

Developer Magic™: Debugger User's Guide

Document Number 007-2579-005

St. Peter's Basilica image courtesy of ENEL SpA and InfoByte SpA. Disk Thrower image courtesy of Xavier Berenguer, Animatica.

Copyright © 1996, 1999 Silicon Graphics, Inc. All Rights Reserved. This manual or parts thereof may not be reproduced in any form unless permitted by contract or by written permission of Silicon Graphics, Inc.

LIMITED AND RESTRICTED RIGHTS LEGEND

Use, duplication, or disclosure by the Government is subject to restrictions as set forth in the Rights in Data clause at FAR 52.227-14 and/or in similar or successor clauses in the FAR, or in the DOD, DOE or NASA FAR Supplements. Unpublished rights reserved under the Copyright Laws of the United States. Contractor/manufacturer is Silicon Graphics, Inc., 1600 Amphitheatre Pkwy., Mountain View, CA 94043-1351.

IRIS, IRIX, and Silicon Graphics are registered trademarks and Developer Magic, and the Silicon Graphics logo are trademarks of Silicon Graphics, Inc. Open Software Foundation, Motif, OSF, OSF/Motif are trademarks of the Open Software Foundation, Inc. PostScript is a trademark of Adobe Systems, Inc. UNIX is a registered trademark in the United States and other countries, licensed exclusively through X/Open Company Limited. X/Open is a trademark of X/Open Company Ltd. The X device is a trademark of the Open Group.

New Features

Developer Magic™: Debugger User's Guide

007–2579–005

This revision of the *Developer Magic: Debugger User's Guide* supports the 2.8 release of MIPSpro WorkShop tools. New features include the following:

- Improved pthread support
 - stability
 - ability for user to continue a single POSIX 6.5 pthread
- Fix and continue for -n32 abi
- Trap manager improvements
 - ability to save breakpoints across debugging sessions
 - ability to set a breakpoint at a line number in a non-loaded DSO — an enhancement to pending traps, introduces in WorkShop 2.7
- Modules-based installation

Record of Revision

<i>Version</i>	<i>Description</i>
2.7	June 1998 Revised to reflect changes for the ProDev WorkShop 2.7 release.
2.8	April 1999 Revised to reflect changes for the MIPSpro WorkShop 2.8 release.

Contents

	<i>Page</i>
About This Guide	xxi
Related Publications	xxii
Obtaining Publications	xxii
Conventions	xxii
Reader Comments	xxiv
WorkShop Debugger Overview [1]	1
Main Debugger Features	1
The Debugger Main View Window	2
About Traps	3
Viewing Program Data	4
Integrating the Debugger with Other WorkShop Tools	5
Accessing the Performance Analyzer from the Main View Window	5
Accessing the Static Analyzer from the Main View Window	6
Accessing Editors from the Main View Window	6
Accessing Configuration Management Tools	7
Recompiling from the Main View Window	7
Debugging with Fix+Continue	8
Debugging with the X/Motif Analyzer	8
Customizing the Debugger	8
Additional Information	9
Basic Debugger Usage [2]	11
Getting Started with the Debugger	11
Basic Tips and Features	11
007-2579-005	iii

	<i>Page</i>
Fortran 90 Code Example and Short Tutorial	13
Example 1: Fortran 90 Example	13
C Code Example and Short Tutorial	15
Example 2: C Code Example	16
Options for Controlling Program Execution	19
Setting Traps	20
Options for Viewing Variables	20
Viewing Variables Using the <code>cvd></code> Command Line	21
Viewing Variables Using Click To Evaluate	22
Viewing Variables Using the Variable Browser	23
Viewing Variables Using the Expression View Window	23
Viewing Variables Using the Array Browser	25
Viewing Variables Using the Structure Browser Window	27
Searching	27
Using the Call Stack View	28
Stopping at Functions or Subroutines	30
Suggestions for Debugging for Serial Execution of Scientific Programs	31
Step 1: Use <code>lint</code>	32
Step 2: Check for Out-of-Bounds Array Accesses	32
Step 3: Check for Uninitialized Variables Being Used in Calculations	34
Step 4: Find Divisions by Zero and Overflows	35
Step 5: Perform Core File Analysis	36
Step 6: Troubleshoot Incorrect Answers	37
Selecting Source Files [3]	39
How to Load Source Files	39
Load Directly into the Main View Window	39
Load from the File Browser Dialog Box	40

	<i>Page</i>
Load from the Open Dialog Box	41
Path Remapping	43
Case Example for Path Remapping	45
Tutorial: The jello Program [4]	47
Starting the Debugger	47
Run the jello Program	48
Perform a Search	51
Edit Your Source Code	53
Setting Traps	54
Examining Data	58
Exiting the Debugger	67
Setting Traps [5]	69
Traps Terminology	70
Trap Triggers	70
Trap Types	71
Setting Traps	72
Setting Traps with the Mouse	73
Setting Traps Using the <code>cvd></code> Command Line	73
Setting Traps Using the Traps Menu in the Main View Window	73
Setting Traps in the Trap Manager Window	76
Setting Single-Process and Multiprocess Traps	77
Syntaxes	78
Setting a Trap Condition	82
Setting a Trap Cycle Count	83
Setting a Trap with the Traps Menu	83
Moving around the Trap Display Area	84
Enabling and Disabling Traps	84

	<i>Page</i>
Saving and Reusing Trap Sets	84
Setting Traps by Using Signal Panel and System Call Panel	84
Controlling Program Execution [6]	87
The Main View Window Control Panel	87
Features of the Main View Window Control Panel	87
Execution Control Buttons	88
Controlling Program Execution Using the PC Menu	91
Execution View	92
Viewing Program Data [7]	93
Traceback through Call Stack View	93
Options for Viewing Variables	96
Using the cvd> Command Line	97
Using Click to Evaluate	97
Using the Array Browser	97
Using the Structure Browser	97
Using the Variable Browser	98
Using the Expression View	98
Evaluating Expressions	99
Expression View Window	99
Assigning Values to Variables	102
Evaluating Expressions in C	103
C Function Calls	104
Evaluating Expressions in C++	104
Limitations	104
Evaluating Expressions in Fortran	105
Fortran Variables	106
Fortran Function Calls	106

	<i>Page</i>
Debugging with Fix+Continue [8]	107
Introduction to Fix + Continue	107
Fix+Continue Functionality	107
Fix+Continue Integration with Debugger Views	108
How Redefined Code Is Distinguished from Compiled Code	108
The Fix+Continue Interface	109
Debugger with Fix+Continue Support	109
Change ID, Build Path, and Other Concepts	109
Restrictions on Fix+Continue	110
Fix+Continue Tutorial	111
Setting up the Sample Session	111
Redefining a Function: <code>time1</code> Program	112
Editing a Function	113
Changing Code	115
Saving Changes	118
Setting Breakpoints in Redefined Code	118
Viewing Status	121
Comparing Original and Redefined Code	121
Switching between Compiled and Redefined Code	121
Comparing Function Definitions	122
Comparing Source Code Files	124
Ending the Session	124
Detecting Heap Corruption [9]	125
Typical Heap Corruption Problems	125
Finding Heap Corruption Errors	126
Compiling with the Malloc Library	126
Setting Environment Variables	127

	<i>Page</i>
Trapping Heap Errors Using the Malloc Library	128
Heap Corruption Detection Tutorial	129
Multiple Process Debugging [10]	135
Using the Multiprocess View Window	136
Starting a Multiprocess Session	137
Viewing Process Status	137
Using Multiprocess View Control Buttons	138
Multiprocess Traps	138
Viewing Pthreaded Applications	138
Adding and Removing Processes	139
Multiprocess Preferences	139
Bringing up Additional Main View Windows	139
Debugging a Multiprocess C Program	139
Launch the Debugger in Multiprocess View	141
Using Multiprocess View to Control Execution	142
Using the Trap Manager to Control Trap Inheritance	144
Debugging a Multiprocess Fortran Program	146
General Fortran Debugging Hints	146
Fortran Multiprocess Debugging Session	146
Debugging Procedure	148
Debugging a Pthreaded Program	152
User-Level Continue of Single 6.5 POSIX Pthread	153
Scheduling Anomalies	153
Blocking Anomalies	155
How to Continue a Single POSIX 6.5 Pthread	155
Other Pthread Debugging Hints	156
Differences between 6.4 and 6.5 Pthreads	156

	<i>Page</i>
Pthread Debugging Session	157
X/Motif Analyzer [11]	159
Introduction to the X/Motif Analyzer	159
Examiners Overview	159
Examiners and Selections	160
Inspecting Data	160
Inspecting the Control Flow	160
Tracing the Execution	161
Restrictions and Limitations	161
X/Motif Analyzer Tutorial	161
Setting up the Sample Session	162
Launching the X/Motif Analyzer	163
Navigating the Widget Structure	163
Examining Widgets	166
Setting Callback Breakpoints	168
Using Additional Features of the Analyzer	172
Ending the Session	175
Customizing the Debugger [12]	177
Customizing the Debugger with Scripts	177
Using a Startup File	177
Implementing User-Defined Buttons	178
Changing X Window System Resources	180
DUMPCORE Environment Variable	182
Appendix A Debugger Reference	183
Main View Window	183
Admin Menu	190
Views Menu	192

	<i>Page</i>
Query Menu	194
Source Menu	194
Display Menu	197
Perf Menu	197
Traps Menu	199
PC Menu	201
Fix+Continue Menu	201
Show Difference Submenu	202
View Submenu	203
Preferences Submenu	203
Keyboard Accelerators	205
Help Menu	205
Some Additional Views	206
Execution View	206
Multiprocess View	206
Status of Processes	207
Multiprocess View Control Buttons	208
Multiprocess View Administrative Functions	208
Controlling Preferences	209
Source View	210
Process Meter	213
Charts Menu	214
Scale Menu	215
Controlling Multiple Processes	215
Ada-specific Windows	216
Task View	216
Admin Menu	217
Config Menu	218
Layout Menu	218
Display Menu	218

	<i>Page</i>
Exception View	219
X/Motif Analyzer Windows	222
Global Objects	223
Admin Menu	224
Examine Menu	224
Examiner Tabs	225
Return Button	226
Breakpoints Examiner	226
Callback Breakpoints Examiner	229
Event-Handler Breakpoints Examiner	231
Resource-Change Breakpoints Examiner	233
Timeout-Procedure Breakpoints Examiner	234
Input-Handler Breakpoints Examiner	235
State-Change Breakpoints Examiner	236
X-Event Breakpoints Examiner	237
X-Request Breakpoints Examiner	239
Trace Examiner	240
Widget Examiner	242
Tree Examiner	243
Callback Examiner	246
Window Examiner	247
Event Examiner	248
Graphics Context Examiner	248
Pixmap Examiner	249
Widget Class Examiner	250
Trap Manager Windows	251
Trap Manager	252
Config Menu	253
Traps Menu	254

	<i>Page</i>
Display Menu	254
Signal Panel	254
Syscall Panel	255
Data Examination Windows	256
Array Browser Window	257
Spreadsheet Menu	261
Format Menu	262
Render Menu	263
Color Menu	263
Scale Menu	265
Examiner Viewer Controls	266
Examiner Viewer Menu	268
Call Stack View Window	271
Config Menu	273
Display Menu	273
Expression View Window	273
Config Menu	274
Display Menu	275
Language Pop-up Menu	275
Format Pop-up Menu	275
File Browser Window	276
Structure Browser Window	277
Using the Structure Browser Overview Window to Navigate	279
Entering Expressions	279
Working in the Structure Browser Display Area	279
Structure Browser Display Menu	280
Node Menu	282
Formatting Fields	284
Variable Browser Window	288
Entering Variable Values	288

	<i>Page</i>
Changing Variable Column Widths	289
Viewing Variable Changes	289
Machine-level Debugging Windows	290
The Disassembly View Window	290
Similarities with Main View Window	291
The Disassemble Menu	292
The Config Menu Preferences Dialog	294
The Register View Window	296
The Register View Window	297
Changing the Register View Display	298
The Memory View Window	299
Viewing a Portion of Memory	300
Changing the Contents of a Memory Location	301
Changing the Memory Display Format	301
Moving around the Memory View Display Area	301
Fix+Continue Windows	301
Fix+Continue Status Window	302
Admin Menu	304
View Menu	304
Fix+Continue Menu	305
Fix+Continue Error Messages Window	307
Admin Menu	307
View Menu	307
Fix+Continue Build Environment Window	308
Changes to Debugger Views	309
Main View	309
Command Line Interface	310
Call Stack View	311
Trap Manager	311
Debugger Command Line	312

	<i>Page</i>
Syntax for dbx-style Commands	312
Blocking Kernel System Calls	326
Appendix B Using the Build Manager	329
Build View Window	329
Build Process Control Area	330
Transcript Area	331
Error List Area	332
Build View Admin Menu	332
Build View Preferences	332
Build Options	333
Using Build View	334
Build Analyzer Window	335
Build Specification Area	336
Build Graph Area	336
Build Graph Control Area	338
Build Analyzer Overview Window	339
Build Analyzer Menus	340
Admin Menu	340
Build Menu	340
Filter Menu	340
Query Menu	341
Index	343
Figures	
Figure 1. The WorkShop Debugger Main View Window	3
Figure 2. Evaluation Pop-Up Menu	22
Figure 3. Expression View Window	24
Figure 4. Array Browser Window	25

	<i>Page</i>
Figure 5. Search Window	28
Figure 6. Call Stack View Window	29
Figure 7. Trap Manager Window	31
Figure 8. File Browser Window	40
Figure 9. Open Dialog Box	42
Figure 10. Path Remapping Dialog Box	44
Figure 11. The Main View Window with jello Source Code	50
Figure 12. The jello Window	51
Figure 13. The Search Dialog	52
Figure 14. Search Target Indicators	53
Figure 15. Stop Trap Indicator	56
Figure 16. Trap Manager Window	57
Figure 17. Call Stack View at spin Stop Trap	59
Figure 18. Variable Browser at spin	60
Figure 19. Variable Browser after Changes	61
Figure 20. Expression View with Language and Format Menus Displayed	62
Figure 21. Structure Browser Window with jello_conec Structure	63
Figure 22. Structure Browser Window with Next Pointer De-referenced	64
Figure 23. Array Browser Window for shadow Matrix	65
Figure 24. Subscript Controls Panel in Array Browser Window	66
Figure 25. Traps Menu in the Main View Window	74
Figure 26. Typical Trap Icons	75
Figure 27. Trap Manager Config, Traps, and Display Menus	76
Figure 28. Trap Examples	78
Figure 29. Signal Panel and System Call Panel	85
Figure 30. The Main View Window Control Panel	87
Figure 31. Pop-up Menu and Step Into Dialog	89

	<i>Page</i>
Figure 32. Pop-up Menu and Step Over Dialog	90
Figure 33. Call Stack View Window	94
Figure 34. Tracing through Call Stack View	96
Figure 35. Expression View with Major Menus Displayed	100
Figure 36. Change Indicators in Expression View	101
Figure 37. Program Results in Execution View	112
Figure 38. Selecting a Function for Redefinition	114
Figure 39. Redefined Function	115
Figure 40. Stopping after Breakpoints in Redefined Code	120
Figure 41. Comparing Compiled and Redefined Function Code	123
Figure 42. Heap Corruption Warning Shown in Execution View	131
Figure 43. Call Stack at Boundary Overrun Warning	131
Figure 44. Main View at Bus Error	133
Figure 45. Multiprocess View	137
Figure 46. Examining Process State Using Multiprocess View	144
Figure 47. Comparing Variable Values from Two Processes	152
Figure 48. First View of the X/Motif Analyzer (Widget Examiner)	164
Figure 49. Widget Hierarchy Displayed by the Tree Examiner	165
Figure 50. Adding a Breakpoint for a Widget	167
Figure 51. Setting Breakpoints for a Widget Class	169
Figure 52. Callback Context Displayed by the Callback Examiner	170
Figure 53. Window Attributes Displayed by the Window Examiner	171
Figure 54. Selecting the Breakpoints Tab from the Overflow Area	173
Figure 55. Breakpoint Results Displayed by the Call Stack View	174
Figure 56. User-Defined Button Example	178
Figure 57. Major Areas of the Main View Window	184
Figure 58. Show/Hide Annotations Button in the Main View Window	189

	<i>Page</i>
Figure 59. Perf Menu and Subwindows	198
Figure 60. Process Menu	207
Figure 61. Source View Window	211
Figure 62. Process Meter	214
Figure 63. Process Menu	215
Figure 64. Task View Window	217
Figure 65. Exception View	220
Figure 66. Launching the X/Motif Analyzer Window	223
Figure 67. Examiner Tabs	225
Figure 68. Breakpoints Examiner Display in the X/Motif Analyzer Window	228
Figure 69. Callback Breakpoints Examiner	230
Figure 70. Event-Handler Breakpoints Examiner	232
Figure 71. Timeout-Procedure Breakpoints Examiner	234
Figure 72. Input-Handler Breakpoints Examiner	235
Figure 73. State-Change Breakpoints Examiner	236
Figure 74. X-Event Breakpoints Examiner	238
Figure 75. X-Request Breakpoints Examiner	239
Figure 76. Request Type Selection Dialog	240
Figure 77. Trace Examiner	241
Figure 78. Widget Examiner	242
Figure 79. Tree Examiner	244
Figure 80. Tree Examiner Window Graphical Buttons	245
Figure 81. Callback Examiner	246
Figure 82. Window Examiner	247
Figure 83. Event Examiner	248
Figure 84. Graphics Context Examiner	249
Figure 85. Pixmap Examiner	250

	<i>Page</i>
Figure 86. Widget Class Examiner	251
Figure 87. Trap Manager Window	252
Figure 88. Signal Panel	255
Figure 89. Syscall Panel	256
Figure 90. Array Browser with Display Menu Options	258
Figure 91. Subscript Controls Area in the Array Browser	259
Figure 92. Array Browser Spreadsheet Area	261
Figure 93. Example of Wrapped Array	262
Figure 94. Color Exception Portion of Array Browser Window	264
Figure 95. Array Browser Graphic Modes	265
Figure 96. Examiner Viewer with Controls and Menus	267
Figure 97. Examiner Viewer Preference Sheet Dialog	270
Figure 98. Call Stack View	272
Figure 99. Expression View	274
Figure 100. Expression View Format Popup with Submenus	275
Figure 101. File Browser Window	277
Figure 102. Structure Browser with Menus Displayed	278
Figure 103. Tree and Linked List Arrangements of Structures	281
Figure 104. Node Menu	282
Figure 105. Structure Browser Preferences Dialog	284
Figure 106. Structure Browser Type Formatting Dialog	286
Figure 107. Variable Browser with Menus Displayed	289
Figure 108. Typical Variable Change Indicators	290
Figure 109. The Disassembly View Window with the Disassembly Menu Displayed	291
Figure 110. The Disassemble From Address Dialog	292
Figure 111. The Disassemble Function Dialog	293
Figure 112. The Disassemble File Dialog	294

	<i>Page</i>
Figure 113. The Disassembly View Preferences Dialog with Display Format Menu	295
Figure 114. The Register View Window	297
Figure 115. The Register View Preferences Dialog	299
Figure 116. The Memory View Window with the Mode Submenu Displayed	300
Figure 117. Fix+Continue Status Window	303
Figure 118. Fix+Continue Status Window Menus	304
Figure 119. Fix+Continue Build Environment Window	308
Figure 120. Debugger Main View Window	310
Figure 121. Command Line Interface with Redefined Function	311
Figure 122. Call Stack View	311
Figure 123. Trap Manager Window with Redefined Function	312
Figure 124. Build Process Control Area in Build View Window	330
Figure 125. Build View Window with Typical Data	331
Figure 126. Build View Preferences Dialog	333
Figure 127. Build Options Dialog	334
Figure 128. Build Analyzer Window	336
Figure 129. Build Graph Icons	338
Figure 130. Build Graph Control Area	338
Figure 131. Build Analyzer Overview Window with Build Analyzer Graph	339
 Tables	
Table 1. Valid C Operations	103
Table 2. Valid Fortran Operations	105
Table 3. Fix+Continue Keyboard Accelerators	205

About This Guide

This publication documents the Developer Magic Debugger, released with the 2.8 version of MIPSpro WorkShop tools running on IRIX systems.

The Debugger is a source-level debugging tool that allows you to see program data, monitor program execution, and fix code for Ada, C, C++, Fortran 77, and Fortran 90.

This manual contains the following chapters:

- Chapter 1, page 1, gives you an introductory functional overview of the Developer Magic Debugger.
- Chapter 2, page 11, outlines principles and procedures of the debugging process and how to approach them using the WorkShop Debugger.
- Chapter 3, page 39, describes how to manage source files.
- Chapter 4, page 47, presents a short Debugger tutorial based on demonstration programs provided with your WorkShop tools.
- Chapter 5, page 69, describes how to set various types of traps.
- Chapter 6, page 87, describes methods for controlling process execution.
- Chapter 7, page 93, explains how to examine Debugger data.
- Chapter 8, page 107, presents a short tutorial using Fix and Continue.
- Chapter 9, page 125, describes heap corruption problems and how to detect them.
- Chapter 10, page 135, describes debugging multiprocess programs.
- Chapter 11, page 159, presents a short tutorial using the X/Motif Analyzer
- Chapter 12, page 177, gives you tips on how you can customize the Debugger to the requirements of your working environment.
- Appendix A, page 183, describes all of the Debugger windows, menus, and other features in detail.
- Appendix B, page 329, describes the use of the Build Manager.

Related Publications

The following documents contain additional information that may be helpful:

- *C++ Language System Library*
- *C++ Language System Overview*
- *C++ Language System Release 3.0.1 Product Reference Manual*
- *C++ Programmer's Guide*
- *Developer Magic: Performance Analyzer User's Guide*
- *Developer Magic: ProDev WorkShop Overview*
- *Developer Magic: Static Analyzer User's Guide*
- *Fortran 77 Language Reference Manual*
- *dbx User's Guide*

Obtaining Publications

Silicon Graphics maintains publications information at the following web site:

<http://techpubs.sgi.com/library>

This library contains information that allows you to browse documents online, order documents, and send feedback to Silicon Graphics.

To order a printed Silicon Graphics document, call 1-800-627-9307.

Customers outside of the United States and Canada should contact their local service organization for ordering and documentation information.

Conventions

The following conventions are used throughout this document:

<u>Convention</u>	<u>Meaning</u>
command	This fixed-space font denotes literal items such as commands, files, routines, path names, signals, messages, and programming language structures.

<code>manpage(x)</code>	<p>Man page section identifiers appear in parentheses after man page names. The following list describes the identifiers:</p> <table><tr><td>1</td><td>User commands</td></tr><tr><td>1B</td><td>User commands ported from BSD</td></tr><tr><td>2</td><td>System calls</td></tr><tr><td>3</td><td>Library routines, macros, and opdefs</td></tr><tr><td>4</td><td>Devices (special files)</td></tr><tr><td>4P</td><td>Protocols</td></tr><tr><td>5</td><td>File formats</td></tr><tr><td>7</td><td>Miscellaneous topics</td></tr><tr><td>7D</td><td>DWB-related information</td></tr><tr><td>8</td><td>Administrator commands</td></tr></table> <p>Some internal routines (for example, the <code>_assign_asgcmd_info()</code> routine) do not have man pages associated with them.</p>	1	User commands	1B	User commands ported from BSD	2	System calls	3	Library routines, macros, and opdefs	4	Devices (special files)	4P	Protocols	5	File formats	7	Miscellaneous topics	7D	DWB-related information	8	Administrator commands
1	User commands																				
1B	User commands ported from BSD																				
2	System calls																				
3	Library routines, macros, and opdefs																				
4	Devices (special files)																				
4P	Protocols																				
5	File formats																				
7	Miscellaneous topics																				
7D	DWB-related information																				
8	Administrator commands																				
<i>variable</i>	<p>Italic typeface denotes variable entries and words or concepts being defined.</p>																				
user input	<p>This bold, fixed-space font denotes literal items that the user enters in interactive sessions. Output is shown in nonbold, fixed-space font.</p>																				
[]	<p>Brackets enclose optional portions of a command or directive line.</p>																				
...	<p>Ellipses indicate that a preceding element can be repeated.</p>																				

Reader Comments

If you have comments about the technical accuracy, content, or organization of this document, please tell us. Be sure to include the title and part number of the document with your comments.

You can contact us in any of the following ways:

- Send electronic mail to the following address:

`techpubs@sgi.com`

- Send a facsimile to the attention of "Technical Publications" at fax number +1 650 932 0801.
- Use the Suggestion Box form on the Technical Publications Library World Wide Web page:

`http://techpubs.sgi.com/library/`

- Call the Technical Publications Group, through the Technical Assistance Center, using one of the following numbers:

For Silicon Graphics IRIX based operating systems: 1 800 800 4SGI

For UNICOS or UNICOS/mk based operating systems or CRAY Origin2000 systems: 1 800 950 2729 (toll free from the United States and Canada) or +1 651 683 5600

- Send mail to the following address:

Technical Publications
Silicon Graphics, Inc.
1600 Amphitheatre Pkwy.
Mountain View, California 94043-1351

We value your comments and will respond to them promptly.

WorkShop Debugger Overview [1]

The Silicon Graphics MIPSpro WorkShop Debugger is a UNIX source-level debugging tool for Silicon Graphics MIPS systems. It displays program data and execution status in real time. This tool can be used to debug Ada, C, C++, FORTRAN 77, and Fortran 90.

Note: Currently, the Debugger does not accommodate all Fortran 90 derived syntax, but it does accommodate Fortran 90 array syntax.

For an introductory tutorial to the principles of debugging, particularly with the WorkShop Debugger, see Chapter 2, page 11.

This chapter presents an overview of the WorkShop Debugger and is divided into the following sections:

- Main Debugger Features, Section 1.1, page 1
- Debugging with Fix+Continue, Section 1.2, page 8
- Debugging with the X/Motif Analyzer, Section 1.3, page 8
- Customizing the Debugger, Section 1.4, page 8

1.1 Main Debugger Features

The following sections outline the primary features and functions of the WorkShop Debugger and include references to comprehensive information found throughout this manual:

- The Debugger Main View Window, Section 1.1.1, page 2
- About Traps, Section 1.1.2, page 3
- Viewing Program Data, Section 1.1.3, page 4
- Integrating the Debugger with Other WorkShop Tools, Section 1.1.4, page 5

1.1.1 The Debugger Main View Window

When you start the Debugger with an executable file, the Main View window displays, loaded with source code, ready to execute your program with your specified arguments. Most of your debugging work takes place in the Main View window, which includes the following:

- A **menu bar** for performing debugger functions
- A **control panel** for specifying and controlling program execution
- A **source code display area** which displays the code for the program you are debugging.
- A **source filename** field which tells gives you the path to the file displayed in the source code display area
- A **status area** for viewing the current status of the program
- The Debugger **command line** in which to enter debugging commands (see Section A.9, page 312, for command syntax)

The major areas of the Main View window are shown in Figure 1, page 3.

Note: For a comprehensive description of the Main View window, see Section A.1, page 183.

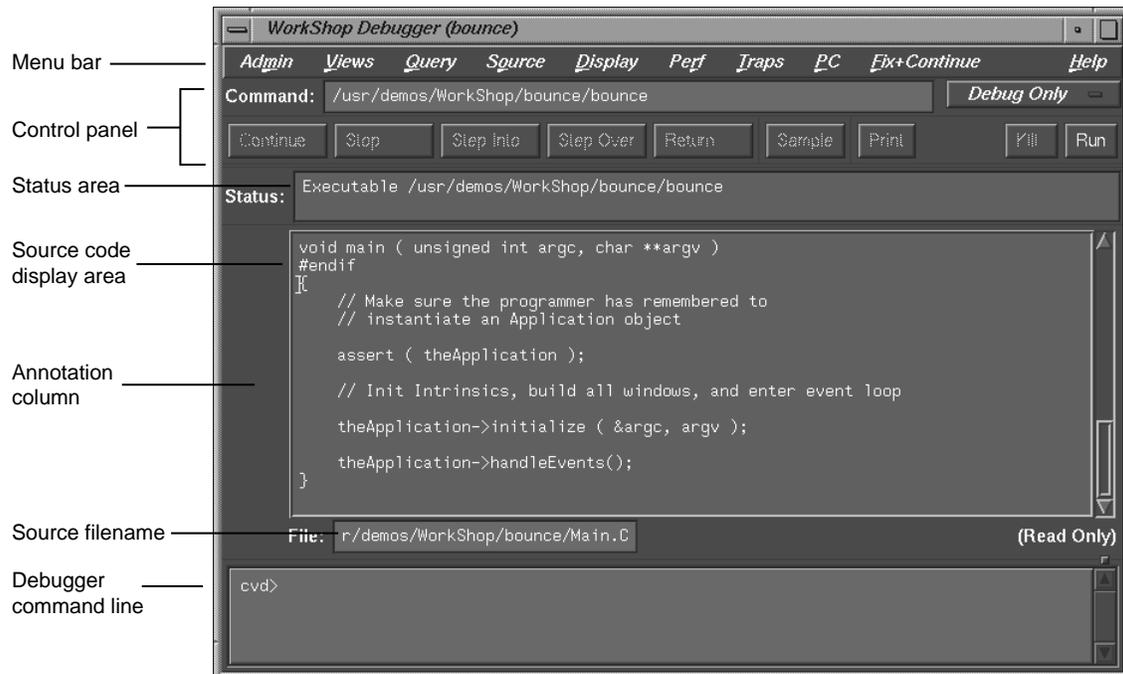


Figure 1. The WorkShop Debugger Main View Window

1.1.2 About Traps

Part of the debugging process requires that you inspect data at various points during program execution. A trap is a mechanism for gathering this data. There are two categories of traps:

- A *stop trap* halts a process so that you can manually examine data.
- A *sample trap* collects specific performance data without stopping.

The Debugger lets you set traps at the following points:

- At a line in a file (a *breakpoint*). These are the most commonly used stop traps. You can set them by clicking in the annotation column to the left of an executable statement in the source code display panel. (See the annotation column in Figure 1.)
- At an instruction address.
- On entry to or exit from a function.

- When a signal is received.
- When a system call is made, at either the entry or exit point.
- When a given variable or address is written to, read from, or executed (a *watchpoint*).
- At set time intervals (a *pollpoint*).

For more information on traps, refer to Chapter 5, page 69.

1.1.3 Viewing Program Data

When you stop a process, the Views menu at the Main View window menu bar provides several options for viewing your data. These Views menu options allow you to inspect the following types of data:

- Call Stack

This option allows you to inspect the call stack at the breakpoints.

- Expression View

This option allows you to inspect the value of specified expressions.

- Variable Browser

This option allows you to inspect the values, types, or addresses of variables.

- Structure Browser

This option allows you to inspect data structures.

- Multiprocess View

This option allows you to inspect the values of multiple and / or pthreaded processes.

- Array Browser

This option allows you to inspect the values of an array variable.

- Memory View
This option allows you to inspect the values in specified memory locations.
- Register View
This option allows you to inspect registers.
- Disassembly View
This option allows you to inspect the disassembled code.

1.1.4 Integrating the Debugger with Other WorkShop Tools

MIPSpro WorkShop tools are designed so that you can move easily between them in a work session.

1.1.4.1 Accessing the Performance Analyzer from the Main View Window

You can run the Performance Analyzer from the Main View window as follows:

Note: If you are in the middle of a run, you must first terminate it.

1. Select the following from the menu bar:

```
Perf
  ->Select Task
      ->Task in List
```

2. Click on the RUN button in the Control Panel.

The executable is run.

3. Select the following from the menu bar:

```
Perf
  ->Examine Results
```

The Performance Analyzer window will display your results. For more information on the Performance Analyzer, see the *Developer Magic: Performance Analyzer User's Guide*.

1.1.4.2 Accessing the Static Analyzer from the Main View Window

The Static Analyzer displays information that makes it easier to determine where to set traps in your source code. To launch the Static Analyzer, select the following from the menu bar:

```
Admin
  ->Launch Tool
    ->Static Analyzer
```

For more information on the Static Analyzer, see the *Developer Magic: Static Analyzer User's Guide*.

1.1.4.3 Accessing Editors from the Main View Window

Once you have isolated your code problem with the WorkShop tools, you will want to correct and recompile your source. WorkShop offers several ways to do this:

1. You can select the following from the Source Menu to make code changes in the source pane of the Main View window:
 - a. Select Source > Make Editable. Make your changes accordingly. (Make Editable toggles with Make Read Only.)
 - b. Select Source > Save to save your changes.
 - c. Select Source > Recompile to recompile your changed code.
2. You can invoke an editor from the Views menu that is a separate window in which to edit your code:

```
Views
  ->Source View
```

Note: You must select Make Editable from the File menu at this point to proceed.

3. You can call up a fork editor window to launch your own editor as follows (first path):

```
Source
  ->Fork Editor
```

4. You can call up a fork editor window to launch your own editor as follows (second path):

```
Views  
  ->Source View
```

Then, from the Source View window select the following:

```
File  
  ->Fork Editor
```

1.1.4.4 Accessing Configuration Management Tools

If you use ClearCase (a Silicon Graphics product), RCS, or SCCS for configuration management, you can integrate the tool into the WorkShop environment by entering the following command:

```
cvconfig [clearcase | rcs | sccs]
```

This will allow you to use Versioning source control (from the Source menu) to check files in and out.

1.1.4.5 Recompiling from the Main View Window

You can recompile your code from the Build View window, accessible from the menu bar as follows:

```
Source  
  ->Recompile
```

For more on the Build View window, see Section B.1, page 329.

To examine build dependencies for your code, launch the Build Analyzer from the menu bar as follows:

```
Admin  
  ->Launch Tool  
    ->Build Analyzer
```

For more on the Build Analyzer window, see Section B.6, page 335.

For general information on the Build Manager tools, see Appendix B, page 329.

1.2 Debugging with Fix+Continue

Fix+Continue gives you the ability to make changes to a program written in C or C++ without having to recompile and link the entire program before continuing to debug the code. With Fix+Continue, you can edit a function, parse the new functions, and continue execution of the program being debugged. Fix+Continue commands may be issued from the Fix+Continue window, which you launch from the `Fix+Continue` item on the Main View menu bar.

You can also issue Fix+Continue commands from the Debugger's command line.

See Chapter 8, page 107, for a comprehensive description of the theory and operation of Fix+Continue, including a short tutorial.

1.3 Debugging with the X/Motif Analyzer

The X/Motif Analyzer provides specific debugging support for X/Motif applications. The X/Motif analyzer is integrated with the Debugger. You issue X/Motif analyzer commands graphically from the X/Motif analyzer subwindow of the Debugger Main View. Select the following from the Main View window to access this subwindow:

```
Views
  ->X/Motif Analyzer
```

See Chapter 11, page 159 for a comprehensive description of the theory and operation of the X/Motif Analyzer, including a short tutorial.

1.4 Customizing the Debugger

WorkShop provides you with a number of ways that you can customize your Debugger as best suited to the needs of your development environment.

See Chapter 12, page 177, for more tips on customizing the Debugger to your specific needs.

1.5 Additional Information

In addition to this manual and the manuals listed in "About This Guide," page 8, you should consult the following sources regarding the WorkShop Debugger:

- The *WorkShop 2.8 Release Notes* are installed on any system running WorkShop 2.8. These notes provide the following information:
 - Installation instructions
 - Changes and additions since the last release
 - Bugfixes
 - Known problems and their workarounds

There are two ways to view the release notes:

- Enter the following command to view the relnotes as simple text:

```
% relnotes WorkShop
```

See the `relnotes(1)` man page for additional information.

- Enter the following command to view the relnotes through the graphical release notes viewer:

```
% grelnotes WorkShop
```

See the `grelnotes(1)` man page for additional information.

- The online Help menu includes a FAQ option. Consult this for a list of frequently asked questions and operating tips.

Basic Debugger Usage [2]

The WorkShop Debugger can be used with the following compilers: C, C++, Ada, FORTRAN 77, and Fortran 90. Currently, the Debugger does not handle all Fortran 90 syntax (that is, derived types), but it does handle Fortran 90 array syntax.

This chapter includes information regarding principles and procedures of the debugging process and how these are to be approached with the WorkShop Debugger.

2.1 Getting Started with the Debugger

Before starting a Debugger session from a remote workstation, you must first enter the following from a window:

```
xhost + machine_name
```

Where *machine_name* is the name or IP address of the machine where the program that you would like to debug will be run.

On the machine where your program will be running, enter the following:

```
echo $DISPLAY
```

The *machine_name* of your workstation should appear, followed by :0.0. If it does not, enter the following on the machine where your program will run (if you are using the `cs`h or `tc`sh shells):

```
setenv DISPLAY machine_name:0.0
```

For other shells, see their respective man pages.

2.1.1 Basic Tips and Features

Note: You can find short Debugger tutorials in Section 2.1.2, page 13 and Section 2.1.3, page 15 for Fortran 90 and C, respectively.

See also Appendix A, page 183 for a comprehensive description of Debugger functions.

To provide debugging information to the Debugger, compile your program with the `-g` option (this disables optimization and produces information for symbolic debugging).

To begin a Debugger session enter:

```
%cvd executable &
```

The Debugger Main View window will automatically appear along with an icon for the Execution View window.

If your program requires data to be read from a file named `input`, for example, then type the following in the command (`cvd>`) pane of the Main View window:

```
cvd> run < input
```

(See Figure 1, page 3.)

The Execution View window receives all the output from running your program that would normally go directly to your screen. The Main View window controls your Debugger session, displaying your source code in its center pane. If you want to see the line numbers for your program, select the following from the Main View window menu bar:

```
Display  
->Show Line Numbers
```

At the bottom of the Main View window there is a window with the `cvd>` prompt from which you can enter all `dbx` commands to the Debugger.

The Debugger allows you to run your program and stop at selected places so you can view current values of program variables to help you find bugs in your program. To stop at a selected statement in your program, you may either set a breakpoint (also called a stop trap) at the desired statement or set a breakpoint prior to the desired statement and then use either the `Step Into` or `Step Over` buttons to reach the desired stopping point. The statement where your program has stopped is indicated in green. A statement highlighted in red indicates that a breakpoint has been set on this line. When the Debugger causes your program to stop at a breakpoint that you have set, the executable statement immediately prior to the breakpoint has been executed, and the executable statement on which the breakpoint has been set has yet to be executed.

Note: Programs featured in this chapter are located in the following directory: `/usr/demos/WorkShop/getstarted`.

2.1.2 Fortran 90 Code Example and Short Tutorial

Use the Fortran 90 files `prog.f` and `dot.f` in `/usr/demos/WorkShop/getstarted` to demonstrate the Debugger features in the following tutorial.

Example 1: Fortran 90 Example

- The Fortran 90 code in the file `prog.f` is as follows:

```
program prog
parameter ( n=3 )
double precision A(n,n), x(n), y(, ) sum xydot

! initialize arrays
x = 2.0d0
do i = 1, n
  y(i) = i
  do j = 1, n
    A(i,j) = i*j - i
  enddo
enddo

! compute the dot product of x and y
call dot(x,y,n,xydot)
print *, 'dot product of x & y = ', xydot

! compute y = Ax
do i = 1, n
  sum = 0.0
  do j = 1, n
    sum = sum + A(i,j)*x(j)
  enddo
  y(i) = sum
enddo
print *, 'y = ', y
stop
end
```

- It includes subroutine dot in file dot.f, as follows:

```
subroutine dot(a,b,m,answer)
double precision a(m), b(m), answer
integer m
answer = 0.0d0
do i = 1, m
    answer = answer + a(i)*b(i)
enddo
end
```

Perform the following steps with these files to demonstrate Debugger features:

1. Enter the following command in the /usr/demos/WorkShop/getstarted directory to produce the executable program:

```
% f90 -g -o progf prog.f dot.f
```

This produces the executable progf.

2. Launch the WorkShop Debugger with your newly-compiled executable as follows:

```
% cvd progf &
```

The WorkShop Debugger Main View displays the source for your prog.f file (see Example 1, page 13).

3. Select the following from the Main View menu bar to turn on file line numbering:

```
Display
->Show Line Numbers
```

The line numbers will display to the left of the source code.

4. Enter a breakpoint at line 15 as follows at the cvd> prompt at the bottom of the Main View window. This will enable you to execute through the end of the initialization of the y array for the sample code:

```
cvd> stop at 15
```

Line 15 is highlighted in red and a stop icon appears in the left column.

5. Run the program. There are two ways that you can do this:
 - a. Click on the Run button at the top of the Main View window.

OR

- b. Enter the following at the `cvd>` prompt:

```
cvd> run
```

The program executes up to Line 15 and waits for further instruction.

6. Enter the following command at the `cvd>` prompt to print the `y` array for this example:

```
cvd> print y
```

The following displays in the `cvd>` command pane:

```
y =  
  (1) = 1.0  
  (2) = 2.0  
  (3) = 3.0  
cvd>
```

Note: You should expand this pane (or use the slider at the right side of the pane) if you do not see the printout.

7. At this point, you can experiment with other command described in this chapter, notably the execution control buttons described in Section 2.1.4, page 19.
8. Select the following to end this tutorial for the Fortran 90 demo program.

```
Admin  
->Exit
```

2.1.3 C Code Example and Short Tutorial

Use the C file `prog.c` and `dot.c` in `/usr/demos/WorkShop/getstarted` to demonstrate the Debugger features in the following tutorial.

