new series opened. This allows objects from one series to be rendered with objects from another. If this feature is not desired simply *Clear* the scene whenever a new series is opened. Scenes are not preserved when the application is terminated. Closing the program requires the scene to be recreated by adding the objects again.

Changing Object Appearance

After an object has been added to the scene, certain aspects of the object's appearance can be modified by opening the *Color/transparency...* dialog from the Scene menu of the Object List window (Fig. 12.1). Only the highlighted objects in the Object List will be modified. The number of highlighted objects is indicating in the title bar of the Changing Scene Objects dialog. Every object will receive the same changes specified in the dialog when the Apply button is selected. Note that the colors of the trace data are not modified by changing the appearance of the object in the scene.

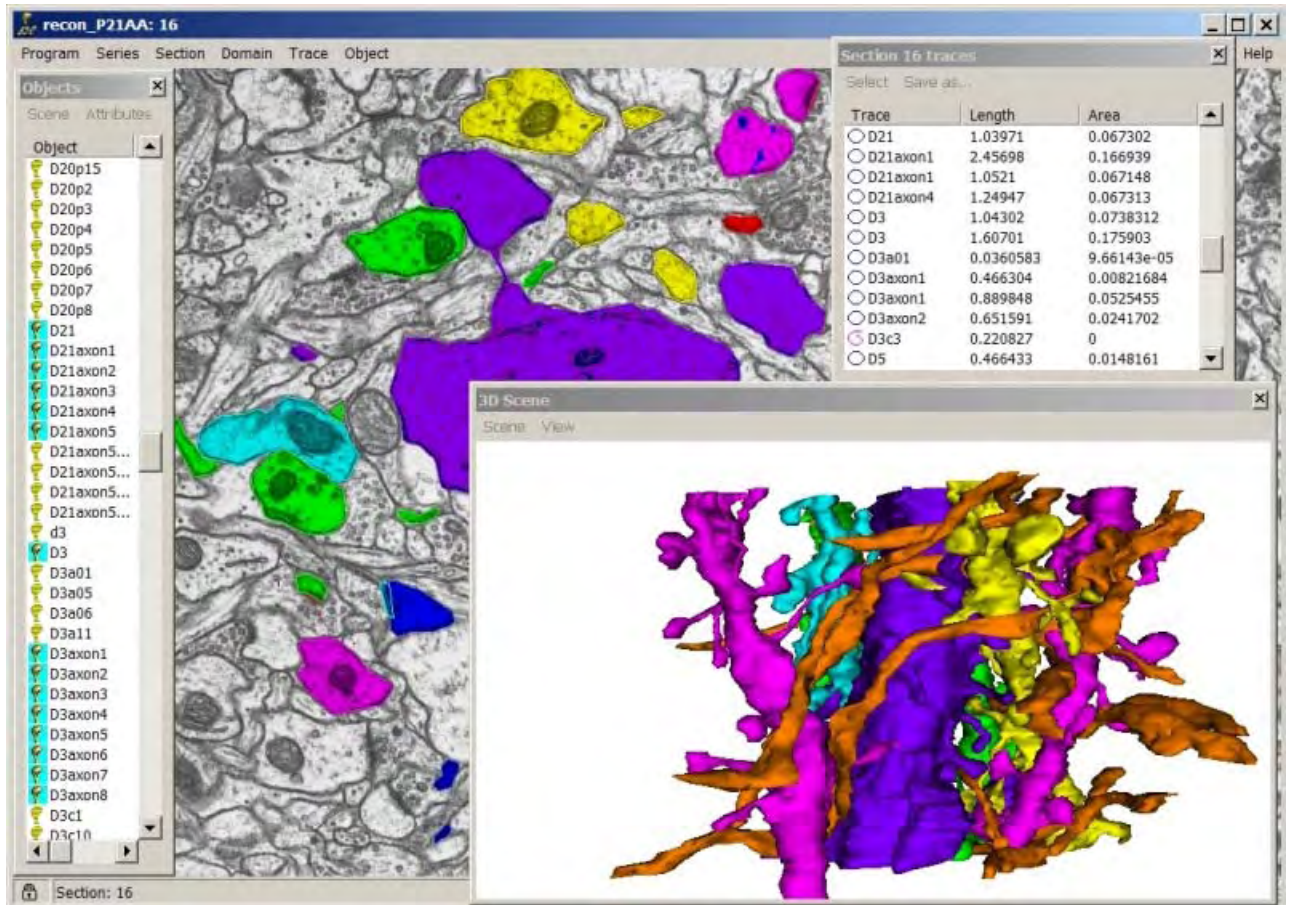


Figure 12.1. A screen shot of RECONSTRUCT™ with a 3D Scene window open. The objects highlighted in the Object List with blue icons reside in the scene. These objects include several dendrites (various colors) and axons (orange) reconstructed from a series of electron micrographs from the hippocampus. The color of the axons was changed from their trace colors using the In Scene Attributes dialog.

Diffuse Color

The parameters in the Change Scene Objects dialog determine object color and reflectance properties. Diffuse Color is the main reflective color value. Surfaces perpendicular to the view point strongly reflect with this color, while surfaces pointing away from the view point reflect very little of the Diffuse Color and instead take on the Emissive and Ambient colors.

Emissive Color

Emissive Color is emitted from the object in the absence of reflected light. The object simply "glows" this color. Emissive color, lacking variations in reflectance, does not give a sense of the 3D shape of the object. Emissive color is used by default for Traces and Trace Midpoint representations that do not have 3D surface structure.

Ambient Intensity

The Ambient Intensity specifies how much ambient light a surface reflects. Ambient light is not directional and so comes from all surface orientations. The ambient color of an object is the Diffuse Color scaled by the Ambient Intensity value between 0 and 1. A value of 0 means no ambient light color, while 1 gives the maximum reflectance for ambient light.

Transparency

The transparency value determines whether light from other sources passes through the object. A value of 1.0 makes the object completely transparent, while 0.0 makes the object completely opaque. Note that the implementation of transparency in the scene is not completely correct, although it is often reasonable (Fig. 12.2). For better transparency export the scene to another application for rendering.

Specular Color

Specular Color and Shininess together determine the specular highlights (e.g., the shiny spot on the sphere in Fig. 12.2). The areas of a surface perpendicular to the view direction receive the Specular Color in addition to any diffuse and ambient colors. To make a surface shiny, choose a whiter (more intense) Specular Color. To make a surface non-shiny, set the Specular Color to black.

Shininess

Shininess is a real value from 0 to 1 that determines the spread of the specular reflection. Lower shininess values produce broad glows, while higher values result in sharper highlights.

Filled or Wireframe Faces

By default the triangular surface patches are filled with the color of the object, both when seen from the front and when seen from the back. Generally speaking, the front surfaces are those seen from the exterior of the object, while the back surfaces are those exposed to the interior of the object (Fig. 12.2). To display the front or back surfaces as colored wireframes, uncheck the Fill checkboxes in the Change Scene Object dialog.

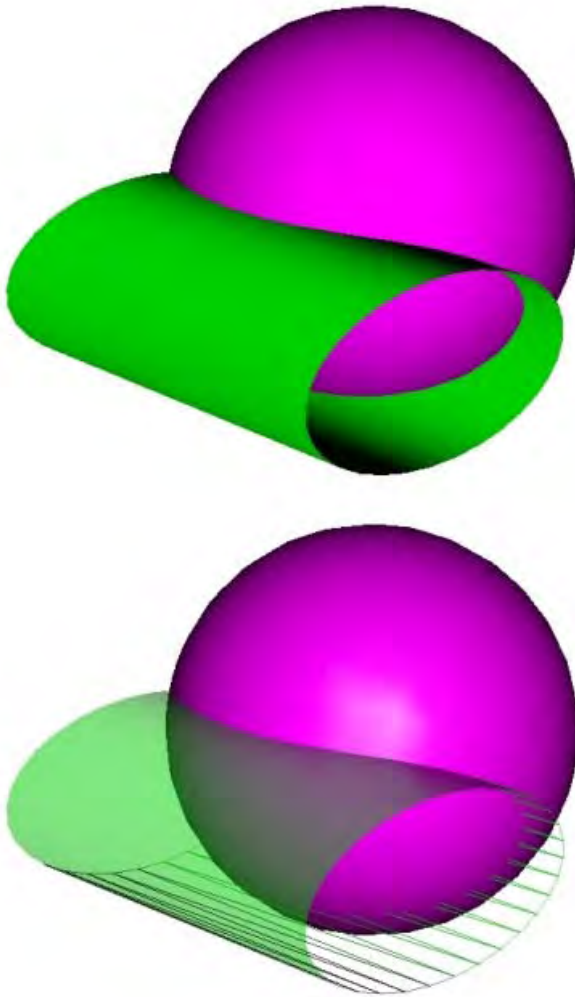


Figure 12.2. The effect of in-scene modifications of two objects. On the left, the scene of an intersecting sphere and cylinder immediately after being added to the scene. On the right is the scene after two modifications. First, the cylinder object was modified to $\text{transparency}=0.5$ and back faces were not filled. Where the back faces are visible inside the cylinder they are seen as wireframe. Then the sphere was made shiny by increasing the intensity of the specular color. Details of the in-scene modification options are given below.

Viewing a Scene

The viewing parameters of the scene are modified by using the mouse to rotate, pan, or zoom, or by using the View menu at the top of the 3D Scene window. The initial or *Reset* view position is from a point beyond the last section of the scene looking down toward section 0. In effect, looking in the minus z-direction. For an object oriented perpendicular to the plane of sectioning, the *Reset* view will be looking at the end of the object from the highest section number. The *Reset* view places the box that bounds all of the objects in the scene in the center of the window. It may be desirable to *Reset* the scene when objects are added.

The center of the scene's bounding box also determines the point around which rotations are performed. Adding objects to the scene expands the bounding box, but removing

objects from the scene does not reduce the bounding box. This allows arbitrary centers of rotation to be defined by adding appropriately positioned objects to the scene and then removing them. To reset the bounding box after removing objects it is necessary to clear the scene, add any desired objects, and then reset the scene.

The scene can be viewed in two different modes: perspective or orthographic. In perspective mode (the default) the scene can be magnified by zooming in until the scene passes through the front view plane. Similarly, the scene can be demagnified by zooming out until the scene passes the rear of the viewing volume. In orthographic mode, the scene can only be magnified by enlarging the window, not by zooming. Zooming in orthographic mode only moves the scene relative to the viewing volume. This feature may be useful for creating a cutting plane to reveal the interior of the object.

Objects in a scene with zigzag appearance are usually due to misalignments between sections. Realignment of sections and traces can correct this problem.

Rotating, Panning, and Zooming

To rotate the scene, press and hold the left mouse button in the 3D Scene window. Drag the cursor across the scene in the desired direction of rotation. The scene will begin to rotate at a speed proportional to the speed of the cursor movement across the window. Releasing the mouse button results in the scene continuing to rotate at this speed in "spin" animation. To stop this rotational spin press the right mouse button.

Scene rotation with the mouse is in two directions left-to-right horizontal rotation or up-and-down vertical rotation. The final motion can be a combination of these rotation directions but not a rotation in the plane of the scene window. To roll the scene in the plane of the scene window it is necessary to use a combination of rotations. First rotate the scene 90 degrees horizontally, then apply the desired vertical rotation and horizontally rotate the scene back to the original viewpoint.

To pan the scene press and hold both left and right mouse buttons simultaneously. Then drag the cursor to shift the scene in a particular direction. To avoid activating rotational spin, release the left button before releasing the right button.

The scene can be zoomed by using the right mouse button alone. Dragging the cursor upward with the right mouse moves the viewpoint away from the scene, effectively zooming out. Dragging downward zooms in, bringing the scene closer. As parts of the scene get very close to the viewpoint they will disappear from view. This feature can be used to look through the exterior surface of an object to see the inside.

Spin

An important aspect of viewing scenes is dynamic rotation to allow the 3D structure to be fully appreciated on a 2D computer screen. As mentioned above, the scene may be animated by using the mouse to spin the scene. Spin can be completely disabled by unchecking the *Spin* item in the View menu at the top of the 3D Scene window.

To produce a particular horizontal or vertical spin, check the *Spin* menu item and use the *Rotation...* dialog to set the rotational velocity. Horizontal spin can be produced by setting the vertical spin rate to zero. Likewise, purely vertical spin can be realized by setting the horizontal spin rate to zero. Setting both rates to zero stops all movement. Spin can also be zeroed using the right mouse button in the 3D Scene window. When *Spin* is

completely disabled, the *Rotation...* dialog allows a specific rotational increment to be applied to the scene.

Background Color

The background color of the scene is white by default, but can be changed to any color supported by the graphics system. To change the background color select the *Background...* item from the View menu. Select a color from the Color dialog and click the OK button. Custom Colors in this dialog reflect the last custom colors accessed by the program and will not be saved with the series.

Saving a Scene

The scene in the 3D Scene window may be saved either as a bitmap image or as a 3D model. In neither case is it possible to reopen and edit the saved scene using RECONSTRUCT™. A saved scene can only be viewed or edited with other programs.

VRML 2.0

To export the scene as a 3D model, use the *VRML 2.0...* menu item from the Export As submenu of the Scene menu. Select the desired directory and filename. Virtual Reality Modeling Language (VRML) files use the .wrl file name extension. This file extension will be automatically added to the filename.

The VRML 2.0 output adheres to the VRML97 file format specification (www.web3d.org). All objects are represented as indexed line or face sets, and are easily imported into other 3D graphics packages for editing or display. The VRML file can be displayed in most web browsers by using a plug-in for VRML. To find out if a VRML plug-in is currently installed, direct the web browser to a VRML plug-in detector (<http://cic.nist.gov/vrml/vbdetect.html>). A number of free VRML plug-ins are available for popular web browsers.

When the scene is opened with the VRML viewer, it should appear as it did in the 3D Scene window at the time it was saved. The background and object colors, including transparency, are preserved. However, the Front and Back faces wireframe setting will not be retained.

The VRML file is a simple text file. Comments in the file indicate the version of RECONSTRUCT™ used to generate the file, the series name and units. Each object in the scene is delineated with comments at the start and end of the object description. The starting object comment, in addition to the object name and type, contains the following information.