Example 2: C Code Example

The following is the same example as the Fortran 90 example in Section 2.1.2, page 13, but it is written in C. Use this example to see how the Debugger can be used to view C structures.

- The C code in the file `prog.c` is as follows:

```
#include <stdio.h>
#define N 3
double dot(double v[],double w[], int m);
void main(){
    int i,j;
    double a[N][N],x[N],y[N],sum,xydot;
    struct node
    {
        int value;
        struct node *next;
    } *list,start;

    /* Initialize arrays */
    for(i=0;i<N;i++){
        x[i]=2;
        y[i]=i;
        for(j=0;j<N;j++){
            a[i][j]=i*j-i;
        }
    }

    /* Compute the dot product x and y */
    xydot=dot(x,y,N);
    printf("dot product of x & y: %f \n",xydot);

    /* Compute y=ax */
    for(i=0;i<N;i++){
        sum=0;
        for(j=0;j<N;j++){
            sum+=a[i][j]*x[j];
        }
        y[i]=sum;
    }
    printf("y = ");
    for(i=0;i<N;j++){
        printf("%f ",y[i]);
    }
}
```

```

    }
    printf("\n");

    /* Built list*/
    start.value=1;
    list=&start
    for(i=1;i<N;i++){
    list->next=(struct node *) malloc(sizeof(struct node));
    list=list->next;
    list->value=i;
    }
    list->next=NULL;
    printf("list: ");
    list=&start;
    for(i=0;i<N;i++){
        printf("%d ",list->value);
        list=list->next;
    }
    printf("\n");
}

```

- It includes function dot in file dot.c, as follows:

```

double dot(double v[],double w[], int m){
    int i;
    double sum;
    for(i=0;i<m;i++){
        sum+=v[i]*w[i];
    }
    return(sum);
}

```

Perform the following steps with this file to demonstrate Debugger features:

1. Enter the following command in the /usr/demos/WorkShop/getstarted directory to produce the executable program:

```
% cc -g -o progc prog.c dot.c
```

This produces the executable progc.

2. Launch the WorkShop Debugger with your newly-compiled executable as follows:

```
% cvd prog.c &
```

The WorkShop Debugger Main View displays the source for your `prog.c` file (see Example 2, page 16).

3. Select the following from the Main View menu bar to turn on file line numbering:

```
Display  
  ->Show Line Numbers
```

The line numbers will display to the left of the code source window.

4. Enter a breakpoint at line 23 as follows at the `cvd>` prompt at the bottom of the Main View window. This will enable you to execute up to the end of the `y` array for the sample code:

```
cvd> stop at 23
```

Line 23 is highlighted in red and a stop icon appears in the left column.

5. Run the program. There are two ways that you can do this:
 - a. Click on the Run button at the top of the Main View window.

OR

 - b. Enter the following at the `cvd>` prompt:

```
cvd> run
```

The program executes up to Line 23 and waits for further instruction.

6. Enter the following command at the `cvd>` prompt to print the `y` array for this example:

```
cvd> print y
```

The following displays in the `cvd>` command pane:

```
y = {  
  [0] 0.0000000000000000e+00  
  [1] 1.0  
  [2] 2.0  
}  
  
cvd>
```

7. At this point, you can experiment with other command described in this chapter, notably the execution control buttons described in Section 2.1.4, page 19.
8. Select the following to end this tutorial for the C demo program.

```
Admin  
->Exit
```

2.1.4 Options for Controlling Program Execution

There are a number of buttons in the Main View window that allow you to control the execution of your program. The following summarizes their functions:

- **Run** creates a new process to execute your program and starts execution. It can also be used to rerun your program.
- **Kill** kills the active process that is executing your program.
- **Stop** stops execution of your program. The first executable statement after the statement where your program has stopped is highlighted.
- **Continue** continues program execution until a breakpoint or some other event stops execution, or program execution terminates. (See also Section 10.4.1.3, page 155.)
- **Step Into** steps to the next executable statement and into function and subroutine calls. Thus, if you set a breakpoint at a subroutine call, click on the **Run** button so the call to the subroutine is highlighted in green, then click on the **Step Into** button to step into this subroutine — source code for this subroutine will automatically be displayed in the Main View window.

By clicking the right mouse button on the **Step Into** button you can select the number of steps the Debugger takes. Left-click on the **Step Into** button to take one step.

- **Step Over** steps to the next executable statement and steps over function and subroutine calls. Thus, if you set a breakpoint at a subroutine call, click on the **Run** button so the call to the subroutine is highlighted in green, then click the **Step Over** button to step over this subroutine to the next executable statement displayed in the source pane of the Main View window.

Right-click on the **Step Over** button to select the number of steps the Debugger takes. Left-click on the **Step Over** button to take one step.

- **Return** executes the remaining instructions in the current function or subroutine. Execution stops upon return from that function or subroutine.

2.1.5 Setting Traps

A trap (also called a *breakpoint*) can be set if you click your cursor in the area to the left of the statement in the area underneath the word **Status** in the Main View window. When you do this, the line in your program will be highlighted in red. To remove the breakpoint, click in the same place and the red highlighting will disappear. Clicking on the **Run** button will cause your program to run and stop at the first breakpoint encountered. To continue program execution, click on the **Continue** button. Breakpoints can only be set at executable statements.

2.1.6 Options for Viewing Variables

There are many ways to view current values of variables with the Debugger. Before you can view current values of variables with the Debugger. Before you can do this, you must first run your program under the Debugger and stop execution at some point. The following sections list ways of viewing current values of variables with the Debugger. It is suggested that you try each of the following methods to determine which you would prefer to use.

2.1.6.1 Viewing Variables Using the `cvd>` Command Line

At the bottom of the Main View window is the command/message pane. Here, you can give the Debugger various instructions at the `cvd>` command line.

Example 1: Value of array x

If you want to know the current value of array x , enter either of the following two commands in the command/message pane:

```
cvd> print x
```

or

```
cvd> p x
```

The current values for x will be printed.

Example 2: Value of $x(2)$

If your program is written in Fortran and you only want the value of $x(2)$, enter:

```
cvd> print x(2)
```

For a C program, enter:

```
cvd> print x[2]
```

Example 3: Change value of $x(2)$ to 3.1

To change the value of $x(2)$ to 3.1 in a Fortran program, enter:

```
cvd> assign x(2) = 3.1
```

For a C program, enter:

```
cvd> assign x[2] = 3.1
```

The value of $x(2)$ will now be 3.1 when execution is resumed.

Such changes are only active during the current Debugger run of your program. In the preceding examples, if x is a large array, you may want to use the `Array Browser` window (see Section 2.1.6.5, page 25).

To view the components of the structure `start` in the `prog.c` example, enter:

```
cvd> print start
```

The current values of each component of `start` will be printed.

To view what the pointer `list` points to in the `prog.c` example, enter:

```
cvd> print *list
```

Note: This pointer must be initialized before you can perform this function.

A complete list of the instructions that can be entered at the `cvd>` command line can be found in Section A.9, page 312.

2.1.6.2 Viewing Variables Using Click To Evaluate

Perform the following to view variables with `Click To Evaluate`:

1. Right-click in the window that contains your source code to bring up the following pop-up menu:



Figure 2. Evaluation Pop-Up Menu

2. Select `Click To Evaluate` from this menu. You can now click on any variable and its value will appear. For example:
 - If you click on the `x` in `x(i)`, the address of `x` will appear.
 - If you click-drag to highlight `x(i)`, the current value of `x(i)` displays.
 - If you highlight an expression, the current value of the expression displays.

2.1.6.3 Viewing Variables Using the Variable Browser

Perform the following to view variables with the Variable Browser:

1. Select the following from the Main View window to call up the Variable Browser:

```
Views  
  ->Variable Browser
```

The Variable Browser automatically displays values of all variables valid within the routine in which you have stopped as well as the address location for each array.

Values of variables can be changed by typing in their new value in the Result column and then hitting the ENTER key (or RETURN key, for some keyboards).

2.1.6.4 Viewing Variables Using the Expression View Window

The Expression View window allows you to enter the variables and/or expressions for which you would like to know values.

1. Select the following from the Main View window menu bar to display the Expression View window:

```
Views  
  ->Expression View
```

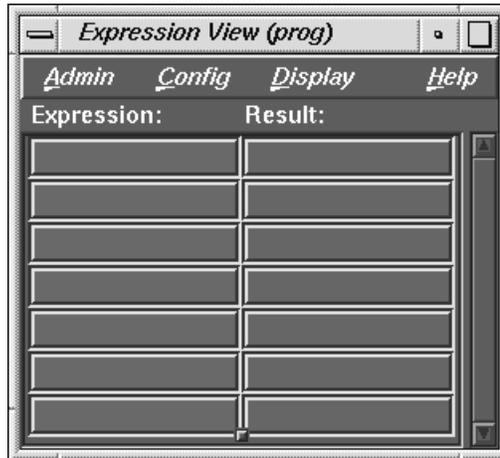


Figure 3. Expression View Window

2. To view, for example, the value component of the structure `start` in `prog.c`, enter the following in the Expression column:

`start.value`

As you step through your program, for example with `Step Over`, the values of all the entries in this window will be updated to their current values.

Values of variables can be changed by typing in their new value in the Result column and then pressing `ENTER`.

To enter an expression from your source code into the Expression View window:

1. Left-click and drag on the expression in your source code.
The expression is highlighted.
2. Left-click in a field in the Expression column.
The cursor appear in the field.
3. Middle-click your mouse.
The desired value appears in the field.

2.1.6.5 Viewing Variables Using the Array Browser

To view array values, select the following from the Main View window menu bar to call up the Array Browser:

Views
->Array Browser

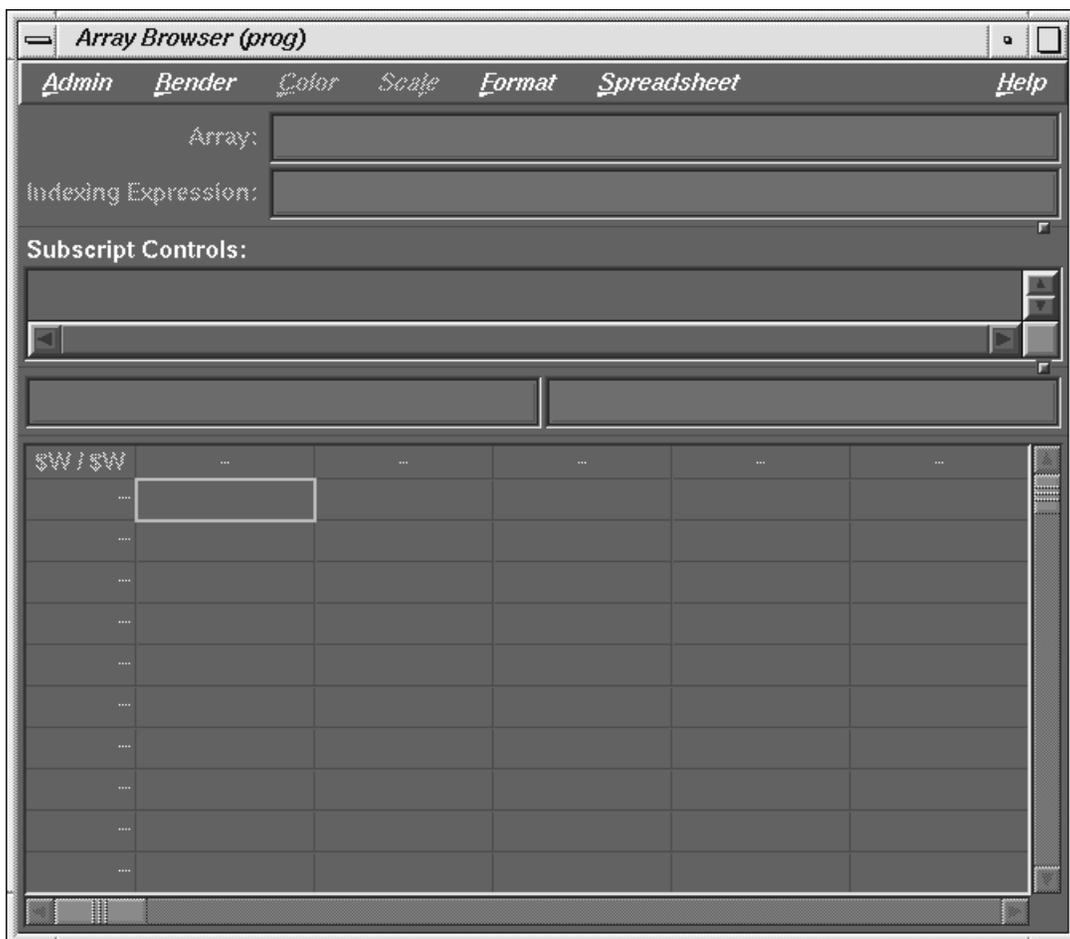


Figure 4. Array Browser Window

To view the values of an array in the Array Browser:

1. Enter the name of this array in the Array field
2. Press ENTER

Here, the current values of, at most, two dimensions of the array will display in the lower pane of the Array Browser window. The values of the array will be updated to their current values as you step through your program.

2.1.6.5.1 Subscript Controls Field

If the array is large or has more than two dimensions, the Subscript Controls panel in the middle of the Array Browser window allows you to specify portions of the array for viewing. You may also use the slide bars at the bottom and right of the window to view hidden portions of an array.

2.1.6.5.2 How to Change Values

Perform the following to change values of array elements:

1. Click on the box with the array value in the lower portion of the Array Browser window.

Note: Once the element is selected, the array index and value appear in the two fields below the Subscript Controls panel in the center of the Array Browser window.

For example, if you click on the value for array element A(2,3), then A(2 , 3) appears in the box above the display of the array values, and its value appears in the box to the right of A(2,3). Simply click in this box, and enter a new value for A(2,3). (Press ENTER to change the value of A(2,3) to the new value.)

2. Enter your change into the Value field described in the note above.
3. If you would like to view a second array at the same time, select the following in the Array Browser window that you have already opened:

```
Admin  
  ->Clone
```

This will bring up a second Array Browser window.

4. Select the following from this new window:

```
Admin  
->Active
```

You can now enter the name of the second array you would like to view.

2.1.6.6 Viewing Variables Using the Structure Browser Window

Perform the following to view values of a C structure:

1. Select the following from the Main View window menu bar to call up the Structure Browser window:

```
Views  
->Structure Browser
```

2. Enter the name of the structure in the Expression pane.

This will bring up a window listing the name of the structure, the names of its components in the left column, and their values in the right column.

3. If one of these components is a pointer, you can see what is being pointed to by double-clicking on its value to the right of the pointer name. This will bring up a new window showing what is being pointed at and an arrow will appear showing this relationship. This can aid in debugging linked lists, for example.

2.1.7 Searching

Often it is useful to search for the occurrences of a variable or some character string in your program so you can set breakpoints. Perform a search as follows:

1. Call up the search utility from the Main View window menu bar as follows to call up the Search window:

```
Source  
->Search
```

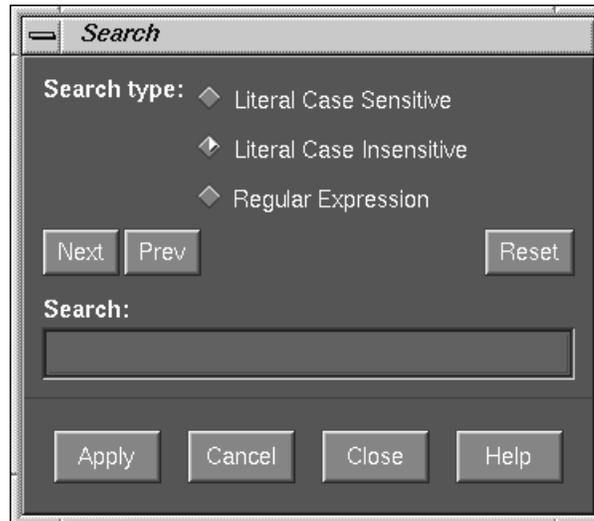


Figure 5. Search Window

2. Enter the character string for which you would like to search in this window, then click on the **Apply** button.

All occurrences of this string are highlighted.

3. Click on **Cancel** to remove highlighting.

2.1.8 Using the Call Stack View

As the Debugger executes, it may stop during a program function or subroutine, or in a system routine. When this happens, you can locate where the Debugger is in your program by examining the call stack.

There are two ways to examine the call stack:

- You can enter the following in the command/message pane:

```
cvd> where
```

This will list the functions and subroutines that were called to bring the Debugger to this place. This call list will include not only your program functions/subroutines but may also include system routines that have been used. You can now move up or down the call tree by issuing, for example:

```
cvd> up [n]
```

In this case, the source code for a function or subroutine that is 'up' *n* items in the call stack will appear in the Main View window. If you omit [n] you move up one item in the call stack.

- You can select the following from the Main View window menu bar:

```
Views  
->Call Stack View
```

This brings up the Call Stack View window.

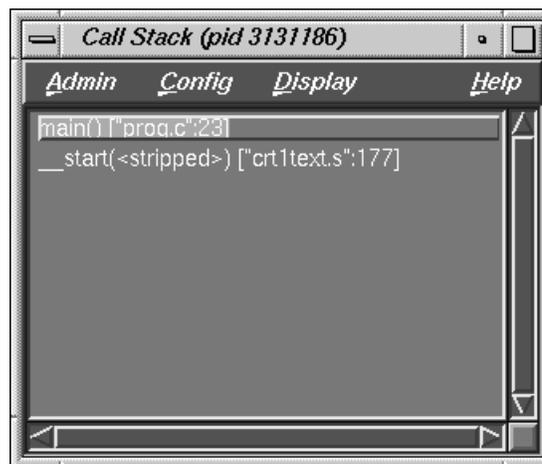


Figure 6. Call Stack View Window

If you double-click on an item here, the source code for the function or subroutine, if available, will display in the Main View window.

2.1.9 Stopping at Functions or Subroutines

In the debugging process, it is sometimes useful to stop at each occurrence of a function or subroutine. The Debugger permits you to do this in either of two ways:

- Using the `cvd` command/message pane.

1. Enter the following in the command/message pane of the Main Window:

```
cvd> stop in name
```

For *name*, specify the name of the function or subroutine in your program where you would like the Debugger to stop.

2. Click on the Run button and the Debugger will stop each time it encounters this function or subroutine.
3. To remove this stopping condition enter the following in the command/message pane:

```
cvd> status
```

For the `Stop in` command above, the trap in the list would appear as:

```
[n] Stop entry name...
```

Here, the value of *n* is a positive integer and *name* is the name of the function or subroutine where the stop has been set.

4. To delete this stop, enter:

```
cvd> delete n
```

- Using the Trap Manager.

1. Select the following from the Main View window menu bar to use the trap manager:

```
Views  
->Trap Manager
```

This will call up the Trap Manager window.

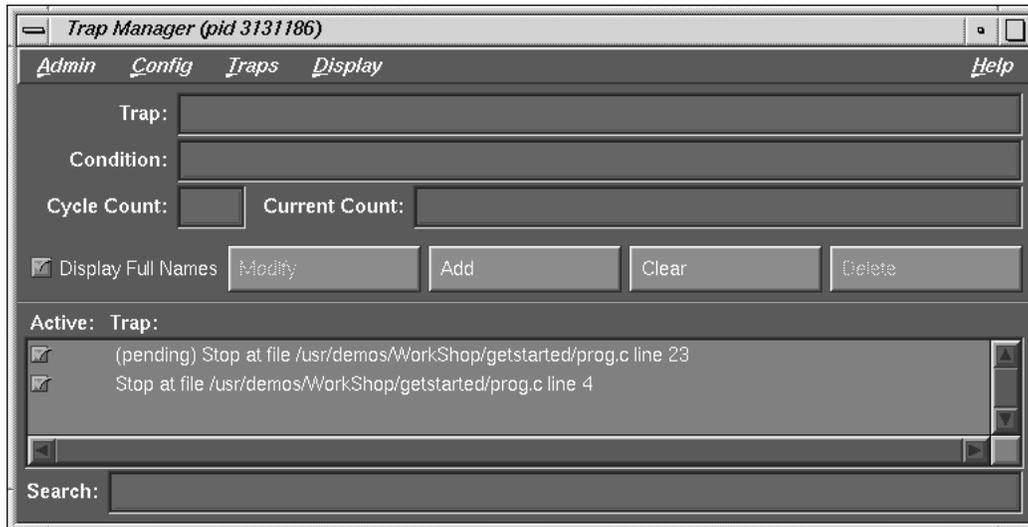


Figure 7. Trap Manager Window

2. In the pane to the right of the word `Trap` enter:

stop in name

3. Click on the `Add` button or press `Enter`. This will add your breakpoint at your desired function or subroutine.

Note: To remove the stopping condition so the Debugger will not stop at each occurrence of *name*, click on the `Delete` button in the Trap Manager window.

If you have multiple traps displayed, click on the trap that you wish to delete before you click on the `Delete` button.

2.2 Suggestions for Debugging for Serial Execution of Scientific Programs

This section offers tips and suggestions for debugging programs written for scientific applications; but many of the suggestions will apply to debugging other types of applications as well.

Note: This section deals only with debugging programs that are running *serially* and not in parallel.

Programs can sometimes appear to have no bugs with some sets of data because all paths through the program may not be executed. To debug your program, therefore, it is important to test it with a variety of different data sets so that, one would hope, all paths in your program can be tested for errors.

Now, assume that your program compiles and produces an executable, but the program execution either does not complete, or it completes but produces wrong answers. In this case, go through the following steps to find many of the commonly occurring bugs.

Note: All compiler options mentioned in the following sections are valid for FORTRAN 77, Fortran 90, C and C++ compilers unless indicated otherwise.

2.2.1 Step 1: Use `lint`

If your program is written in C, you should use the `lint` utility. This will help you identify problems with your code at the compile step. For example, if your C program is in a file named `prog.c`, invoke `lint` with:

```
> lint prog.c
```

The output from this command is directed to your screen.

Note: There is a public domain version of `lint` for FORTRAN 77 called `ftnchek`. You can get `ftnchek` from the `/pub` directory at the anonymous ftp site `ftp.dsm.fordham.edu` at Fordham University. Also, there is a `lint` for Fortran 90 for SGI/Cray PVP machines and the Cray T3E called `cflint`. You can invoke `cflint` if you compile with the `-r1` option.

2.2.2 Step 2: Check for Out-of-Bounds Array Accesses

A common programming error is the use of array indices outside their declared limits. For help finding these errors, compile your program as follows:

```
> -g -DEBUG:subscript_check=ON
```

Then run the generated executable under `cvd` and click on the RUN button. (See the `DEBUG_group(5)` man page for more information on this option.)

The following list explains compiling dependencies when working with out-of-bounds array accesses:

- If you are running a C or C++ program, your program will stop at the first occurrence of an array index going out of bounds. You can now examine the value of the index that caused the problem by using any of the methods

described in Section 2.1.6, page 20. If you compile with the `-g` option, the compiler generates symbolic debugging information so your program will execute under `cvd`. It also disables optimization. Sometimes, disabling optimization will cause the bug to disappear. If this happens, you should still carefully go through each of these steps as best as you can.

- If you are using the Fortran 90 compiler, after compiling with the preceding options and after running the generated executable under `cvd`, enter the following in the `cvd` pane:

```
cvd> stop in __f90_bounds_check
```

Note: `__f90` in this command has two lead “`_`” characters.

Now, click on the RUN button. Next select the following from the Main View window menu bar:

```
Views  
->Call Stack View
```

Next, double click on the function or subroutine immediately below `__f90_bounds_check`. This will cause the source code for this function or subroutine to display in the Main View window, and the line where `cvd` has stopped will be highlighted. You can now find the value of the index that caused the out-of-bounds problem.

- If you are using the FORTRAN 77 compiler, after compiling with the preceding options and after running the generated executable under `cvd`, enter the following in the `cvd`> pane:

```
cvd> stop in s_rnge
```

Now, click on the RUN button. Next, select the following from the Main View menu bar:

```
Views  
->Call Stack View
```

Double-click on the function or subroutine immediately below `s_rnge`. This will cause the source code for this function or subroutine to display in the Main View window and the line where the Debugger stopped will be highlighted. You can now find the value of the index that caused the out-of-bounds problem.

Note: For Fortran programs, bounds checking cannot be performed in subprograms if arrays passed to a subprogram are declared with extents of 1 or * instead of passing in their sizes and using this information in their declarations. An example of how the declarations should be written to allow for bounds checking is:

```
SUBROUTINE SUB(A,LDA,N, ...)  
  INTEGER LDA,N  
  REAL A(LDA,N)
```

2.2.3 Step 3: Check for Uninitialized Variables Being Used in Calculations

To find uninitialized REAL variables being used in floating-point calculations, compile your program with the following:

```
-g -DEBUG:trap_uninitialized=ON
```

This will force all uninitialized stack, automatic, and dynamically allocated variables to be initialized with 0xFFFA5A5A. When this value is used as a floating-point variable involving a floating-point calculation, it is treated as a floating-point NaN and it will cause a floating-point trap. When it is used as a pointer or as an address a segmentation violation may occur. For example, if x and y are real variables and the program is compiled as described previously, $x = y$ will not be detected when y is uninitialized since no floating point calculations are being done. However, the following will be detected:

```
x = y + 1.0
```

After you compile your program with the preceding options, enter the following:

```
cvd executable
```

Then click the RUN button. To find out where your program has stopped, select the following from the Main View window menu bar:

```
Views  
  ->Call Stack View
```

Here, you will see that many routines have been called. Double-click on the routine closest to the top of the displayed list that is not a system routine. (You will probably recognize the name of this source file.) This will bring up the source code for this routine and the line where the first uninitialized variable (subject to the above-mentioned conditions) was used. You can now examine

the values of the indices which caused the problem using any of the methods described in Section 2.1.6, page 20.

Note: You cannot use `cvd` to detect the use of uninitialized `INTEGER` variables.

2.2.4 Step 4: Find Divisions by Zero and Overflows

If you are using a `csh` or `tsch` shell, perform the following to find floating-point divisions by zero and overflows.

1. Enter the following:

```
setenv TRAP_FPE ON
```

Note: For other shells, see their man pages.

2. Compile your program using the following options:

```
-g -lfpe
```

3. Enter the following:

```
cvd executable
```

4. In the `cvd` command/message pane enter:

```
cvd> stop in _catch
```

5. Click on the `RUN` button.

6. Select the following from the Main View window:

```
Views  
->Call Stack View
```

7. Double-click on the routine closest to the top of the displayed list that is not a system routine. (You will probably recognize the name of this source file.)

The line where execution stopped is highlighted in the source code display area of the Main View window.

You may now use any of the methods to find variable values, described in Section 2.1.6, page 20, to discover why the divide-by-zero or overflow occurred.

For more information on handling floating-point exceptions, see the `handle_sigfpe(3)` and `fsigfpe(3F)` man pages.

Perform the following to find integer divisions by zero:

1. Compile your program using the following options:

```
-g -DEBUG:div_check=1
```

2. Enter the following:

```
cvd executable
```

3. Click the Run button.

The program will automatically stop at the first line where an integer divide-by-zero occurred. You may now use any of the methods to find variable values, described in Section 2.1.6, page 20, to discover why the divide-by-zero occurred.

2.2.5 Step 5: Perform Core File Analysis

Sometimes during program execution a core file is produced and the program does not complete execution. The file is placed in your working directory and named `core`.

Some machines are configured to not produce a core file. To find out if this is the case on the machine you are using, enter the following:

```
% limit
```

If the limit on `coredumpsize` is zero, no corefile will be produced. If the limit on `coredumpsize` is not large enough to hold the program's memory image, the core file produced will not be usable.

To change the configuration to allow core files to be produced enter the following:

```
% unlimited coredumpsize
```

Once you have a core file, you can perform the following analysis:

1. You can find the place in your program where the execution stopped and the core file was produced by entering:

```
cvd executable core
```

Here, *executable* is the executable that you were running.

The Main View window will come up and the source line where execution stopped may be highlighted in green.

2. If the source line is not highlighted in green, select the following from the Main View window menu bar:

```
Views
  ->Call Stack View
```

3. Double-click on the routine closest to the top of the displayed list that is not a system routine. (You will probably recognize the name of this source file.) This will bring up the source code for this routine, and the last line executed will be highlighted in green.

If the executable was formed by compiling with the `-g` option, then you can view values of program variables when program execution stopped.

To find the assembly instruction where execution stopped, select the following from the Main View window menu bar:

```
Views
  ->Disassembly View
```

Remember that this is the last statement executed before the `core` file was produced. It therefore does not necessarily mean that the bug in your program is in this line of code. For example, a program variable may have been initialized incorrectly; but the core was not produced until the variable was used later in the program.

2.2.6 Step 6: Troubleshoot Incorrect Answers

Assume that the preceding steps have been taken and that all detected problems have been corrected. Your program now completes execution, but obtains incorrect answers. What you do at this point will likely depend on special circumstances. The following is a list of some commonly used debugging tips that may or may not apply to your situation.

1. Try running your program on a very small size problem where you can easily obtain intermediate results. Run your program under `cvd` on this small problem and compare with the known correct results.
2. If you know that a certain answer being calculated is not correct, set breakpoints in your program so you can monitor the value of the answer at various points in your program.

3. You may want to set breakpoints on each call to a selected function or subroutine where you suspect there may be problems. (See Section 2.1.6, page 20 for suggested methods.)
4. Debug COMMON blocks and EQUIVALENCE statements in Fortran. Variables used in these statements must have the same type and dimension everywhere they appear and they must occur in the same order. Normally `ftnchek`, for FORTRAN 77 programs, and `cflint`, for Fortran 90 programs, will find these errors. However, for FORTRAN 77 programs it is best to use an include statement for each COMMON block. For Fortran 90 programs, it is best to use a module for each COMMON block. It is best not to use EQUIVALENCE statements.
5. Save local data that is otherwise not saved. In Fortran, values of local variables are not guaranteed to be saved from one execution of the subprogram to the next unless they are either initialized in their declarations or they are declared to have the `SAVE` attribute. Some compilers and machines automatically give all local variables the `SAVE` attribute, so moving a working program from one compiler or machine to a compiler or machine that does not do this may cause this bug to manifest. The Fortran standards require that you give all uninitialized local variables the `SAVE` attribute if you would like their values saved.

Selecting Source Files [3]

This chapter shows you how to select source files for the source pane of the Main View window of the Debugger (see Figure 1, page 3). It covers the following topics:

- How to load source files, Section 3.1, page 39
- Path Remapping, Section 3.2, page 43

3.1 How to Load Source Files

The following sections show you the three ways you can load source files for debugging.

Note: For demonstration purposes, before you begin this section, perform the following from your shell:

```
% mkdir demos
% mkdir demos/jello
% cd demos/jello
% cp /usr/demos/WorkShop/jello/* .
% make
% cvd &
```

3.1.1 Load Directly into the Main View Window

Perform the following steps to load your source file directly into, or run your executable from, the Debugger Main View window:

- Enter the source file directly.

Enter the name (or full pathname, if necessary) of the source file in the File: field.

For example enter:

```
File: jello.c
```

if you launched `cvd &` from within the `jello/demos` directory.

- Enter the executable directly.

Enter the name (or full pathname, if necessary) of the executable in the Command: field.

For example, enter:

Command: `jello`

if you launched `cvd &` from within the `jello/demos` directory.

3.1.2 Load from the File Browser Dialog Box

You can load source files from the File Browser dialog box, available as follows from the menu bar of the Main View window:

```
Views  
->File Browser
```

The File Browser window is shown in Figure 8.

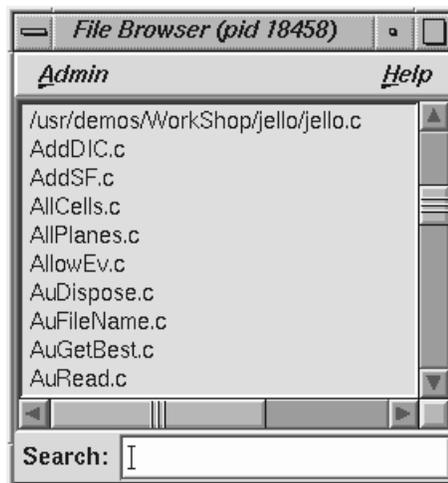


Figure 8. File Browser Window

This dialog box provides you with a list of the source files that your executable file can use, including any files in linked libraries.

locate a file To locate a file, enter your desired filename in the Search field.

load a file To load a file directly into the Main View window from the File Browser dialog box, simply double-click on the file name.

Note: You may be unable to locate some files because the source supports system routines. Source for these routines may not be available on your system.

3.1.3 Load from the Open Dialog Box

You can load source files from the Open . . . dialog box, available as follows from the menu bar of the Main View window:

```
Source  
->Open . . .
```

The standard dialog box lists all available files and the currently selected directory in the Selection field. You can change this directory as you wish.

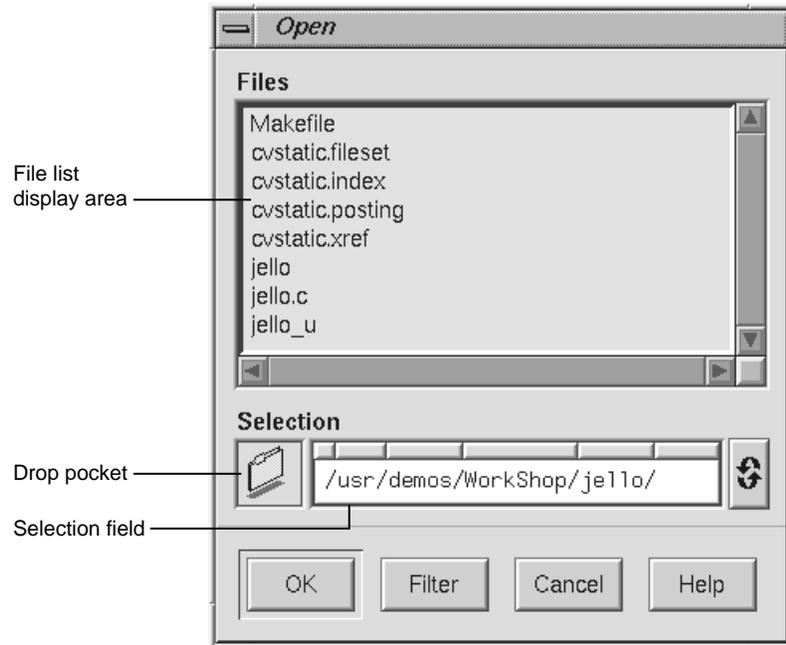


Figure 9. Open Dialog Box

There are several ways to load a file. You can:

- double-click on the file name.
- type the full pathname of the file in the Selection field and click the OK button.
- drag the file icon into the drop pocket. (Use an application like `fm` to produce file icons.)

If you specify a file name without a full path, the Debugger uses the current path remapping information to try to locate the file (see Path Remapping, Section 3.2, page 43).

3.2 Path Remapping

Path remapping allows you to modify mappings to redirect file names, located in your executable file, to their actual source locations on your file system. Since WorkShop uses full (that is, absolute) path names, path remapping generally is not necessary. However, if you have mounted executable files on a different tree from the one on which they were compiled, you will need to remap the root prefix to get access to the source files in that hierarchy.

The most basic remapping is for ".", which allows you to specify the directories to be searched for files. This basic function works just like `dbx` and can be modified by using `use/full_path_name` and `dir/full_path_name` in the command line.

Open the Path Remapping window as follows from the menu bar of the Main View window:

```
Admin  
  ->Remap Paths...
```

The following displays:

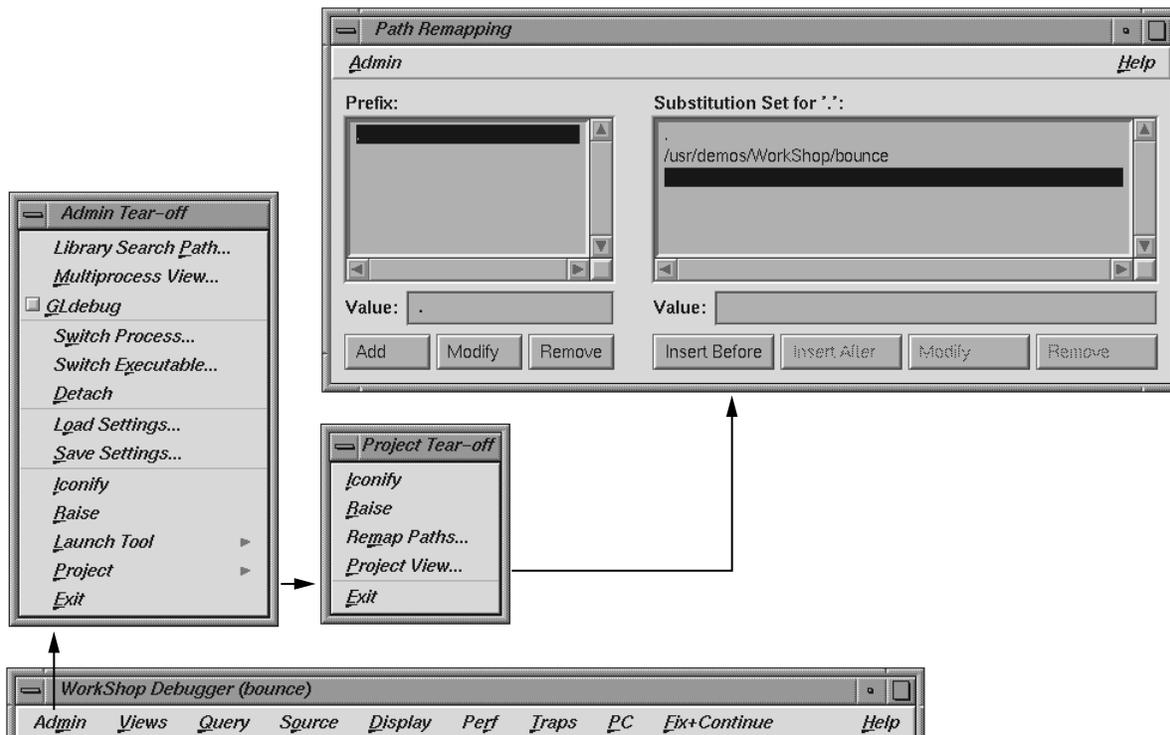


Figure 10. Path Remapping Dialog Box

For each prefix listed in the `Prefix` list, there is an ordered set of substitutions used to find a real file. By default, path remapping is initialized so that `."` is mapped to the current directory. The `Substitution Set for \".:'` list shows the substitution list for the currently highlighted item in the `Prefix` list. You can perform the following operations through the `Path Remapping` dialog box:

- To view the substitution set for a different prefix, click that prefix.
- To add a new prefix, enter the new value in the `Value:` field below the `Prefix` list and click the `Add` button. A new substitution set is created with the prefix name as the first element. Click on this element to highlight it.

Next, type the desired substitution in the `Value` field below the `Substitution Set` list and insert it by clicking on either the `Insert Before` button or the `Insert After` button.

- To modify the currently selected prefix, edit the string in the `Value` field and click the `Modify` button.
- To remove the current prefix and its substitution set, select the prefix and click the `Remove` button.

3.2.1 Case Example for Path Remapping

In some cases, if source files have been moved to new locations, path remapping is required to help the Debugger find the source files again.

The following tutorial shows you a case for remapping. It includes demo files bundled with your WorkShop Debugger:

1. Create a new directory in *your_home_directory*:

```
% mkdir jellodemos
```

2. Change to the new directory:

```
% cd jellodemos
```

3. Copy the Jello demo files from the Workshop demo directory into your new directory:

```
% cp /usr/demos/WorkShop/jello/* .
```

4. Enter the following to ensure that the `jello` executable contains the `jello` demos source path:

```
% make clobber  
% make
```

5. Create another new directory in your `jellodemos` directory:

```
% mkdir ./newdir
```

6. Move the Jello source files to a new location:

```
% mv /*.c ../newdir
```

7. Start the WorkShop Debugger:

```
% cvd ./jello &
```

The Main View window will display with no source in the source pane.
The following message appears:

```
Unable to find file <your_home_directory/newdir/jello.c
```

8. Choose the following from the menu bar in the Main View window:

```
Admin  
  ->Remap Paths...
```

The Path Remapping window displays.

9. In the Substitution Set for `'.'`: dialog box:

- a. Select `your_home_directory/jellodemos/newdir/jello.c`

The *path/filename* will appear in the Value: field.

- b. Enter the following:

```
.
```

10. Enter the following in the Value: field below the Substitution Set for `'.'`: dialog box:

```
newdir
```

11. Click on the Insert Before button.

The directory is inserted before the highlighted empty line in the Substitution Set for `'.'`: dialog box and after the first element, which was not highlighted.

Now, the source appears in the Main View source pane as
`your_home_directory/newdir/jello.c`

Tutorial: The `jello` Program [4]

This chapter presents a short tutorial with the demonstration program `jello`, provided with your MIPSpro WorkShop software package. This tutorial walks you through commonly encountered debugging situations.

The tutorial is divided into four parts:

- Start the Debugger, Section 4.1, page 47
- Run the `jello` Program, Section 4.2, page 48
- Perform a Search, Section 4.3, page 51
- Edit Your Source Code, Section 4.4, page 53
- Set Traps, Section 4.5, page 54
- Examine Your Data, Section 4.6, page 58

Before you begin this tutorial, you should be aware of the following:

- This tutorial must be run on an SGI workstation.
- WorkShop identifies files with the path names in which they were compiled. The path names in the tutorial may not match the ones on your system.

4.1 Starting the Debugger

Use the following syntax to start the Debugger:

```
cvd [-pid pid] [-host host] [executable_file [corefile]] [&]
```

Note: The `cvd` command should be invoked in the same directory as your program source code.

- The `-pid` option lets you attach the Debugger to a running process. You can use this to determine why a live process is in a loop.

- The `-host` option lets you specify a remote host on which to run your program while the Debugger runs locally. This option is seldom used, except under the following circumstances:
 - You do not want the Debugger windows to interfere with the execution of your program on the remote host.
 - You are supporting an application remotely.
 - You do not want to use the Debugger on the target machine for any other reason.

Note: The host and local machines must be running the same version of WorkShop. Also, the `.rhost` files on the machines must allow `rsh` commands to operate between them.

- The `executable_file` argument is the name of the executable file for the process or processes you want to run. This file is produced when you compile with the `-g` option, which disables optimization and produces the symbolic information necessary to debug with WorkShop. The `-g` option is most commonly used, but it is optional; if you wish, you can invoke the Debugger first and specify the name of the executable file later.
- Sometimes when a file is being executed, a core file is produced (its default name is `core`). Use the following command to determine why a program crashed and produced the core file:

```
cvd executable_file core
```

See Section 4.2 for more information.

4.2 Run the `jello` Program

In this part of the tutorial, you invoke the Debugger and start a typical process running. The `jello` program simulates an elastic polyhedron bouncing around inside of a revolving cube. The program's functionality is mainly contained in a single loop that calculates the acceleration, velocity, and position of the polyhedron's vertices.

Enter the following commands to run the `jello` program:

1. Go to the directory with the `jello` demo program:

```
> cd /usr/demos/WorkShop/jello
```

2. List the contents of this directory:

```
>ls
```

3. Enter the following to make the program if jello is not listed (from Step 2):

```
> make jello
```

4. Invoke the Debugger with the jello program:

```
> cvd jello &
```

The Main View window appears, as shown in Figure 11, page 50, and scrolls automatically to the main function.

In addition to the Main View window, the Execution View icon also appears. When you run the jello program, the command you used to invoke jello is displayed in this window.

The Execution View window is the interface between the program and the user for programs that use standard input/output or generate stderr messages.

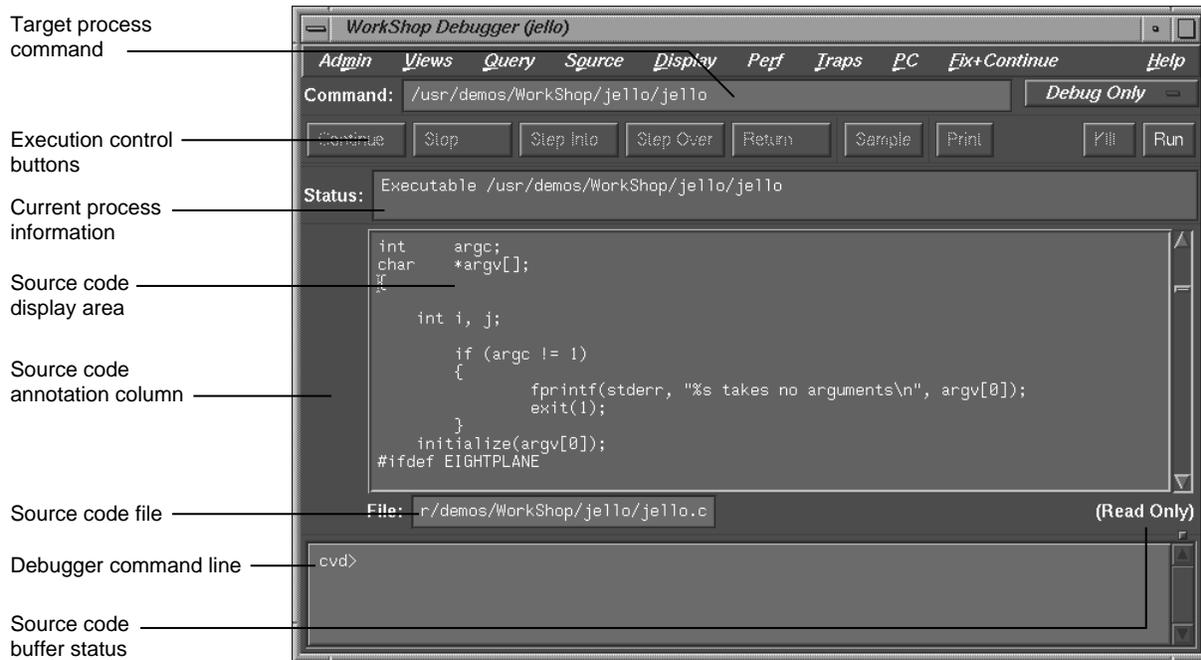


Figure 11. The Main View Window with jello Source Code

Note: The Main View window brings up the source file in read-only mode. You can change this to read/write mode if you select the following from the Main View window menu bar (provided you have the proper file access permissions):

```
Source
->Make Editable
```

5. Click the Run button in the upper-right corner of the Main View window to run the jello program.

The jello window opens on your display (see Figure 12, page 51). Enlarge this window to watch the program execute. The polyhedron is initially suspended in the center of the cube.

If you wish, you can perform the following with the jello program:

- a. Click the left mouse button anywhere inside the jello window.

The polyhedron drops to the floor of the cube.

- b. Hold down the right mouse button to display the pop-up menu and select `spin`.

The cube rotates and the polyhedron bounces inside the cube.

- c. Hold down the right mouse button to display the pop-up menu and select the `Display` option.

This opens a submenu that allows you to change the appearance of the `jello` polyhedron.

- d. Feel free to select (right-click) from this menu to see how the `jello` display changes.

Note: You may encounter flashing colors inside windows while running `jello`. This is normal.

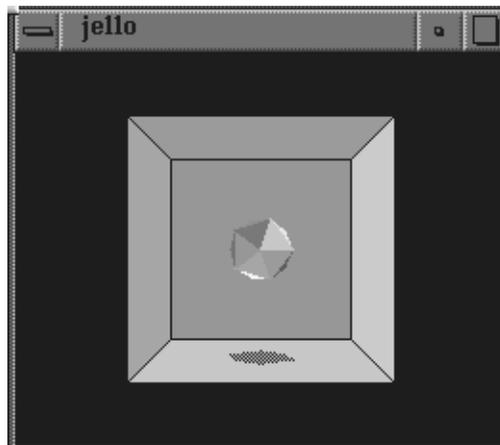


Figure 12. The `jello` Window

6. Right-click and select `Exit` from the pop-up menu to exit this demonstration executable.

4.3 Perform a Search

This part of the tutorial covers the search facility in the Debugger. You will search through the `jello` source file for a function called `spin`. The `spin` function recalculates the position of the cube.

1. Select the following from the Main View menu bar:

Source
->Search

The Search dialog box appears.

2. Type **spin** in the Search field in the dialog box, as shown in Figure 13.

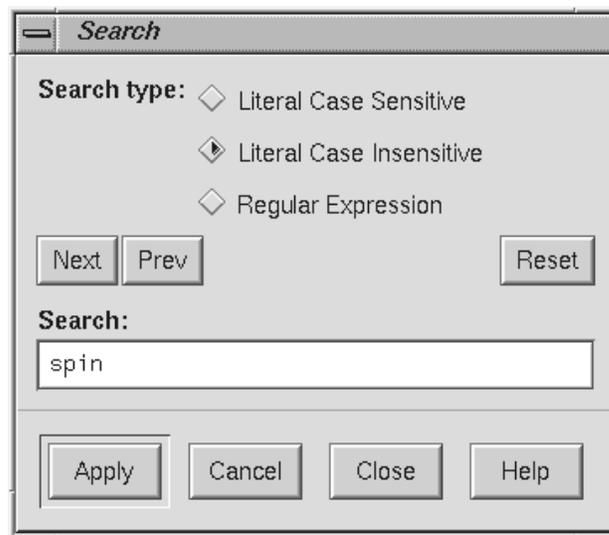


Figure 13. The Search Dialog

3. Click the **Apply** button.

The search takes place on the displayed source files. Each instance of `spin` is highlighted in the source code and flagged with target indicators in the scroll bar to the right of the display area. (See search target indicators in Figure 14, page 53.) The **Next** and **Prev** buttons in the **Search** dialog box let you move from one occurrence to the next in the order indicated.

For more information on **Search**, see Section A.1.4, page 194.

4. Click the **Close** button and the dialog box disappears.

- Click the middle mouse button on the last search target indicator at the right side of the source code pane (see the figure below). This scrolls the source code down to the last occurrence of `spin`, the location of the `spin` function.

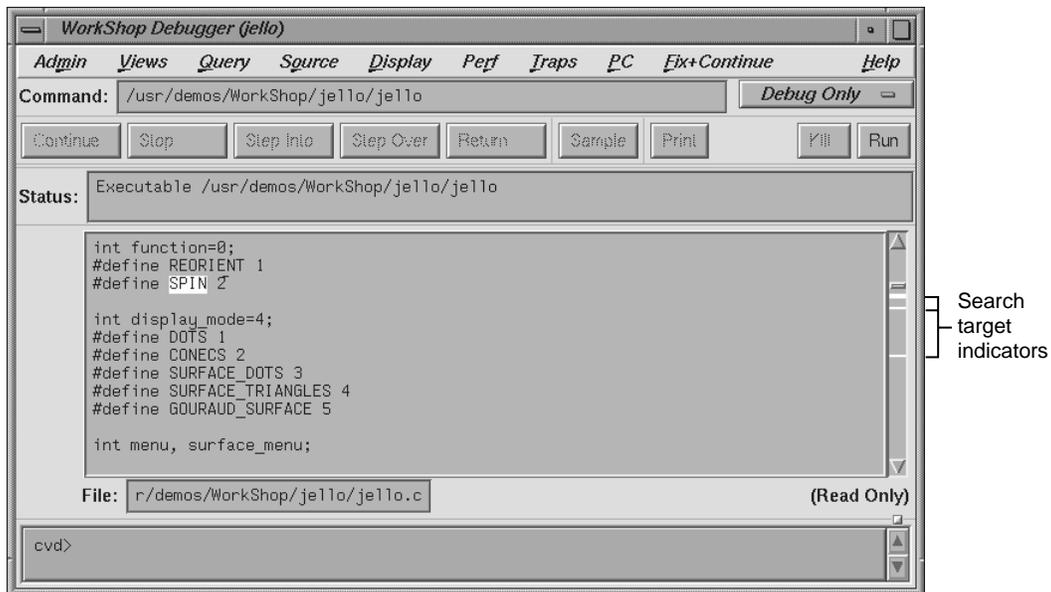


Figure 14. Search Target Indicators

- Proceed to Section 4.4 to edit your code.

4.4 Edit Your Source Code

To edit and recompile your source code, follow these steps:

- Select the following from the Main View menu bar:

```
Source
  ->Fork Editor
```

A text editor will appear. (In this case, if you are proceeding from Section 4.3, page 51, notice that the `spin` function is displayed.)

Note: If you use source control, you can check out the source code through the configuration management shell by selecting the following from the menu bar:

```
Source
  ->Versioning
    ->CheckOut
```

2. Edit the source code as follows:

Change all occurrences of 3600 to 3000 in the following code:

```
if ((a+=1)>3600) a -= 3600;
if ((b+=3)>3600) b -= 3600;
if ((c+=7)>3600) c -= 3600;
```

Save your changes.

3. Select the following from the menu bar to recompile your code:

```
Source
  ->Recompile
```

The `Build View` window displays and starts the compile. Your `makefile` will determine which files need to be recompiled and linked to form a new executable.

Any compile errors are listed in the window, and you can access the related source code by clicking the errors.

Note: This does not apply to warnings generated by the compiler.

For more information on compiling, see *Using the Build Manager*, Appendix B, page 329.

When the code is successfully rebuilt, the new executable file reattaches automatically to the Debugger and the Static Analyzer. Previously set traps are intact unless you have traps triggered at line numbers and have changed the line count.

4.5 Setting Traps

Stop traps (also called breakpoints) stop program execution at a specified line in the code. This allows you to track the progress of your program and to check the values of variables at that point. Typically, you set breakpoints in your

program prior to running it under the Debugger. For more information on traps, refer to Chapter 5, page 69.

In this part of the tutorial, you set a breakpoint at the `spin` function.

1. Click the Run button to run the `jello` executable.

The demo window will display.

2. Click the left mouse button in the Main View source code annotation column next to the line containing `if ((a+=1)>3600) a -= 3600;`, or `if ((a+=1)>3000) a -= 3000;`, if you are proceeding from the previous section).

A stop trap indicator appears in the annotation column as shown in Figure 15, page 56.

3. Right-click on the `jello` window and select `spin` from the pop-up menu.

The program runs up to your stop trap and halts at the beginning of the next call to the `spin` function. When the process stops, an icon appears and the line is highlighted. Note that the `status` field indicates the line at which the stop occurs.

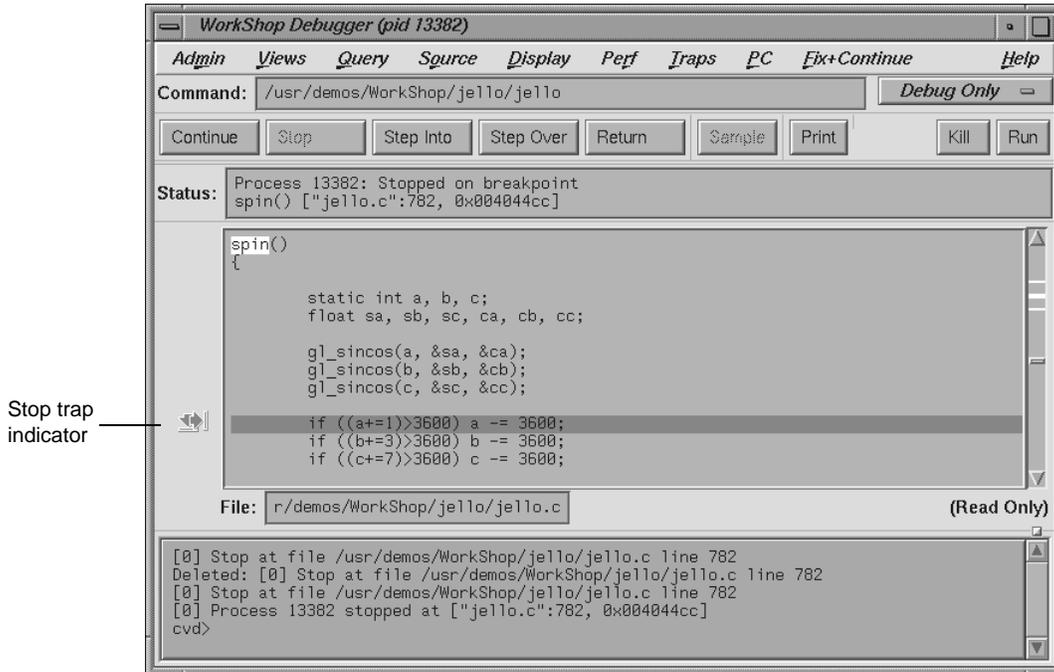


Figure 15. Stop Trap Indicator

4. Click the Continue button at the upper-left corner of the Main View window several times. Observe that the jello window goes through a spin increment with each click.
5. Select the following from the Main View window menu bar:

```
Views
  ->Trap Manager
```

The Trap Manager window appears as shown in Figure 16, page 57.

The Trap Manager window lets you list, add, edit, disable, or remove traps in a process. In Step 2, you set a breakpoint in the spin function by clicking in the source code annotation column. The trap now displays in the trap display area of the Trap Manager window.

The Trap Manager window also permits you to do the following:

- Define other traps.

- Set conditional traps in the Condition field near the top of the window.
- Specify the number of times a trap should be encountered before it activates by using the Cycle Count field.
- Manipulate traps by using trap controls (Modify, Add, Clear, Delete).
- View all traps (active and inactive) in the trap display area.

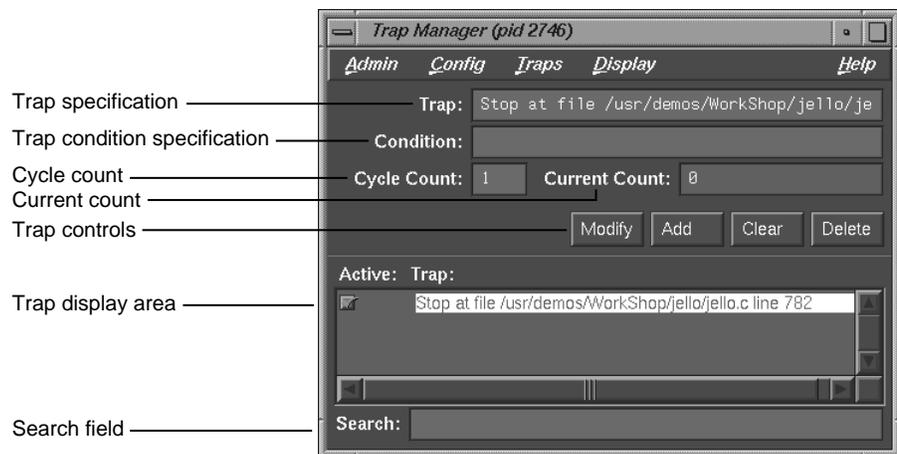


Figure 16. Trap Manager Window

6. Click the button to the left of the stop trap in the trap display area.
The trap is temporarily disabled. (If you click again in this box, the trap will be re-enabled.)
7. Click the Clear button, move the cursor to the Trap field, then type:
watch display_mode
Click the Add button.
This sets a watchpoint for the display_mode variable. A *watchpoint* is a trap that causes an interrupt when a specified variable or address is read, written, or executed.
8. Click the Continue button in the Main View window to restart the process.

The process now runs somewhat slower but still at a reasonable speed for debugging.

9. Hold down the right mouse button in the `jello` window to display the pop-up menu. From this menu, select `display` and then select the `conecs` option with the right mouse button.

This triggers the watchpoint and stops the process.

10. Go to the `Trap Manager` window and click the button next to the `display_mode` watchpoint to deactivate it. Then, click the button next to the `spin stop trap` to reactivate it.
11. Enter `100` in the `Cycle Count` field and click the `Modify` button. Notice how the trap description changes in the `Trap Manager` window.
12. Click the `Continue` button in the `Main View` window.

This takes the process through the stop trap for the specified number of times (100), provided no other interruptions occur.

The `Current Count` field keeps track of the actual number of iterations since the last stop, which is useful if an interrupt occurs. Note that it updates at interrupts only.

13. Select `Close` from the `Admin` menu to close the `Trap Manager` window.

4.6 Examining Data

This part of the tutorial describes how to examine data after the process stops.

1. Select the following from the `Main View` window menu bar:

```
Views  
->Call Stack
```

The `Call Stack View` window appears as shown in Figure 17, page 59. The `Call Stack View` window shows each frame in the call stack at the time of the breakpoint, with the calling parameters and their values. Through the `Call Stack View->Display` menu, you can also display the calling parameters' types and locations, as well as the program counter (PC). The program counter is the address at which the program has stopped. For more information about the program counter, see Section 7.1, page 93.

In this example, the `spin` and `main` stack frames are displayed in the `Call Stack View` window, and the `spin` stack frame is highlighted, indicating that it is the current stack frame.

2. Select the following from the `Call Stack` window menu bar:

```
Admin  
  ->Active
```

Notice that the `Active` toggle button is turned on. Active views are those that have been specified to change their contents at stops or at call stack context changes. If the toggle is on, the call stack is updated automatically whenever the process stops.

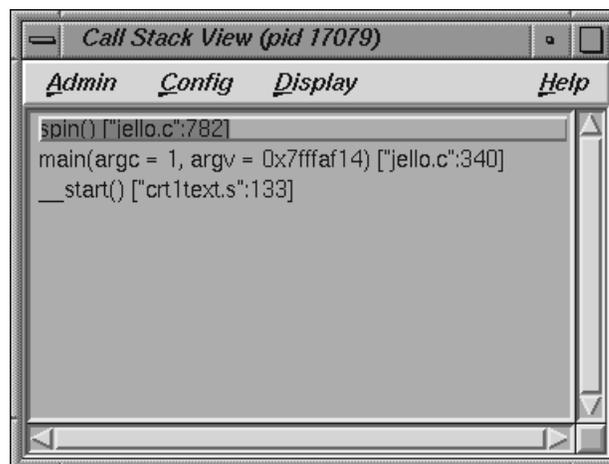


Figure 17. `Call Stack View` at `spin` Stop Trap

3. Double-click the main stack frame.

This shifts the stack frame to the main function, scrolls the source code in the `Main View` window (or `Source View`) to the place in main where `spin` was called, and highlights the call. Any active views are updated according to the new stack frame.

4. Double-click the `spin` stack frame.

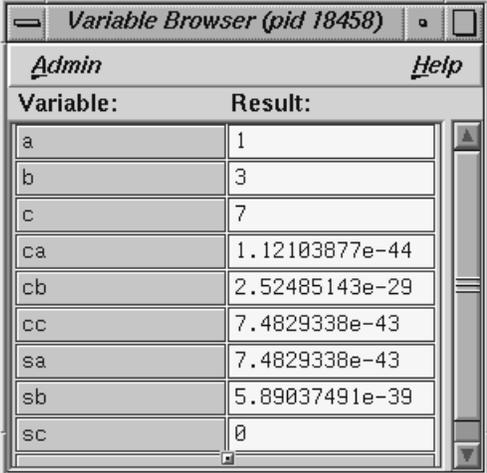
This returns the stack frame to the `spin` function.

Select `Variable Browser` from the `Views` menu in the `Main View` window.

The `Variable Browser` window appears. This window shows you the value of local variables at the breakpoint. The variables appear in the left column (read-only), and the corresponding values appear in the right column (editable).

Your `Variable Browser` window should resemble the one in Figure 18, although you may need to enlarge the window to see all the variables (the values will be different).

The `jello` program uses variables `a`, `b`, and `c` as angles (in tenths); `ca`, `cb`, `cc` as their corresponding cosines; and `sa`, `sb`, `sc` as their sines. Whenever you stop at `spin`, these values change.



The screenshot shows a window titled "Variable Browser (pid 18458)". It has a menu bar with "Admin" and "Help". Below the menu bar is a table with two columns: "Variable:" and "Result:". The table contains the following data:

Variable:	Result:
a	1
b	3
c	7
ca	1.12103877e-44
cb	2.52485143e-29
cc	7.4829338e-43
sa	7.4829338e-43
sb	5.89037491e-39
sc	0

Figure 18. `Variable Browser` at `spin`

5. Double-click some different frames in `Call Stack View` and observe the changes to `Variable Browser` and the `Main View` window.

These views update appropriately whenever you change frames in `Call Stack View`. Notice also the change indicators in the upper-right corners of the `Result` fields in Figure 19, page 61. These appear if the value has changed. If you click the *folded* corner, the previous value displays (and the indicator appears *unfolded*). You can then toggle back to the current value.

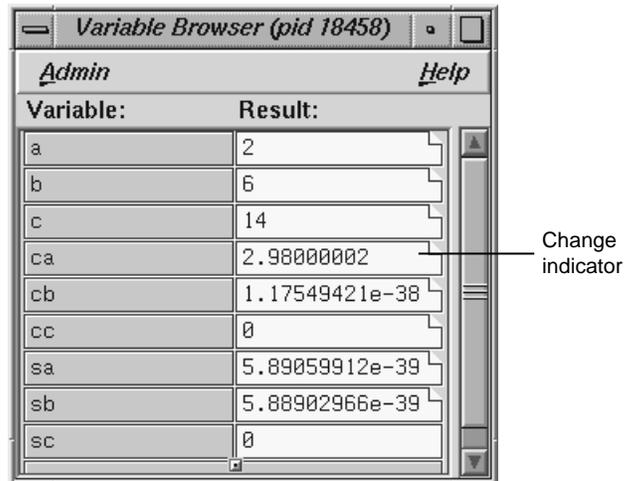


Figure 19. Variable Browser after Changes

6. Select Close from the Admin menu in Variable Browser and Close from the Admin menu in Call Stack View to close them.
7. Select Expression View from the Views menu in the Main View window.

The Expression View window appears. It lets you evaluate an expression involving data from the process. The expression can be typed in, or more simply, cut and pasted from your source code. You can view the value of variables (or expressions involving variables) any time the process stops. Enter the expression in the left column, and the corresponding value appears in the right column. For more information, see Section 7.3, page 99.

8. Hold down the right mouse button in the Expression column to bring up the Language menu. Then hold down the right mouse button in the Result column to display the Format menu.

The Language menu (shown on the left side of Figure 20, page 62) lets you apply the language semantics to the expression.

The Format menu (shown on the right side of Figure 20, page 62) lets you view the value, type, address, or size of the result. You can further specify the display format for the value and address.

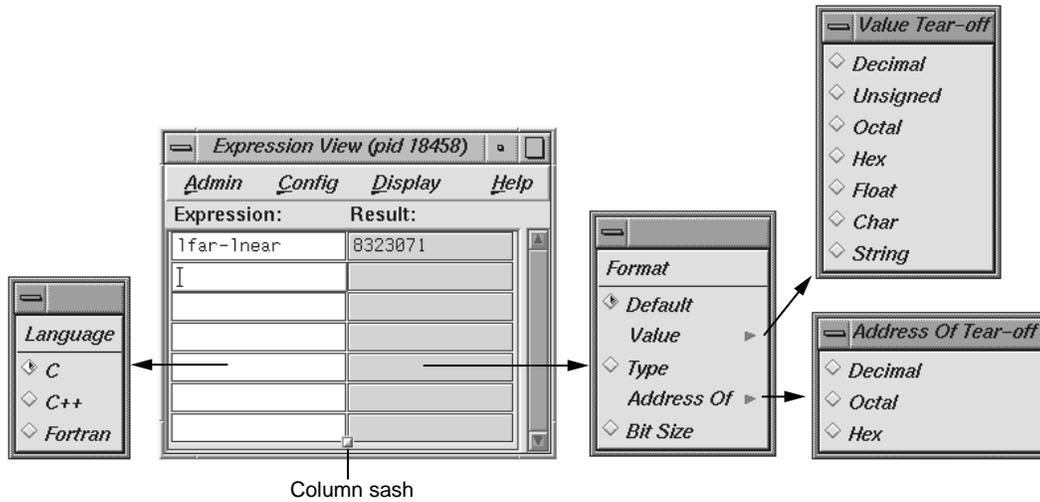


Figure 20. Expression View with Language and Format Menus Displayed

9. Click on the first Expression field in the Expression View window. Then enter **(a+1)>3600** in the field and press Enter.

This is a test performed in jello to ensure that the value of *a* is less than 3600. This uses the variable *a* that was displayed previously in Variable Browser. After you press Enter, the result is displayed in the right column; 0 signifies FALSE.

10. Select the following from the Expression View window menu bar to close that window:

```
Admin
    ->Close
```

11. Select the following from the Main View window to open the Structure Browser:

```
Views
    ->Structure Browser
```

12. Enter `jello_conec` in the Expression field and press Enter.

The Structure Browser window displays the structure for the given expression; field names are displayed in the left column, and values in the right column. If only pointers are available, the Structure Browser will de-reference the pointers automatically until actual values are encountered. You can then perform any further de-referencing by double-clicking pointer addresses in the right column of the data structure objects. A window similar to the one shown in Figure 21 now appears.

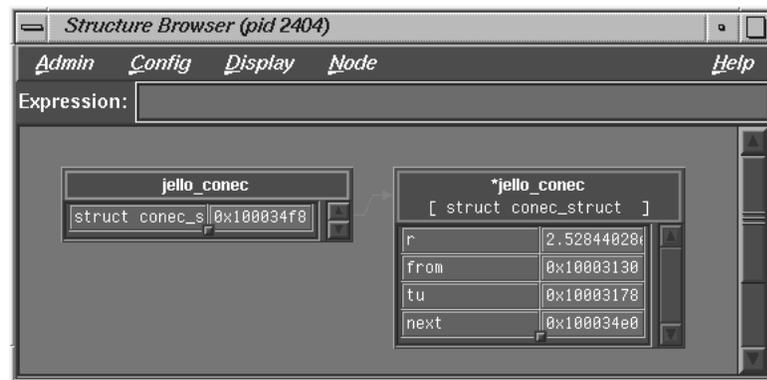


Figure 21. Structure Browser Window with `jello_conec` Structure

13. Click once to focus, then double-click the address of the next field (in the right column of the `jello_conec` structure).

Double-clicking the address corresponding to a pointer field de-references it. Double-clicking the field name displays the complete name of the field in the Expression field at the top of the Structure Browser window. (See Figure 22, page 64.)

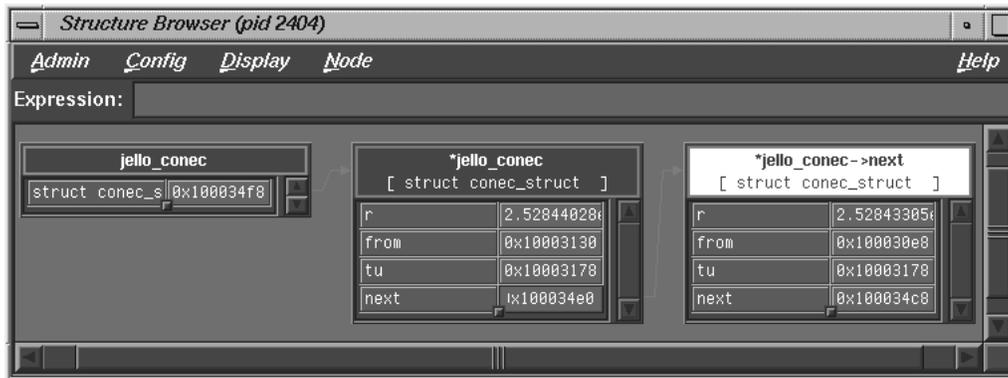


Figure 22. Structure Browser Window with Next Pointer De-referenced

14. Select Close from the Admin menu of Structure Browser window to close it.
15. Select the following from the Main View window menu bar:

Views
 ->Array Browser

The Array Browser lets you see or change values in an array variable. It is particularly valuable for finding bad data in an array or for testing the effects of values you enter.

16. Type **shadow** in the Array field and press Enter.

You can now see the values of the shadow matrix, which displays the polyhedron's shadow on the cube. The Array Browser template should resemble Figure 23, page 65, but with different data values. If any areas are hidden, hold down the left mouse button and drag the sash buttons at the lower right of the array specification and subscript control areas to expose the area.

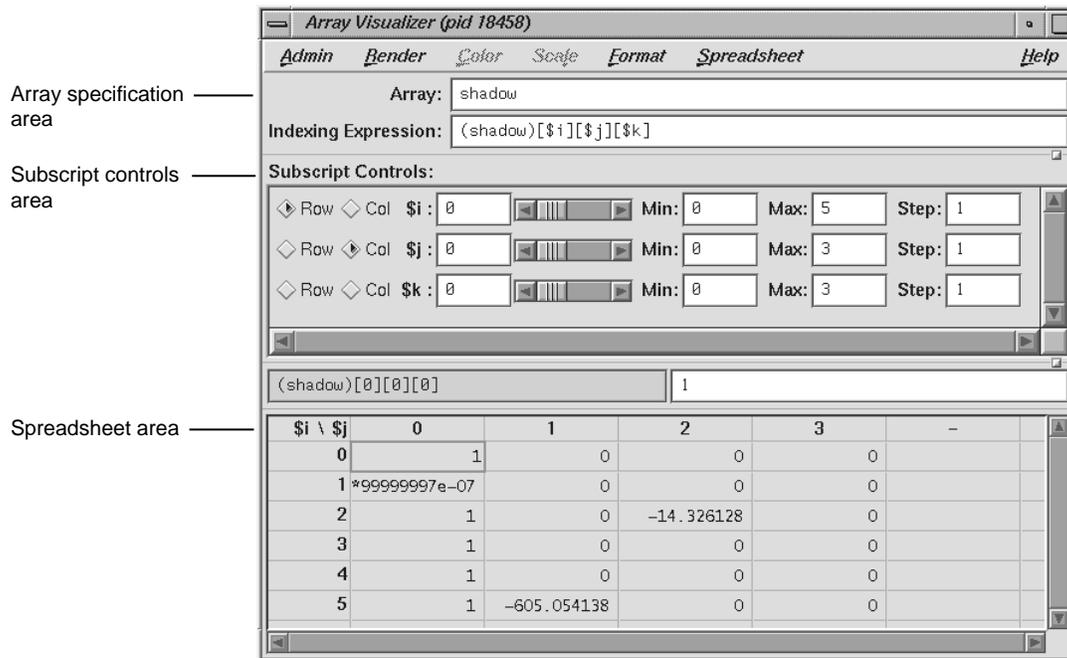


Figure 23. Array Browser Window for shadow Matrix

17. Select the Col button next to the $\$k$ index in the Subscript Controls pane (you may need to scroll down to it).

The Array Browser can handle matrices containing up to six dimensions but displays only two dimensions at a time. Selecting the Col button for $\$k$ has the effect of switching from a display of $\$i$ by $\$j$ to a display of $\$i$ by $\$k$.

Figure 24, page 66, shows a close-up view of the subscript control area.

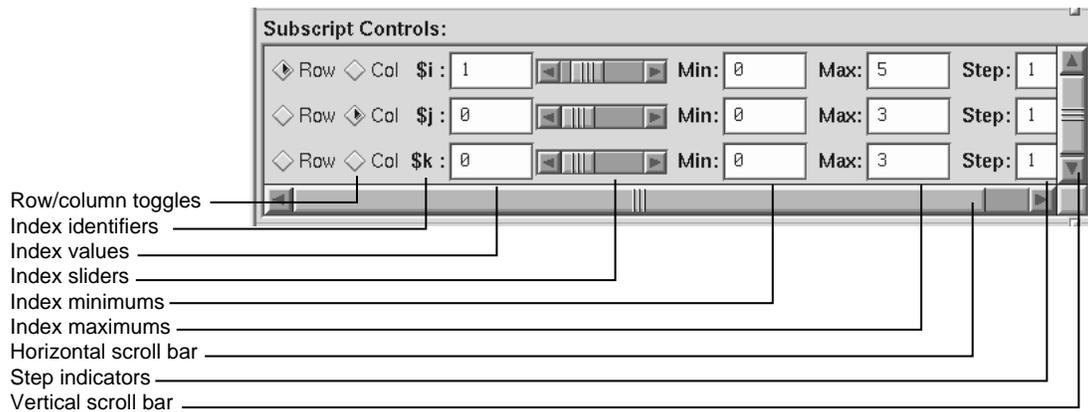


Figure 24. Subscript Controls Panel in Array Browser Window

The row and column toggles indicate whether a particular dimension of the array appears as a row, column, or not at all in the spreadsheet area. Although an array may be of 1 to 6 dimensions, you can view only one or two dimensions at a time. The index values shown as Min and Max initially exhibit the lower and upper bounds of the dimension indicated. The values may be changed to allow the user to display a subset of the available index range in that dimension. The index sliders let you move the focus cell along the particular dimension. The focus cell may also be changed by selecting a cell with the left mouse button. The index slider for a dimension whose row and column toggles are both off may be used to select a different two-dimensional plane of a multidimensional array. Use the horizontal and vertical scroll bars to expose hidden portions of the Array Browser window.

18. Select Surface from the Render menu.

The Render menu displays the data from the selected array variable graphically, in this case as a three-dimensional surface. The selected cell is highlighted by a rectangular prism. The selected subscripts correspond to the x- and y-axes in the rendering with the corresponding value plotted on the z-axis. The data can be rendered as a surface, bar chart, multiple lines, or points.

4.7 Exiting the Debugger

There are several ways to exit the Debugger:

- Select `Exit` from the `Admin` menu.
- Type `quit` at the Debugger command line as follows:
`cvd> quit`
- Double-click on the icon in the upper-left corner of the Main View window.
- Press `Ctrl-c` in the same window where you entered the `cvd` command.

Setting Traps [5]

Setting traps (often referred to as breakpoints) is one of the most valuable functions of a debugger or performance analyzer. A trap enables you to select a location or condition within your program at which you can stop the execution of the process, or collect performance data and continue execution. You can set or clear traps from the Main View window or the Trap Manager window. You can also specify traps in the Debugger command line at the bottom of the Main View window. For signal traps, you can also use the Signal Panel window; and for system call traps, use the Syscall Panel window.

When you are debugging a program, you typically set a trap in your program to determine if there is a problem at that point. The Debugger lets you inspect the call stack, examine variable values, or perform other procedures to get information about the state of your program.

Traps are also useful for analyzing program performance. They let you collect performance data at the selected point in your program. Program execution continues after the data is collected.

This chapter covers the following topics:

- Trap Terminology, Section 5.1, page 70
- Setting traps in the Main View and Source View windows, Section 5.2, page 72
- Setting traps in the Trap Manager, Section 5.3, page 76
- Setting traps with Signal Panel and Syscall Panel, Section 5.4, page 84

For a tutorial on the use of traps, see Section 4.5, page 54.

5.1 Traps Terminology

A *trap* is an intentional process interruption that can either stop a process or capture data about a process. It has two parts:

<i>trigger</i>	Specifies when the trap fires
<i>action</i>	Specifies what happens when the trap fires

5.1.1 Trap Triggers

You can set traps at a specified location in your program or when a specified event occurs. You can set a trigger at any of the following points:

- At a given line in a file (often referred to as a *breakpoint*)
- At a given instruction address
- At the entry or exit for a given function
- After set time intervals (referred to as a *pollpoint*)
- When a given variable or address is read, written, or executed (referred to as a *watchpoint*)
- When a given signal is received
- When a given system call is entered or exited

In addition, you can use an expression to specify a condition that must be met before a trap fires. You can also specify a *cycle count*, which specifies the number of passes through a trap before firing it.

Note: When you set a breakpoint in C++ code that uses templates, that breakpoint will be set in each instantiation of the template.

5.1.2 Trap Types

Traps can effect any of a number of actions in your debugging process. The following is a list of the variety of traps and their functions:

Note: See Section 5.3.1.1, page 78 for trap syntaxes.

- A *stop trap* will cause one or all processes to stop. In single process debugging, a stop trap stops the current process. In multiprocess debugging, you can specify a stop trap to stop all processes or only the current process.
- A *pending trap* is a trap with a destination address that does not resolve at the initial startup of the debugger: it allows you to place breakpoints on names that do not yet exist. By enabling pending traps, you make unknown function names acceptable breakpoint locations or entries. Because of this, the trap manager will not error a pending trap off if it can not be found in the executable. Instead, when a dlopen of a DSO (dynamic shared object) occurs, and the function name is found that matches the trap, the trap is resolved and becomes active.

There are two ways to enable pending traps.

- You can add the following line to your `.Xdefaults` file:

```
*AllowPendingTraps: true
```

- You can enable pending traps by entering the following in the `cvd>` pane of the Main View window:

```
set $pendingtraps=true
```

The default for this variable is `false`.

Note: Only “stop entry” / “stop in” and “stop exit” traps can be pending traps, not “stop at” traps. The debugger will only accept a “stop at” trap on a DSO (dynamic shared object) and line number loaded at startup. (See Section 5.3.1.1, page 78 for trap syntaxes.)

To get around this, set a pending “stop in” trap and run to the trap location. Now you can set “stop at” traps because the DSO is loaded. The debugger will then remember these traps if you run, kill, and re-run *within a single session*.

Pending traps display with the prefix `(pending)` preceding the normal trap display line.

Note: Breakpoints on misspelled names are not flagged as errors when pending traps are enabled. This is because the Debugger has no way of knowing if a pending trap will ever be tripped.

- A *sample trap* collects performance data. Sample traps are used only in performance analysis, not in debugging. They collect data without stopping the process. You can specify sample traps to collect such information as call stack data, function counts, basic block counts, PC profile counts, mallocs/frees, system calls, and page faults. Sample traps can use any of the triggers that stop traps use. Sample traps are often set up as pollpoints so that they collect data at set time intervals.
- An *exception trap* fires when a C++ exception is raised.

You can add a conditional expression to an exception trap through the Trap Manager window. However, the context in which the expression is evaluated is not that of the throw; the context is the exception handling of the C++ runtime library. Therefore, only global variables have unambiguous interpretation in the `if` clause.

You should not include complex expressions involving operators such as `*` and `&` in your type specification for an exception trap. If you create an exception trap with a specific base type, however, you will also stop your program on throws of pointer, reference, `const`, and `volatile` types. For example, if:

```
cvd> stop exception char
```

then your program will stop at type `char`, `char*`, `char&`, `const char&`, and so forth.

5.2 Setting Traps

You can set traps directly in the Main View window by using the Traps menu or by clicking the mouse in the source annotation column. You can also specify traps at the Debugger (`cvd>`) command line.

The following sections describe the ways that traps can be set:

- Section 5.2.1, page 73, Setting Traps Using the Mouse
- Section 5.2.2, page 73, Setting Traps Using the `cvd>` Command Line
- Section 5.2.3, page 73, Setting Traps Using the Traps Menu in the Main View Window

5.2.1 Setting Traps with the Mouse

The following lists ways to set traps by using your mouse:

- The quickest way to set a trap is to click in the source annotation column in the Main View or Source View windows. A subsequent click removes the trap.
- If data collection mode has been specified in the Performance Data window, clicking produces a sample trap; otherwise, a stop trap is entered.