Traces

shifted by: x y z

Trace Slabs

shifted by: x y z

Trace Midpoints

size: a b shifted by: x y z

Box

size: a b c shift by: x y z

Boissonnat Surface

shifted by: x y z

Cylinder

center: x y z axes: x y z, x y z, x y z size: l w h shifted by: x y z

Ellipsoid

center: x y z axes: x y z, x y z, x y z size: l w h shifted by: x y z

Sphere

center: x y z size: a shifted by: x y z

These comment values record the settings of the Shift Object By and Size settings of the 3D Options tab. In the case of the cylinder and ellipsoid, however, the "l w h" values reflect the Size parameter settings only when these are positive. Otherwise, the actual eigenvalues calculated from the scatter matrix (Chapter 11: Objects) are reported, with l the diameter of the principal axis, w the diameter of the secondary axis, and h the diameter of the minor axis. The axes are the eigenvectors of the scatter matrix given in the same order. The center is the centroid of the trace points. For the sphere, the center is calculated midpoint of the object extent and the size is the diameter of the sphere specified by the Size A parameter setting.

VRML 1.0

Since some older programs accept VRML 1.0 input but not VRML 2.0, a provision is provided to export the scene as VRML 1.0. Output is similar in all respects to the VRML 2.0, except the .wrl file created adheres to the 1.0 format specification (www.web3d.org). For the most part, this only means that the Background color will not be saved with the model.

DXF

Some even older programs accept AutoCAD's DXF input but not VRML. For this reason support is provided to export the scene as DXF, but this representation is less rich than VRML. The .dxf output does not include the scene view point or color and shading information. An Object consisting of line data, will be output simply as a series of LINE entities, while an Object consisting of triangular faces will be output as a series of 3DFACE entities.

Bitmap

A simple method of capturing the RECONSTRUCT™ display as an image is to use the Print Screen key on the keyboard. This copies the entire screen to the clipboard where it can be imported into image editing software for refinement. The Export As *Bitmap...* operation expedites this process by storing just the image of the 3D Scene window. The image data is stored in a Windows Bitmap file.

JPEG

The Export As *JPEG...* operation stores the image of the 3D Scene window in a JPEG file. Since JPEG is a compressed image format, the user is prompted for the amount of compression to use before the file is written.

360° Bitmaps

The Export As *360° Bitmap...* operation stores a sequence of images of the 3D Scene window into a set of files. The image data is stored as Windows Bitmap files, with each filename modified by a three-digit number before the .bmp extension. Each image in the sequence is generated by rotating the scene by the Horizontal and Vertical increments specified by the user. When either the total Horizontal angle or the total Vertical angle exceeds 360° the image sequence is terminated and the resulting final orientation of the 3D scene is displayed on screen. Any positive or negative Horizontal or Vertical increment can be specified, but no more than 999 image files will be generated.

Scene Information

The number of data elements that compose a scene determines the speed of rendering, the size of VRML files, and the complexity of future scene processing. Therefore it is often desirable to reduce the total number of elements in a 3D Scene. The total number of elements, including vertices, line segments, triangles and normal vectors, is tallied in the Scene Info dialog box (Fig. 12.3). To reduce the number of vertices and triangles it is necessary to reduce the number of trace segments prior to generating the 3D

representation. This can be done using the *Simplify* command (Chapter 8: Traces), but note that such changes are permanent so you may want to make copy of any objects before you *Simplify* them.

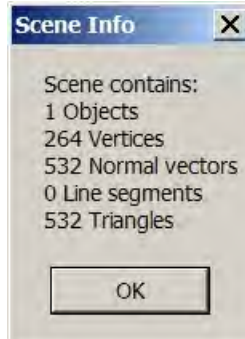


Figure 12.3. The Scene Info dialog box reports the total number of scene elements, including those not visible due to transparency, occlusion, or view point.



Bibliography

Boissonnat J-D (1988) Shape reconstruction from planar cross sections. *Computer Vision, Graphics, and Image Processing* 44:1-29.

Fiala JC (2005) Reconstruct: A free editor for serial section microscopy. *J Microscopy* 218:52-61.

Fiala JC, Harris KM (2001a) Cylindrical diameters method for calibrating section thickness in serial electron microscopy. *J Microscopy* 202:468-472.

Fiala JC, Harris KM (2001b) Extending unbiased stereology of brain ultrastructure to three-dimensional volumes. *J Amer Med Informatics Assoc* 8:1-16.

Fiala JC, Harris KM (2002) PC-based alignment and reconstruction of serial sections. *Microscopy and Analysis USA Edition*. 52:5-7.

Howard CV, Reed MG (1998) Unbiased stereology. *Three-Dimensional Measurement in Microscopy*. Springer-Verlag: New York.

Schneider PJ, Eberly DH (2003) *Geometric Tools for Computer Graphics*. Morgan Kaufman Publishers: Amsterdam.

Copyright © 2006 by John C. Fiala
Revised by K. Harris and K. Sorra; 2009



Appendix 1

RECONSTRUCT™ Frequently Asked Questions

General

1. [Is it possible to run RECONSTRUCT™ on a Mac?](#)
2. [Can I use RECONSTRUCT™ for stereological measurements such as unbiased counts?](#)
3. [Can I use RECONSTRUCT™ with light microscopic sections?](#)
4. [Can I use RECONSTRUCT™ with confocal microscopy?](#)
5. [How many different colors/objects can I have?](#)
6. [Is it possible to organize data in a hierarchical tree?](#)
7. [How can I create a local copy of the Users Manual?](#)
8. [How should I reference the software in my paper?](#)

Images

1. [What image file formats does RECONSTRUCT™ accept?](#)
2. [After Importing images, why do I see only part of them or none at all?](#)
3. [Why does Export Images produce such a huge image?](#)
4. [Why can't I Export Images from my series?](#)
5. [How can I speed up the paging and display of sections?](#)

Moving and Aligning

1. [Why don't the Arrow keys move my section images?](#)
2. [Can RECONSTRUCT™ align my sections all by itself?](#)
3. [How can I align one series to another?](#)
4. [My first 50 images align well, then they seem rotationally off. Can I adjust?](#)

Traces and Trace Editing

1. [Can RECONSTRUCT™ automatically trace my images?](#)
2. [How do I draw without having the trace filled with a color?](#)
3. [Is it possible to trace a single structure on a single section? In my serial sections I have very small structures that are embedded in a single section. I tried to trace the single structure but no image of the structure was observed in the reconstructed scene.](#)
4. [How can I place text on my sections/reconstructions?](#)

5. [Can I "personalize" RECONSTRUCT™ to create and save my own palette?](#)
6. [Can I switch tools using the keyboard?](#)
7. [Why do my traces expand/shift when I edit them?](#)
8. [How do I know if my tracing is correct for Boissonnat surfacing?](#)
9. [Why is Auto-simplify always checked in the AutoTracing tab when I start RECONSTRUCT™?](#)
10. [How can I fix the orientation of my traces?](#)
11. [Why doesn't Simplify alter my traces?](#)
12. [How can I trace "donut" shapes with the black center region taken into account instead of being ignored?](#)
13. [Can I trace/reconstruct an involuted object with holes inside holes?](#)
14. [How are traces stored in XML format?](#)
15. [What is the reference point for the max X and max Y measurements in the Trace lists dialogue? Are X and Y with respect to the position of the trace on the domain image, or are X and Y coordinates derived with respect to the centroid of individual traces?](#)

Measurements

1. [How do I measure size/diameter/distance?](#)
2. [Can I compute the length of a region of interest by using the maximum and minimum X and Y coordinates of the trace?](#)
3. [How can I make a measurement on a piece of an object?](#)
4. [What are the units of my traces/measurements?](#)
5. [Can I superimpose multiple traces to obtain a "Total Occupied Area" of a region of interest over time?](#)