To determine if data collection is on, look at the upper-right corner of the Main View window to see which debugging option is selected (Debug Only, Performance, or Purify).

When the trap is set, a trap icon appears.

5.2.2 Setting Traps Using the `cvd>` Command Line

The `cvd>` command line is discussed in . `stop` commands similar to those found in `dbx`, or as documented in Section 5.3.1, page 77, may be entered into the command line portion of the Main View window as an alternative way to set traps.

5.2.3 Setting Traps Using the Traps Menu in the Main View Window

To set a trap using the Traps menu, you first need to know which type of trap you wish to set, then select the location in your program at which to set the trap.

To set a stop trap or sample trap at a line displayed in the Main View window (or the Source View window), click in the source code display area next to the desired line in the source code, or click-drag to highlight the line. Then, select either of the following from the menu bar:

```
Traps
  ->Set Trap
  ->Stop
```

or

```
Traps
  ->Set Trap
  ->Sample
```

The Traps menu in the Main View window is shown in Figure 25.

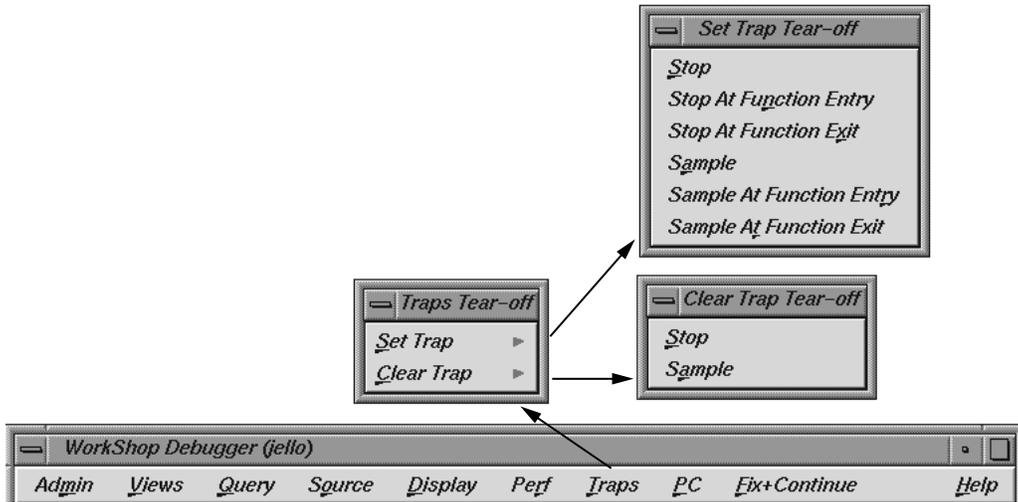


Figure 25. Traps Menu in the Main View Window

For a trap at the beginning or end of a function, highlight the function name in the source code display area and select one of the following from the Traps->Set Traps submenu as is appropriate to your needs:

- Stop At Function Entry
- Stop At Function Exit
- Sample At Function Entry
- Sample At Function Exit

Traps are indicated by icons in the source annotation column (and also appear in the Trap Manager window if you have it open). Figure 26, page 75, shows some typical trap icons. Sampling is indicated by a dot in the center of the icon. Traps appear in normal color or grayed out, depending on whether they are active or inactive. A transcript of the trap activity appears in the Debugger command line area. The active/inactive nature of traps is discussed in Section 5.3.6, page 84.

The Clear Trap selection in the Traps menu deletes the trap on the line containing the cursor. You must designate a Stop or Sample trap type, since both types can exist at the same location appearing superimposed on each other.

When the Group Trap Default toggle is checked (ON), the pgrp option is added into the resulting trap when a trap is set. This option causes the trap to apply to all processes/threads in the group of which the current process is a member.

When the Stop All Defaults toggle is checked (ON), the all option is added into the resulting trap when a trap is set. This option causes the trap to apply to all processes/threads in the current debugging session.

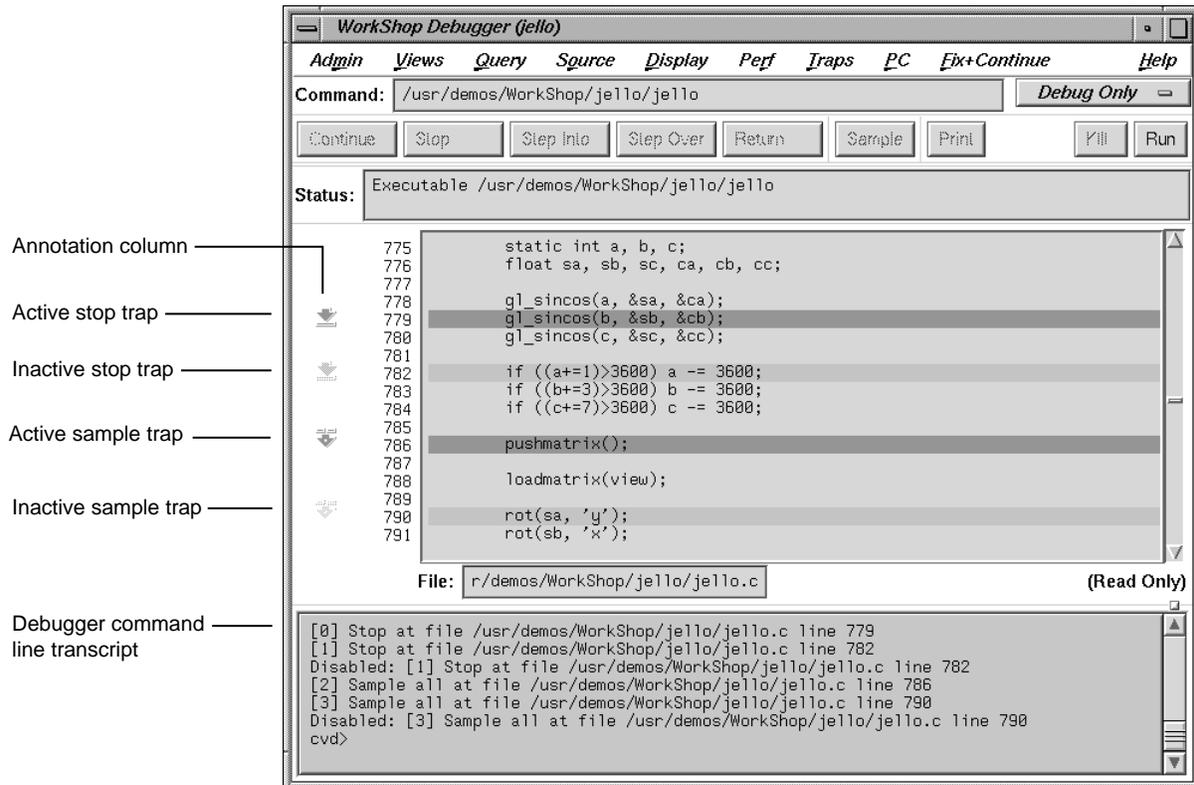


Figure 26. Typical Trap Icons

5.3 Setting Traps in the Trap Manager Window

The Trap Manager window is brought up by selecting the following from the Main View window menu bar:

```
Views
  ->Trap Manager
```

This tool helps you manage all traps in a process. Its two major functions are to:

- List all traps in the process (except signal traps)
- Add, delete, modify, or disable the traps listed

The Trap Manager window appears in Figure 27 with the Config, Traps, and Display menus shown.

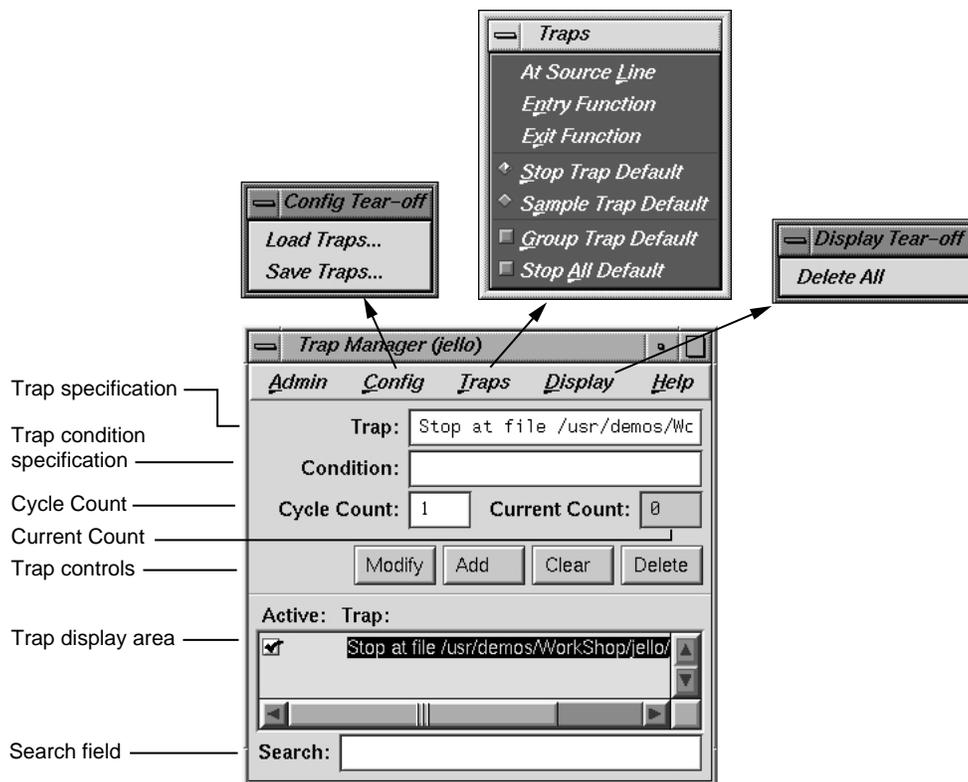


Figure 27. Trap Manager Config, Traps, and Display Menus

5.3.1 Setting Single-Process and Multiprocess Traps

New or modified traps are entered in the Trap field. Traps have the following general form:

```
[stop | sample] [all] [pgrp] location | condition
```

The entry [stop | sample] refers to the trap action. You can set a default for the action by using the Stop Trap Default or Sample Trap Default selections of the Traps menu and omitting it on the command line.

The entries [all] and [pgrp] are used in multiprocess analysis. The [all] entry causes all processes in the process group to stop or sample when the trap fires. The [pgrp] entry sets the trap in all processes within the process group that contains the code where the trap is set. You can set a default for the action by setting the Stop All Default or Group Trap Default toggles in the Traps menu.

After you enter the trap (by using the Add or Modify button or by pressing Enter), the full syntax of the specification appears in the field. The Clear button clears the Trap and Condition fields and the cycle fields.

Some typical trap examples are provided in Figure 28, page 78. The entries made in the Trap field are shown in the left portion of the figure, the trap display in the Trap Manager window resulting from these entries is shown on the right, and the trap display shown at the command line in the Main View window is shown at the bottom.

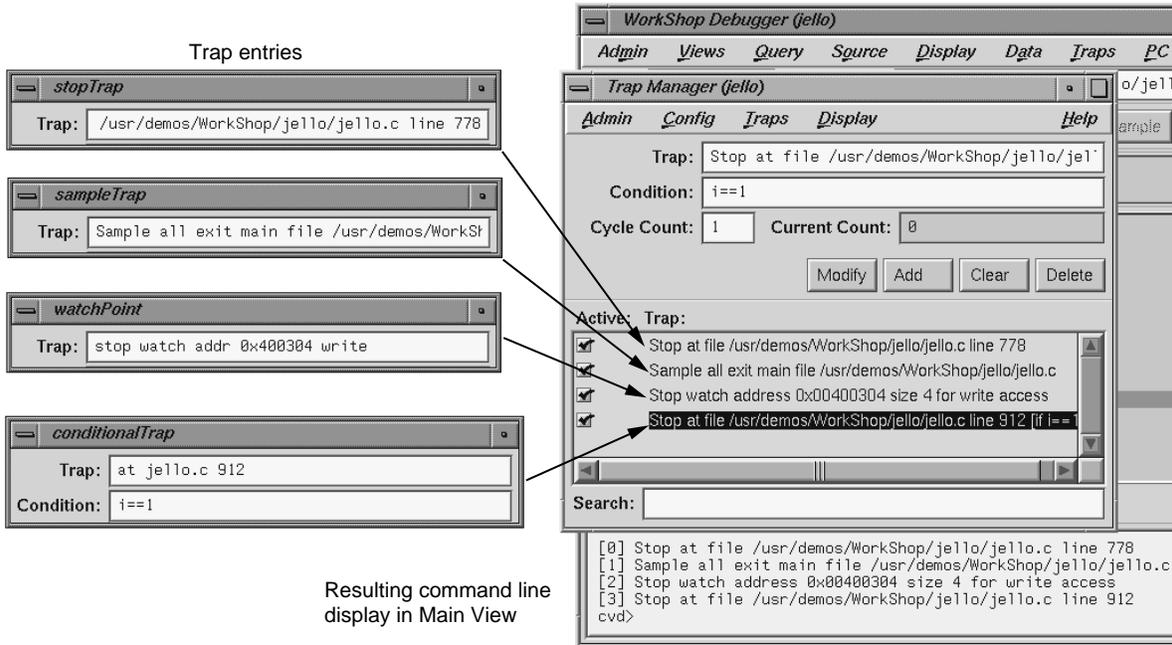


Figure 28. Trap Examples

5.3.1.1 Syntaxes

Specific command syntaxes that may be entered in the Trap text field are shown in the following sections:

Note: A backslash (\) shown at the end of a line in these examples indicates that the command is continued. Do not enter a backslash in the text field.

5.3.1.1.1 Setting aTrap in filename at line-number

```
[stop | sample] [all] [pgrp] at [{file} filename]\
  [[line] line-number]
```

This command sets a trap at the specified line in the specified file, for example:
Trap: stop at 1449.

5.3.1.1.2 Setting a Trap on *instruction-address*

```
[stop | sample] [all] [pgrp] addr instruction-address
```

This command sets a trap on the specified instruction address. Instruction addresses may be obtained from the Disassembly View window, which is brought up from the Views submenu on the Main View window menu bar. The addresses may be entered as shown, such as '0af8cbb8' X, or as 0x0af8cbb8. For example:

```
Trap: stop addr 0x0af8cbb8
```

5.3.1.1.3 Setting a Trap on Entry to *function*

```
[stop | sample] [all] [pgrp] entry function \  
[[file] filename]
```

```
[stop | sample] [all] [pgrp] in function \  
[[file] filename]
```

This command sets a trap on entry to the specified function. For example:

```
Trap: stop entry anneal
```

or

```
Trap: stop in anneal
```

If the filename is given, the function is assumed to be in that file's scope.

5.3.1.1.4 Setting a Trap on Exit from *function*

```
[stop | sample] [all] [pgrp] exit function \  
[[file] filename]
```

This command sets a trap on exit from the specified function. For example:

```
Trap: stop exit anneal file generic.c
```

If the filename is given, the function is assumed to be in that file's scope.

5.3.1.1.5 Setting a Watchpoint on Specified *expression*

```
[stop | sample] [all] [ pgrp] watch expression \ [  
[for] read | write | execute [access]]
```

This command sets a watchpoint on the specified expression (using the address and size of the expression for the watchpoint span). The watchpoint may be specified to fire on write, read, or execute (or some combination thereof). If not specified, the write condition is assumed. This syntax has no provision for looking on only a portion of an array. The next syntax item can handle such a request. For example:

```
Trap: stop watch x for write
```

5.3.1.1.6 Setting a Watchpoint for *address* and *size*

```
[stop | sample] [all] [pgrp] watch addr[ess] address \  
[[size] size] [[for] read | write | execute \  
[access]]
```

This command sets a watchpoint for the specified address and size in bytes. Typically the expression is the name of a variable. The watchpoint may be specified to fire on write, read, or execute (or some combination thereof) of memory in the given span. If not specified, the size defaults to 4 bytes. Also, if not specified, the write condition is assumed. Addresses may be found by choosing the following from the Main View window menu bar:

```
Views  
->Variable Browser
```

or

```
Views  
->Structure Browser
```

The window displays addresses for arrays and has options to display addresses for other variables. Addresses may be entered in the form 0x0123fabc or '0123fabc'X. For example:

```
Trap: stop watch addr 0x0123fabc 16
```

If the array you are watching at address 0x0123fab0 is defined such that each element is 4 bytes long, this example trap would be watching 4 adjacent array elements and would cause a stop if any of those 4 elements were updated, since a size of 16 is specified.

5.3.1.1.7 Setting a Trap at *signal-name*

```
[stop | sample] [all] [pgrp] signal signal-name
```

This command sets a trap upon receipt of the given signal. This is the same as the `dbx(1) catch` subcommand. (For a list of signals, or an alternative way to set traps involving signals, see Section 5.4, page 84.) For example:

```
Trap: stop signal SGIFPE
```

5.3.1.1.8 Setting a Trap on Entry to *sys-call-name*

```
[stop | sample] [all] [pgrp] syscall entry sys-call-name
```

This command sets a trap on entry to the specified system call. This is slightly different from setting a trap on entry to the function by the same name. A `syscall entry` trap sets a trap on entry to the actual system call. A `function entry` trap sets a trap on entry to the stub function that calls the system call. (For a list of system calls, or an alternative way to set traps involving system calls, see Section 5.4, page 84.) For example:

```
Trap: stop syscall entry write
```

5.3.1.1.9 Setting a Trap on Exit from *sys-call-name*

```
[stop | sample] [all] [pgrp] syscall exit sys-call-name
```

This command sets a trap on exit from the specified system call. This is slightly different from setting a trap on exit from the function by the same name. A `syscall exit` trap sets a trap on exit from the actual system call. A `function exit` trap sets a trap on exit from the stub function that calls the system call. (For a

list of signals, or an alternative way to set traps involving signals, see Section 5.4, page 84.) For example:

```
Trap: stop syscall exit read
```

5.3.1.1.10 Setting a Trap at *time* Interval

```
[stop | sample] pollpoint [interval] time [seconds]
```

This command sets a trap at regular intervals of seconds. This is typically used only for sampling. For example:

```
Trap: stop pollpoint 3
```

5.3.1.1.11 Setting a Trap for C++ Exception

```
[stop | sample] exception [all | item] itemname
```

This command sets a trap on all C++ exceptions, or exceptions that throw the base type *item*.

```
[stop | sample] exception unexpected [all | [item [, item]]]
```

Stops on all C++ exceptions that have either no handler or are caught by an unexpected handler. If you specify *item*, stops on executions that throw the base type *item*.

5.3.2 Setting a Trap Condition

The Condition field in the Trap Manager window lets you specify the condition necessary for the trap to be fired. A condition can be any legal expression and is considered to be true if it returns a nonzero value when the corresponding trap is encountered.

The expression must be valid in the context in which it will be evaluated. For example, a Fortran condition like `a .gt. .2` cannot be evaluated if it is tested while the program is stopped in a C function.

There are two possible sequences for entering a trap with a condition:

1. Define the trap.
2. Define the condition.
3. Click Add.

or

1. Define the trap.
2. Click Add.
3. Define the condition.
4. Click Modify (or press Enter).

An example of a trap with a condition is shown in Figure 28, page 78. The expression `i==1` has been entered in the `Condition` field. (If you were debugging in Fortran, you would use the Fortran expression `i .eq. 1` rather than `i==1`.) After the trap has been entered, the condition appears as part of the trap definition in the display area. During execution, any requirements set by the trigger must be satisfied first for the condition to be tested. A condition is true if the expression (valid in the language of the program you are debugging) evaluates to a nonzero value.

5.3.3 Setting a Trap Cycle Count

The `Cycle Count` field in the `Trap Manager` window lets you pass through a trap a specific number of times without firing. If you set a cycle count of n , the trap will fire every n th time the trap is encountered. The `Current Count` field indicates the number of times the process has passed the trap since either the cycle count was set or the trap last fired. The current count updates only when the process stops.

5.3.4 Setting a Trap with the Traps Menu

The `Traps` menu of the `Trap Manager` window lets you specify traps in conjunction with the `Main View` or `Source View` windows. Clicking `At Source Line` sets a trap at the line in the source display area that is currently selected. To set a trap at the beginning or end of a function, highlight the function name in the source display and click `Entry Function` or `Exit Function`.

5.3.5 Moving around the Trap Display Area

The trap display area displays all traps set for the current process. There are vertical and horizontal scroll bars for moving around the display area. The `Search` field lets you incrementally search for any string in any trap.

5.3.6 Enabling and Disabling Traps

Each trap has an indicator to its left for toggling back and forth between active and inactive trap states. This feature lets you accumulate traps and turn them on only as needed. Thus, when you do not need the trap, it will not be in your way. When you do need it, you can easily activate it.

5.3.7 Saving and Reusing Trap Sets

The `Load Traps` selection in the `Config` menu lets you bring in previously saved trap sets. This is useful for reestablishing a set of traps between debugging sessions. The `Save Traps...` selection of the `Config` menu lets you save the current traps to a file.

5.4 Setting Traps by Using `Signal Panel` and `System Call Panel`

You can trap signals by using the `Signal Panel` and set system calls by using `System Call Panel` (see Figure 29, page 85).

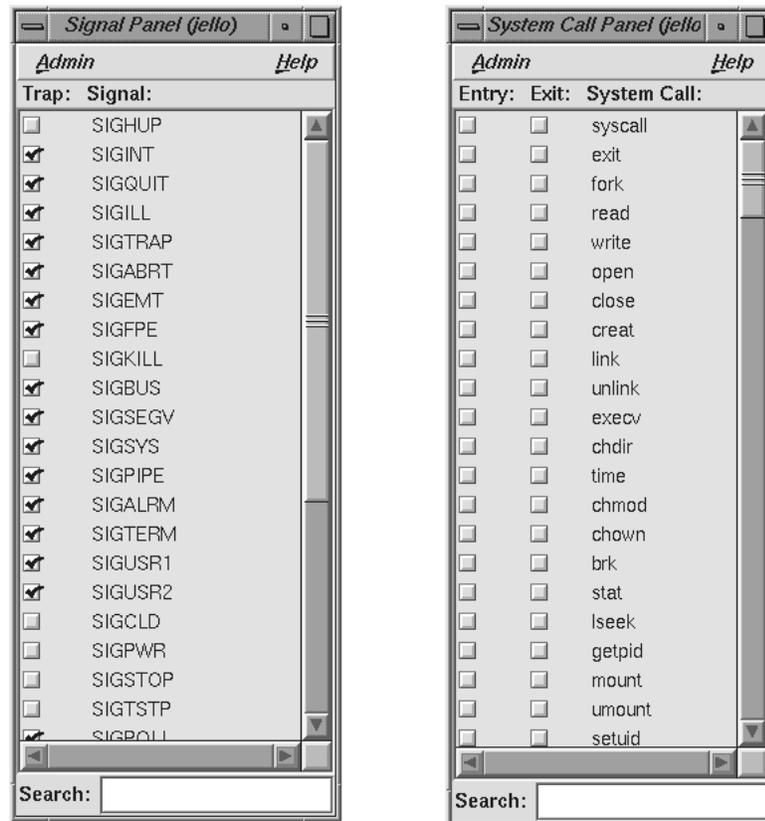


Figure 29. Signal Panel and System Call Panel

You can select either panel from the Views menu of the Main View window menu bar. The Signal Panel sets a trap on receipt of the signal(s) selected. The System Call Panel sets a trap at the selected entry to or return from the system call.

Note: When debugging IRIX 6.5 pthreads, the Signal Panel is inaccessible if more than one thread is active.

Controlling Program Execution [6]

This chapter shows you how to control the execution of your program with Workshop Debugger. It includes the following topics:

- Main View Control Panel, Section 6.1, page 87
- Controlling Program Execution Using the PC Menu, Section 6.2, page 91
- Execution View, Section 6.3, page 92

6.1 The Main View Window Control Panel

The Main View window control panel allows you to choose an executable, and control its execution:

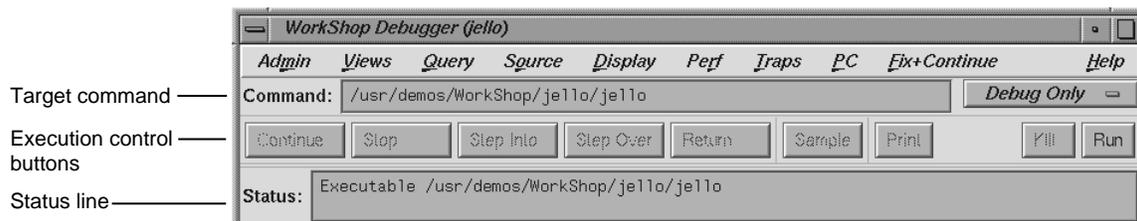


Figure 30. The Main View Window Control Panel

6.1.1 Features of the Main View Window Control Panel

The control panel includes the following:

- | | |
|---------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Command field | Use this field to enter commands (with arguments) to run your program. |
| Status field | Displays information about the execution status of your program. The top line in this box indicates whether the program is running or stopped. The message No executable displays if no executable is loaded. When your program stops at a breakpoint, an additional status line lists the current stack frame. |

Note: To see all of the stack frames, select the following from the Menu Bar:

Views->Call Stack View

6.1.2 Execution Control Buttons

The execution control buttons enable you to control program execution. Most of these buttons are not active until the Run button has been selected and the program is executed.

Note: The Print button does not affect program execution. It is described in Appendix A, page 183.

Run	Creates a new process for your program and starts its execution. The Run button is also used to re-run a program.
Kill	Kills the active process.
Continue	Resumes program execution after a halt and continues until a breakpoint or other event stops execution.
Stop	Stops execution of your program. When program execution stops, the current source line is highlighted in the Main View window and annotated with an arrow.
Step Into	Steps to the next source line and into function or subroutine calls. To step a specific number of lines, right-click on the button. This displays the pop-up menu shown in Figure 31, page 89. You can select one of the fixed values or enter your own number of steps by selecting N. . . . Selecting N. . . displays the dialog box shown at the right in Figure 31, page 89.

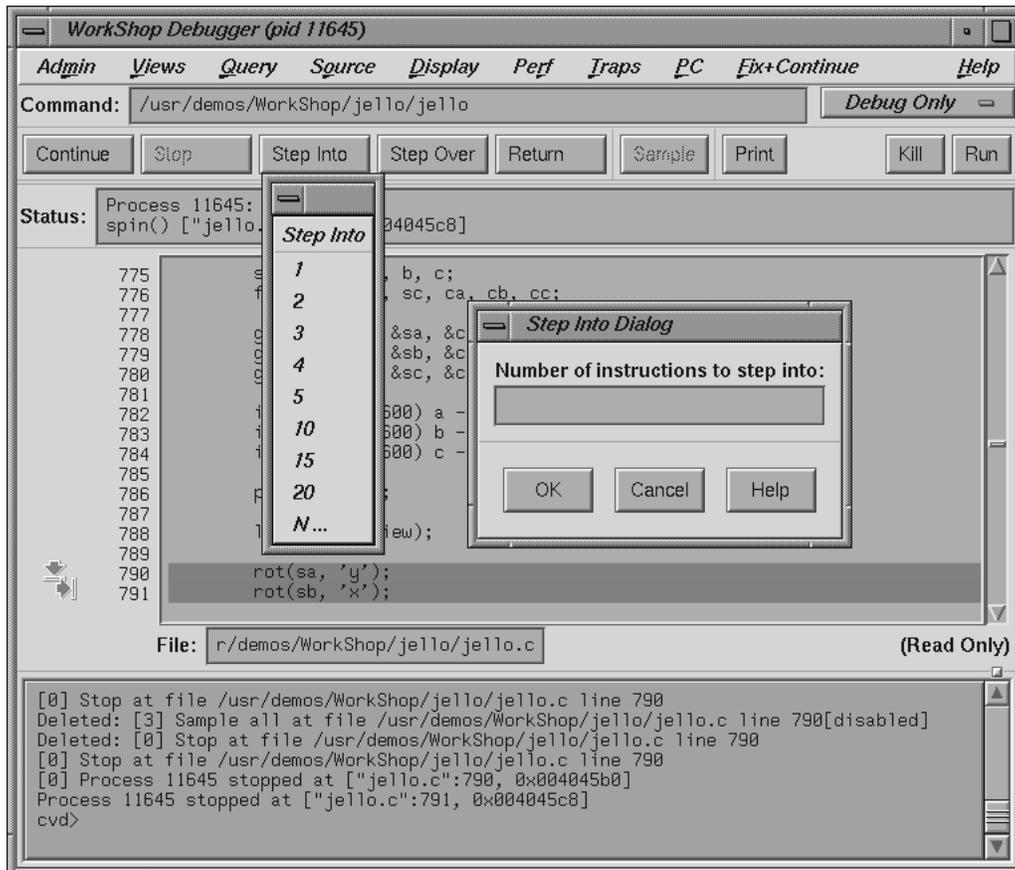


Figure 31. Pop-up Menu and Step Into Dialog

Step Over

Steps over function or subroutine calls to the next source line. To step a specific number of lines, right-click on this button to display the pop-up menu shown in Figure 32, page 90. You can select one of the fixed values or enter your own number of steps by selecting *N . . .*. If you select *N . . .*, the dialog box shown at the right in Figure 32, page 90 displays.

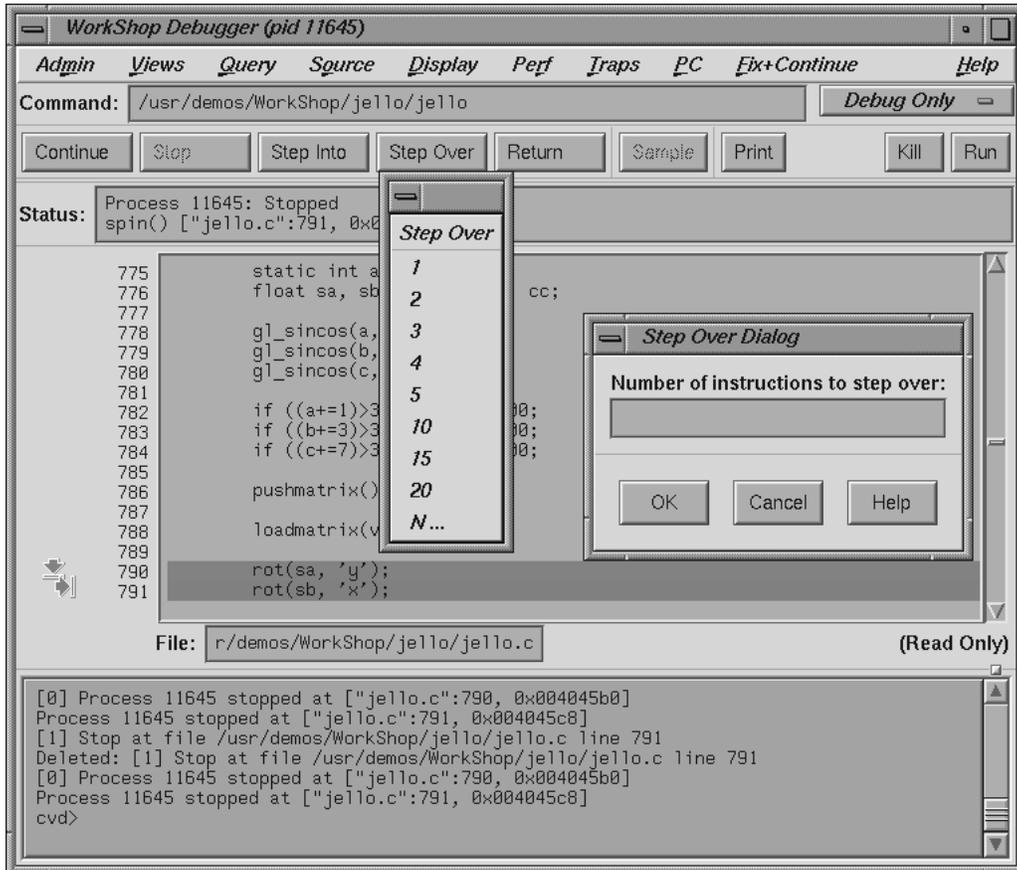


Figure 32. Pop-up Menu and Step Over Dialog

Return

Executes the remaining instructions in the current function or subroutine, and stops execution at the return from this subprogram.

Sample

Collects performance data. Before this button is operative, a performance task must have been previously specified in the Performance Task window and data collection must have been enabled.

For further information about using the Performance Analyzer, see *Developer Magic: Performance Analyzer User's Guide*.

6.2 Controlling Program Execution Using the PC Menu

The PC (program counter) menu in the Main View window includes two tools, `Continue To` and `Jump To`, which allow you to control program execution without setting breakpoints.

These tools are inoperative until a process has been executing and is stopped. At that point, you must place your cursor on a source line you wish to target, then select from the PC menu depending on your requirements. The program counter tools are described as follows:

<code>Continue To</code>	This tool lets you select a target location in the current program (by placing the cursor in the line). The process proceeds from the current program counter to that point, provided there are no interruptions. It then stops there, as it would for a stop trap. <code>Continue To</code> is equivalent to setting a one-time trap. If the process is interrupted before reaching your target location, then the command is cancelled.
<code>Jump To</code>	This tool lets you select a target location in the current program (by placing the cursor in the line). This location must be in the same function. Instead of starting from the current program counter, <code>Jump To</code> skips over any intervening code and restarts the process at your target. This is particularly useful if you want to get around bad code or irrelevant portions of the program. It also lets you back up and reexecute a portion of code.

6.3 Execution View

The `Execution View` window is a simple shell that lets you set environment variables and inspect error messages. If your program is designed to be interactive using standard I/O, this interaction will take place in the `Execution View` window. Any standard I/O that is not redirected by your `Target Command` is displayed in the `Execution View` window.

Note: When you launch the debugger, the `Execution View` window is launched in iconified form.

Viewing Program Data [7]

After you set traps (breakpoints) in your program, use the Run button to execute your program. When a trap stops a process, you can view your program data using the tools described in this chapter. This chapter covers:

- Tracing Through Call Stack View, Section 7.1, page 93.
- Evaluating Expressions, Section 7.3, page 99.

The Debugger also lets you examine data at the machine level. The tools for viewing disassembled code, machine registers, and data by specific memory location are described in Appendix A, page 183.

7.1 Traceback through Call Stack View

The Views menu may be used to bring up the Call Stack View window. This window displays the functions/subroutines (that is, the “frames”) in the call stack when the process associated with your program has stopped. This display provides a traceback of subprograms from the system routine which starts your executable, found at the bottom of the list, to the routine in which you are currently stopped, found at the top of the list. If the trap/breakpoint at which you are currently stopped is located in your source code, that code is displayed in the source pane of the Main View window. If the trap/breakpoint is located in a system routine, the Call Stack View allows you to double-click on another routine’s name to bring up that routine’s source and associated data. A typical Call Stack View window is shown in Figure 33, page 94.

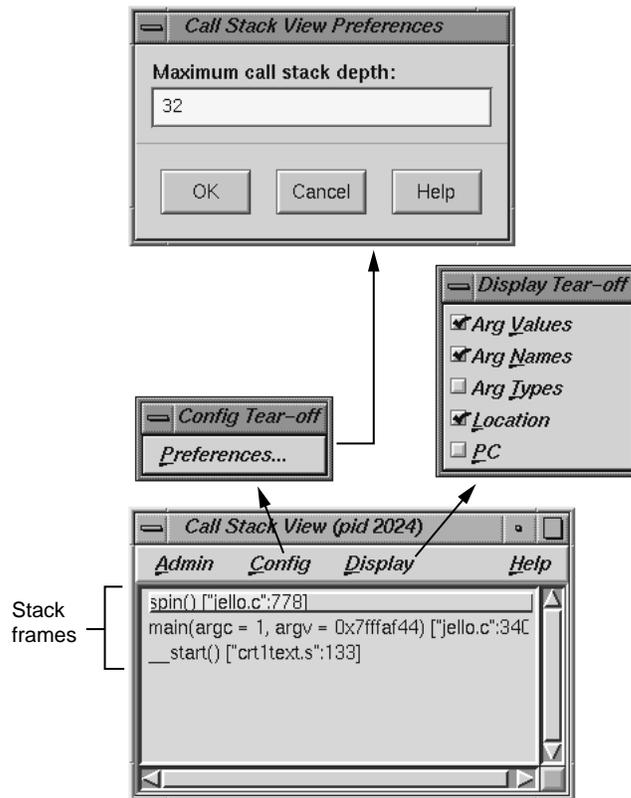


Figure 33. Call Stack View Window

The Call Stack View window lets you see the argument names, values, types, and locations of functions, as well as the program counter (PC).

If symbolic information for the arguments has been stripped from the executable file, the label <stripped> appears in place of the arguments. By default call stack depth is set to 10, but you can reset the depth of the Call Stack View by selecting the following from the Main View window:

```
Config
->Preferences...
```

To move through the call stack, double-click a frame in the stack. The frame becomes highlighted to indicate the current context. The source display in the

Main View or Source View windows scrolls automatically to the location where the function was called and any other active views update.

The source display has two special annotations:

- The location of the current program state is indicated by a large arrow. This represents the PC (program counter).
- The location of the call to the function selected in the Call Stack View window is indicated by a smaller arrow. This represents the current context, and the source line is highlighted.

Figure 34, page 96, illustrates the correspondence between a frame and the source code when a frame is clicked in the Call Stack View window. In this example, the stack frame `spin` has been selected; the Main View display scrolls to the place where the trap occurred. If the second stack (`main`) had been selected, the window would have scrolled to the place where the function `main` calls `spin`.

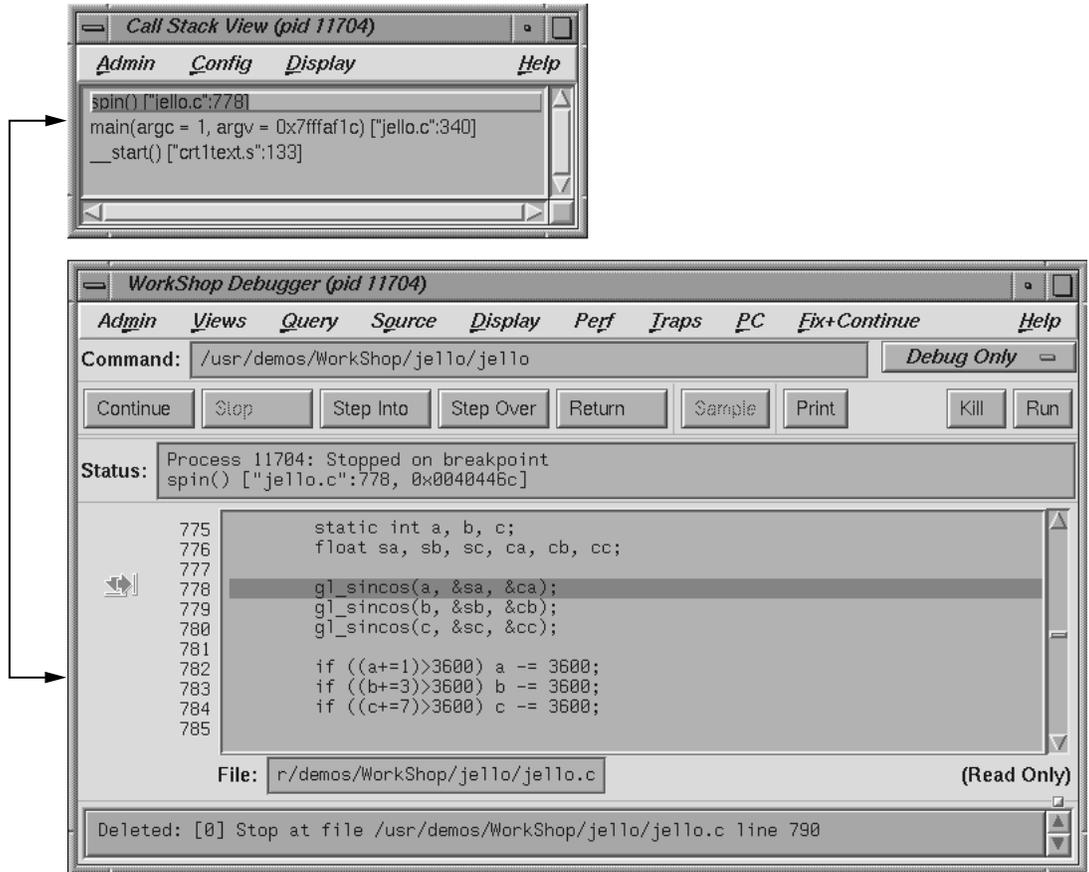


Figure 34. Tracing through Call Stack View

7.2 Options for Viewing Variables

The WorkShop Debugger provides several options for viewing variables or expressions involving these variables. In this section only brief descriptions of these options will be provided along with references where further information may be obtained elsewhere in this document.

7.2.1 Using the `cvd`> Command Line

At the bottom of the Main View window is the `cvd` command/message pane. Here, you can enter `print xxx` commands to obtain the current value of variable `xxx`.

For examples of these commands, see Section 2.1.6.1, page 21.

For the syntax of these commands (such as `assign`, `print`, `printd`, `printo`, `printx`, and so on), see Section A.9.1, page 312.

7.2.2 Using Click to Evaluate

In the Main View window's source pane, a click of the right mouse button will bring up a pop-up menu from which you can select `Click to Evaluate`. When this option is on, a click on a variable will cause its value to appear. If you click on a subscript variable, its address will appear. If you hold down the left mouse button and wipe across a variable and its subscripts to highlight them, the current value will be displayed. The same is true if an expression in the source is highlighted.

7.2.3 Using the Array Browser

To view values of arrays, select the following from the Main View menu bar:

```
Views  
->Array Browser
```

This calls up the `Array Browser` window. When the name of an array is entered into the `Array` field, the values of 1-dimensional arrays or the values in a selected plane of a multi-dimensional array are displayed.

See Section 2.1.6.5, page 25 for short description of the `Array Browser`. Also, Section A.6.1, page 257, "Array Browser Window", provides a more detailed description and graphic of the window and associated submenus.

7.2.4 Using the Structure Browser

To view C structures, select the following from the Main View menu bar:

```
Views  
->Structure Browser
```

This calls up the `Structure Browser` window. When the name of a structure is entered into the `Expression` field, the objects in the structure display in the lower portion of the window.

See Section 2.1.6.6, page 27 for a short description of the `Structure Browser`. Also, Section A.6.5, page 277, "`Structure Browser Window`", provides a more detailed description and graphic of the window and associated submenus.

7.2.5 Using the `Variable Browser`

To call up the `Variable Browser` window, select the following from the `Main View` menu bar:

```
Views  
->Variable Browser
```

This window lists the names and values of variables associated with the current routine in which the process has stopped.

See Section 2.1.6.3, page 23 for a short description of the `Variable Browser`. Also, Section A.6.6, page 288, "`Variable Browser Window`", provides a more detailed description and graphic of the window and associated submenus.

7.2.6 Using the `Expression View`

Select the following from the `Main View` menu bar to bring up the `Expression View` window:

```
Views  
->Expression View
```

Expressions may be entered into this window in the `Expression` column; and the corresponding `Results` entry will display the value of the evaluated expression using values of variables associated with the current routine in which the process has stopped.

See Section 2.1.6.4, page 23 for a short description of the `Expressions View` window. Also, Section A.6.3, page 273, `Expression View Window`, provides a more detailed description and graphic of the window and associated submenus.

The next section of this chapter, "`Evaluating Expressions`", gives further details on how to create acceptable expressions for the `Expression View`.

7.3 Evaluating Expressions

You can evaluate any valid expression at a stopping point and trace it through the process. Expressions are evaluated by default in the frame and language of the current context. Expressions may contain data names or constants; however, they may not contain names known only to the C preprocessor, such as in a `#define` directive or a macro.

To evaluate expressions, you can use `Expression View`, which lets you evaluate multiple expressions simultaneously, updating their values each time the process stops.

Note: You can also evaluate expressions from the command line. See Section A.9, page 312 for more information.

7.3.1 Expression View Window

The `Expression View` window is shown in Figure 35, page 100, with its major menus displayed. The `Expression View` window has two pop-up menus, the `Language` menu and the `Format` menu:

- The `Language` menu is invoked by holding down the right mouse button while the cursor is in the `Expression` column.
- The `Format` menu is displayed by holding down the right mouse button in the `Result` column.

To specify the expression to be evaluated, click in the `Expression` column and then enter the expression in the selected field. It must be a valid expression in the current or selected language: Ada, C, C++, or Fortran. To change languages, display the `Language` menu and make your selection. When you press `Enter`, the result of the expression is displayed in the `Result` column.

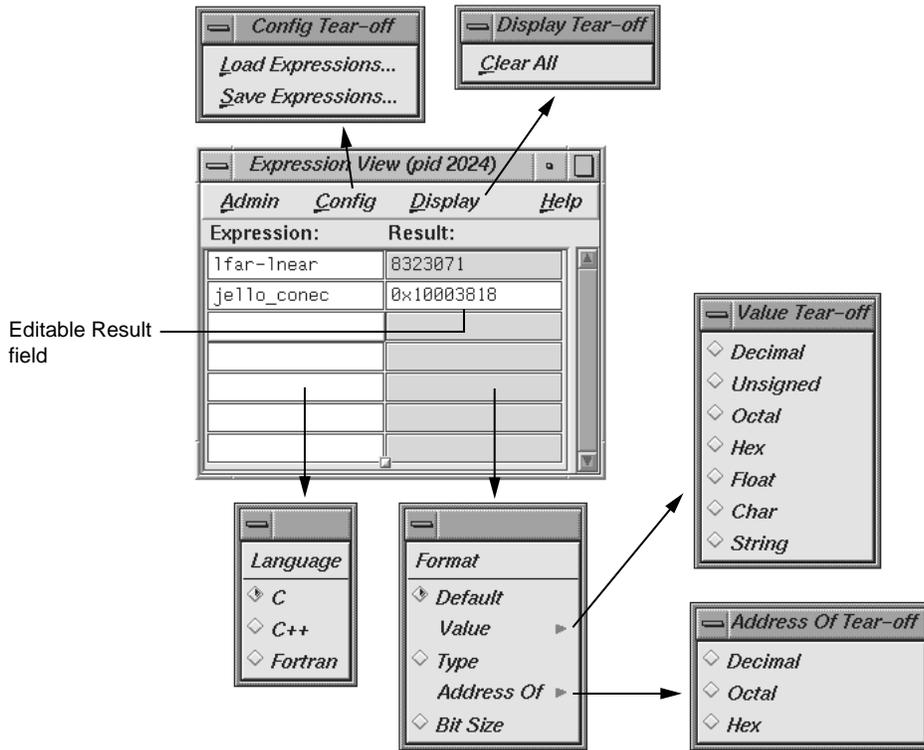


Figure 35. Expression View with Major Menus Displayed

To change the type of result information displayed in the right column, hold down the right mouse button over the right column. This displays the *Format* menu. From here you can select the following:

- Select the *Default Value* menu to see the value as decimal, unsigned, octal, hex, float, char, or string characters.
- Select the *Type Address Of* menu to display the address in decimal, octal, or hexadecimal.
- Select *Bit Size* to specify the size of the result, in bits.



Caution: The Debugger uses the symbol table of the target program to determine variable type. Some variables in libraries, such as `errno` and `_environ`, are not fully described in the symbol table. As a result, the Debugger may not know their types. When the Debugger evaluates such a variable, it assumes that the variable is a fullword integer. This gives the correct value for fullword integers or pointers, but the wrong value for non-fullword integers and for floating-point values.

To see the value of a variable of unknown type, use C type cast syntax to cast the address of the variable to a pointer that points to the correct type. For example, the global variable `_environ` should be of type `char**`. You can see its value by evaluating `*(char***)&_environ`.

After you display the current value of the expression, you may find it useful to leave the window open so that you can trace the expression as it changes value from trap to trap (or when you change the current context by double-clicking in the call stack). Like other views involved with variables, `Expression View` has variable change indicators for value fields that let you see previous values, as shown in the following figure:

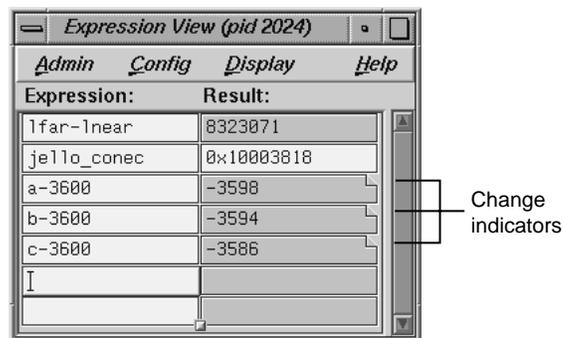


Figure 36. Change Indicators in Expression View

Another useful technique is to save your expressions to a file for later reuse. To save expressions, select the following from the Main View menu bar:

```
Config  
  ->Save Expressions
```

To load expressions, select the following from the Main View menu bar:

```
Config  
  ->Load Expressions...
```

7.3.2 Assigning Values to Variables

To assign a value to a variable, click the left column of the Expression View window and enter the variable name. The current value appears in the right column. If this Result field is editable (highlighted), you can click it and enter a new value or legal expression. Press Enter to assign the new value. You can perform an assignment to any expression that evaluates to a legal lvalue (in C). The C operator "=" is not valid in Expression View. Valid expression operations are shown in the following paragraphs.

7.3.3 Evaluating Expressions in C

The valid C expressions are shown in Table 1.

Table 1. Valid C Operations

Operation	Symbol
Arithmetic ¹	+ - ++ --
Arithmetic (binary)	+ - * / %
Logical	&& !
Relational	< > <= >= == !=
Bit	& ^ << >> ~
Dereference	*
Address	&
Array indexing	[]
Conditional	? :
Member extraction ²	. ->
Assignment ³	= += -= /= %= >>= <<= &= ^= =
Sizeof	
Type-cast	
Function call	

¹ Unary - increment and decrement do not have side-effects

² These operations are interchangeable.

³ A new assignment is made at each stepping point. Use Assignments with caution to avoid inadvertently modifying variables.

7.3.3.1 C Function Calls

Function calls can be evaluated in expressions, as long as enough actual parameters are supplied. Arguments are passed by value. Following the rules of C, each actual parameter is converted to a value of the same type as the formal parameter, before the call. If the types of the formal parameters are unknown, integral arguments are widened to full words, and floating-point arguments are converted to doubles.

Functions may return pointers, scalar values, unions, or structs. Note that if the function returns a pointer into its stack frame (rarely a good programming practice), the value pointed to will be meaningless, since the temporary stack frame is destroyed immediately after the call is completed.

Function calls may be nested. For example, if your program contains a successor function `succ`, the Debugger will evaluate the expression `succ(succ(succ(3)))` to 6.

7.3.4 Evaluating Expressions in C++

C++ expressions may contain any of the C operations. You can use the word `this` to explicitly reference data members of an object in a member function. When stopped in a member function, the scope for `this` is searched automatically for data members. Names may be used in either mangled or de-mangled form. Names qualified by class name are supported (for example, `Symbol::a`).

If you wish to look at a static member variable for a C++ class, you need not specify the variable with the class qualifier if you are within the context of the class. For example, you would specify `myclass::myvariable` for the static variable `myvariable` outside of class `myclass` and `myvariable` inside `myclass`.

7.3.4.1 Limitations

Constructors may be called from `Expression View`, just like other member functions. To call a constructor, you must pass in a first argument that points to the object to be created. C++ function calls have the same possibility of side effects as C functions.

7.3.5 Evaluating Expressions in Fortran

You can enter any Fortran expression under the `Expression` heading and its current value will appear in the same row under the `Results` column. Fortran expressions may contain any of the arithmetic, relational, or logical operators. Relational and logical operator keywords may be spelled in upper case, lower case, or mixed case.

The usual forms of Fortran constants, including complex constants, may be used in expressions. String constants and string operations, however, are not supported. The operators in Table 2 are supported on data of integer, real, and complex types.

Table 2. Valid Fortran Operations

Operation	Symbol
Arithmetic (unary)	- +
Arithmetic (binary)	- + * / **
Logical	.NOT. .AND. .OR. .XOR. .EQV .NEQV.
Relational	.GT. .GE. .LT. .LE. .EQ. .NE.
Array indexing	()
Intrinsic function calls (except string intrinsics)	
Function subroutine calls	
Assignment ⁴	=

⁴ A new assignment is made at each stepping point. Use assignments with caution to avoid inadvertently modifying variables.

7.3.5.1 Fortran Variables

Names of Fortran variables, functions, parameters, arrays, pointers, and arguments are all supported in expressions, as are names in common blocks and equivalence statements. Names may be spelled in upper case, lower case, or mixed case.

7.3.5.2 Fortran Function Calls

The Debugger evaluates function calls the same way that compiled code does. If it can be, an argument is passed by reference; otherwise, a temporary expression is allocated and passed by reference. Following the rules of Fortran, actual arguments are not converted to match the types of formal arguments. Side effects can be caused by Fortran function calls. A useful technique to protect the value of a parameter from being modified by a function subroutine is to pass an expression such as `(parameter + 0)` instead of just the parameter name. This causes a reference to a temporary expression to be passed to the function rather than a reference to the parameter itself. The value is the same.

Debugging with Fix+Continue [8]

Fix+Continue allows you to make changes to a C/C++ program you are debugging without having to recompile and link the entire program. With Fix+Continue you can edit a function, parse the new function, and continue execution of the program being debugged.

Fix+Continue is an integral part of the Debugger. You issue Fix+Continue commands graphically from the `Fix+Continue` submenu of the Main View window, or from the `cvd>` command line prompt in the Command/Message pane of the Main View window.

This chapter provides an introduction to the Fix+Continue functionality as well as a tutorial to demonstrate many of the Fix+Continue functions.

8.1 Introduction to Fix + Continue

The following sections give you an overview to Fix+Continue functionality.

8.1.1 Fix+Continue Functionality

Fix+Continue lets you perform the following activities:

- Redefine existing function definitions
- Disable, re-enable, save, and delete redefinitions
- Set breakpoints in redefined code.
- Single-step within redefined code.
- View the status of changes
- Examine differences between original and redefined functions

A typical Fix+Continue cycle proceeds as follows:

1. You redefine a function with Fix+Continue. When you continue executing the program, the Debugger attempts to call the redefined function. If it cannot, an information pop-up window appears and the redefined function will be executed the next time the program calls that function.

2. You redefine other functions, alternating between debugging, disabling, re-enabling, and deleting redefinitions. You might save function redefinitions to their own files, or save files to a different name, to be used later with the present or with other programs.

During debugging you can review the status of changes by listing them, showing specific changes, or looking at the `Fix+Continue Status View`. You can compare changes to an individual function or to an entire file with the compiled versions. When you are satisfied with the behavior of your application, save the changed file as a replacement for the compiled source.

8.1.2 Fix+Continue Integration with Debugger Views

Fix+Continue interacts with the following Views:

- The Views main view, the Source View, and Fix+Continue Status windows distinguish between compiled and redefined code, and allow editing in redefined code.
- The following status windows elements work with redefined code:
 - Call Stack window
 - Trap Manager
 - Debugger command line

8.1.3 How Redefined Code Is Distinguished from Compiled Code

Redefined functions have an identification number and special line numbers. They are color-coded according to their state (that is, edited, parsed, and so on).

Line numbers in the *compiled* file stay the same, no matter how redefined functions change. However, when you begin editing a function, the line numbers of the function body are represented in decimal notation ($n.1, n.2, \dots, n.m$), where n is the compiled line number where the function body begins, and m is the line number relative to the beginning of the function body, starting with the number 1.

The Call Stack window and the Trap Manager functions both use function-relative decimal notation when referring to a line number within the body of a redefined function.

The Debugger command line reports ongoing status. In addition to providing the same commands available from the menu, edit commands allow you to

add, replace, or delete lines from files. Therefore, you can operate on several files at once.

8.1.4 The Fix+Continue Interface

You can access Fix+Continue through the `Fix+Continue` menu. It includes three supporting windows: `Status`, `Message`, and `Build Environment`. These windows are part of Fix+Continue, and do not operate unless it is installed.

8.1.5 Debugger with Fix+Continue Support

Without Fix+Continue, the Debugger source views are `Read-Only` by default. That is so you can examine your files with no risk of changing them. When you select `Edit` from the `Fix+Continue` menu, the Debugger source code status indicator (in the lower-right corner of the Debugger window) remains `Read-Only`. This is because edits made using Fix+Continue are saved in an intermediate state. Instead, you must choose `Save File->Fixes As ...` to save your edits.

When you edit a function, it is highlighted in color; and if you switch to the compiled version of your code, the color changes to show that the function has been redefined. If you try to edit the compiled version of your code, the Debugger beeps indicating `Read-Only` status.

When you have completed your edits and want to see the results, select `Parse` and `Load`. When the parse and load has executed successfully, the color changes again. If the color does not change, there may be errors: check the `Message Window`.

8.1.6 Change ID, Build Path, and Other Concepts

The Fix+Continue features finding files and accessing functions through ID numbers as follows:

- Each redefined function is numbered with a change ID. Its status may be shown as `redefined`, `enabled`, `disabled`, `deleted`, or `detached`.
- Fix+Continue needs to know the location of `include` files and other parameters specified by compiler build flags. You can set the build environment for all files or for a specific file. You can display the current build environment from the `Fix+Continue` menu, the command line, or

the `Fix+Continue Status Window`. When you finish a `Fix+Continue` session, you can unset the build environment.

- Output from a successful run is displayed in the `Execution View`. This functionality is the same as it is in the Debugger without `Fix+Continue`.

8.2 Restrictions on `Fix+Continue`

`Fix+Continue` has the following restrictions:

- When you work with C code, you must use the `-o32` compiler option.
- `Fix+Continue` does not support C++ templates.
- You may not add, delete, or reorder local variables in a function.
- You may not change the type of a local variable.
- You may not change a local variable to be a register variable and vice-versa.
- You may not add any function calls that increase the size of the parameter area.
- You may not add an `alloca` function to a frame that did not previously use an `alloca` function.
- Both the old and new functions must be compiled with the `-g` option.

In other words, the layout of the stack frames of both the old and new functions must be identical for you to continue execution in the function that is being modified. If not, execution of the old function continues and the new function is executed the next time the function is called.

- If you redefine functions that are in but not on top of the call stack, the modified code will not be executed when they combine. Modified functions will be executed only on their next call or on a rerun.

For example, consider the following call stack:

```
foo()  
bar()  
foobar()  
main()
```

- If you redefine `foo()`, you can continue execution, provided that the layout of the stack frames are the same.

- If you redefine `main()` after you have begun execution, the redefined `main()` will be executed only when you rerun.
- If you redefine `bar()` or `foobar()`, the new code will not be executed when `foo()` returns. It will be executed only on the next call of `bar()` or `foobar()`.

8.3 Fix+Continue Tutorial

This tutorial illustrates several features of Fix+Continue. The demo files included in `/usr/demos/WorkShop/time1` contain the complete C++ source code for the program `time1`. Use this program for your tutorial. Here you will see how Fix+Continue can modify functions without recompiling and linking the entire program.

This section contains the following subsections:

- Setting up the Sample Session, Section 8.3.1, page 111.
 - Redefining a C/C++ Function, Section 8.3.2, page 112
- Setting Breakpoints in Redefined Code, Section 8.3.3, page 118.
- Viewing Status, Section 8.3.4, page 121.
- Comparing Original and Redefined Code, Section 8.3.5, page 121.
- Ending the Session, Section 8.3.6, page 124.

8.3.1 Setting up the Sample Session

For this tutorial, use the demo files in the `/usr/demos/WorkShop/time1` directory that contains the complete source code for the C++ application `time1`. To prepare for the session, you must create the fileset and launch Fix+Continue from the Debugger as shown below:

1. Enter the following commands:

```
% mkdir demos/time1
% cd demos/time1
% cp /usr/demos/WorkShop/time1/* .
% make time1
% cvd time1 &
```

The `cvd` command brings up the Debugger, from which you can use the Fix+Continue utility. The Execution View icon and the Main View window appear. Note that the Debugger shows a source code status indicator of (Read Only).

2. Open the Execution View window and position it next to the Main View window.
3. Click Run to run `time1`.

The Execution View shows the program output (see Figure 37).

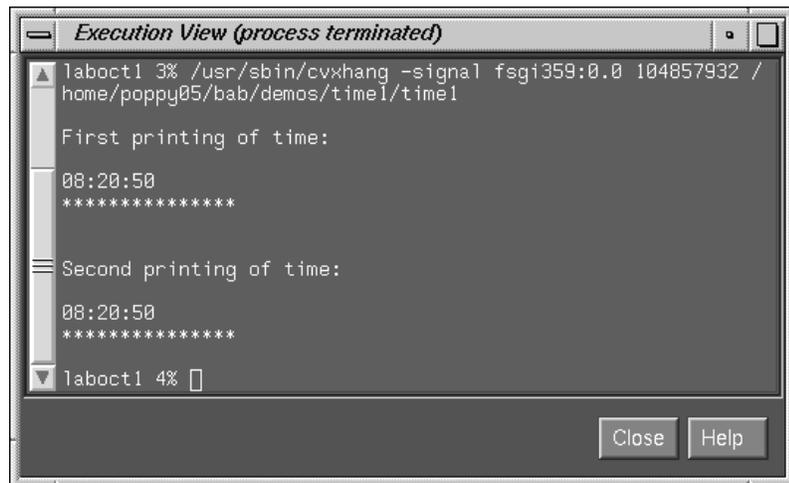


Figure 37. Program Results in Execution View

8.3.2 Redefining a Function: `time1` Program

In this section, you will do the following in the `time1` program:

- Edit a C/C++ function
- Change the code of an existing C/C++ function and then parse and load the function, rebuilding your program to see the effect of your changes on program output (without recompiling)
- Save the changed function to its own separate file.

8.3.2.1 Editing a Function

Perform the following in the Debugger Main View window (the `time1` program should be displayed) to edit a function:

1. Selecting the following from the Main View window menu bar to show line numbers:

```
Display  
  ->Show Line Numbers
```

2. Click on the source annotation column to the left of line 17 to set a breakpoint at that point.
3. Set a breakpoint at line 20 from the `cvd>` prompt as follows:

```
cvd> stop at 20
```

4. Click on the Run button to execute the program.

The following output appears in the Execution View window:

```
First printing of time:  
  
08:20:50  
*****  
  
Second printing of time:
```

5. Enter the following in the `cvd >` field to choose a class member function to edit (in this case, `printTime`, a C++ member function of class `Time`):

```
cvd> func Time::printTime
```

This command opens the `time1/time1.c` file, which contains the implementation of class `Time`. The cursor is placed at the beginning of the `printTime` function. (See Figure 38, page 114.)

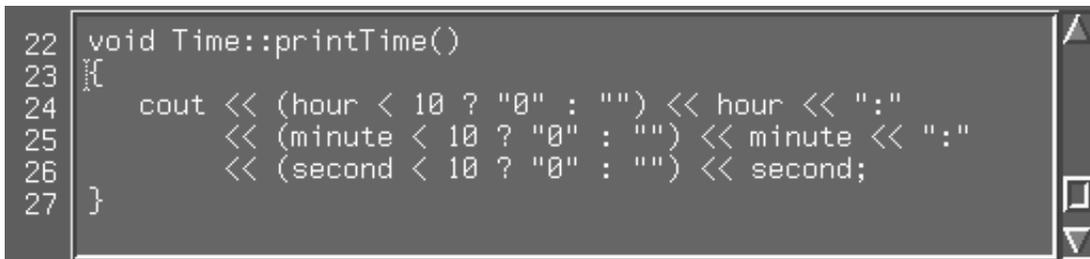
Note: The syntax of the `func` command is as follows:

- For C++ class member functions:

```
cvd> func className::classMemberFunction
```

- For all other C/C++ functions:

```
cvd> func functionName
```



```
22 void Time::printTime()
23 {
24     cout << (hour < 10 ? "0" : "") << hour << ":"
25         << (minute < 10 ? "0" : "") << minute << ":"
26         << (second < 10 ? "0" : "") << second;
27 }
```

Figure 38. Selecting a Function for Redefinition

6. Select the following option from the menu bar to highlight the function to be edited:

```
Fix+Continue
->Edit
```

Note: You can also use `Alt-Ctrl-e` to perform this task.

Note the results as shown in Figure 39, page 115. Line numbers changed to a decimal notation based on the first line number of the function body. The function body highlights to show that it is being edited. The line numbers of the rest of the file are not affected.

```

22 void Time::printTime()
23.1 {
23.2     cout << (hour < 10 ? "0" : "") << hour << ":"
23.3         << (minute < 10 ? "0" : "") << minute << ":"
23.4         << (second < 10 ? "0" : "") << second;
23.5 }

```

Figure 39. Redefined Function

8.3.2.2 Changing Code

1. To change the time output as shown in Step 5 in Section 8.3.2.1, page 113, delete the 0 from "0" in line 23.2.
2. Select the following option from the menu bar to parse the modified function and load it for execution:

```

Fix+Continue
    ->Parse and Load

```

An icon for the Fix+Continue Error Messages window displays; and the following message appears in the `cvd>` field:

```
Change id: 1 modified
```

Note: You can also use `Alt-Ctrl-x` to perform this task.

If there are errors:

- a. Go to the error location(s) by double-clicking the related message line in the Fix+Continue Error Messages window.
- b. Correct the errors.
- c. Repeat steps 1 and 2.

Continue to step 3 when you see the change ID and the following messages:

```

Change id: 1 redefined
Change id: 1 saved func
Change id: 1 file not saved
Change id: 1 modified

```

Note: The new function value is not active until the function is called.

3. Click on the `Continue` button to continue program execution.

The following output appears in the `Execution View` window:

```
Second printing of time:

8:20:50
*****
```

Notice how the time printout has changed from `08:20:50` to `8:20:50`.

8.3.2.2.1 Deleting Changed Code

To cancel any of your changes, you must bring up the source file in which the change was made and perform the following steps:

1. Enter the following at the `cvd>` prompt:

```
cvd> func Time::printTime
```

2. Select the following option from the menu bar to delete your changes:

```
Fix+Continue
  ->Delete Edits
```

The `Verify before deleting...` dialog displays.

3. Click `OK` in the `Verify before deleting...` dialog.

Your deletion is complete.

8.3.2.2.2 Changing Code from the Debugger Command Line

You can redefine a C++ class member function from the Debugger command line as follows:

1. Click on the `Kill` button in the `Main View` window.
2. Click on the `Run` button to re-run the program.
3. Choose a class member function (in this case, `printTime`) to edit by entering the following at the `cvd>` command line:

```
cvd> func Time::printTime
```

- Use the `redefine` command to edit the function:

```
cvd> redefine Time::printTime
```

Note the results as shown in Figure 39, page 115. Here, line numbers have changed to decimal notation based on the first line number of the function body. Note also that the command line prompt has changed.

- Enter the following at the prompt:

```
"/path/name/time1.C":23.1> .
```

- Change the function source by entering the following at the command line:

```
cvd> replace_source "time1.C":23.2
"time1.C:23.2> cout << (hour < 10 ? "" : "") << hour << ":"
"time1.C:23.3> .
```

Parse and Load is executed at this point. You exit out of the function edit mode and return to the main source code. The following messages appear in the command/message pane:

```
Change id: 2 redefined
Change id: 2 save func
Change id: 2 file not saved
Change id: 2 modified
Change id: 2 , build results:
      2      enabled .../time1.C Time::printTime(void)
```

Note: 2 in these messages is the redefined function ID. You will use this ID in the procedure in Section 8.3.5.1, page 121. The new function value is not active until the function is called.

- Continue execution by entering the following command at the `cvd>` prompt:

```
cvd> continue
```

The following displays in the Execution View window:

```
Second printing of time:

8:20:50
*****
```

If you prefer to use the command line, experiment with `add_source` and other commands that give you the same functionality described for the menu commands. For details on each command, see Section A.9, page 312.

8.3.2.3 Saving Changes

Your original source files are not updated until the changed source file is saved. You could save redefined function changes to the `time1.C` file. However, if you did, the file would not match the tutorial. So perform the following steps:

1. Enter the following command:

```
cvd> func Time::printTime
```

2. Select the following from the menu bar:

```
Fix+Continue  
->Save As
```

A *file_name* dialog box opens.

3. The dialog box enables you to save your file changes back to the original source files or save them to a different file. However, since you do not want to save your changes, press the `Cancel` button on the bottom of the dialog box.

Note: You usually want to wait until you are finished with `Fix+Continue` before you save your changes. In addition to the method described above, you can also save your changes by selecting the following:

```
Fix+Continue  
->Save All Files...
```

8.3.3 Setting Breakpoints in Redefined Code

To see how the Debugger works with traps in redefined code, this section shows you how to set breakpoints, run the Debugger, and view the results (Figure 40, page 120).

1. Reset to the beginning of program execution by entering the following at the `cvd>` prompt:

```
cvd> kill  
cvd> run
```

2. Bring up the `time1.C` source file by entering the following at the `cvd>` prompt:

```
cvd> func Time::printTime
```

3. Select the following from the menu bar:

```
Fix+Continue  
->Edit
```

Note: You can also use the `Alt-Ctrl-e` accelerator to perform this task.

4. Enter the following after line 23.4 in the source pane of the Main View window:

```
cout << " AM"<< endl;
```

5. Select the following from the menu bar:

```
Fix+Continue  
->Parse and Load
```

6. Bring up the `time1.C` source file again by entering the following command at the `cvd>` prompt:

```
cvd> func Time::printTime
```

7. Set a breakpoint at line 23.6 by entering the following message at the `cvd>` prompt:

```
cvd> stop at 23.6
```

The following message appears in the Command/Message pane:

```
[2] Stop at file /path/name/time1.C line 23.6
```

8. Click on the Continue button in the Main View window.

The following figure shows how the breakpoint has been reached in the redefined code:

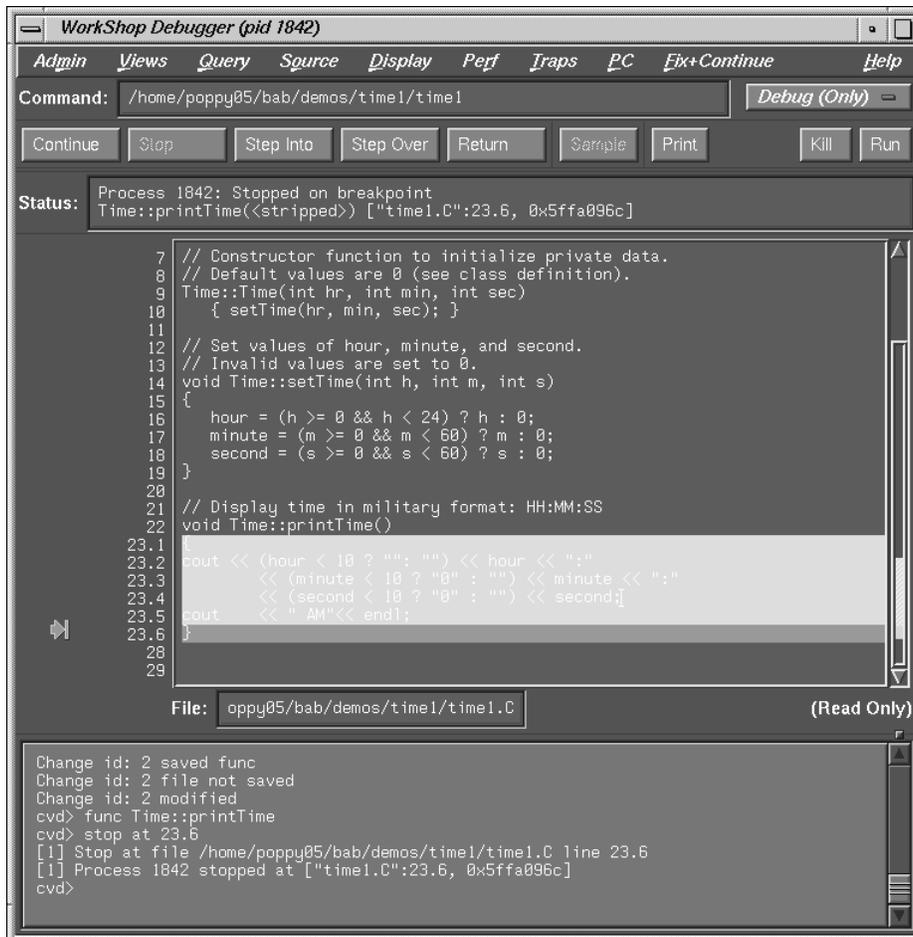


Figure 40. Stopping after Breakpoints in Redefined Code

9. Select the following from the menu bar to see the results of continuing to the breakpoint:

Views
->Call Stack

10. Select the following from the menu bar to view the locations of the breakpoints:

```
Views
  ->Trap Manager
```

11. Remove the breakpoint by clicking on the source annotation column to the right of line 23.6.

8.3.4 Viewing Status

To view status, select the following from the menu bar:

```
Fix+Continue
  ->View
    ->Status Window
```

The Fix+Continue Status window opens.

8.3.5 Comparing Original and Redefined Code

You can use Fix+Continue to compare original code with and modified code. This section shows you several ways to view your changes.

8.3.5.1 Switching between Compiled and Redefined Code

Follow these steps to see how the redefined code affects your executable:

1. Click the Run button to view your redefined code. If you are following procedures in order in this section, you should see the following display:

```
8:20:50 AM
```

2. Enter the following at the `cvd>` prompt:

```
cvd> func Time::printTime
```

3. Select the following from the menu bar:

```
Fix+Continue
  ->Parse and Load
```

4. Select the following from the menu bar to disable your changes:

```
Fix+Continue  
->Edit<-->Compiled
```

5. Click the Continue button to see the printing of Time as in the original executable. The following displays:

```
08:20:50
```

6. Re-enter the following at the `cvd>` prompt:

```
cvd> func Time::printTime
```

7. Select the following from the menu bar to re-enable your changes:

```
Fix+Continue  
->Edit<-->Compiled
```

8. Click on the RUN button to see the changed printout of Time. The following displays:

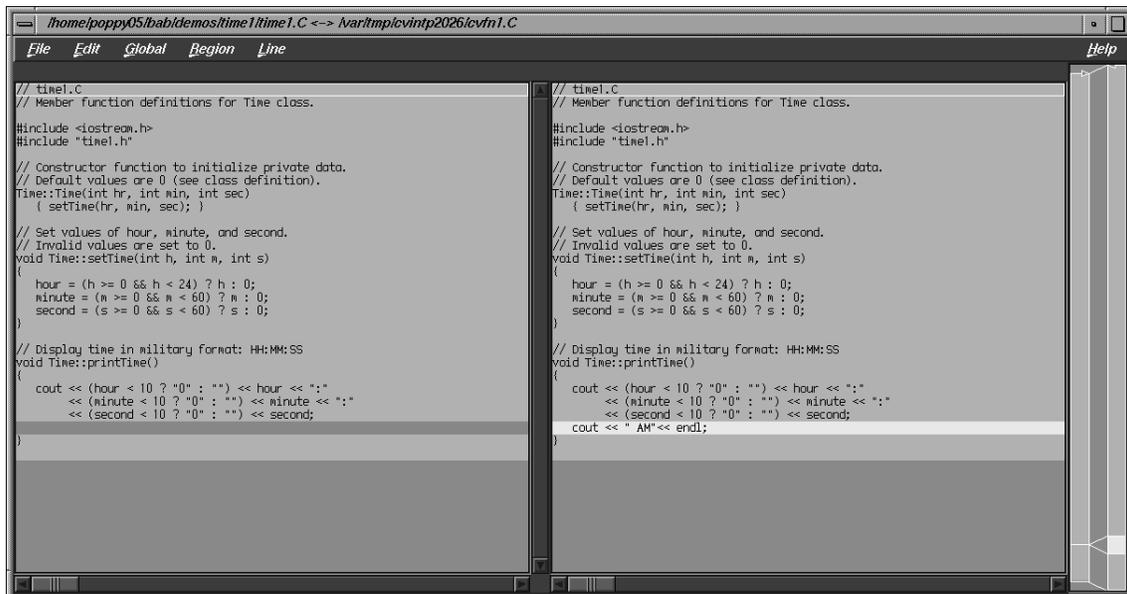
```
08:20:50 AM
```

8.3.5.2 Comparing Function Definitions

1. Place the cursor in the `time1.C` function.
2. Select the following from the menu bar:

```
Fix+Continue  
->Show Difference  
->For Function
```

A window opens to display an `xdiff` comparison of the files as follows:



```

// time1.C
// Member function definitions for Time class.
#include <iostream.h>
#include "time1.h"

// Constructor function to initialize private data.
// Default values are 0 (see class definition).
Time::Time(int hr, int min, int sec)
{ setTime(hr, min, sec); }

// Set values of hour, minute, and second.
// Invalid values are set to 0.
void Time::setTime(int h, int m, int s)
{
    hour = (h >= 0 && h < 24) ? h : 0;
    minute = (m >= 0 && m < 60) ? m : 0;
    second = (s >= 0 && s < 60) ? s : 0;
}

// Display time in military format: HH:MM:SS
void Time::printTime()
{
    cout << (hour < 10 ? "0" : "") << hour << ":" <<
        << (minute < 10 ? "0" : "") << minute << ":" <<
        << (second < 10 ? "0" : "") << second;
}

// time1.C
// Member function definitions for Time class.
#include <iostream.h>
#include "time1.h"

// Constructor function to initialize private data.
// Default values are 0 (see class definition).
Time::Time(int hr, int min, int sec)
{ setTime(hr, min, sec); }

// Set values of hour, minute, and second.
// Invalid values are set to 0.
void Time::setTime(int h, int m, int s)
{
    hour = (h >= 0 && h < 24) ? h : 0;
    minute = (m >= 0 && m < 60) ? m : 0;
    second = (s >= 0 && s < 60) ? s : 0;
}

// Display time in military format: HH:MM:SS
void Time::printTime()
{
    cout << (hour < 10 ? "0" : "") << hour << ":" <<
        << (minute < 10 ? "0" : "") << minute << ":" <<
        << (second < 10 ? "0" : "") << second;
    cout << " AM" << endl;
}

```

Figure 41. Comparing Compiled and Redefined Function Code

You can get the same result by entering the `show_diff #` command from the Debugger command line, where `#` is the redefined function ID.

If you do not like `xdiff`, you can change the comparison tool by selecting the following from the menu bar:

```

Fix+Continue
->Show Difference
    ->Set Diff Tool...

```

8.3.5.3 Comparing Source Code Files

If you have made several redefinitions to a file, you may need a side-by-side comparison of the entire file. To see how changes to the entire file look, select the following from the menu bar:

```
Fix+Continue  
  ->Show Difference  
    ->For File
```

This opens an `xdiff` window that displays the entire file rather than just the function.

You can get the same comparison results from the Debugger command line if you enter the following command:

```
show_diff -file time1.C
```

8.3.6 Ending the Session

Exit the Debugger by selecting the following from the menu bar:

```
Admin  
  ->Exit
```

Detecting Heap Corruption [9]

The heap is a portion of memory used to support dynamic memory allocation/deallocation via the `malloc` and `free` function. This chapter describes heap corruption detection and covers the following topics:

- Typical Heap Corruption Problems, Section 9.1, page 125
- Finding Heap Corruption Problems, Section 9.2, page 126
- Heap Corruption Detection Tutorial, Section 9.3, page 129

9.1 Typical Heap Corruption Problems

Due to the dynamic nature of allocating and deallocating memory, the heap is vulnerable to the following typical corruption problems:

<i>boundary overrun</i>	A program writes beyond the <code>malloc</code> region.
<i>boundary underrun</i>	A program writes in front of the <code>malloc</code> region.
<i>access to uninitialized memory</i>	A program attempts to read memory that has not yet been initialized.
<i>access to freed memory</i>	A program attempts to read or write to memory that has been freed.
<i>double frees</i>	A program frees some structure that it had already freed. In such a case, a subsequent reference can pick up a meaningless pointer, causing a segmentation violation.
<i>erroneous frees</i>	A program calls <code>free()</code> on addresses that were not returned by <code>malloc</code> , such as static, global, or automatic variables, or other invalid expressions. See the <code>malloc(3F)</code> man page for more information.

9.2 Finding Heap Corruption Errors

To find heap corruption problems, you must relink your executable with the `-lmalloc_ss` library instead of the standard `-lmalloc` library. By default, the `-lmalloc_ss` library catches the following errors:

- `malloc` call failing (returning `NULL`)
- `realloc` call failing (returning `NULL`)
- `realloc` call with an address outside the range of heap addresses returned by `malloc` or `memalign`
- `memalign` call with an improper alignment
- `free` call with an address that is improperly aligned
- `free` call with an address outside the range of heap addresses returned by `malloc` or `memalign`

If you also set the `MALLOC_FASTCHK` environment variable, you can catch these errors:

- `free` or `realloc` calls where the words prior to the user block have been corrupted
- `free` or `realloc` calls where the words following the user block have been corrupted
- `free` or `realloc` calls where the address is that of a block that has already been freed. This error may not always be detected if the area around the block is reallocated after it was first freed.

9.2.1 Compiling with the Malloc Library

You can compile your executable from scratch as follows:

```
% cc -g -o targetprogram targetprogram.c -lmalloc_ss
```

You can also relink it by using:

```
% ld -o targetprogram targetprogram.o -lmalloc_ss ...
```

An alternative to rebuilding your executable is to use the `_RLD_LIST` environment variable to link the `-lmalloc_ss` library. See the `rld(1)` man page.

9.2.2 Setting Environment Variables

After compiling, invoke the Debugger with your executable as the target. In *Execution View*, you can set environment variables to enable different levels of heap corruption detection from within the `malloc` library, as follows:

`MALLOC_CLEAR_FREE`

Clears data in any memory allocation freed by `free`. It requires that `MALLOC_FASTCHK` be set.

`MALLOC_CLEAR_FREE_PATTERN` *pattern*

Specifies a pattern to clear the data if `MALLOC_CLEAR_FREE` is enabled. The default pattern is `0xcafebabe` for the 32-bit version, and `0xcafebabe0xcafebabe` for the 64-bit versions. Only full words (double words for 64-bits) are cleared to the pattern.

`MALLOC_CLEAR_MALLOC`

Clears data in any memory allocation returned by `malloc`. It requires that `MALLOC_FASTCHK` be set.

`MALLOC_CLEAR_MALLOC_PATTERN` *pattern*

Specifies a pattern to clear the data if `MALLOC_CLEAR_MALLOC` is enabled. The default pattern is `0xfacebabe` for the 32-bit version, and `0xfacebabe0xfacebabe` for the 64-bit versions. Only full words (double words for 64-bits) are cleared to the pattern.