Objects and The Object List

1. [Why is the object list so slow to fill?](#)
2. [Can tips be provided for dividing an Object into section ranges and for Object Renaming?](#)
3. [Why does Add to Scene never finish and then crash when I close the series?](#)

Surfacing and 3D Output

1. [Why does my Boissonnat surface have irregularities?](#)
2. [Can I fix surface irregularities with Simplify?](#)
3. [How do I generate a 3D representation for an object that appears on a single section?](#)
4. [I copied the same trace on adjacent sections. Why is the 3D surface so bad?](#)
5. [How can I reduce the number of lines/faces in my VRML output?](#)
6. [How can I bring multiple series together accurately within the 3D Scene?](#)
7. [How can a 3D scene window be reset?](#)
8. [Why does the program crash when I open or interact with the 3D Scene?](#)
9. [Why does my object appear as only lines in the 3D Scene?](#)

11. [Why can't I smooth my reconstruction in 3D Studio MAX?](#)
12. [How can I record/save a video sequence of my rotating 3D Scene?](#)
13. [How can I view a saved VRML scene?](#)

ANSWERS

General

Is it possible to run RECONSTRUCT™ on a Mac?

RECONSTRUCT™ will run under other operating systems provided a Win32 emulator/translator is installed. RECONSTRUCT™ binary may run on Virtual PC, Apple's Pentium/Windows emulator, but this has not been systematically tested by the developers of RECONSTRUCT™. Anecdotal feedback suggests RECONSTRUCT™ runs in Bootcamp and in Parallels. Bootcamp is a tool from Apple that lets you partition your computer's hard drive and install Windows on one of the partitions. Once you've done this, it's basically like running Windows on any other machine, with full 3D hardware acceleration, etc. Parallels is a 3rd party application that lets you use the Bootcamp partition to run Windows within the Mac environment, much like any other application.

Can I use RECONSTRUCT™ for stereological measurements such as unbiased counts?

(See Chapter 8) Yes. RECONSTRUCT™ works well for making stereological measurements on serial sections and includes facilities for generating sampling grids.

Can I use RECONSTRUCT™ with light microscopic sections?

Yes. RECONSTRUCT™ allows you to align and reconstruct from any serial section images. If you have optical sections which are already aligned, using them will be no problem. For sections mounted and imaged separately, the challenge in realigning them depends on the amount of warping/distortion in the sections.

Can I use RECONSTRUCT™ with confocal microscopy?

Yes. individual section images or TIFF image stacks from confocal microscopy can be imported into RECONSTRUCT™. The Wildfire Tool has been used successfully to semiautomatically reconstruct confocal image data.

How many different colors/objects can I trace?

RECONSTRUCT™ offers more than 16 million colors for tracing/shading and more than 10^{35} different trace names, so you can create more objects than you will ever need.

Is it possible to organize data in a hierarchical tree?

(See Chapter 5) RECONSTRUCT™ does not have built-in object classes or a data hierarchy, but data can be organized to the extent that name strings limit the view of the data, e.g. show only the objects associated with dendrite D01 and their measurements. This is possible in RECONSTRUCT™ through object names and their limit strings in the Lists tab of the Series Options.

For example, if one names objects as follows:

```
N01D01
N01D01Synapse01
N01D01Synapse02
:
N01D01Synapse22
N01D02
N01D02Synapse01
:
N01D03
```

N01D03Synapse01

:

N02D01

:

N02D02

:

N03D01

:

Then if the Limit string for the Object List is set to "N01*", the Object List will only show the objects associated with neuron N01. To see all synapses on all neurons, one can use the limit string "N##D##Synapse*". To see all the dendrites on Neuron N02, one can use limit string "N02D##". And so on...

Limiting the view of the traces to relevant Objects is also possible. Select everything you don't want to see in the Object list and then Modify>Attributes... to set the Hidden flag for these objects. Then when you browse the sections, the traces of these objects will not be displayed, only the selected objects will be visible. To unhide, check the "Unhide all traces when loading section" option in the General tab. More information on limit strings can be found in Chapter 5 of the Users Manual.

If your series is becoming to complicated, you can use file system subfolders to break your series into manageable subsets. For example, suppose your original aligned series is located in the folder "myStudy" and is named "aligned.ser". Create a new series in a subfolder "myStudy\N01\N01.ser", and then import a subset of your original dataset.

Select Import>Series... from the series menu in the new series. Use the Browse... button to select myStudy\aligned.ser as the series to import from. Check all the boxes and verify the section range incorporates the entire series. Check the domains named box and leave the wildcard "*" in the name field to import all images, but uncheck the box next to "Copy image files...". This will force the new series to display images by accessing the files in the parent folder rather than unnecessarily copying the image files into the subfolder. Check the trace named box and use a wildcard name string to select a subset of the data you want to be included in the new series, e.g. "N01*" to include all traces that begin with the string "N01". Click Import and verify the traces are correctly in the new series.

You can repeat the Import dialog to grab additional subsets of data using different wildcard strings, but be carefully not to copy the same traces twice as this will create problems for your reconstructions.

You can repeat the process of creating subfolder series and importing part of your dataset until the entire series is divided up into manageable units. To avoid confusion you may want to delete the original traces from the parent series aligned.ser. Later it will be possible to reverse the process and copy the subfolder traces back into the main folder if you want to reconstitute the whole dataset as one series.

How can I create a local copy of the Users Manual?

Upgrade to the latest version of RECONSTRUCT™ and download the manual.chm file from the Users Group (http://tech.groups.yahoo.com/group/reconstruct_users/) into the RECONSTRUCT™ program files subfolder. If you're looking for the old PDF manuals, see the Older Versions folder in the files section of the Users Group.

How should I reference the RECONSTRUCT™ software in my paper?

Please cite the paper: Fiala JC (2005) Reconstruct: a free editor for serial section microscopy. J Microscopy 218:52-61.

Images

What image file formats does RECONSTRUCT™ accept?

(See Chapter 3) RECONSTRUCT™ currently supports several common single image file formats including Windows Bitmaps, JPEG, TIFF, and GIF. The complete set of acceptable file formats is determined by the image input/output library developed by Andy Key (see Credits, below). Acceptable formats include Windows and OS/2 bitmaps (.BMP .VGA .BGA .RLE .DIB .RL4 .RL8) except for CCITT G3 MH Fax compressed and 24 bit RLE files; CompuServe GIF files (.GIF); JPEG File Interchange Format files (.JPG .JPEG .JPE); TIFF files (.TIF .TIFF) with 1, 4, 8 and 24-bits per pixel; Truevision Targa/Vista bitmaps (.TGA .VST .AFI); ZSoft PC Paintbrush Image format (.PCX .PCC); GEM Raster (.IMG); Portable Greyscale-map (binary P5 type) (.PGM); and Portable Pixel-map (binary P6 type) (.PPM).

After Importing or calibrating images, why do I see only part of them or none at all?

(See Chapter 6) Use the Section List to make sure your sections exist and to Page To one of them. Then hit the "Home" key to bring the section image into view. If you still can't see anything, check the Domain List to verify that there is actually a domain image present.

Why does *Export Images* produce such a huge image?