`MALLOC_FASTCHK`

Enables additional corruption checks when you call the routines in this library. Error detection is done by allocating a space larger than the requested area, and putting specific patterns in front of and behind the area returned to the caller. When `free` or `realloc` is called on a block, the patterns are checked, and if the area was overwritten, an error message is printed to `stderr` using an internal call to the routine `ssmalloc_error`. Under the Debugger, a trap may be set at exit from this routine to catch the program at the error.

`MALLOC_MAXMALLOC` *n*

Where *n* is an integer in any base, sets a maximum size for any `malloc` or `realloc` allocation. Any request exceeding that size is flagged as an error, and returns a NULL pointer.

`MALLOC_NO_REUSE`

Specifies that no area that has been freed can be reused. With this option enabled, no actual free calls are made and process space and swap requirements can grow quite large.

`MALLOC_TRACING`

Prints out all `malloc` events including address and size of the `malloc` or `free`. When running a trace in the course of a performance experiment, you need not set this variable because running the experiment automatically enables it. If the option is enabled when the program is run independently, and the `MALLOC_VERBOSE` environment variable is set to 2 or greater, trace events and program call stacks are written to `stderr`.

`MALLOC_VERBOSE`

Controls message output. If set to 1, minimal output displays; if set to 2, full output displays.

For further information, see the `malloc_ss(3)` man page.

9.2.3 Trapping Heap Errors Using the Malloc Library

If you are using the `-lmalloc_ss` library, you can use the Trap Manager to set a stop trap at the exit from the function `ssmalloc_error` that is called when an error is detected. Errors are detected only during calls to heap management routines, such as `malloc()` and `free()`. Some kinds of errors, such as overruns, are not detected until the block is freed or reallocated.

When you run the program, the program halts at the stop trap if a heap corruption error is detected. The error and the address are displayed in Execution View. You can also examine the Call Stack View at this point to get stack information. To find the next error, click the Continue button.

If you need more information to isolate the error, set a watchpoint trap to detect a write at the displayed address. Then rerun your program. Use `MALLOC_CLEAR_FREE` and `MALLOC_CLEAR_MALLOC` to catch problems from attempts to access uninitialized or freed memory.

Note: You can run programs linked with the `-lmalloc_ss` library outside of the Debugger. The trade-off is that you have to browse through the `stderr` messages and catch any errors through visual inspection.

9.3 Heap Corruption Detection Tutorial

This tutorial demonstrates how to detect corruption errors by using the `corrupt` program. The `corrupt` program has already been linked with the SpeedShop `malloc` library (`libmalloc_ss`). The `corrupt` program listing is as follows:

```
#include <string.h>
void main (int argc, char **argv)
{
    char *str;
    int **array, *bogus, value;

    /* Let us malloc 3 bytes */
    str = (char *) malloc(strlen(`bad`));

    /* The following statement writes 0 to the 4th byte */
    strcpy(str, `bad`);

    free (str);

    /* Let us malloc 100 bytes */
    str = (char *) malloc(100);
    array = (int **) str;

    /* Get an uninitialized value */
    bogus = array[0];

    free (str);
    /* The following is a double free */
    free (str);
    /* The following statement uses the uninitialized value as a pointer */
    value = *bogus;
}
```

To start the tutorial:

1. Enter the following:

```
% mkdir demos
% mkdir demos/mallocbug
% cd demos/mallocbug
% cp /usr/demos/WorkShop/mallocbug/* .
```

2. Invoke the Debugger by typing:

```
% cvd corrupt &
```

The Main View window displays with `corrupt` as the target executable.

3. Open the Execution View window (if it is minimized) and set the `_SSMALLOC_FASTCHK` and `_SSMALLOC_CLEAR_MALLOC` environment variables.

If you are using the C shell, type:

```
% setenv _SSMALLOC_FASTCHK
% setenv _SSMALLOC_CLEAR_MALLOC
```

If you are using the Korn or Bourne shell, type:

```
$ _SSMALLOC_FASTCHK=
$ _SSMALLOC_CLEAR_MALLOC=
$ export _SSMALLOC_FASTCHK _SSMALLOC_CLEAR_MALLOC
```

4. To trap any `malloc` corruption problems, you must enter the following at the `cvd>` prompt:

```
cvd> set $pendingtraps=true
cvd> stop exit ssmalloc_error
```

A stop trap is set at the exit from the `malloc` library `ssmalloc_error`.

5. Enter the following at the `cvd>` prompt:

```
cvd> run
```

The program executes. Observe Execution View as the program executes.

A heap corruption is detected and the process stops at one of the traps. The type of error and its address display in Execution View (see example in Figure 42, page 131.)

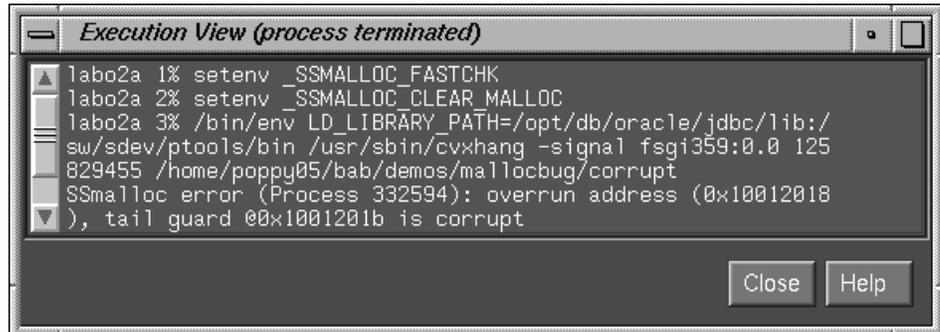


Figure 42. Heap Corruption Warning Shown in Execution View

6. Select the following from the Main View window menu bar:

Views

->Call Stack

Call Stack View opens displaying the call stack frame at the time of the error (see Figure 43).

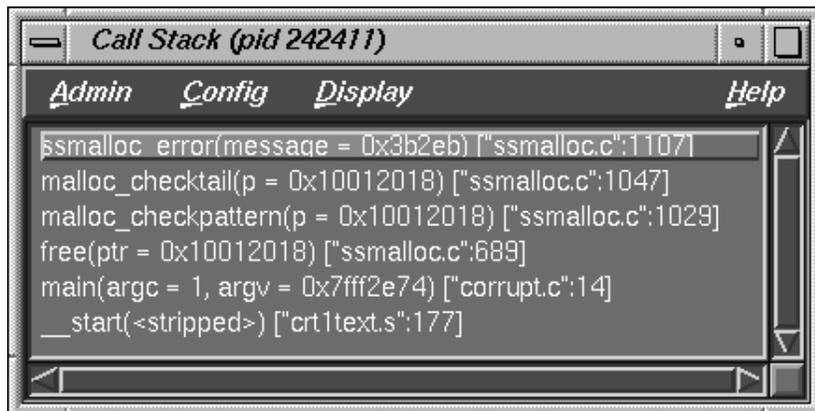


Figure 43. Call Stack at Boundary Overrun Warning

7. Click the Continue button in the Main View window's control panel. Watch the Execution View and Call Stack View windows.

The process continues from the stop at the boundary overrun warning until it hits the next trap where an erroneous free error occurs.

8. Click the Continue button again and watch the Execution View and Call Stack View windows.

This time the process stops at a bus error or segmentation violation. The PC stops at the following statement because `bogus` was set to an uninitialized value:

```
value=*bogus
```

9. Enter `p &bogus` on the Debugger command line at the bottom of the Main View window.

This gives us the address for the `bogus` variable and has been done in Figure 44, page 133. We need the bad address so that we can set a watchpoint to find out when it is written to. (This example has an address of `0x7ffaef4`; your address will be different.)

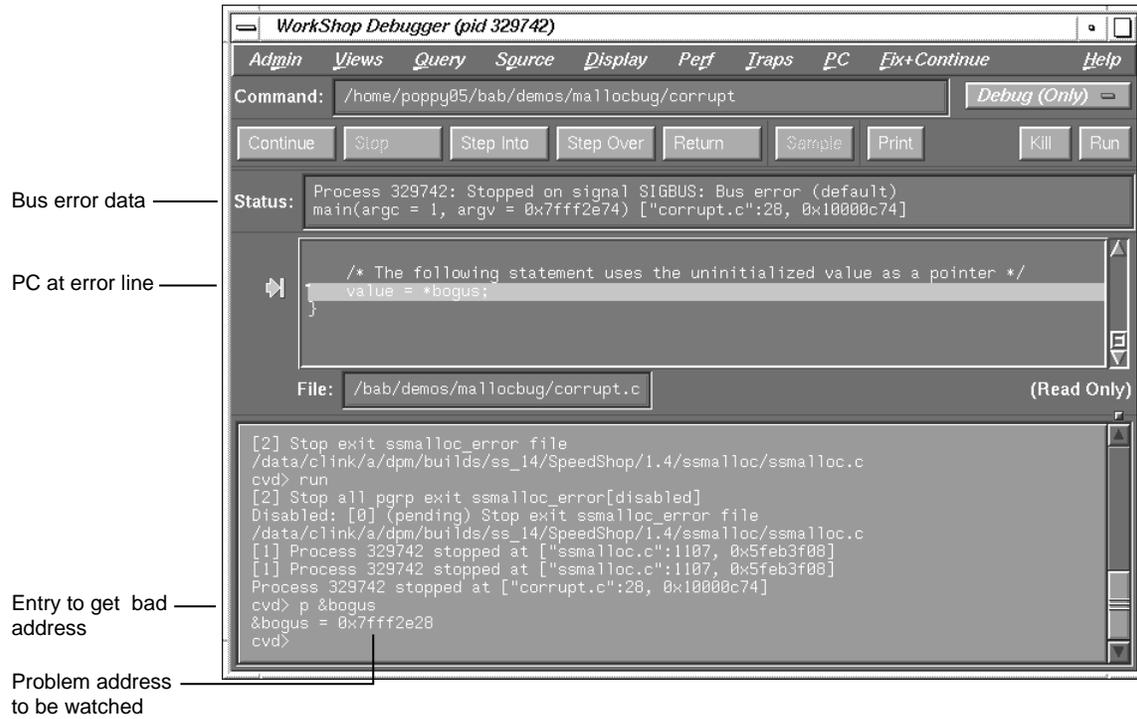


Figure 44. Main View at Bus Error

10. Deactivate the stop trap by clicking the toggle button next to the trap description in the Trap Manager window, and click the Kill button in the Main View window to kill the process.
11. Click on the Clear button in the Trap Manager window.
12. Type the following command in the Trap field. This includes the address you obtained from the Debugger command line (see Figure 44, page 133). This sets a watchpoint that is triggered if a write is attempted at that address.

Note: Use the address from your system, not the one shown here.

stop watch address 0x7fffaef4 for write

13. Click the Add button.
14. Click the Run button and observe the Main View window.

The process stops at the point where the bogus variable receives a bad value. Details of the error display in the Main View window's Status field.

Multiple Process Debugging [10]

The WorkShop Debugger supports the debugging of programs that use multiple processes spawned by `fork` or `sproc`, as well as threaded applications. The Debugger allows you to control a single process at a time or all members of a process group. You can attach child processes automatically. You can also specify that spawned processes inherit traps from the parent process. The Trap Manager provides special commands to facilitate debugging multiple processes by setting traps that apply to the entire process group.

WorkShop supports performance analysis and debugging of multiprocess applications, including processes spawned either with `fork` or `sproc` and threaded applications. You can perform process control operations on a single process or on all members of a process group. You can attach WorkShop automatically to child processes. You can also specify spawned processes to inherit traps. The Trap Manager provides special trap commands to facilitate debugging multiple processes simultaneously.

Note: The `Multiprocess View` window is for use by C, C++, and Fortran users. If you are debugging Ada code, you should use the `Task View` window available through the `View` menu of the `Main View` window (see Section A.3.1, page 216).

Currently, `Multiprocess View` handles the following multiple process situations:

- *True multiprocess program*, which refers to a tightly integrated system of `sproc'd` processes, generated by the MIPSpro Automatic Parallelization Option. For more information on parallel processing, see the `auto_p(5)` man page, or the *MIPSpro Automatic Parallelizer Programmer's Guide*.
- *Auto-fork application*, which is a process that spawns a child process and then runs in the background.
- *Fork application*, which is a process that spawns child processes and can interact with them.
- *Locally distributed application*, which is an application that involves two different executables running in different processes on the same host coordinated by a rendezvous mechanism.

This chapter discusses the details of multiprocess debugging in WorkShop and includes the following topics:

- Using the `Multiprocess View` window, Section 10.1, page 136
- Controlling Execution and Setting Traps in a Multiprocess Program, Section 10.2, page 139
- Debugging a Multiprocess Fortran Program, Section 10.3, page 146
- Debugging pthreaded applications, Section 10.4, page 152

10.1 Using the Multiprocess View Window

The `Multiprocess View` window is brought up by selecting the following from the menu bar of the Main View window:

```
Admin  
  ->Multiprocess View
```

This window can display individual processes or operate on a process group. By default, a *process group* includes the parent process and all descendants spawned by `sproc`. Processes spawned with `fork` during the session can be added to the process group automatically when they are created. For a program compiled with the MIPSpro Automatic Parallelization Option, a process group includes all threads generated by the option. Any process to which you have read/write access can also be added to the process group. All `sproc'd` processes must be in the same process group, since they share information.

Note: Any child process that performs an `exec` with `setuid` (set user ID) enabled will not become part of the process group.

Each process in the session can have a standard main view window session associated with it. However, all processes in a process group appear on a single `Multiprocess View` window.

When debugging multiprocess applications, you should disable the `SIGTERM` signal by selecting the following from the Main View window menu bar:

```
Views  
  ->Signal Panel
```

Note: Although multiprocessing debugging is possible with `SIGTERM` enabled, the multiprocess application may not terminate gracefully after execution is complete.

10.1.1 Starting a Multiprocess Session

The first step in debugging multiple processes is to invoke the Debugger with the parent process. Then select the following from the menu bar:

```
Admin
  -> Multiprocess View
```

The following figure shows a typical Multiprocess View window.

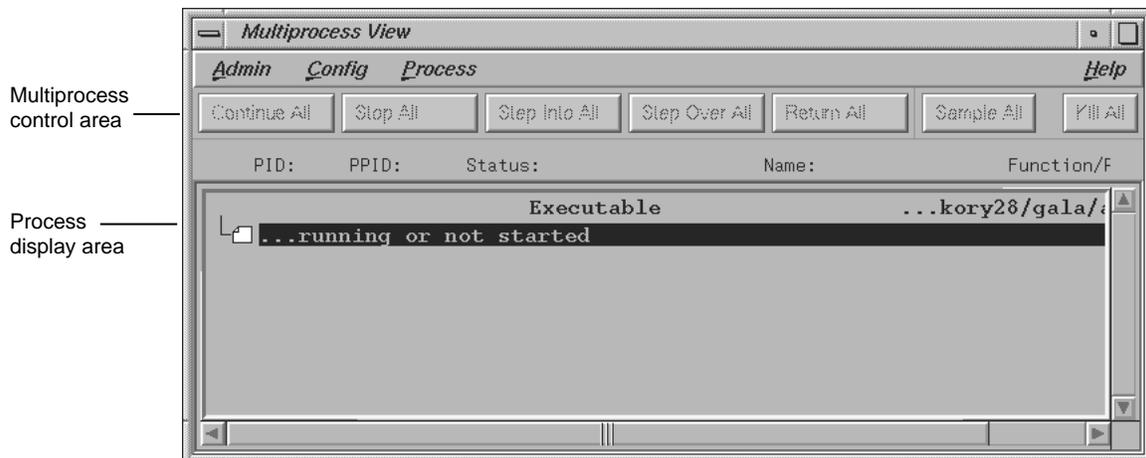


Figure 45. Multiprocess View

10.1.2 Viewing Process Status

The process display area of the Multiprocess View lists the status of all processes and threads in the process group. For definitions of the various status and states, see Section A.2.2.1.

To get more information about a process or thread displayed in the process display area, right-click on the process or thread entry. A Process menu pops up which is applicable to the selected entry. From this menu you can do the following:

- change entry focus
- create a new window
- focus attention to a user-entered thread

- show thread-specific details
- add or remove an entry

For complete details about the `Process` menu, see Section A.2.5, page 215.

10.1.3 Using Multiprocess View Control Buttons

The `Multiprocess View` window uses the same control buttons as the `Main View` window with the following exceptions:

- Buttons are applied to all processes as a group.
- There are no `Return`, `Print`, or `Run` buttons.

Control buttons in the `Multiprocess View` window have the same effect as clicking the corresponding button in the `Main View` window of each individual process. For definitions of the buttons, see Section A.2.2.2, page 208.

10.1.4 Multiprocess Traps

As discussed in Chapter 5, page 69, “Setting Traps”, the trap qualifiers `[all]` and `[pgrp]` are used in multiprocess analysis. The `[all]` entry stops or samples all processes when a trap fires. The `[pgrp]` entry sets the trap in all processes within the process group that contains the trap location. The qualifiers can be entered by default by using the `Stop All Default` and `Group Trap Default` selections, respectively, in the `Traps` menu of `Trap Manager`. The `Trap Manager` is brought up from the `Views` menu of the `Main View` window.

10.1.5 Viewing Pthreaded Applications

The `Multiprocess View` supports a hierarchical view of your pthreaded applications. Select the folder icons of your choosing to get more information about a process or thread.

Perform the following from within the `Multiprocess View` window to get additional information about a process or thread:

1. Double-click on a folder icon.
The process display expands to show its pthreads.
2. Double-click to select the pthread of your choosing.
The call stack for that thread displays.

10.1.6 Adding and Removing Processes

To add a process, select `Add . . .` from the `Process` menu. In the `Switch Dialog` dialog window, select one of the listed processes or enter a process ID in the `Process ID` field and click the `OK` button.

To remove a process, click on the process name in the `Multiprocess View` window and select `Remove` from the `Process` menu. Be aware that a process in a `sproc` process group cannot be removed. Likewise, you cannot remove a `pthread` from a `pthread` group.

10.1.7 Multiprocess Preferences

The `Preferences . . .` option in the `Config` menu brings up the `Multiprocess View Preferences` dialog. The preferences on this dialog lets you determine when a process is added to the group, specify process behavior, specify the number of call stack levels to display, and so forth.

For details about `Multiprocess View Preference` options, see Section A.2.2.4, page 209.

10.1.8 Bringing up Additional Main View Windows

To create a `Main View` window for a process, highlight that process in the `Multiprocess View` window. Then, select the following in the `Multiprocess View` window:

```
Process  
->Create new window
```

10.2 Debugging a Multiprocess C Program

This section uses a C program that generates numbers in the Fibonacci sequence to demonstrate the following tasks when using the debugger to debug multiprocess code:

- Stopping a child process on a `sproc`
- Using the buttons in the `Multiprocess View` window
- Setting traps in the parent process only
- Setting group traps

The `fibonacci` program uses `spawn` to split off a child process, which in turn uses `spawn` to split off a grandchild process. All three processes generate Fibonacci numbers until stopped. You can find the source for `fibonacci.c` in the `/usr/demos/WorkShop/mp` directory. A listing of the `fibonacci.c` source code follows:

```
#include <stdio.h>
#include <sys/types.h>
#include <sys/prctl.h>

int NumberToCompute = 100;
int fibonacci();
void run(),run1();

int fibonacci(int n)
{
    int f, f_minus_1, f_plus_1;
    int i;

    f = 1;
    f_minus_1 = 0;
    i = 0;

    for (; ; ) {
        if (i++ == n) return f;
        f_plus_1 = f + f_minus_1;
        f_minus_1 = f;
        f = f_plus_1;
    }
}

void run()
{
    int fibon;
    for (; ; ) {
        NumberToCompute = (NumberToCompute + 1) % 10;
        fibon = fibonacci(NumberToCompute);
        printf("%d'th fibonacci number is %d\n",
            NumberToCompute, fibon);
    }
}

void run1()
```

```
{
int grandChild;

    errno = 0;
    grandChild = sproc(run,PR_SADDR);

    if (grandChild == -1) {
        perror("SPROC GRANDCHILD");
    }
    else
        printf("grandchild is %d\n", grandChild);
    run();
}

void main ()
{
int second;

    second = sproc(run1,PR_SADDR);
    if (second == -1)
        perror("SPROC CHILD");
    else
        printf("child is %d\n", second);

    run();
    exit(0);
}
```

10.2.1 Launch the Debugger in Multiprocess View

Perform the following to start, compile the program, and run the Debugger:

1. Copy the program source from the demo directory as follows:

```
% cp /usr/demos/WorkShop/mp .
```

2. Compile `fibonacci.c` by entering the following command:

```
% cc -g fibonacci.c -o fibonacci
```

3. Invoke the Debugger on `fibonacci` as follows:

```
% cvd fibonacci &
```

4. Call up the `Multiprocess View` by selecting the following from the Main View menu bar:

```
Admin  
    ->Multiprocess View...
```

The next section uses the `fibonacci` program to illustrate some of the functionality of the `Multiprocess` window.

10.2.2 Using `Multiprocess View` to Control Execution

To examine each process as it is created, you must set preferences so that each child process created stops immediately after being created. The following steps show how this can be done:

1. Select the following from the menu bar in the `Multiprocess View` window:

```
Config  
    ->Preferences...
```

2. Toggle off `Resume child after attach on sproc` in the `Multiprocess View Preferences` window.
3. Toggle off `Copy traps to sproc'd processes` so you can experiment with setting traps later.
4. Click on the `OK` button to accept the changes.
5. Click on the `Run` button in the Main View window to execute the `fibonacci` program.

Watch the `Multiprocess View` window, You will see the main process appear and spawn a child process, which stops as soon as it appears. This is because you turned off the `Resume child after attach on sproc` option. Notice also that the Main View window switched to the stopped child process.

6. Click on the `Stop All` button in the `Multiprocess View` window.

The control buttons on the `Multiprocess View` window may be used to control all processes simultaneously, or the control buttons on any Main View window may be used to control that individual process separately.

7. Click on the first line (that is, the main process) in the process pane of the `Multiprocess View` window to highlight this line.

8. Select the following from the menu bar of the Multiprocess View window:

```
Process
  ->Create new window...
```

A new Main View window displays with a debug session for the main process.

Note: You may get a warning that `.../write.s` is missing. This refers to assembly code and can be ignored. The new Main View window will not have source in its source pane.

9. Select the following from the menu bar of the Main View window you just created to create a Call Stack View window:

```
Views
  ->Call Stack View
```

10. Double-click on the line in the Call Stack View window that contains `run ()`. This brings up the `fibonacci.c` source for the main process in the Main View window.

11. Select the following from within the Call Stack View window to close it:

```
Admin
  ->Close
```

12. Click on Continue All in the Multiprocess View window. The first child, created in Step 5, now spawns a grandchild process that stops in `_nsproc`, as shown here:

Note: A Main View window switches to the new stopped process.

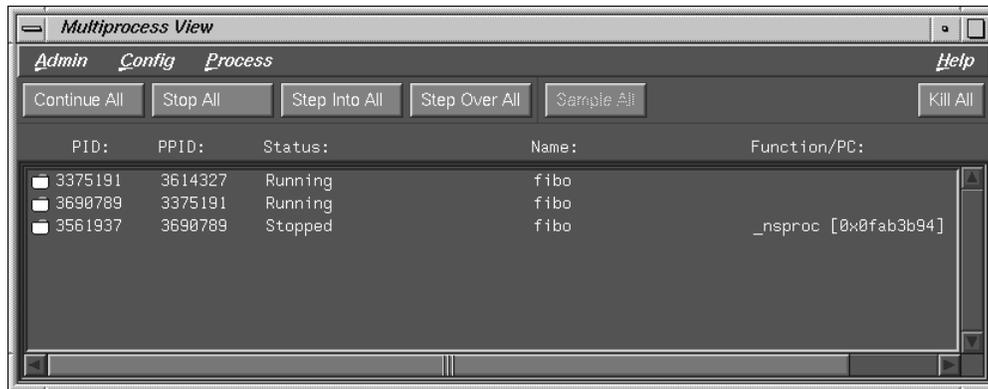


Figure 46. Examining Process State Using Multiprocess View

13. Click on Stop All in the Multiprocess View window.
14. Repeat steps 7 through 11 to bring up a Main View window for the parent process.

10.2.3 Using the Trap Manager to Control Trap Inheritance

Note: The instructions in this section assume that you have just run the tutorial in Section 10.2.2, page 142.

This section shows you how to use the Trap Manager to set traps that affect one or all processes in the fibc process group. For complete information on using the Trap Manager, refer to Chapter 5, page 69.

1. Select the following from the Main View window for the parent process:

```
Views
->Trap Manager
```

Note: Traps are specific to the processes in the Main View window in which they are set.

2. Select the following (from the same Main View window) to turn on line numbering in the source pane, if not already showing:

```
Display
->Show Line Numbers
```

- Click to the left of line 32, to set a breakpoint/stop trap for the parent process. Line 32 reads as follows:

```
32 Number to Compute = (NumberToCompute + 1) % 10
```

Line 32 highlights in red to indicate that a breakpoint has been set. A corresponding trap command appears in the Trap text box in the Trap Manager window; and the trap is added to the list on the Active Traps list of the same window. Remember, this trap will affect only the parent process.

- Click on the Continue All button in the Multiprocess View window.

Note: The parent process has stopped, but the other processes are probably still running.

- Insert the word `pgrp` (that is, “process group”) after the word `stop` in the Trap field of the Trap Manager window.

The trap should now read `Stop pgrp at . . .`. As the command suggests, `pgrp` affects the whole process group.

- Click on the Modify button.

The trap now affects two child processes. Watch the Multiprocess View window to see the running processes in the process group stop at the trap on line 32.

- Select the following from the Trap Manager window:

```
Traps
->Group Trap Default
```

Note: Now, any additional traps that you set using the Trap Manager will affect the entire process group. Any previously set traps will not be affected.

- Select the text of line 23, found in the source pane of the Main View window associated with the parent process. This line reads as follows:

```
23 f_minus_1 = f;
```

- Select the following from the menu bar of the Trap Manager window:

```
Traps
->At Source Line
```

The trap you have just set includes the modifier `pgrp`.

10. Select the following from any Main View window to close your session and end this tutorial:

```
Admin  
->Exit
```

10.3 Debugging a Multiprocess Fortran Program

The section of this chapter presents a few standard techniques to assist you in debugging a parallel program. This section shows you how to debug the sample program.

See also Chapter 2, page 11 for important related information.

10.3.1 General Fortran Debugging Hints

Debugging a multiprocessed program is more involved than debugging a single-processor program. Therefore, you should debug a single-processor version of your program first and try to isolate the problem to a single parallel DO loop.

Once you have isolated the problem to a specific DO loop, change the order of iterations in a single-processor version. If the loop can be multiprocessed, then the iterations can execute in any order and produce the same answer. If it cannot be multiprocessed, you will see that changing the order in which the loops execute will cause the single-processor version to produce wrong answers. If wrong answers are produced, you can use standard single-process debugging techniques to find the problem. (See Chapter 2, page 11 for important related information.)

If this technique fails, you must debug the multiprocessed version. To do this, compile your code with the `-g` and `-FLIST:=ON` flags. The `-FLIST:=ON` flags save the file containing the multiprocessed DO loop Fortran code in a file called `total.w2f.f` and a file `tital.rii` and an `rii_files` directory.

10.3.2 Fortran Multiprocess Debugging Session

This section shows you how to debug a small segment of multiprocessed code.

The source code for this tutorial, `total.f`, can be found in the directory `/usr/demos/WorkShop/mp`.

A listing of this code is as follows:

```

program driver
  implicit none
  integer iold(100,10), inew(100,10),i,j
  double precision aggregate(100, 10),result
  common /work/ aggregate
  result=0.
  call total(100, 10, iold, inew)
  do 20 j=1,10
    do 10 i=1,100
      result=result+aggregate(i,j)
10    continue
20  continue
  write(6,*)' result=',result
  stop
  end

  subroutine total(n, m, iold, inew)
  implicit none
  integer n, m
  integer iold(n,m), inew(n,m)
  double precision aggregate(100, 100)
  common /work/ aggregate
  integer i, j, num, ii, jj
  double precision tmp

  C$DOACROSS LOCAL(i,ii,j,jj,num)
  do j = 2, m-1
    do i = 2, n-1
      num = 1
      if (iold(i,j) .eq. 0) then
        inew(i,j) = 1
      else
        num = iold(i-1,j) +iold(i,j-1) + iold(i-1,j-1) +
&      iold(i+1,j) + iold(i,j+1) + iold(i+1,j+1)
        if (num .ge. 2) then
          inew(i,j) = iold(i,j) + 1
        else
          inew(i,j) = max(iold(i,j)-1, 0)
        end if
      end if
      ii = i/10 + 1
    
```

```
        jj = j/10 + 1
        aggregate(ii,jj) = aggregate(ii,jj) + inew(i,j)
    end do
end do
end
```

In the program, the local variables are properly declared. The `inew` always appears with `j` as its second index, so it can be a share variable when multiprocessing the `j` loop. The `iold`, `m`, and `n` are only read (not written), so they are safe. The problem is with `aggregate`. The person analyzing this code deduces that, because `j` is always different in each iteration, `j/10` will also be different. Unfortunately, since `j/10` uses integer division, it often gives the same results for different values of `j`.

While this is a fairly simple error, it is not easy to see. When run on a single processor, the program always gets the right answer. Sometimes it gets the right answer when multiprocessing. The error occurs only when different processes attempt to load from and/or store into the same location in the `aggregate` array at exactly the same time.

10.3.2.1 Debugging Procedure

Perform the following to debug this code:

1. Create a new directory for this exercise:

```
% mkdir demos/mp
```

2. `cd` to the new directory and copy the following program source into it:

```
% cp /usr/demos/WorkShop/mp .
```

3. Edit the `total.f` file in a shell editor, such as `vi`:

```
% vi total.f
```

4. Reverse the order of the iterations for demonstration purposes.

Replace

```
do j = 2, m-1
```

with

```
do j = m-1, 2, -1
```

This still produces the right answer with one process running, but the wrong answer when running with multiple processes. The local variables look right, there are no equivalence statements, and `inew` uses only simple indexing. The likely item to check is `aggregate`. Your next step is to look at `aggregate` with the Debugger.

5. Compile the program with `-g` option as follows:

```
% f77 -g -mp total.f -o total
```

6. If your debugging session is not running on a multiprocessor machine, you can force the creation of two threads, for example purposes, by setting an environment variable. If you use the C shell, type:

```
% setenv MP_SET_NUMTHREADS 2
```

If you use the Korn or Bourne shell, type:

```
$ MP_SET_NUMTHREADS=2  
$ export MP_SET_NUMTHREADS
```

7. Enter the following to start the Debugger:

```
% cvd total &
```

The Main View window displays.

8. Select the following from the Main View menu bar to show the line numbers:

```
Display  
->Show Line Numbers
```

9. Select the following from the Main View menu bar:

```
Source  
->Go To Line...
```

And enter **44**.

Line 44 is as follows:

```
aggregate(ii,jj) = aggregate(ii,jj) + inew(i,j)
```

10. You will now set a stop trap at this line, so you can see what each thread is doing with `aggregate`, `ii`, and `jj`. You want this trap to affect all threads of the process group. One way to do this is to turn on trap inheritance in the Multiprocess View Preferences dialog. To open this dialog,

select the following from the Main View menu bar to open the Multiprocess View window:

```
->Admin
  ->Multiprocess View
```

Then, select the following from within the Multiprocess View window:

```
Config
  ->Preferences
```

Another way is to use the Trap Manager to specify group traps, as follows.

- a. Select the following from the Main View window menu bar to open the Trap Manager:

```
Views
  ->Trap Manager
```

- b. Select the following from the Trap Manager window:

```
Traps
  ->Group Trap Default
```

11. Click-drag to select line 44 in the Main View window.
12. Open the Trap Manager window from the Main View window menu bar as follows:

```
Views
  ->Trap Manager
```

Then select the following from the the Trap Manager window:

```
Traps
  ->At Source Line
```

This sets a stop trap that reads as follows in the `cvd>` pane of the Main View window:

```
Stop pgrp at file /usr/demos/WorkShop/mp/total.f line 44
```

13. Select the following from the menu bar in the Main View window to monitor status of the two processes:

```
Admin  
  ->Multiprocess View
```

You are now ready to run the program.

14. Click the Run button in the Main View window.

As you watch the `Multiprocess View`, you see the two processes appear, run, and stop in the function `_mpdo_total_1`. It is unclear, however, if the Main View window is now relative to the master process, or that it has switched to the slave process.

15. Right-click on the name of the slave process in the `Multiprocess View` window and select the following:

```
Process  
  ->Create a new window
```

A new window is displayed that launches a debug session for the process. Now, both master and slave processes should display in respective Main View windows.

16. Invoke the Variable Browser as follows from the Menu Bar of each process:

```
Views  
  ->Variable Browser
```

17. Look at `ii` and `jj` in the following figure.

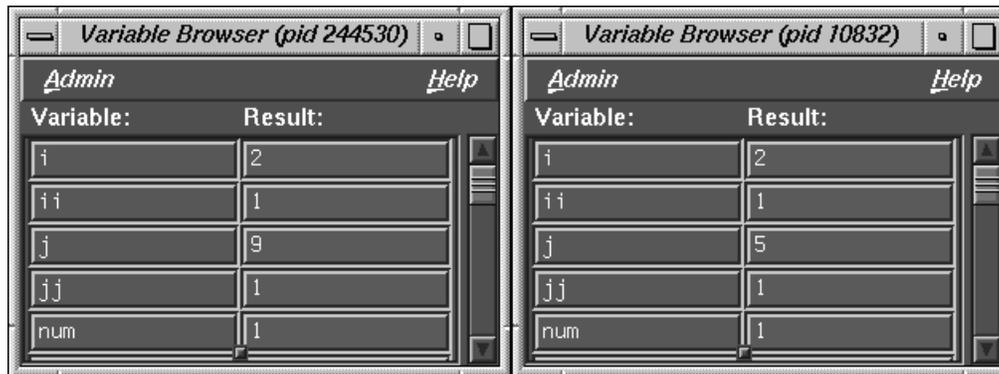


Figure 47. Comparing Variable Values from Two Processes

They have the same values in each process; therefore, both processes may attempt to write to the same member of the array aggregate at the same time. So aggregate should not be declared as a share variable. You have found the bug in your parallel Fortran program.

10.4 Debugging a Pthreaded Program

`cvd` supports the debugging of IRIX 6.4 and 6.5 pthreads. You can view pthread creation and execution through the `Multiprocess View` window. Through this window you can:

- View a hierarchal display of a threaded application
- View a process/pthread relationship
- Expand individual call stacks

C, C++, and Fortran users should use the `Multiprocess View` window when debugging pthreads. Ada users should use the `Task View` window.

The next sections give hints on debugging pthreaded programs, list differences between 6.4 and 6.5 threads, and illustrate how to debug a program that uses IRIX 6.5 pthreads.

10.4.1 User-Level Continue of Single 6.5 POSIX Pthread

The ability to “continue” or “free run” a single POSIX pthread under IRIX 6.5 is available at the user level with WorkShop release 2.8. However, use of this new debugging feature can, in certain specific circumstances, lead to anomalous and possibly confusing behavior. Such behavior occurs when the single thread that is continued or free run encounters either a “blocking” or “scheduling” situation in the operating system or the pthreads library. When such situations arise, the operating system (or, in some cases, the pthreads library) must take action to dispose of the single continued or free run thread and, possibly, newly created threads. In the course of this action the debugging user will see things occur, with both the single continued or free run thread as well as all other threads, that will be perplexing because complex thread scheduling algorithms are invoked by both the operating system and the pthreads library to recover from the original blocking or scheduling incident. This note attempts to provide information which might lessen frustration on the part of the debugging user in the face of blocking or scheduling incidents. Debugging true POSIX pthreads is tricky business, and users of this new feature, allowing a continue or free run of a single 6.5 POSIX pthread, will gain even more appreciation of this fact.

This feature has been used internally for some time by the WorkShop debugger. The continue or free run of a single 6.5 pthread is used each time a user requests a single thread step-over of a function. The single thread is allowed to free run through the function which is being stepped over. Thus, if any blocking or scheduling situations occur in the course of this stepping over and associated free run of a single thread, then anomalous behavior can, and does, occur. For example, when attempting to step over a function call which eventually calls the routine `pthread_create`, the user will see the step over succeed or not succeed in non-deterministic fashion due to complex scheduling algorithms in the pthreads library. This is explained in greater detail below in Section 10.4.1.1, page 153.

As another example, when a single thread step over is requested of a routine which does a `printf` somewhere down the call tree, the blocking operating system call `writew` is eventually invoked and this can, and does, lead to anomalous behavior in the eyes of the debugging user. This is explained in detail in Section 10.4.1.2, page 155.

10.4.1.1 Scheduling Anomalies

Scheduling anomalies *may* occur when the single 6.5 POSIX pthread which is being continued or free run creates a new pthread via call to routine `pthread_create`. At the time of the call to `pthread_create` the OS kernel and the pthreads library get into a complex algorithm in deciding how to give

life to the new child pthread. An actual OS kernel thread (OS kernel threads are not available at the user level — they differ from the user level pthread) must be either created anew or found elsewhere to support the user's new child pthread.

First, let's assume the OS kernel thread is to be found elsewhere, depending on a vast number of things (for example, number of cpu's, environment variables, and so on). The OS kernel *may* (this is very non-deterministic) decide to just put the child pthread on a ready queue, in need of an OS kernel thread. Thus the child does nothing right away. Meanwhile if the parent pthread, via a call to routine `pthread_cond_wait`, monitors the child pthread's struggle for life, it (the parent) gets parked on a mutex because the child obviously has not been born yet (remember, it's on the ready queue). The parent pthread's OS kernel thread becomes available, which causes the OS scheduler to check for work for this newly freed OS kernel thread. Behold, it finds the child sitting in the ready queue and assigns the parent's OS kernel thread to the new child pthread. The child then runs to completion and releases its (parent's old) OS kernel thread. The parent, checking for the child's new life via `pthread_cond_wait`, now recaptures its OS kernel thread and things work fine in the eyes of the debugging user.

Now, let's assume the OS kernel thread required by the new child pthread must be created anew. The child is not placed on the 'ready queue'. Again, this is a non-deterministic decision which depends on a large number of variables (that is number or cpu's, and so on). The OS kernel creates a new OS kernel thread for the child pthread and "engages" it (the child) to that new OS kernel thread. However, "marriage" of the OS kernel thread and the new child pthread cannot occur until the new OS kernel thread actually runs. This never occurs because, in allowing the single parent 6.5 pthread to continue or run free, it was requested that only one user pthread be run - the parent. If the parent, however, is using `pthread_cond_wait` to monitor the new life of its child, then it (the parent) is parked on a mutex waiting for the child to run. This is a Catch-22. The parent awaits the child but the child cannot run because only one pthread, the parent, has been requested to run. The debugger user, meanwhile, is not amused. The debugger displays "running" as the overall status and this is because no events of interest are occurring. Everybody is waiting on everybody else. Things are not working.

10.4.1.2 Blocking Anomalies

Blocking anomalies occur when the single 6.5 POSIX thread which is being continued or free run encounters a blocking condition in the course of its running. The notion of blocking has three distinct flavors:

- Blocking syscalls in the OS kernel (see Section A.10, page 326 for a list). When one of these kernel syscalls is blocked by another thread's usage, the OS kernel decides what the next move is regarding the OS kernel thread attached to the user pthread making the call. Control could just transfer to another application, to disk I/O, or whatever. These syscalls are all I/O-related. The OS kernel thread says, in effect, "I'm waiting for something to do -> I'm blocked", and it is immediately available for reassignment. The best example of a blocking kernel syscall is `writev`, which is used by the common library routine `printf`. Thus, if a single 6.5 POSIX thread is continued or free run and it eventually calls `printf`, the debugging user may see the anomalous behavior of another pthread started running because the continued or free run 6.5 pthread was blocked in `writev`. The `printf` pthread's OS kernel thread is reassigned and another pthread suddenly takes off.
- Various lock blocking in the pthread library, such as Mutex (mutual exchange lock). This occurs in user space (libc, user code, and so on). The pthread library sense that a pthread is going to block due to another pthread's usage. Control transfers to the routine `usync_control` which eventually calls a blocking kernel syscall (see preceding item). Again, the OS kernel decides the fate of the associated OS kernel thread. Unexpected things could start running.
- Other lock blocking in the pthread library, whereby the pthread library senses that a user pthread is going to block but does not go off to `usync_control`. Instead it goes to the `pthread_scheduler` in the pthread library for the disposition of the associated OS kernel thread. The `pthread_scheduler` then reassigns the associated OS kernel thread to another user pthread and unexpected things could start running.

10.4.1.3 How to Continue a Single POSIX 6.5 Pthread

To continue (or free run) a single POSIX 6.5 pthread, simply click on the `Continue` button in the Main View window. Note that this is different from the function of the `Continue` button in the `Multiprocess View` window, which continues all threads.

10.4.2 Other Pthread Debugging Hints

Observe the following guidelines when debugging pthreaded programs:

- C++ exception handling works per process not per thread.
- Using the `step over` function on a `pthread_exit` may produce unexpected results.

10.4.3 Differences between 6.4 and 6.5 Pthreads

The following differences exist between the IRIX 6.4 and the IRIX 6.5 implementation of pthreads:

- In the IRIX 6.5 implementation, many user threads are serviced by a few kernel micro-threads. This can result in the following problems:
 - After stopping at a breakpoint, a pthreaded program may abort with a `Trace/BPT trap error` when you try to use the `continue` function.
 - Data race problems may occur as kernel micro-threads switch from one user pthread to another.