(See Chapter 6) *Export Images* at full resolution takes the smallest pixel size of any image and the total area occupied by the entire series and renders at that resolution. So if you had a minimum pixel size of 0.00254 microns and a section touched (-22,-22) and another section touched (21.7,21.7) microns in x and y, full resolution is $(43.7/0.00254)^2$ pixels times 3 bytes per pixel = 847 Mb! You'll obviously want to type in a more reasonable pixel size in the Export Images dialog. The real underlying problem may be that your aligned series is really very big because there's a shift in the sections as you go through the series. This is why RECONSTRUCT™ transforms your images in real-time rather than creating new bitmaps, to save on this huge (wasted) memory overhead and speed up paging of sections. NOTE: *Export Images* always exports RGB color.

Why can't I Export Images from my series?

(See Chapter 6) If you get a uniformly black image from Export, check that the area you want is what you're specifying in the Export dialog. Try Exporting just one image, setting the area you want in the main window and then using Make Images Resemble Current View. Use the default pixel size and verify the output. There was a bug prior to version 1.0.5.3 that caused a uniformly black image to be created when files were large. This was caused by a failure in a Windows API call (under XP) when the image size got too big (around 18 Mpixels). This API call is used to draw traces onto the image, so now if the image is too big and traces are requested As Displayed, the program will export the image but may not show the traces. Even when images are too big to use the API call, you can still export traced areas using the Fill option in the Export Images dialog.

How can I speed up the paging and display of sections?

Two main factors determine how quickly sections are displayed: the time it takes to read image files from the storage device, and the number of screen pixels that contain image data. The time it takes to read image files is influenced by the type of storage device, the size of the files, and the number of files that have to be read to display the visible portion of the section. You can greatly increase the speed of paging through section by optimizing these factors.

To optimize file access times, store image files on the local hard drive of the computer (usually the C:drive). Avoid accessing files from CD or DVD optical drives, external hard drives, or network drives. The speed of the drive will mostly determine how fast the image data can be loaded into memory. Smaller image files can generally be loaded into memory faster than large image files, but keep in mind that any data compression (e.g. JPEG) will have to be decompressed before display, so there is increased processing overhead for compressed versus uncompressed image files. The relative speed of the processor/memory versus the speed of the hard drive will determine whether compressed file formats offer any speed improvements over uncompressed formats.

When image files contain high resolution data, RECONSTRUCT™ has a built-in mechanism for taking advantage of the reduced resolution requirements of displaying the image on a lower resolution computer screen. By loading low resolution proxy image files instead of the full resolution images, the section can be more quickly displayed at low resolution. Proxy images are created from the Proxies tab of the Series Options. See the program Users Manual for more details about creating and using proxy images. The main point is that it is possible to speed up the display of sections by creating reduced resolution proxy images. If the full resolution images have lower resolution than the screen (~1 megapixels), proxy images are probably not needed. If each section contains multiple images, a significant speed up can be obtained by changing the view and/or window size so that only one of the images is visible. RECONSTRUCT™ only loads an image file when it is visible on screen. Zooming-in until just one image domain fills the main window results in only one image file being read from the drive during section display. When a section only contains one image, however, zooming-in does not improve performance. In fact, it may slow down the section display because more screen pixels have to be rendered.

Section images are rendered onto a bitmap that is sized to the main window, but only the pixels that contain image data require processing. The size of the black background has little effect on the speed of section display. To speed up display, zoom-out to reduce the number of pixels that contain image data on screen. This is always effective for sections that contain a single image. When sections contain multiple images, use proxies to get the speed benefit of zooming-out to low magnification. Setting the zoom to Actual Pixels may offer minor improvements over displaying at a higher magnification, but the real speed factor is the total number of screen pixels that need to be rendered. Keep this as small as possible for speed.

Although image file storage and the size of the image on the screen are the two main factors influencing the speed of section display, a third factor is the complexity of the transformation required to position the image on the section. Image domains without movements will be displayed somewhat faster than rotated, flipped, or scaled images. For images that are aligned prior to using RECONSTRUCT™, such as confocal stacks, the ideal scenario is not to introduce any transformations (movements) of the images. Lock the sections immediately after importing the images to prevent the inadvertent introduction of transformations.

Moving and Aligning

Why don't the Arrow keys move my section images?

(See Chapter 10) The size of the movement generated by pressing the arrow keys may depend on the pixel size of your section. Try increasing the value of the movement parameters in the Series Options Movement tab and see if you can get a movement with the arrow keys. Also, make sure your sections are "Unlocked" so that movements are enabled.

Can RECONSTRUCT™ align my sections all by itself?

(See Chapter 10) No, it is not possible to execute an alignment without user input, but in some cases it may be possible to automate some aspects of the process such as creating traces for point correspondences. Then, assuming you use the same fiducials on every section, it is just a matter of repeating Ctrl-S and Alt-T-N-R on every section.

How can I align one series to another?

(See Chapter 10) If you have a rotational offset between series, you can bring them into registration by propagating a corrective rotation throughout one of the series. This will keep the alignment, and just rotate everything.

Start by making a copy of the series you plan to rotate by creating a new series and then using Import Series (everything but the "Copy image files into series folder") from the series to be rotated. Import a section from the other series that can be used to determine the rotation angle. This section should be numbered either before or after the current series. You could put one in the middle but you would need to shift the sections first to make room.

Rotate an adjacent section to match the imported section by using any method such as correspondence points.

Then use Propagate... (from the Movement submenu) to apply this same rotation to all the remaining sections.

(See also: How can I bring multiple series together accurately within the 3D Scene? Below.)

My first 50 images align well, then they seem rotationally off. Can I adjust?

(See Chapter 10) Yes. Since you can detect the rotational offset it is also possible to correct it though manual rotations using the F3 and F4 keys. The size of keyboard movements can be modified by holding the Control or Shift key down. See the Options...Movement page to set the keystroke movement sizes and modifiers. The center of the movement can be set in the Movement Type-In dialog under the Section dialog.

Traces and Trace Editing

Can RECONSTRUCT™ automatically trace my images?

(See Chapter 8) The general problem of automatic segmentation is unsolved, and very likely it cannot be solved to the satisfaction of an expert microscopist, since the program would have to be an expert microscopist itself. You may be able to obtain a degree of help in the segmentation process by using the Wildfire Tool. The Wildfire Tool works well for some types of images, but it does not work well for others (e.g transmission electron micrographs). Also, you may have to spend some time trying different settings for the Stop Criteria in the Autotracing tab. Even then, don't expect wildfire to segment perfectly with one mouse click. But even if you have to click at multiple points in your region and then Merge the results, Wildfire can sometimes still save time and effort.

How do I draw without having the trace filled with a color?

(See Chapter 8) Set the Fill Mode to None in the Names/Colors options before drawing.

Is it possible to trace a single structure on a single section? In my serial sections I have very small structures that are embedded in a single section. I tried to trace the single structure but no image of the structure was observed in the reconstructed scene.

(See Chapter 8) Yes. If you generate the object as a Sphere in 3D (or as Traces, Trace Midpoints, or Box) it will be visible. Go to the 3D tab in the Series Options.

How can I place text on my sections/reconstructions?

A way to get text on sections is to create text traces and paste them where desired. An arial font set of traces is available as a RECONSTRUCT™ series "letters.ser" in the Files section of the RECONSTRUCT™ Users Group. (http://tech.groups.yahoo.com/group/reconstruct_users/). You can copy the font into your series and then copy and paste or stamp letters using the Stamp or Grid tools. It's not as convenient as typing text, but it may be enough if you have very simple labels.

Can I "personalize" RECONSTRUCT™ to create and save my own palette?

(See Chapter 8) Yes. Use the Names/Colors tab to set the individual palette entries. You can change the shape of the first 10 palette entries by 1) pressing the palette button, 2) picking up the Stamp Tool, 3) going to the trace whose shape you want to copy and pressing the left mouse button while holding the Ctrl key down, and 4) pressing Ctrl-1 (or any other digit) where 1 is the second button on the palette (the first button is zero). The new palette entries will be saved when the Series is saved. Each series saves its own palette and you have to copy the palette from a previous series to get it into a new one using the Import Series dialog.

Can I switch tools using the keyboard?

(See Chapter 4) When you're tracing and need to switch to the Arrow Tool, press Ctrl-Tab and then Tab followed by Enter. To get back to another tool using the keyboard just add more Tabs before pressing Enter.

Why do my traces expand/shift when I edit them?

(See Chapter 8) The expansion of traces is caused by the *Simplify* algorithm that is applied to your trace to eliminate loops. You can turn this off in the series options AutoTracing tab, but this is not recommended as you may introduce loops that will seriously affect your measurements and 3D reconstruction. For best results keep automatic simplification on and set the associated Simplify on Demand value to a negative number (which forces simplification to be at screen resolution.) Then edit your traces when fully zoomed in on them, so that the expansion of only one screen pixel will have a minimal effect on the trace. In addition you can set the Shrink-back option in the Autotracing tab to prevent expansion.

How do I know if my tracing is correct for Boissonnat surfacing?

(See Chapter 11) A proper trace for Boissonnat surfacing should be listed next to a blue circle icon and a positive area. This will always be the case when you use one of the drawing tools for closed traces with automatic simplification turned on.

Why is Auto-simplify always checked in the AutoTracing tab when I start RECONSTRUCT™?

(See Chapter 8) You'll want Auto-simplify to be on all the time, to prevent creating traces with loops when you're drawing. Because it's important not to make loops, Auto-simplify comes on everytime you restart RECONSTRUCT™. This is just to let you know that you shouldn't turn it off under normal circumstances.

How can I fix the orientation of my traces?

(See Chapter 8) You can verify the orientation by opening the Trace List and paging through the sections. If the Area column for a trace has a negative value, the trace has an orientation that defines it as an inner surface of the object (a hole). You can fix this by selecting the Trace and applying Reverse from the trace menu. If the Area is zero, then the trace does not define a closed surface which is also a problem. You can also verify this by the icon next to Trace name which will be open instead of closed. To fix an open trace, select it and use the Attributes dialog to change it to closed. Then use *Simplify* from the Trace menu to set the proper orientation and remove any loops. Refresh the Trace List to see the changes.

Why doesn't *Simplify* alter my traces?

(See Chapter 8) First you have to mark the trace as "unsimplified". Select the trace and open the Trace Attributes dialog. Uncheck the Simplified flag for the trace and press OK. Now the trace can be simplified again by using the *Simplify* command from the Trace menu. You can control the amount of simplification by setting the "Simplify on demand" resolution in the AutoTracing tab of the Series Options.

How can I trace "donut" shapes with the black center region taken into account instead of being ignored?

(See Chapters 8 & 11) To define an interior boundary of an object you'll need to use a Reverse trace. This is done by drawing a normal trace and then using the Reverse item from the Trace menu with the trace selected to reverse it's orientation. You can verify that a trace has a reversed orientation by checking the area measurement in the Trace List. Reversed traces have negative areas. Assuming your donut's profile is white, on a black background, you should be able to use Wildfire equally well for the black interior as for the white object. Change the Wildfire stop limits in the Autotracing tab of the series options to Brightness "differs by" instead of "is less than" and then just click on the black interior instead of the white. To generate the traces I used regional wildfires with a Stop at Brightness less than 75. Then, switch to Stop at Brightness greater than 75 and generated reverse traces in any interior region(s).

Can I trace/reconstruct an involuted object with holes inside holes?

You can have any number of traces interior to other traces, but they shouldn't touch. Such nested traces must follow a pattern: positive, negative, positive, negative, positive, negative -- the outermost trace must be positive (have a positive area) and the innermost trace must be negative (have a negative area).

How are traces stored in XML format?

(See Chapter 3) For details on the XML format for the section files that contain the trace data, see the User Manual.

The z-traces are stored in a similar format in the series file. In general the trace data is stored as an ordered list of x,y points. The coordinates and units for these points are what they are on the section (i.e. they are in section units, see below). But be aware that the values in the XML files are subject to the transformations encoded there. In other words, to get the actual section coordinates in microns you may need to apply a transformation. This is done for you, along with calculation of the z-distance, when you export VRML from the 3D Scene or export Lines from the Series Export list. For details on the basis of the transformation represented by the xcoef and ycoef attributes, see: Fiala JC, Harris KM (2001) Extending unbiased stereology of brain ultrastructure to three-dimensional volumes. J Am Med Inform Assoc. Jan-Feb;8(1):1-16.

What is the reference point for the max X and max Y measurements in the Trace lists dialogue? Are X and Y with respect to the position of the trace on the domain image, or are X and Y coordinates derived with respect to the centroid of individual traces?

The X and Y coordinates of traces and objects are defined relative to an arbitrary origin around which data are positioned. In other words, there is a point (0,0) on every section that is defined prior to the addition of any image or trace data. By default, images are imported such that the lower left corner is at (0,0), but this is entirely optional. You can import or move images to any other point in space, e.g. by setting the X,Y offset in the import dialog.

Traces are not drawn or stored in image coordinates, since any number of images at different magnifications can be placed in the section. Rather traces are drawn and stored in section coordinates relative to the (0,0) origin around which the images are placed. The Extent (min and max) values in the Trace List are also given relative to this origin, in whatever units that were defined for the series, and the values delineate the extent of the trace relative to the X,Y axes of the series. If the images are not rotated after importing them, then the X,Y axes are aligned with the horizontal and vertical dimensions of the original images.

Measurements

How do I measure size/diameter/distance?

(See Chapter 9) RECONSTRUCT™ has at least 4 ways to make distance measurements:

Use the line or open point-by-point tools within a section, and then list the length in the Trace List

Use the z-trace tool within or across sections, and then list the length in the Z-trace list.

While using any of these tools, look at the status bar to read off the distance

Create 2 or more objects that define the end points of the distance and then use the Distances list (with limits!).

Of course the trick is to determine where to make the diameter measurement, especially if your dendrite is not cross-sectioned. You can check the orientation of a z-trace measurement relative to the dendrite using the 3D scene window and transparency. Alternatively, if your dendrites are obliquely sectioned, you may be able to create a central axis object by stamping small point traces at the center of each section profile. Then the Distance List can be used to report the distance from the synapse object to the central axis, giving effectively the radius of the dendrite at the level of the synapse.

Can I compute the length of a region of interest by using the maximum and minimum X and Y coordinates of the trace?

No. The Extent does not give the maximum width of each trace, but instead gives the size of the bounding box aligned with the image. An accurate length measurement requires a line of measurement to be drawn along the curved interior either in a single section and across serial sections as needed. The Z-trace can be used for this, and the length measurements are displayed in the Z-trace list.

How can I make a measurement on a piece of an object?

If the piece is defined by a range of sections, you can rename your object using the "\$" special character to create a separate object for each section. These can be again renamed in groups to get pieces of your object.