You can alleviate these problems with the following command:

```
ccall pthread_setconcurrency n
```

Where *n* is the number of kernel micro-threads made available to service user threads. In this manner, you can increase the number of kernel micro-threads available. However, you must not specify a value for *n* that is greater than the number of physical CPUs on your system.

- The `Multiprocess View` window shows different sets of threads / processes depending on which version of IRIX you are running. (This is because in IRIX releases before IRIX 6.5, threads are handled as sprocs, not as threads.)

10.4.4 Pthread Debugging Session

Pthread debugging is highly variable not only from environment to environment from from IRIX release to IRIX release. Because of this, it is not possible to provide a representative pthread debugging tutorial that can be used by all users.

See Section 10.4.1, page 153, Theory of 6.5 Pthreads, for an in depth description of current pthread implementation in IRIX.

X/Motif Analyzer [11]

This chapter provides an introduction to the X/Motif Analyzer as well as a tutorial to demonstrate most of the Analyzer functions. The chapter contains the following sections:

Motif is a library of routines that enable you to create user interfaces in an X-environment. The Motif libraries handle most of the low-level event handling tasks common to any GUI system. This way, you can create sophisticated interfaces without having to contend with all of the complexity of X.

The X/Motif Analyzer helps you debug code that calls Motif library routines. The X/Motif Analyzer is integrated with the Debugger so you can issue X/Motif Analyzer commands graphically. To access X/Motif analyzer subwindow select the following from the Main View window menu bar:

```
Views  
->X/Motif Analyzer
```

11.1 Introduction to the X/Motif Analyzer

The Analyzer contains X/Motif objects (for example, widgets and X graphics contexts) that can be difficult or impossible to inspect through ordinary debugging procedures. It also allows you to set widget-level breakpoints and collect X-event history information in the same manner as using `xscope`. See the `xscope(1)` man page for more information.

11.1.1 Examiners Overview

When the X/Motif Analyzer first displays, it is set to examine widgets. At this point, the window may be blank, or it may display a widget found in the call stack of a stopped process.

At the bottom of the X/Motif Analyzer window is a tab panel that shows the current set of examiners. In addition to this tab, the `Widget Examiner`, `Breakpoints`, `Trace`, and `Tree Examiners` tabs are at the bottom of the window. These four tabs are always present. Other examiners are available from the `Examine` menu of the X/Motif Analyzer window.

Some examiners cannot be manually selected. They appear only when appropriate to the call stack context. For example, the `Callback Examiner` appears only when a process is stopped somewhere in a widget callback.

11.1.2 Examiners and Selections

If you select text in one examiner and then choose another examiner by using the `Examine` menu, the new examiner is brought up and the text is used as an expression for it. If you selected text that is an inappropriate object for the new examiner, an error is generated.

Alternatively, you can select text, pull down the `Examine` menu, and choose `Selection`. Here, the X/Motif Analyzer attempts to select an appropriate examiner for the text. If the type of text is unknown, the following message displays:

```
Couldn't examine selection in more detail
```

Otherwise, the appropriate examiner is chosen and the text is evaluated.

You can also accomplish this by triple-clicking on a line of text. If the type of text is unknown, nothing happens. Otherwise, the appropriate examiner is chosen and the text is evaluated.

11.1.3 Inspecting Data

X/Motif applications consist of collections of objects (that is, Motif widgets) and make extensive use of X resources such as windows, graphics context, and so on. The construction model of an X window system hinders you from inspecting the internal structures of widgets and X resources because you are presented with ID values. The X/Motif Analyzer lets you to see the data structures behind the ID values.

11.1.4 Inspecting the Control Flow

Traditional debuggers enable you to set breakpoints only in source lines or functions. With the X/Motif Analyzer, you can set breakpoints for specific widgets or widget classes, for specific control flow constructs like callbacks or event handlers, and for specific X events or requests.

11.1.5 Tracing the Execution

The X/Motif Analyzer can trace `Xlib`-level server events and client requests, `Xt`-level event dispatching information, widget life cycle, and widget status information.

11.2 Restrictions and Limitations

The X/Motif Analyzer has the following restrictions and limitations:

- The `Breakpoints Examiner` is active only after you have stopped a process and if you have changed `$LD_LIBRARY_PATH`. See Section 11.3.2, page 163 for more information regarding the correct `$LD_LIBRARY_PATH`.
- Sometimes, gadget names may be unavailable and are displayed as `<object>`. You can minimize this condition by first loading the widget tree.
- `editres` requests (such as, widget selection and widget tree) work only if the process is running or if the process is stopped outside of a system call. This can be annoying when the process is stopped in `select()`, waiting for an X server event.
- The process state and appearance of the Main View window flickers while the X/Motif Analyzer tries to complete an `editres` request when the process is stopped.
- `editres` requests may be unreliable if the process is stopped.

11.3 X/Motif Analyzer Tutorial

This section illustrates several features of the X/Motif Analyzer. The demo files in the `/usr/demos/WorkShop/bounce` directory are used to demonstrate the debugging of a running X-application. These files contain the complete C++ source code for the `bounce` program.

This section includes the following subsections:

- Setting up a Sample Session, Section 11.3.1, page 162.
- Launching the X/Motif Analyzer, Section 11.3.2, page 163.
- Navigating the Widget Structure, Section 11.3.3, page 163.
- Examining Widgets, Section 11.3.4, page 166.

- Setting Callback Breakpoints, Section 11.3.5, page 168.
- Using Additional Features of the Analyzer, Section 11.3.6, page 172.
- Ending the Session, Section 11.3.7, page 175.

11.3.1 Setting up the Sample Session

Perform the following to prepare for this session:

1. Enter the following commands:

```
% mkdir demos/bounce
% cd demos/bounce
% cp /usr/demos/WorkShop/bounce .
% make clean
% make bounce
% cvd bounce &
```

The Debugger is launched, from which you can use the X/Motif Analyzer. Upon invocation, you see the Execution View icon and the Main View window.

2. Double-click the Execution View icon to open the window. Then, tile your windows so you can clearly see all windows.
3. Click on the Run button in the Main View window to run the bounce program.

The Execution View window will update with the command that cvd is executing.

4. Click Run in the Bounce window.

Note: You will see no action until you have finished the next steps.

5. Select the following from the menu bar of the Bounce program window:

```
Actors
  ->Add Red Ball
```

6. Click on the Kill button in the Main View window to terminate the process that has been running.
7. The Execution View shows the program output.

11.3.2 Launching the X/Motif Analyzer

Once the bounce fileset is built and the debugger is active, you need to launch the X/Motif Analyzer as follows:

1. Select the following from the menu bar of the Main View window:

```
Views  
  ->X/Motif Analyzer
```

2. Click OK when asked if you want to change your `$LD_LIBRARY_PATH` environment variables to include `.../usr/lib/WorkShop/Motif`.

This includes instrumented versions of the Silicon Graphics libraries `Xlib`, `Xt`, and `Xm`. These libraries provide debugging symbols and special support for the X/Motif Analyzer.

Note: There are no instrumented MIPS/ABI versions of the libraries.

You are now ready to begin the sample session.

Note: Follow the steps in this tutorial precisely as written.

11.3.3 Navigating the Widget Structure

When the X/Motif Analyzer is launched, it brings up the X/Motif Analyzer window with an empty `Widget Examiner` tab panel. The tab panels also show the `Breakpoints`, `Trace`, and `Tree Examiner` tab panels (see Figure 48, page 164).

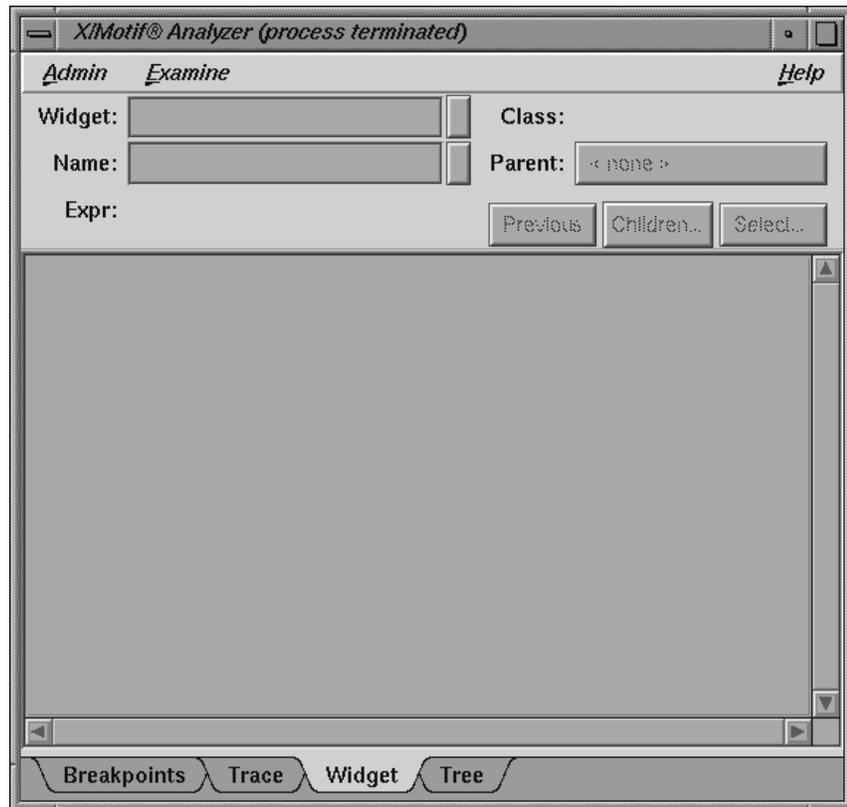


Figure 48. First View of the X/Motif Analyzer (Widget Examiner)

1. Widen the X/Motif Analyzer window shown in Figure 48, page 164. This will make it easier to understand what you will be asked to do in this tutorial.
2. Click Run in the Main View window to re-run the bounce program.
The instrumented versions of the Motif libraries will now be used.
3. When the Bounce window appears, re-size it to make it taller.
4. Click Run in the Bounce window.
Note: You will not see any action until you perform the next steps.
5. Position the X/Motif Analyzer and Bounce windows side-by-side.

- Click on the **Select** button in the X/Motif Analyzer window.

This brings up an information dialog and changes the cursor to a plus sign (+).

Note: Do not click on the **OK** button in this dialog.

- Select the **Step** widget by clicking on the **Step** button in the **Bounce** window with the (+) cursor, as described in the `cvmotif` information dialog.

The **Widget Examiner** displays the **Step** widget structure.

- Click the **Tree** tab in the X/Motif window.

The **Tree Examiner** panel displays the widget hierarchy of the target object (see Figure 49, page 165).

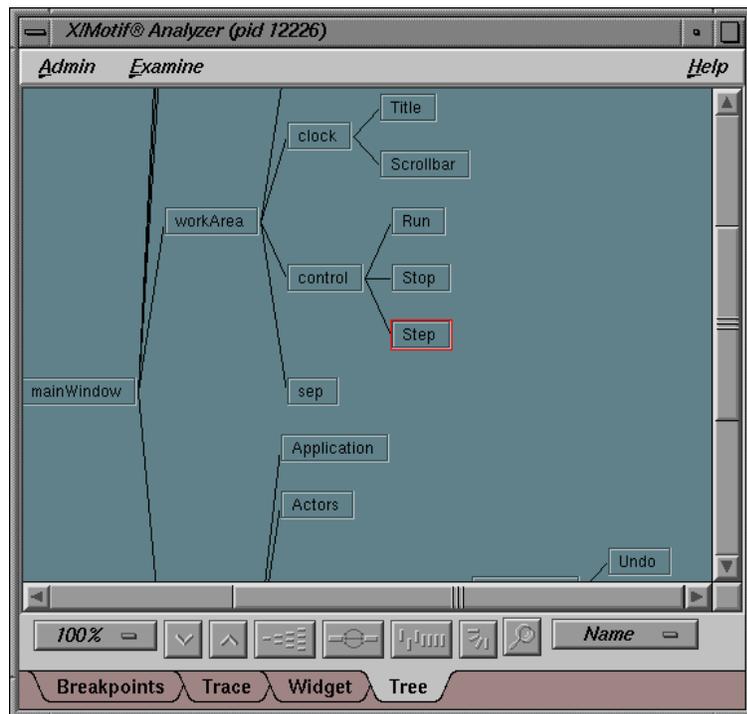


Figure 49. Widget Hierarchy Displayed by the Tree Examiner

9. Double-click the Run node in the tree. (Run is in the upper-right area of the window).

This brings up the `Widget Examiner` that displays the Run widget structure. Notice that the `parent` text area displays the name of the current widget's parent, which is `control`.

10. Click on the word `control` displayed on the `Parent` button at the top of the `Widget Examiner` in the `X/Motif Analyzer` window.

This switches the view to the Run widget's parent, the `control` object, as shown in the `Name` field. And, the `Widget Examiner` displays the `Control` widget structure.

You can now navigate through the widget hierarchy using either the `Widget Examiner` or the `Tree Examiner`.

11.3.4 Examining Widgets

Note: This section follows from the previous section.

1. In the `Widget Examiner`, click on the `Children...` button to see the menu, and select `Run` from that menu.

The Run widget structure displays in the examiner.

2. Select the following from the `Bounce` window:

```
Actors
  ->Add Red Ball
```

You should see a bouncing red ball.

3. Enter `stop in Clock::timeout` at the `cvd>` prompt in the `Main View` window.

After you press `ENTER`, the red ball will stop bouncing.

4. Select `Continue` in the `Main View` window a few times to observe the behavior of `bounce` with this breakpoint added.
5. Select the `Breakpoints` tab in the `X/Motif Analyzer` window. This calls up the `Breakpoints Examiner` which allows you to set widget-level breakpoints.

- In the Callback Name text field, enter `activateCallback`, then click on the Add button to add a breakpoint for the `activateCallback` object of the RUN button widget. The result is displayed in Figure 50, page 167.

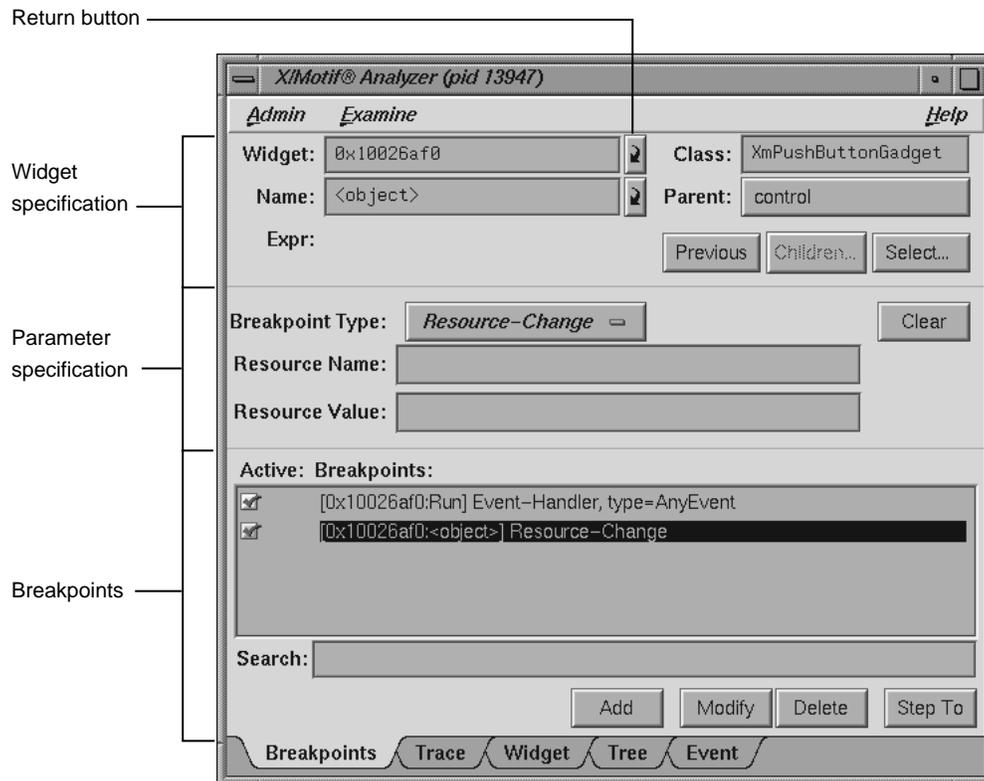


Figure 50. Adding a Breakpoint for a Widget

- Click on the red breakpoint down arrow in the Annotation Column of the Main View window to remove the `Clock::timeout` breakpoint.

Note: If you click on the line, but not the down arrow, the breakpoint will be deleted; but the source pane will still display the arrow.

- Click on the Continue button in the Main View window.
- Click on the Stop button in the Bounce window.

10. Click on the Run button in the Bounce window. The process stops in the Run button's registered `activateCallback`. This is the routine that was passed to `XtAddCallback` routine. Notice that the Callback tab (for the Callback Examiner) is added to the tab list.

11.3.5 Setting Callback Breakpoints

1. Click on the Breakpoints list item (the Active box will be checked) to highlight the breakpoint in the X/Motif Analyzer Breakpoint Examiner window.
2. Delete the widget address in the Widget text field by backspacing over the text.
3. Click on the Modify button to change the `activateCallback` breakpoint to apply to all push-button gadgets `XmPushButtonGadget` (see in the Class text field) rather than just the Run button. See Figure 51.

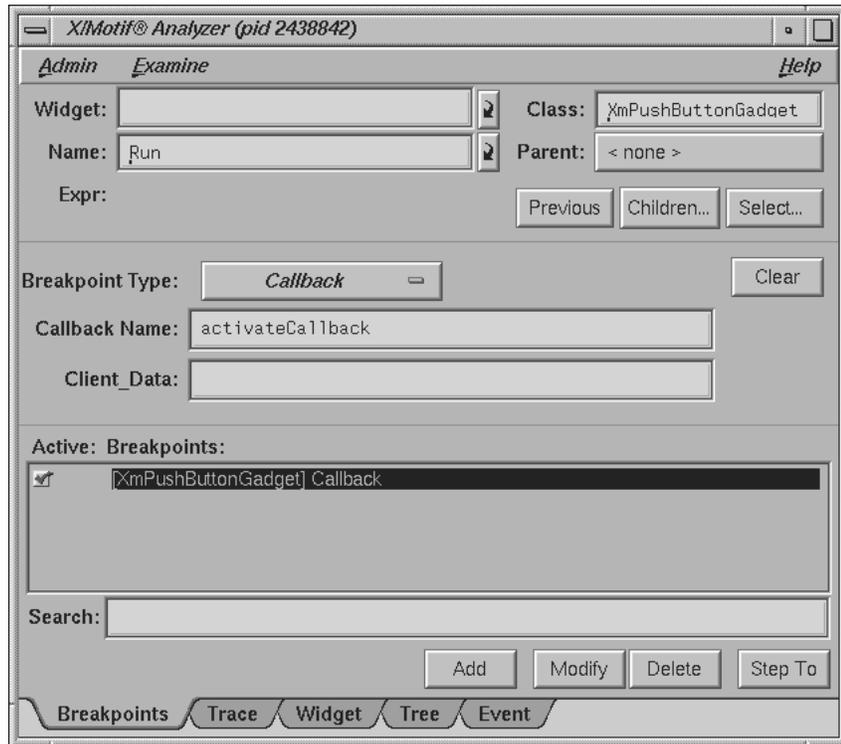


Figure 51. Setting Breakpoints for a Widget Class

4. Click Continue in the Main View window.
5. Click Stop in the Bounce window.

The process now stops in the Stop button's activateCallback routine.

6. Click the Callback tab in the X/Motif Analyzer window to go to the Callback Examiner. This examiner displays the callback context and the appropriate call_data structure (see Figure 52).

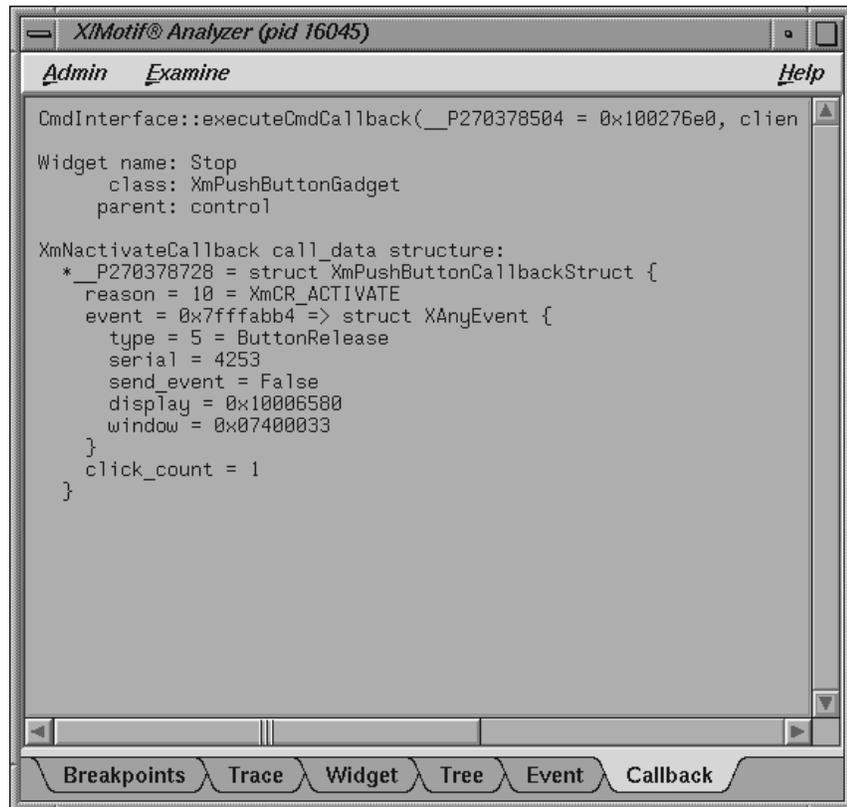


Figure 52. Callback Context Displayed by the Callback Examiner

7. Double-click the window value in the callback structure, fourth line from bottom.
8. Select the following in the Callback Examiner

```

Examine
->Window
    
```

The X/Motif Analyzer displays the window attributes for that window (the window of the `Stop` button). Notice that the `Window` tab (for the Window Examiner) is added to the tab list. See Figure 53.

Note: You can also accomplish the same action by triple-clicking the window value in the callback structure of the Callback Examiner (Step 7). In general, triple-clicking on an address brings you to that object in the appropriate examiner.

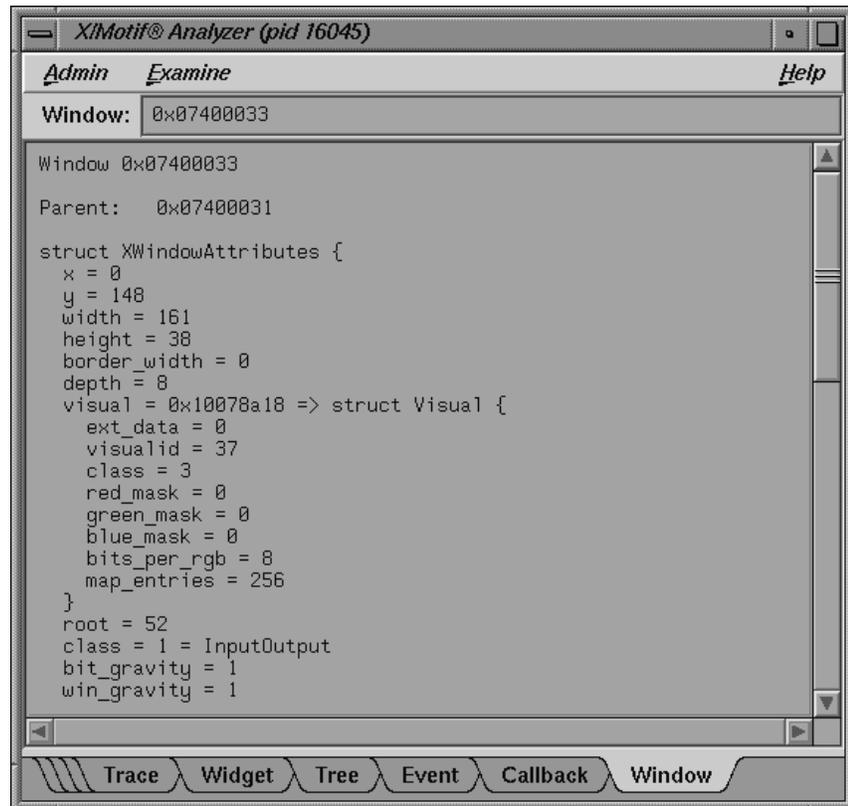


Figure 53. Window Attributes Displayed by the Window Examiner

11.3.6 Using Additional Features of the Analyzer

The following steps demonstrate additional Analyzer features.

1. In the X/Motif Analyzer window, click the `Widget` tab.
2. Double-click the `widget_class` value (on the fourth line) to highlight it.
3. Pull down the following:

```
Examine  
  ->Widget Class
```

The X/Motif Analyzer window displays the class record for the `XmPushButtonGadget` routine. Notice that the `Widget Class` tab (for the widget class examiner) is added to the tab list.

(Again, the same action can be accomplished by triple-clicking the `widget_class` value in the Widget Examiner.)

4. (This time we used the triple-click shortcut to select the correct examiner.) Triple-click the superclass value on the third line. The X/Motif Analyzer window displays the class record for `XmLabelGadget`, the superclass of `XmPushButtonGadget`.
5. Triple-click the superclass value on the right side of the third line. The X/Motif Analyzer window displays the class record for `XmGadget`, the superclass of `XmLabelGadget`.
6. Select the `Widget` tab to change to the Widget Examiner.
7. Triple-click the parent value on the fifth line. The X/Motif Analyzer window displays the control widget, the parent of `Run`. This action produces the same results as selecting the `control` text in the `Parent` text box.
8. Right-click on the tab overflow area (the area where the tabs overlap, to the far left of the tab list) as labeled in Figure 54, and select the `Breakpoints` tab.

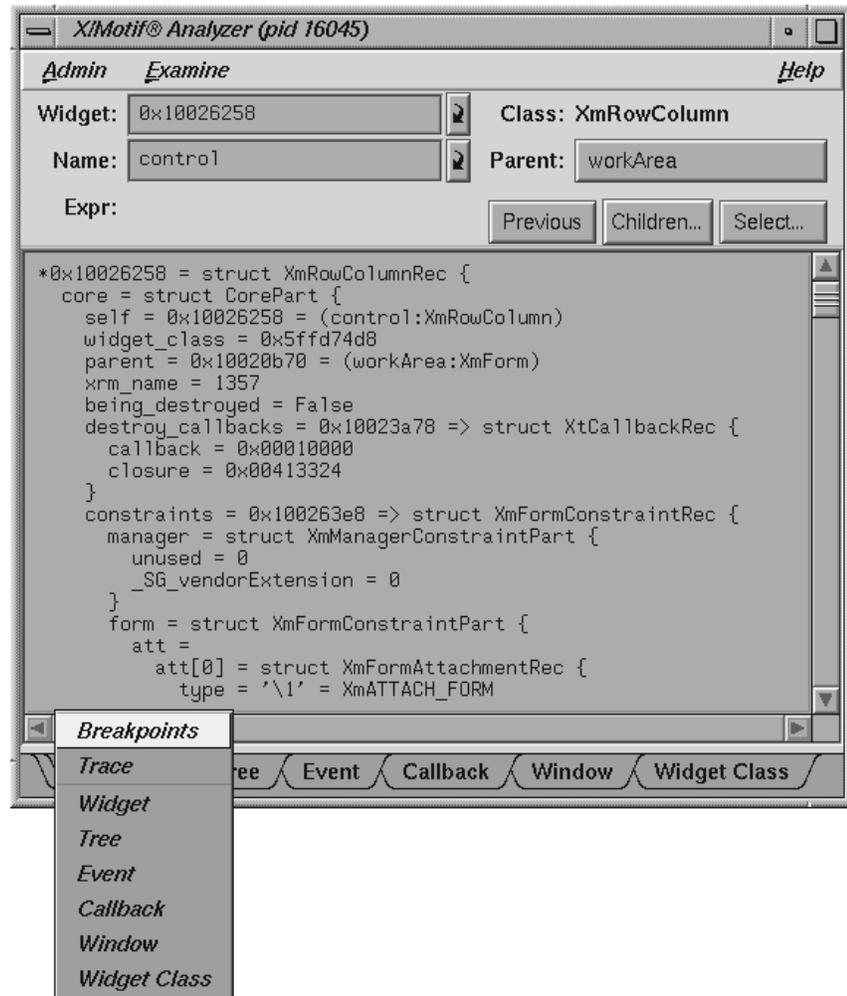


Figure 54. Selecting the Breakpoints Tab from the Overflow Area

9. Click on the word `Callback` in the Breakpoint Type text field of the Breakpoints Examiner window to bring up a submenu, and select `Resource-Change`.
10. In the Class text field, enter: **Any**.
11. In the Resource Name text field, enter: **sensitive**.

12. Click Add. This adds a breakpoint. The Active: Breakpoints: list should now include the following text:

```
[Any] Resource-Change,name=sensitive
```

13. Click Continue. The status will update to Stopped in the SetValues routine, since the breakpoint set in the previous step was reached.
14. Select the following in the Main View window:

```
Views
  ->Call Stack
```

Notice the call to XtSetValues on the second line (see Figure 55).

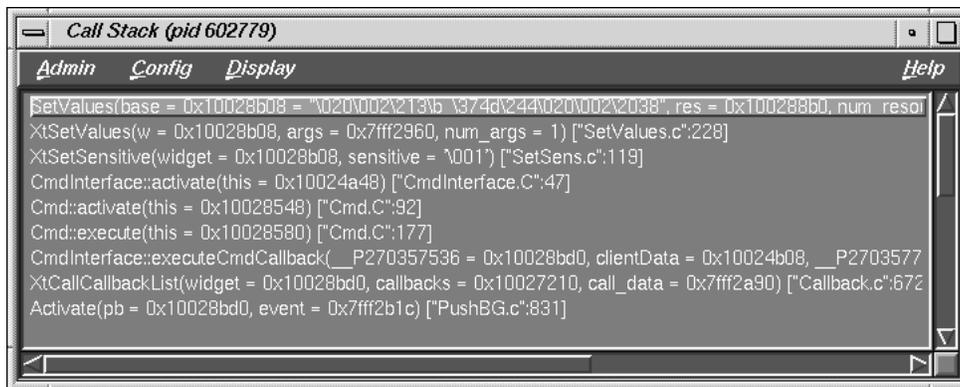


Figure 55. Breakpoint Results Displayed by the Call Stack View

15. In the Call Stack View, double-click the Cmdinterface::activate line (just below XtSetSensitive). This is where the sensitive resource was changed.
16. In the Widget Examiner window, double-click the widget address in the Widget text field, press backspace, enter `_w`, and press ENTER. The X/Motif Analyzer displays the Run widget, which is the widget currently being changed.
17. Click Continue in the Main View window. The status will update to Stopped in the SetValues routine, since the breakpoint set in the previous step was reached again.

18. In the Call Stack View, double-click on the `Cmdinterface::activate` line (just below `XtSetSensitive`).
19. Perform the following sub-steps:
 - a. Double-click in `Widget` text field of the `Widget Examiner`.
 - b. Press backspace.
 - c. Enter `_w`.
 - d. Press ENTER.

The X/Motif Analyzer window displays the `Step` widget, which is the widget currently being changed.

11.3.7 Ending the Session

Select the following to close the X/Motif Analyzer:

```
Admin  
->Close
```

Select the following to exit the Debugger (from the Main View window):

```
Admin  
->Exit
```

Note: If you exit the Debugger first, you exit the X/Motif Analyzer as well.

For more information on the X/Motif Analyzer, see Section A.4, page 222.

Customizing the Debugger [12]

This section shows you how you may customize the WorkShop Debugger specifically to your environment needs.

12.1 Customizing the Debugger with Scripts

If there are Debugger commands or combinations of Debugger commands that you use frequently, you may find it convenient to create a script composed of Debugger commands. Debugger scripts are ASCII files containing one Debugger command and its arguments per line. A Debugger script can in turn call other Debugger scripts. There are three general methods for running scripts:

- Enter the `source` command and the filename at the Debugger command line. This is useful for scripts that you need only occasionally.
- Include the script in a startup file. This is useful for scripts that you want implemented every time you use the Debugger.
- Define a button in the graphical interface to run the script. Use this method for scripts you use frequently but apply only at specific times during a debugging session.

12.1.1 Using a Startup File

A startup file lets you preload your favorite buttons and aliases in a file that runs when the Debugger is invoked. It also is useful if you have traps that you set the same way each time. The suggested name for the startup file is `.cvdrc`. However, you can select a different name as long as you specify its path in the `CVDINIT` environment variable. The Debugger uses the following criteria when looking for a startup file:

- Checks the `CVDINIT` environment variable.
- Check for a `.cvdrc` file in the current directory.
- Checks for a `.cvdrc` file in the user's home directory.

12.1.2 Implementing User-Defined Buttons

You can implement buttons by providing a special Debugger startup file or by creating them on the fly within a debugging session. Buttons appear in the order of implementation in a row at the bottom of the control panel area. Currently, you can define only one row of custom buttons. Figure 56 is a typical example of the Main View window with user-defined buttons. The definitions for the user-defined buttons display in the Debugger command line area.

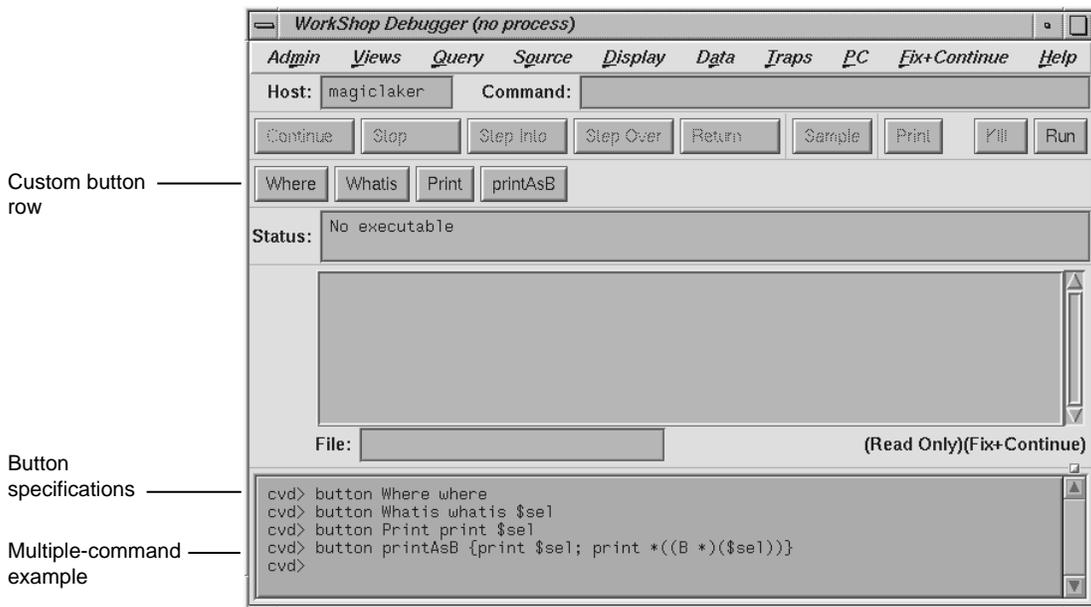


Figure 56. User-Defined Button Example

The syntax for creating a button is as follows:

```
button label command [$sel]
```

The syntax for creating a multiple-command button is as follows:

```
button label {command1 [$sel]; command2 [$sel]; ...}
```

The button command accepts the following options:

<i>label</i>	Specifies the button name. Button labels should be kept short since there is only room for a single row of buttons. There can be no spaces in a label.
<i>command</i>	Specifies one of the Debugger commands, which are entered at the command line at the bottom of the Main View window. See Section A.9, page 312.
<i>\$sel</i>	Provides a substitute for the current cursor selection and should be appropriate as an argument to the selected command.
<i>commandn...</i>	Specifies Debugger commands to be applied in order. Commands must be separated by semicolons (;) and enclosed by braces ({}). The multiple-command button is a powerful feature; it lets you write a short script to be executed when you click the button.

The following command displays a list of all currently defined buttons:

```
button
```

The following command deletes the button corresponding to the label:

```
unbutton label
```

You might use this command if you needed room to create a new button. The effect of `unbutton` is temporary so that subsequently running the startup file reactivates the button.

The following command displays the definition of the specified button, if it exists. If the button does not exist, an error message is displayed:

```
button label
```

12.1.3 Changing X Window System Resources

While there are many X Window System resources that you can change, we recommend that you avoid modifying these resources if at all possible. In some cases, there may be no way within WorkShop to make the desired change. If you must modify resources, the following X Window System resources for the Debugger and its views may be useful:

`*AllowPendingTraps:`

If set to true, enables support for pending traps.

The default value is false.

`*autoStringFormat`

If set to true, sets default format for `*char` results as strings in Expression View, the Variable Browser, and the Structure Browser.

The default format is the hexadecimal address.

`cvmain*sourceView*nameText.columns`

Sets the length of the File field in the Main View window.

The default value is 30 characters.

`Cvmain*disableLicenseWarnings` and `*disableLicenseWarnings`

Disables the license warning message that displays when you start the Debugger and the other tools.

`*editorCommand`

If you prefer to view source code in a text editor rather than in Source View, lets you specify a text editor.

The default value is the vi editor.

`*expressionView*maxNumOfExpr`

Lets you set the maximum number of expressions that can be read from a file by Expression View.

The default value is 25.

`*UseOldExprEval` and `*useOldExprEval`

When set to true, allows you to use the older, dbx-like, expression evaluator rather than the default C++ expression evaluator introduced in WorkShop version 2.6.4. Using the older evaluator may result in faster evaluation of some expressions.

The default value is false.

`*varBrowser*maxSymSize`

Lets you set the maximum number of variables that can be displayed by the Variable Browser.

The default value is 25.

The following resources apply to Source View:

`*svComponent*lineNumbersVisible`

Displays source line numbers by default.

`*sourceView*nameText.columns`

Sets the length of the File field in Source View.

The default value is 30 characters.

`*sourceView*textEdit.scrollHorizontal`

If set to true, displays a horizontal scroll bar in Source View.

`*tabWidth`

Sets the number of spaces for tabs in Source View.

The following resource applies to Build View:

`*buildCommand`

Is used to determine which program to used with `make(1)`, `smake(1)`, `clearmake(1)`, and so forth.).

The default value is `make(1)`.

`*runBuild`

Specifies whether `cvmake(1)` begins its build immediately upon being launched.

The default value is true.

To change these resources, you need to set the desired value in your `.xdefaults` file, and run the `xrdb(1)` command. Then, restart your application so that the resource gets picked up.

12.2 DUMPCORE Environment Variable

The `DUMPCORE` environment variable allows the Debugger to dump a core file in the event that there is a debugger execution problem during the debug session. To enable core files, enter either of the following before you startup your debug session:

- For the C shell, enter:

```
% setenv DUMPCORE
```

- For the Korn shell or Bourne shell, enter:

```
$ DUMPCORE=;export DUMPCORE
```

Debugger Reference [A]

This chapter describes the function of each window, menu, and display in the Debugger's graphical user interface and describes the commands available on the Debugger command line (see Section A.9, page 312).

This chapter contains the following sections:

- Main View Window, Section A.1, page 183
- Basic Windows, Section A.2, page 206
- Ada-Specific Windows, Section A.3, page 216
- X/Motif Analyzer Windows, Section A.4, page 222
- Trap Management Windows, Section A.5, page 251
- Data Examination Windows, Section A.6, page 256
- Machine-Level Debugging Windows, Section A.7, page 290
- Fix+Continue Windows, Section A.8, page 301
- Debugger Command Line, Section A.9, page 312

A.1 Main View Window

The major areas of the Main View window are shown in Figure 57, page 184.