If the piece is more complicated, e.g. the spine head in a dendrite reconstruction, you'll need to separate the traces that define the piece. The easiest way to do this may be to copy the object using the Copy command in the Object List. After you create a Copy, Hide the original object. (Make sure the Unhide flag is not checked in the Series Options, and use the Trace List while paging through sections to verify original object is hidden.) Now you can edit the copied object without fear of modifying the original. Pick up the Scalpel Tool and use the left mouse drag to cut off the piece from the original trace. Use the middle mouse button to select the new piece and use Ctrl-A to open the Trace Attributes dialog. Change the trace name and color. While the new trace is selected use Ctrl-C to put a copy of it on the trace clipboard. Go to the next section, cut off the next piece, and select it with middle mouse button. Press Ctrl-B to paste the attributes from the clipboard to the new trace (it will change name and color). Keep going like this using the Scalpel and Ctrl-B until you have completely cut the piece from the main object. You can verify the result by looking at the new objects in the 3D Scene.

Is there a way I can superimpose multiple traces in order to obtain a "Total Occupied Area" of a region of interest imaged over time?

One approach is to copy the traces into a single section, merge them, and then obtained measurements from the trace list. This will work reliably as long as all the traces overlap; non-overlapping traces won't merge.

What are the units and coordinate system for my traces and measurements?

(See Chapter 9) The units are arbitrary but assumed to be the same for all input values including the image pixel size(s), traces, and measurements. The descriptor for these units is set in the General tab of the Series Options and by default is "microns". If you prefer to work in units of "pixels", for example, it is possible to do that instead. Just set the units descriptor in the General tab of the Series Options to "pixels" and then always import your images using a pixel size of 1.

Objects and The Object List**Why is the object list so slow to fill?**

(See Chapter 9) When you open the Object List the program reads all of the sections to see what traces have been created and to compute the values for the list (e.g. volumes, surface area, etc.). The length of time it takes to read this list will be the time it takes to read one section times the number of sections in the series. You can speed up the Object List by:

making sure your section files are stored locally on your hard drive for fast access

turning off all list parameters like volume, surface area, etc.

turning off all other programs that are actively accessing the hard drive such as any virus scanning software.

Due to the nature of Windows list views, the list display will also be slow (or even stop altogether) if there are tens of thousands of objects in the list. A particularly critical number for Win32 appears to be 65,535. If you are tracing one object through the series make sure you are using the same object name for every trace, not an incrementing name. Names with different numbers in them are interpreted as belonging to different objects.

Can tips be provided for dividing an Object into section ranges and for Object Renaming?

(See Chapter 11) Open the Object List and highlight the object name (e.g. myObject) you want to break into parts. Select Modify->Attributes->Of Traces... from the top of the Object List. This opens the Trace Attributes dialog on the object.

In the Name field of the Attributes dialog change "myObject" to "myObject_ \$" and press OK. (Use caution when renaming to not give the parts of an object the name of another existing object because then they will become part of that object and may be difficult to separate).

Reopen the Object List and see that myObject has become a set of objects: myObject_1, myObject_2, myObject_3,... Each of these objects corresponds to one section of myObject.

Now it is possible to rename a group of sections into one subpart. For example, highlight myObject_10 through myObject_19 and select Modify->Attributes->Of Traces... changing the name field from "*" to "myObject_part2". Reopen the Object List to see the change. You can verify that myObject_part2 is a slice of the object by generating a 3D representation.

To put all the parts back together into one object, highlight them all and Modify Trace Attributes changing the name field from "*" to "myObject". When you reopen the list the original myObject will be all that you see

Why does Add to Scene never finish and then crash when I close the series?

(See Chapter 12) The problem is probably that you are too impatient for the complexity of the object you've created. A complex object made up of lots of traces, each with thousands of little line segments, will take a long time to generate all the triangular faces corresponding to the line segments. You can abort the process by using the Abort button on the dialog but you'll still need to wait for it to finish the section and clean up memory before it's done. Rushing things by closing the series will only result in a program crash, but you can sometimes exit the entire program by clicking on the X in the title bar of the main window.

You can speed up future Add to Scene commands by simplifying the objects to reduce the number of line segments. Here's a specific example:

Download the letters.zip series from the Files section of the Users Group, and open the series.

Open the Object List and add the object to the scene. You'll see the outlines of the letters. If you zoom-in or use the Scene Info... dialog you'll see that they are made up of lots of little line segments.

Open the Series Options and in 3D tab set the representation type to Trace Slabs. If you add the object to the scene now it will take forever due to the complexity of the traces.

In the Autotracing tab of the Series Options, set the Simplify on Demand resolution to 0.004

From the Object List Modify menu, *Simplify* the object. (Note: The Simplified flag is already unchecked.)

Now add the object to the scene again. After just a few seconds the solid letters will appear because now the traces are much simpler.

If the program gets stuck generating a 3D scene as a Boissonnat surface and "Aborting..." goes on forever, the program has gotten stuck trying to surface a single section of trace data. Most likely there is a problem with the trace data within that section, such as overlapping or duplicate traces. You can narrow the problem to the particular section by limiting the range of sections in the 3D Tab of the Series Options. Once you find a section that is not surfacing, carefully check the traces for overlaps or duplicates. One technique to find copied traces is to look for duplicate trace names with equal areas in the Trace List. To verify duplication, select a trace and move it temporarily with the arrow key to see if there is an identical trace underneath it.

Surfacing and 3D Output

Why does my Boissonnat surface have irregularities?

(See Chapter 11) A Boissonnat surface may have irregularities for any of the following reasons:

- There are duplicate traces (even hidden ones) within a section.
- There are extraneous traces within a section.
- Traces have loops or self-intersections.
- Two traces of the object within a section touch or overlap.
- Some exterior traces of the object are counterclockwise.
- Some interior traces are clockwise.
- A branch terminates abruptly leading to a connection to another branch.
- A bug in the surfacing algorithm leaves holes and/or interior triangles.

Most of these problems (and sometimes even the last one) can be fixed by locating the offending traces and eliminating or redrawing them (see below).

Can I fix surface irregularities with *Simplify*?

(See Chapters 8 & 11) The mistakes in your 3D surface might be caused by several things: 1) small loops in your traces, 2) overlapping or duplicated traces (of the same name), 3) a trace is Reversed (counterclockwise), or 4) a failure of the 3D surfacing algorithm. If you traced with Auto-simplify off, then 1) or 3) is the most likely culprit. These can be fixed with *Simplify* (see below). If 2) is the problem, *Simplify* won't help you. In regards to 4), the surface generated by the algorithm occasionally has mistakes that can't be traced to 1), 2) or 3), but sometimes a change in the trace can produce a different result and "fix" the problem so *Simplify* might help.

In tracking down a small surface defect, use the section range values in the 3D tab of the Series Options to progressively reduce the range of sections until just a 2 sections are involved. Then go to those sections and examine the Trace List with the Area column turned on. If any traces have a negative area or are inappropriately duplicated, fix them. You can also examine your trace at high zoom to check for loops or intersections with traces of the same name. Note that it might be the case that 2 traces come too close together rather than touching. You can try redrawing these parts of the trace with the Scissors tools.

To try Simplifying your whole object, open the Object List. Highlight the object and select "copy" from the Create menu. This will create an exact copy of the traces with the name "Copy of...". You can later keep or delete this object without losing your existing work. Highlight only the copy and select Modify>Attributes... and uncheck the Simplified flag in the dialog. This will mark your object copy as not yet simplified. Now set the Simplify on Demand parameters in the Autotracing tab. For a pixelsize of 0.0084 microns, try setting the resolution to +0.02 for starters. If you're worried about your traces expanding by 0.02 microns, then also check the Shrink back... box. Click OK.

Now, with the object copy highlighted in the Object List, run *Simplify* from the Modify menu. You should see an immediate change in your trace on the section. If not, then you may have forgotten to uncheck the Simplified flag for the traces. Finally, try regenerating your surface by adding the copied, simplified object to the 3D Scene. Make sure you Clear the Scene of the original object first. You can repeat the process with a larger resolution parameter to simplify the object even further.

How do I generate a 3D representation for an object that appears on a single section?

If you generate the object as a Sphere (or as Traces, Trace Midpoints, or Box) it will be visible in the 3D Scene. See the 3D tab in the Series Options.

I copied the same trace on adjacent sections. Why is the 3D surface so bad?

(See Chapter 11) You might think that it would be easy for the program to correctly surface identical traces on adjacent sections, but in fact it is harder than non-identical traces. The Boissonnat algorithm is undefined for identical traces so a trick is used to get a surface anyway. Often this trick works, but sometimes it does not. Try this: Start over with a fresh rectangle trace drawn with the rectangle tool. (Note this trace is already as simplified as it can get -- re-Simplifying it would actually make it more complicated, strangely enough.) Copy the rectangle trace with Ctrl-C and then paste it into adjacent sections. Now open the Object List and surface the object by adding it to the 3D scene. You may get a solid, but let's assume you get a surface with a mistake in it. Go to the center section and select the rectangular trace. Use the Type In... dialog from the Movement menu (found under Section dropdown) to translate the trace imperceptibly, by say 0.001 in the X and Y directions. Now add the object back to the scene and your surface should be solid. Shifting the identical trace by the tiniest amount is usually enough to avoid the identical traces dilemma.

How can I reduce the number of lines/faces in my VRML output?

(See Chapters 8 & 11) You can greatly reduce the number of segments in an existing trace (or object) by reapplying the *Simplify* command. You may have to uncheck the "simplified" flag in the Trace Attributes in order to do this. When the "simplify on demand" parameter is zero or negative, the size of the resulting trace segments will be at least 1% of the overall trace extent. When a positive "simplify on demand" parameter is given, *Simplify* will use this value as the smallest segment length and the trace will expand by about this distance all the way around. Provided the "simplify on demand" parameter is large enough

to increase the minimum segment length (which is initially the screen resolution when the trace was made) the result will be a reduced number of triangles in the final surface. You may want to play around with *Simplify* on a single trace before applying to a whole Object. If you want to play around with *Simplify* on a whole Object, I suggest you first make a copy of your object or series. *Simplify* on a whole Object is not undo-able.

How can I bring multiple series together accurately within the 3D Scene?

If your series were parallel, you could just add the objects from each series to the scene in sequence:

Open first series, add objects to scene

WITHOUT EXITING RECONSTRUCT™, open the next series

In the Series Options 3D tab set the offset (Shift object by: X,Y,Z) if necessary to position the objects in the scene. (You can permanently shift the position of each series by the appropriate amount and leave the offsets at zero, as described in the next part for rotation.)

Add objects to the scene, and go to 2 until done with all series. (The object names must not replicate names already in the scene, because this will replace the existing scene object with the added one.)

You can also combine the series using montaging on a section-by-section basis:

Create a new series to hold all the series together

Import everything from the first series

Import everything from the second series

Go through section by section aligning the new domains to the edge of the first ones. Record the domain movements while aligning and then select and repeat these movements on the imported traces as well. (There are a number of ways to do this more efficiently, but ultimately you're going to have to create a fully aligned montage of all images on each section.)

Repeat 3 until done.

How can a 3D Scene window be reset?

(See Chapter 12) The 3D Scene window has a bounding box that surrounds all the objects in the scene. The center of this box is the center of the scene rotation. If the 3D objects are small after a reset, then the bounding box is much larger than the actual scene contents. Usually this happens when objects are added to the scene and later deleted, as the scene bounding box is not resized with deletions.

To adjust:

- 1) Exit RECONSTRUCT™ and restart it.
- 2) Add 1 object to the scene. Scene rotation will now roughly be around the object center.
- 3) Add a 2nd object to the scene that is much further away in 3D space. Scene rotation will now be about the midpoint of the 2 objects.
- 4) Delete the first object. Scene rotation now will still be about the midpoint of the 2 objects, rather than just the one remaining. Reset the scene. Rotation remains at the 2 object midpoint.
- 5) Clear the scene.
- 6) Add the 2nd object to the scene. Scene rotation is still at the midpoint.
- 7) Reset the scene. Scene rotation is now at the center of the 2nd object alone.

Why does the program crash when I open or interact with the 3D Scene?

(See Chapter 12) A common culprit is the OpenGL implementation in your video driver. OpenGL is used by RECONSTRUCT™ to display the 3D Scene. You may need to update your video driver. I recommend first repeating the same operation on a different machine to see if it works there. If it works fine, then the video driver is the prime suspect.

Why does my object appear as only lines in the 3D Scene?

If you want a surfaced object, select that option in the 3D tab of the Series Options dialog.

Why can't I smooth my reconstruction in 3D Studio MAX?

Try generating your 3D Scene using Face Normals (Series Options 3D tab) instead of Vertex Normals. While Vertex Normals generate a smoother picture in the 3D Scene window they prevent 3D Studio MAX from smoothing between faces.

How can I record/save a video sequence of my rotating 3D Scene?

(See Chapter 12) One method is to simply record the screen display using another program such as CamStudio or SnagIt. Another approach for making 3D animations of rotating scenes requires a little more work:

1. Layout the scene you want and size the 3D Scene Window for image size you want in your animation
2. Put the scene in the desired starting orientation and uncheck the Spin item in the View menu.
3. Save the scene image to a file using Export As... Bitmap.
4. Open the View menu Rotation... dialog and enter a horizontal angle, say 36 degrees.
5. Go to step (3) to continue saving and rotating the scene until you've done this 9 times.

You now have a sequence of images showing the scene at different angles. You can put these together into an animation using your favorite image editing software such as Corel Photopaint, and sometimes reduce the images to 8-bit color to save the movie as an animated GIF. You could also save the VRML of your scene and import it into sophisticated graphics software like 3D Studio MAX for rendering and creating animations. This will produce much prettier images, but it works in a similar way. Another alternative is to display your VRML file using a VRML viewer during your presentation. A good VRML viewer can rotate the scene for you.

How can I view a saved VRML scene?

Display your VRML file using a VRML viewer such as FluxPlayer from Media Machines. See also the information on VRML files on the File Formats page.

Flip	90° Rotate	- Rotate	+ Rotate	+ Horz. Slant	- Vertical Slant	+ Vertical Slant	- Horz. Slant	- Horz. Scale	- Vertical Scale	+ Vertical Scale	+ Horz. Scale
F1	F2	F3	F4	F5	F6	F7	F8	F9	F10	F11	F12
Ctrl small movement	Shift large movement	Ctrl-0 1 .. 9 get palette entry	Ctrl-0 1 .. 9 set palette entry	Page Up/Down page sections	Page Up/Down page sections	Backspace undo trace point	Home center	End last zoom	Brightness	Contrast	

Print without Fit to Page setting turned on.

Cut on dotted lines. Place over F1-F12 function keys on keyboard.