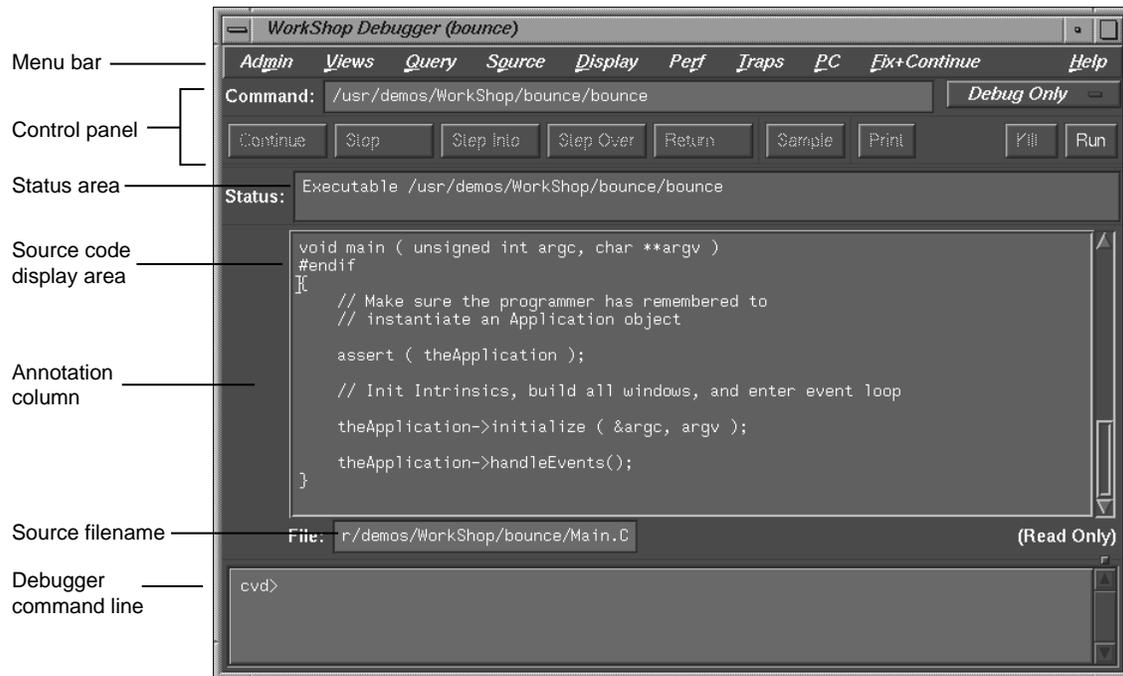


Figure 57. Major Areas of the Main View Window

The Main View window contains a menu bar, from which you can perform a number of actions and launch windows. The menu bar contains the following menus, which are discussed in detail in later pages:

- Admin, Section A.1.1, page 190
- Views, Section A.1.2, page 192
- Query, Section A.1.3, page 194
- Source, Section A.1.4, page 194
- Display, Section A.1.5, page 197
- Perf, Section A.1.6, page 197
- Traps, Section A.1.7, page 199
- PC, Section A.1.8, page 201

- `Fix+Continue`, Section A.1.9, page 201
- `Help`, Section A.1.10, page 205

The Main View window also contains several input fields, a source code display area, and buttons that trigger commonly used actions.

In the Main View window, actions can be applied to a single process or to all processes. To force an action to apply to all processes, click the right mouse button over the `Run`, `Continue`, or `Stop` buttons to set the *All Processes* mode. For more information about multiprocessing and pthreadd programs, refer to Chapter 10, page 135.

The Main View window contains the following items:

Command text field	Displays full pathname of the executable file that you are currently debugging, including any run-time arguments (if any).
Debug option menu	Allows you to toggle among various modes. The following mode choices are available: <ul style="list-style-type: none"> • <code>Debug (Only)</code> runs the Debugger in Debug mode with no performance tools enabled. • <code>Performance</code> mode causes performance data to be gathered and instrumented code to be generated for performance analysis while using the Debugger. • <code>Purify</code> mode activates the Purify memory corruption analysis tool. The code displayed in the Main View window, Source View window, and so forth will be code generated by Purify. (This option appears only if Purify is installed on your system. Purify is not a Silicon Graphics product nor is it part of the WorkShop package. It is a product of Rational Software and is neither available from nor supported by SGI.)
<code>Cont [All]</code>	Continues execution of the current process or all processes in the program. When you click on the <code>Continue</code> button, the program runs either to a breakpoint exception or to termination. This button is active only after the running process(es)

has stopped. If the program has not been run or has been killed, the Continue button is grayed out. If the target program has not yet started executing, use the Run button to start execution.

Stop [All]

Stops execution of the current running process(es). This button is valid only when a process(es) is running; otherwise the button is grayed out. Traps can also be set to stop the program at a specific location or on a particular condition.

Step Into button

Executes the code involving a single source line of the current process. If a function is encountered in the source line, or is the source line is a subroutine call, the process will step into a function or subroutine call and stop at the first executable line in the function or subroutine. The Step Over button can be used to continue to the next source line in the current file. If a trap is encountered while executing the step into command, the process is stopped where the trap was fired. The Step Into button is active only after the running process(es) has stopped; otherwise the button is grayed out.

When you click-right on the Step Into button, a menu pops up to allow you to choose the number of source lines to be stepped. If you choose the N . . . menu entry, a dialog window is opened to allow you to enter a step value.

Step Over

Executes the code involving a single source line of the current process. The current process continues to the next source line in the current file, and does not count any statements in functions that may be encountered in the source line. If the source line is a subrouting call, the process will stop at the next source line in the current file. The Step Into button can be used to step into a function or subroutine call and stop at the first executable line in the function or subroutine. If a trap is encountered while executing the step over command, the process is stopped where the trap was fired. The Step Over button is active only

	after the running process(es) has stopped; otherwise the button is grayed out.
	When you right-click on the <code>Step Over</code> button, a menu pops up to allow you to choose the number of source lines to be stepped. If you choose the <code>N . . .</code> menu entry, a dialog window is opened to allow you to enter a step value.
<code>Return[all]</code>	Continues execution of the process until the current function being executed returns. The process is stopped immediately upon returning to the calling function. All code within the current function is executed as usual. If a breakpoint is encountered, the action is canceled and the process is stopped where the breakpoint was fired. You can use this button only after the running process(es) has stopped; otherwise the button is grayed out. This action is not allowed if the executable is instrumented for performance analysis. Note: For IRIX 6.5, this command button always returns all pthreads.
<code>Sample</code>	Allows you to collect process state data to be used by the Performance Analyzer for program evaluation. You can use this button only when the process(es) is running and the <code>Enable Data Collection</code> mode is set on the Performance panel; otherwise the button is grayed out. <i>See Developer Magic: Performance Analyzer User's Guide</i> for more information.
<code>Print</code>	Prints, in the <code>Source View</code> window, the value of any highlighted text, or the source pane of the <code>Main View</code> window.
<code>Kill [All]</code>	Kills the currently running process or all running processes in your debug session by sending the process(es) the equivalent of a <code>kill -9</code> signal. You can use this button if the process(es) is running or stopped; otherwise the button is grayed out.

Run [All]	Runs the program that you are currently debugging or all programs. After the initial run, allows you to rerun the program(s) while maintaining any breakpoints and command line arguments you have set.
Status area	Displays information about the process that you are debugging, such as process id, thread id, function name, list of arguments, location of the PC, and so forth.
Source Code area	Displays the source code for the program that you are currently debugging.
Annotation column	Clicking in this area displays information specific to a line number, such as breakpoints, location of the PC, and so forth.
File text field	Displays the name of the file shown in the source code area.
Command line area	Area of the Main View window where you can enter Debugger commands and view line-mode debugger message, and in which Debugger messages are displayed.
Show/Hide annotations button	This button (see Figure 58, page 189) is visible only when you run or load a performance experiment (see the <i>Developer Magic: Performance Analyzer User's Guide</i> for more information on the performance tools). This is a toggle button that shows or hides performance related annotations.

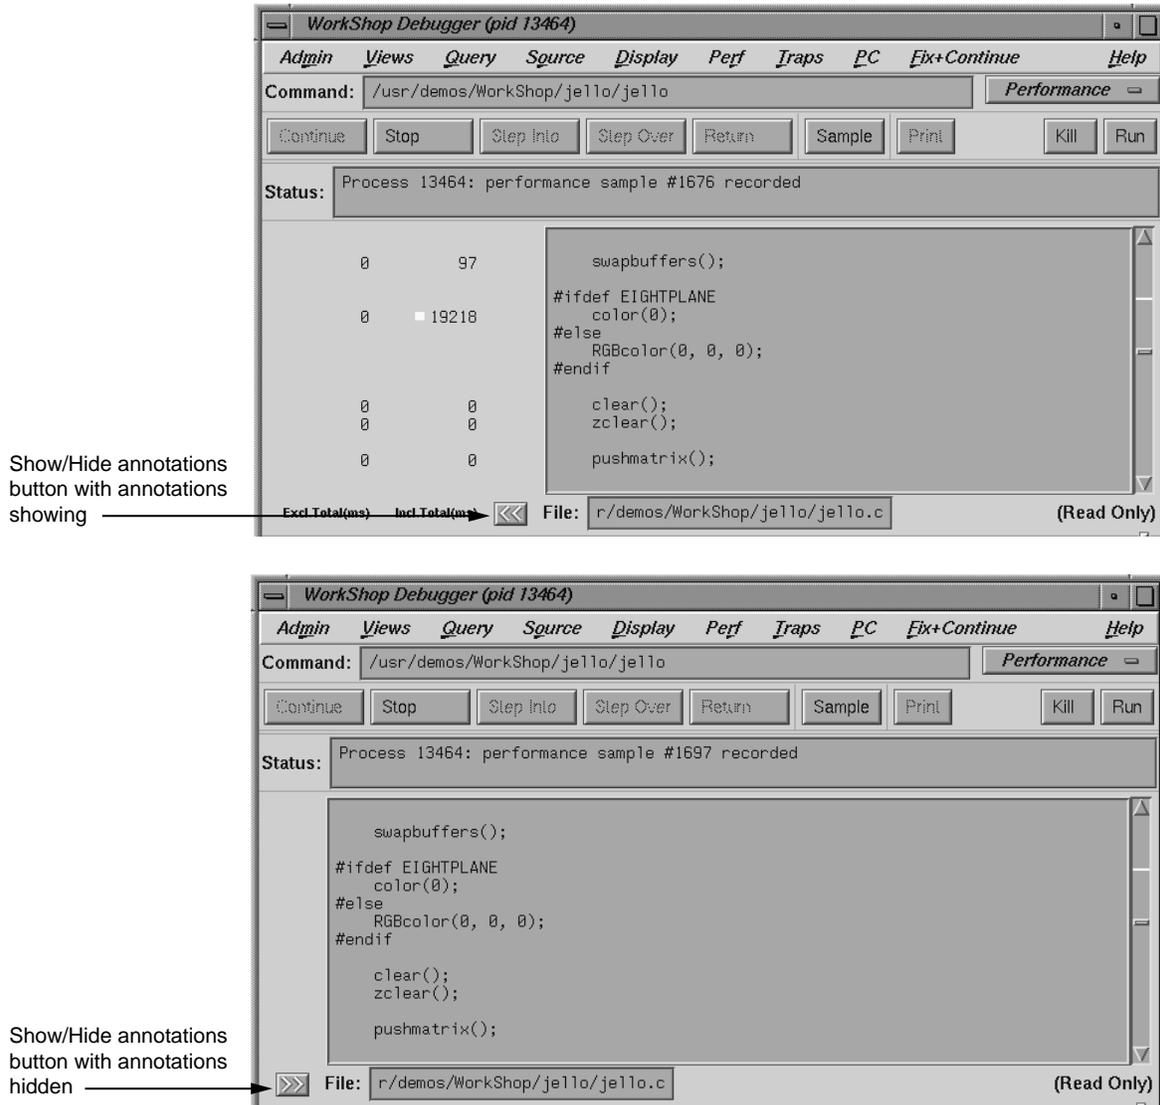


Figure 58. Show/Hide Annotations Button in the Main View Window

A.1.1 Admin Menu

The Admin menu in the Main View window performs administrative and general management tasks dealing with processes, windows, and user preferences. The Admin menu provides the following selections:

Library Search
Path...

Controls where the Debugger looks for DSOs when you invoke the Debugger on an executable or core file. The Library Search Path dialog allows you to reset the LD_LIBRARY_PATH and _RLD_ROOT environment variables. You can also reset _RLD_LIST to control the set of DSOs that will be used by the program. See the rld(1) man page for more information on these variables. Any changes you make to these variables are propagated into Execution View when you run the program.

There are two ways to open this dialog. First, from the Main View window menu bar, as follows:

```
Admin  
  ->Library Search Path...
```

Or, the Library Search Path dialog opens automatically if you invoke the Debugger on an executable or core file and the it is unable to find all of the required DSOs. In this case, an annotated list of required DSOs displays at the top of the dialog box with such status messages as OK, Error: Cannot find library, or Error: Core file and library mismatch (which indicates that the Debugger found a DSO that did not match the core file). Below this list are three to nine fields in which you can modify the value of the corresponding named environment variable.

Remap Path...

Opens a dialog that allows you to enter a new pathname.

Multiprocess
View...

Displays the Multiprocess View window, which allows you to control processes and threads. You should note that if you exit from

	Multiprocess View, you will exit from your debugging session. More additional information about Multiprocess View, refer to Section A.2.2, page 206.
GLdebug	Provides a toggle to turn on GLdebug. GLdebug is a graphical software tool for debugging application programs that use the IRIS Graphics Library (GL). GLdebug locates programming errors in executables when GL calls are used incorrectly. For more information, refer to the <i>GLdebug User's Guide</i> .
Switch Process...	Changes the current process. You will be queried for the new process ID. You can select one from the list of items presented, type one in or paste one in from another window. Switching processes changes the session.
Switch Executable...	Changes the current executable. This option also allows you to debug a different core file.
Detach	Releases the process from the Debugger. This allows you to make changes to the source code. You must detach the process before you recompile the program.
Load Settings...	Allows you to use the previously saved preference settings from a file you choose in the Load Settings dialog.
Save Settings...	Allows you to save the current preference settings to an initialization file used when the Debugger is first started, or any file you choose through this dialog. These can include such items as window sizes, current views, window configurations, and so on.
Iconify	Iconifies all session views.
Raise	Brings all session view windows to the foreground and displays any iconified windows.
Launch Tool	Allows you to run other WorkShop tools. You can switch to the other tools by selecting Build Analyzer, Static Analyzer, Performance Analyzer, or Tester. Selecting Debugger

	allows you to start another debugging session. If you have ProDev ProMP (formerly called WorkShop Pro MPF) installed on your system, the Parallel Analyzer selection is also available.
Project	Allows you to bring up a Project View to manage your work on a project-by-project basis. You can also Iconify or Raise windows, Remap paths, and Exit the project.
Close	Closes the Main View window.
Exit	Closes all views in the session and terminates the session.

A.1.2 Views Menu

The Views menu in Main View window provides the following selections for viewing the process(es) and their corresponding data:

Array Browser	Displays values from an array or array-slice in a two-dimensional spreadsheet and optionally in a three-dimensional representation; that is, a bar graph, surface, multiple lines, or points in space. These help you pick out bad data more readily. Arrays can contain up to 100 x 100 elements.
Call Stack	Displays the call stack along with parameters to the calls. If you double-click an entry in the stack, you switch the current context to that entry and you can check the state of variables.
Disassembly View	Displays assembly code corresponding to the source code.
Exception View	Displays an Ada-specific window used for exception handling.
Execution View	Displays the iconified Execution View window, which handles the input and output of the target process.
Expression View	Evaluates expressions in Fortran, C, or C++. To enter an expression, select it in the source code display and paste it into the Expression View field, using the middle mouse button.

File Browser	Displays a list of source files and library routines used by the current executable. Double-click a source file in the list to load it directly into the source display area in Main View or Source View windows. The Search field allows you to find files in the list quickly.
Memory View	Displays the value at a given memory address.
Process Meter	Monitors the resource usage of a running process without saving the data. (Used with the Performance Analyzer.)
Register View	Displays the values stored in the hardware registers for the target process.
Signal Panel	Displays the signals that can occur. You can specify which signals trigger traps and which are to be ignored.
Source View	Displays source code. Allows you to set traps, perform searches, and inspect source code without losing information in the Main View window.
Structure Browser	Displays data structures in a graphical format. You can de-reference pointers by double-clicking.
Syscall Panel	Allows you to set traps at the entry to or exit from system calls.
Task View	Brings up an Ada-specific view that provides task and callstack information for processes.
Trap Manager	Allows you to set, edit, and manage traps. The Trap Manager is used by both the Debugger and the Performance Analyzer.
Variable Browser	Displays values of local variables and parameters for the current context.

X/Motif Analyzer	Provides you with specific debugging support for X/Motif applications. There are various examiners for different X/Motif objects, such as widgets and X graphics contexts, that might be difficult or impossible to inspect using ordinary Debugger functionality.
------------------	--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------

A.1.3 Query Menu

The *Query* menu allows you to perform some of the queries available in the Static Analyzer. These queries are convenient if you have previously built a *cvstatic* fileset. However, if you need to build the fileset from scratch, the process becomes more involved. For complete information about using the Static Analyzer, refer to the *Developer Magic: Static Analyzer User's Guide* and the *cvstatic(1)* man page

With a current fileset, you can double-click any defined entity in the source code, select the *Where Defined?* option from the submenu appropriate to its type, and the source code display area will scroll to the location where the item is defined.

A.1.4 Source Menu

The *Source* menu in the Main View window provides the following selections to manage source code files:

Open . . .	Loads a source file.
Open Recent . . .	Provides you with a popup dialog that gives you a selection of recently-opened files from which to choose.
Save	Records changes made during the debugging session to the source file. You must first select <i>Make Editable</i> , which appears in the <i>Source</i> menu when the file is read-only.
Save As . . .	Records changes made during the debugging session to the source file under a different filename. You must first select <i>Make Editable</i> , which appears in the <i>Source</i> menu when the file is read-only.

Save As Text...	Records the information in the display area as a text file.
Insert Source...	Inserts the text of a file within your current file. You must first select <code>Make Editable</code> , which appears in the <code>Source</code> menu when the file is read-only. You must first select <code>Make Editable</code> , which appears in the <code>Source</code> menu when the file is read-only.
Fork Editor	Starts your default editor on the current file. The default editor is determined by the <code>editorCommand</code> resource in the <code>app-defaults</code> file. The value of this resource defaults to <code>wsh -c vi +%d</code> , which means run <code>vi</code> in a <code>wsh</code> window and scroll to the current line. If the editor allows you to specify a starting line, enter <code>%d</code> in the resource to indicate the new line number.
Recompile	Displays the <code>Build View</code> window, which allows you to compile the source code associated with the current executable.
Make Read Only / Make Editable	Toggles the source code displayed between read-only and writable states so that you can edit your code.
Search...	Searches for a literal case-sensitive, literal case-insensitive, or regular expression. After you have set your target and clicked <code>Apply</code> (or pressed <code>Enter</code>), each instance is marked by a search target indicator in the scroll bar. You can search forward or backward in the file by clicking the <code>Next</code> and <code>Prev</code> buttons. You can also click an indicator with the middle mouse button to scroll to that point. Clicking <code>Reset</code> removes the search target indicators.
Go to Line...	Allows you to scroll to a position in the source code by specifying a line number. <code>Go to Line...</code> brings up a dialog box. You can enter a line number or use the slider at the top of the box to select a line number. You do not have to display line numbers to use this feature.

Versioning

Provides access to the configuration management tool, if you have designated one.

Type the following at the Execution View prompt:

```
/usr/sbin/cvconfig [rcs | sccs | ptools | clearcase]
```

Note: You must have root permissions to run cvconfig.

The Versioning submenu appears.

Selecting any submenu option displays a shell in which you can access the configuration management tool. The following selections are available on the submenu:

- CheckIn — Saves the source file and checks it into the database as a new version.
- CheckOut — Recalls the source file from the tool's database if you have the proper authority, locks it, and makes it editable.
- UncheckOut — Cancels the checkout, with no changes registered.

A.1.5 Display Menu

The Display menu in the Main View window provides the following selections to annotate the displayed source code:

Show Line Numbers/Hide Line Numbers	Displays or hides line numbers in the annotation column corresponding to the source code.
Preferences...	Displays the Annotations Preferences dialog box, which allows you to show or hide column annotations and menus specific to the different WorkShop tools. If you have purchased ProDev ProMP, you can display and manipulate loop indicators. The Performance toggle displays experiment statistics. The Tester module allows you to see coverage statistics. Turning off the Performance toggle deletes the performance annotations from the Source View.
Hide Icons/Show Icons	Hides or displays the annotation column, which is located to the left of the source code display area.

A.1.6 Perf Menu

The Perf (Performance) menu (see Figure 59, page 198) includes the following menu selections:

Select Task submenu	Allows you to choose the task for your performance analysis. The choices available are shown in Figure 59, page 198. You may select only one task per performance analysis run. If none of the given tasks satisfy your requirements, you can choose Custom. . . , which brings up the configuration dialog open to the General tab. From here, you can design your own task requirements.
---------------------	--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------

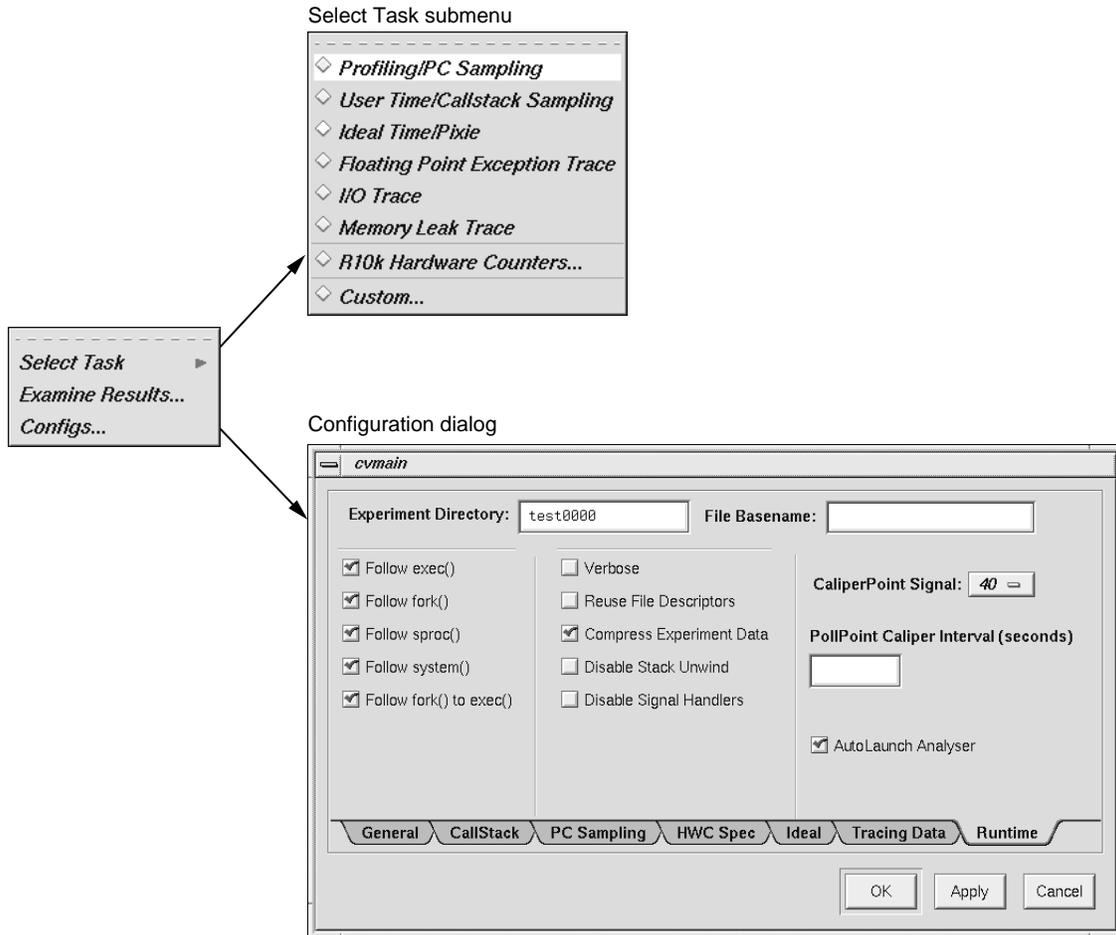


Figure 59. Perf Menu and Subwindows

- Examine Results... Launches the Performance Analyzer. For complete information about the Performance Analyzer, see the *Developer Magic: Performance Analyzer User's Guide*.
- Configs... Brings up the configuration dialog open to the Runtime tab as shown in Figure 59, page 198. The dialog opens with the Experiment Directory text field filled in with a default

value. The Performance Analyzer provides a default directory named `test0000`. If you use the default or any other name that ends in four digits, the four digits are used as a counter and will be incremented automatically for each subsequent experiment

A.1.7 Traps Menu

The Traps menu (see Figure 25, page 74) offers the Set Trap and Clear Trap submenus and the Group Trap Default and Stop All Default menu options.

The Set Trap submenu offers submenus for managing breakpoints and sample points. The following submenu selections are available:

Stop	<p>Sets a breakpoint at a designated line in your source code. To set a breakpoint at a line displayed in the Main View or Source View windows:</p> <ol style="list-style-type: none"> 1. position the cursor on the appropriate line in the source code display area 2. select the Set Trap submenu 3. choose the Stop option <p>Note: The preferred method for setting a breakpoint is to click in the annotations area of the Main View window, across from the line at which you want to set the breakpoint.</p>
Stop At Function Entry	<p>Sets a breakpoint at the beginning of a function. To set a breakpoint at a function, click on the function name in the source code display area and select the Set Trap submenu, then choose the Stop At Function Entry option.</p>
Stop At Function Exit	<p>Sets a breakpoint at the end of a function. To set a breakpoint at a function exit, click on the function name in the source code display area and select the Set Trap submenu, then choose the Stop At Function Exit option.</p>

Sample	<p>Sets a sample trap at a line displayed in the Main View or Source View windows. To set a sample trap:</p> <ol style="list-style-type: none"> 1. highlight on the appropriate line 2. pull down the Set Trap submenu 3. select the Sample option
Sample At Function Entry	<p>Sets a sample trap at the beginning of a function. To set the sample trap, highlight the function name in the source code display area, then pull down the Set Trap submenu and select the Sample At Function Entry option.</p>
Sample At Function Exit	<p>Sets a sample trap at the end of a function. To set the sample trap, highlight the function name in the source code display area, then pull down the Set Trap submenu and select the Sample At Function Exit option.</p>

The Clear Trap submenu contains selections that allow you to delete a trap on the line containing the cursor. You must designate Stop or Sample trap type, since both types can exist at the same location, appearing superimposed on each other. The following submenu selections are available:

Stop	Designates the stop trap type.
Sample	Designates the sample trap type.

The last two menu options allow you to specify the following items:

Group Trap Default	<p>Interacts with Source View. If set to true, all subsequent Source View trap requests will be group traps. That is, all members of the process group will apply this trap. This option is the same as typing <code>stop pgrp in filename</code> from the command line. Default is false unless you are using IRIX 6.5 pthreads, when the implied setting is always true.</p>
Stop All Default	<p>Interacts with Source View. If set to true, all subsequent Source View trap requests will apply the Stop All command to the trap. That is, whenever this trap is encountered, all other members of the process group also will be</p>

stopped. This option is the same as typing `stop all` in *filename* from the command line. Default is false unless you are using IRIX 6.5 pthreads, when the implied setting is always true.

If both of the default options are set to true, it is the same as typing `stop all prgp` in *filename* from the command line.

A.1.8 PC Menu

The PC (program counter) menu in the Main View window provides the following selections for controlling the execution of a process:

`Continue To` Continues the process to the selected point in the program unless some other event interrupts. Select a line by clicking on it.

Note: The process must be stopped before you can use `Continue To`.

`Jump To` Goes directly to a selected point within the same function, jumping over intervening code. Then the Debugger waits for a command to resume execution. Select a line by clicking on it.

A.1.9 Fix+Continue Menu

The Fix+Continue menu offers the following menu selections:

`Edit` Allows you to edit text using the Debugger editor.

`External Edit` Allows you to edit text by using an external editor. The default editor is `vi`, but can be changed by using the `Set Edit Tool...` pop-up menu in the Admin menu of the Status window. See Section A.8.1, page 302, for further information.

`Parse and Load` Compiles your modified program and loads it for execution. You can execute the modified program by clicking on the `Run` or `Continue` buttons in the Main View window.

`Show Difference submenu` Allows you to see the difference between the original code and your modifications. See Section A.1.9.1, page 202, for further information.

Edited<-->Compiled	Enables or disables your changes. This switch allows you to see how your application executed before and after the changes you made.
Save As...	Allows you to save your changes to a file. You can save changes to the current source file (the default) or to a separate file.
Save All Files...	Launches the <code>Save File+Fixes As...</code> dialog that allows you to update the current session and save all the modifications.
View submenu	Allows you to change to different views. <code>Fix+Continue</code> supports status, message, and build environment windows. See Section A.1.9.2, page 203, for further information.
Preferences submenu	Allows you to set your <code>Fix+Continue</code> preferences. See Section A.1.9.3, page 203, for further information.
Cancel Edit	Takes you out of edit mode and cancels any changes you have made.
Delete Edits	Deletes any modifications that you made.

A.1.9.1 Show Difference Submenu

This submenu allows you to view differences between your original and your modified code. It contains the following options:

For Function	Opens a window that shows you the differences between the original source and your modified source.
For File	Opens a window that shows you the differences between the original source file and your modified version.
Set Diff Tool ...	Launches the <code>Fix+Continue Preferences Dialog</code> that allows you to set the tool that displays code differences. The default is <code>xdiff(1)</code> . For further information on the <code>Fix+Continue Preferences Dialog</code> , see Section A.1.9.3, page 203.

A.1.9.2 View Submenu

This submenu allows you to open different Fix+Continue view windows. It contains the following options:

Status Window	Launches the Fix+Continue Status window. See Section A.8.1, page 302, for more information.
Message Window	Launches the Fix+Continue Message window. See Section A.8.2, page 307, for more information.
Build Environment Window	Launches the Fix+Continue Build Environment window. See Section A.8.3, page 308, for more information.

A.1.9.3 Preferences Submenu

This submenu allows you to set various options for the Fix+Continue environment, such as the difference tool, the external editor command, and so on. The menu contains the following options:

Show Preferences	<p>Launches the Fix+Continue Preference Dialog that displays preferences currently enabled for the session, and allows you to change the settings. The following preferences are available through the dialog:</p> <ul style="list-style-type: none">• External Editor Command text field that allows you to choose your text editor. The default is vi.• File Difference Tool text field allows you to choose the tool to use when comparing code. The default is xdiff.• Copy Traps On Previous Definition toggle allows you to edit and parse code. When Fix+Continue copies traps from the old definition to the new one by mapping old lines to new lines. (This mapping is the same as what can be generated using the UNIX diff utility.) If Copy Traps On Previous Definition is on and the mapped line the new definition is modified, then Fix+Continue will look at the switch.
------------------	----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------

- **Copy Traps Even On Changed Lines** toggle causes the debugger to copy traps onto a mapped line.
- **Continue Even If Line Has Changed** toggle allows you to edit and compile code in which your program is currently stopped. **Fix+Continue** can continue in the new definition provided some conditions are satisfied. The line from which the program continues depending on the mapping from the line in which it stopped. In case it can continue in the new definition from a line which you have modified, **Fix+Continue** consults this toggle to determine whether to continue in the new or old definition. This toggle allows you to override the default behavior.
- **Warn Unfinished Edits Before Run** toggle pops up a warning dialog before a run if you have unfinished edits.
- **Warn Unfinished Edits Before Continue** toggle pops up a warning dialog before a continue if you have unfinished edits.
- **Save deactivated code during File Save** toggle save old code. The **Fix+Continue** file save substitutes new definitions in place of old ones. If you want to save your original functions in the same file, this switch allows you to save the old (original or compiled) code under an `#ifdef`. When you compile, the old code will not get compiled. You can manually edit the source to use the old definition in any way you desire.

Reset To Factory Defaults	Sets preferences to the installed defaults.
Save Preferences	Brings up the File dialog that allows you to save your preferences to a file.
Load Preferences...	Brings up the File dialog that allows you to load preferences from a file.

A.1.9.4 Keyboard Accelerators

Use the accelerators in Table 3 to issue Fix+Continue commands directly from the keyboard. The accelerators are listed alphabetically by command.

Table 3. Fix+Continue Keyboard Accelerators

Command	Key Sequence
Cancel Edit	Alt+Ctrl+q
Edit	Alt+Ctrl+e
Parse And Load	Alt+Ctrl+x

A.1.10 Help Menu

The Help menu provides the following options:

Click for Help	Provides information on the selected window or menu.
Overview	Provides general information on the current tool.
Index...	Displays the entire list of help topics, alphabetically, hierarchically, or graphically.
Keys & Shortcuts	Lists the keys and shortcuts for the current tool.
Product Information	Provides copyright and version number information on the tool.

A.2 Some Additional Views

This section discusses some of the additional views available through the Debugger: the Execution View, Multiprocess View, Source View, and Process Meter.

A.2.1 Execution View

The Execution View window is a simple shell that allows you to set environment variables and inspect error messages. If your program is designed to be interactive using standard I/O, this interaction will take place in the Execution View window. Any standard I/O that is not redirected by your Target command is displayed in the Execution View window. Execution View is launched (and iconified) automatically with the Debugger.

A.2.2 Multiprocess View

WorkShop supports debugging of multiprocess applications, including pthreadd programs and processes spawned with either `fork` or `sproc` commands.

Multiprocess debugging is supported primarily through the Multiprocess View window. To display this window, select Multiprocess View from the Admin menu of the Main View window. Multiprocess View displays a hierarchical view of your pthreadd application. Pthreadd processes are marked with a folder icon. Clicking the folder changes the view to show that process's pthreadds. Clicking on a thread opens a call stack for that thread.

For each process or thread, clicking the right mouse button brings up a menu that applies to the selected item.

Note: This menu is a duplicate of the Process menu. See Section A.2.5, page 215 for more information.



Figure 60. Process Menu

A.2.2.1 Status of Processes

When the `Multiprocess View` window comes up, it lists the status of all processes in the process group. This view includes the following information:

PID	Shows the process identifier (PID).
PPID	Lists the parent process PIDs.
Status/State	Shows <code>Status</code> or <code>State</code> depending on how you have set your preferences on the <code>Multiprocess View Preferences</code> menu, brought up by selecting <code>Preferences</code> from the <code>Config</code> menu. <code>Status</code> is user-level status and <code>State</code> is the kernel-level status.
Name	Identifies the process by file name.
Function/PC	Indicates the current function and program counter (PC) for any stopped processes.

The following `Status` and `State` conditions are possible:

<u>Status</u>	<u>State</u>
Running	RUNNING
Stopped	RUNNING
Stopped on breakpoint	RUNNING (but at a trap pc)

Waiting to terminate	JOIN
Thread terminated	DEAD
Waiting on kernel	READY
Waiting on mutex	MUTEX-WAIT
Sleeping in system call	RUNNING

A.2.2.2 Multiprocess View Control Buttons

The `Multiprocess View` window uses the same control buttons as are in the `Main View` window with following exceptions:

- There are no `Run`, `Return`, or `Print` buttons in the `Multiprocess View`.
- The buttons in this view apply to all processes as a group.
- Using a control button in the `Multiprocess View` window has the same effect as clicking the button in each process's `Main View` window.

The following command buttons are available:

`Continue All`
`Stop All`
`Step Into All`
`Step Over All`
`Sample All`
`Kill All`

These buttons operate identically to those described for the `Main View` window, Section A.1, page 183, with the `All` option effect. Refer to that section for descriptions of these buttons.

A.2.2.3 Multiprocess View Administrative Functions

The `Admin` menu in the `Multiprocess View` window lets you perform several administrative functions. Only the `Save as Text...`, `Close`, and `Exit` items are described here. All other options perform as those found in the `Admin` menu of the `Main View` window, described in Section A.1.1, page 190.

Save as Text...	The process status list from in the Multiprocess View window is saved to the file you select using the Save Text dialog.
Launch Tool Project	
Close	Closes the Multiprocess View window only.
Exit	Exits all views in the session and terminates the session.

A.2.2.4 Controlling Preferences

The Preferences... option in the Config menu brings up the Multiprocess View Preferences dialog that allows you to control when processes are added to the group and specifies their behavior. This option also contains a Save option that allows you to save your preferences.

The Multiprocess View preference options are:

Stack Depth	Allows you to set how many lines of the call stack should be displayed when opening the call stack. Default is 10.
Levels to open	Allows you to specify the number of levels of hierarchy to be displayed. For pthreadd programs, you can display information by program, threads, and callstack. For non-pthreadd programs, you can display information by process and callstack. Default is one level.
Attach to forked processes	Automatically attaches new processes spawned by the fork command to the group. (Note that processes spawned by sproc are always attached.) Default is off.
Copy traps to forked processes	Copies traps you have set in the parent process to new forked processes automatically. Alternatively, if you create parent traps with Trap Manager and specify pgrp, then the children inherit these traps automatically, regardless of the state of this flag. Default is off.
Copy traps to sproc'd processes	Copies traps you have set in the parent process to new sproc'd processes automatically. As in the

	previous option, if you create parent traps with the Trap Manager and specify <code>pgrp</code> , the children inherit these traps automatically, whether this flag is set or not. Default is on.
Resume parent after fork	Restarts the parent process automatically when a child is forked. Default is on.
Resume child after attach on fork	Restarts the new forked process automatically when it is attached. If this option is left off, a new process will stop as soon as it is attached. Default is on.
Resume parent after sproc	Restarts the parent process automatically when a child is <code>sproc'd</code> . Default is on.
Resume child after attach on sproc	Restarts the new <code>sproc'd</code> process automatically when it is attached. If this option is left off, a new process will stop as soon as it is attached. Default is on.
Combine threads at same location	Applies a collapsing algorithm to display threads stopped at the same location at the same time. (It is possible for threads to arrive at the same location through different logical routes.) Default is on.
Show Thread Status vs Thread State	Displays thread status, which is the user-level status as opposed to showing thread state, which is the kernel-level status. Default is status.
Show/Hide Buttons	When this option is ON, the Continue All, Stop All, Step Into All, Step Over All, Sample All, and Kill All buttons appear near the top of the Multiprocess View window.
Show/Hide Header Information	When this option is ON, column headings display above the list of processes/threads.

A.2.3 Source View

The Source View window is brought by choosing the following from the Main View window menu bar:

```
Views
  ->Source View
```

By default, a copy of the source on display in the Main View window source pane is displayed in this window.

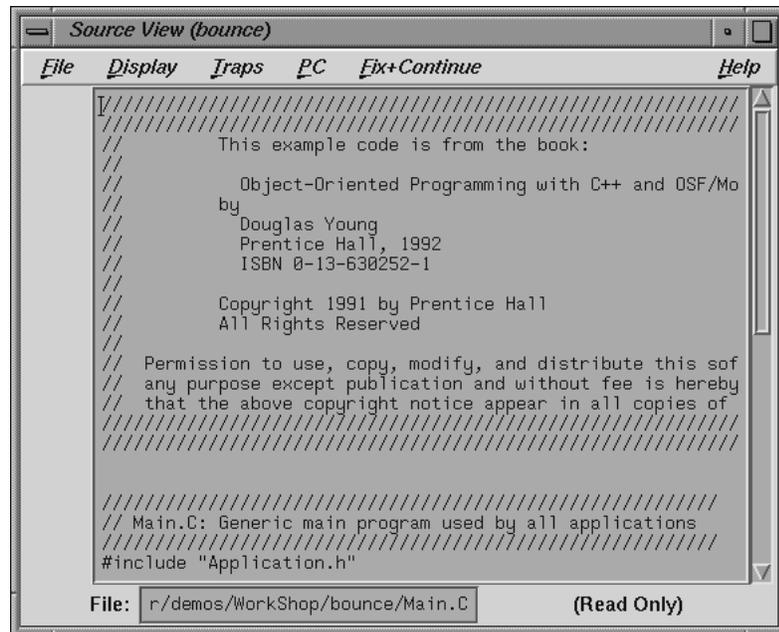


Figure 61. Source View Window

The Source View menu bar contains selections duplicated from the Main View window: Display, Traps, PC, and Fix+Continue. Each of these menus has the same functionality as its counterpart in the Main View window (see Section A.8.4.1, page 309). The only new menu selection is the File menu described below:

Open...	Launches the Open dialog that allows you to choose a file to load into Source View.
Save	Records changes made to the file during the current debugging session. You must first select Make Editable from this File menu when the file is read only.
Save As...	Records changes made during the debugging session to the source file under a different file

	name, the name of which you can enter or select using the <code>Save As Text...</code> dialog.
<code>Save As Text...</code>	Records information in the display area as a text file, the name of which you can enter or select using the <code>Save As Text...</code> dialog.
<code>Open Separate...</code>	Launches the <code>Open Separate</code> dialog that allows you to create a new <code>Source View</code> with the contents of a different source file.
<code>Insert File...</code>	Inserts the text of a file within your current file. This item description is available only if the file is editable. The <code>Make Editable</code> item description from this <code>File</code> menu can be used to switch the file from read only.
<code>Clone</code>	Clones the current window.
<code>Fork Editor</code>	Starts your default editor on the current file. The default editor is determined by the <code>editorCommand</code> resource in the <code>app-defaults</code> file. The value of this resource defaults to <code>wsh -c vi +%d</code> , which means run <code>vi</code> in a <code>wsh</code> window and scroll to the current line. If the editor allows you to specify a starting line, enter <code>%d</code> in the resource to indicate the new line number.
<code>Recompile</code>	Displays the <code>Build View</code> window that allows you to compile the source code associated with the current executable.
<code>Make Editable</code>	Toggles the source code displayed between (<code>Read Only</code>) and (<code>Editable</code>) so that you can edit your code.
<code>Search</code>	Searches for a literal case-sensitive, literal case-insensitive, or regular expression. After you have set your target and clicked <code>Apply</code> (or pressed <code>Enter</code>), each instance is marked by a search target indicator in the scroll bar. You can search forward or backward in the file by clicking the <code>Next</code> or the <code>Prev</code> button. You can also click an indicator with the middle mouse button to scroll to that point. Clicking <code>Reset</code> removes the search target indicators.

Go To Line...	Launches the Go To Line dialog that allows you to go to a specific line in the source. You can type in the line, or select the line number via the slider bar.
Versioning	Provides access to the configuration management tool, if you have designated one. The <code>cvconfig</code> script allows you to designate ClearCase, RCS or SCCS. Type the following: <code>/usr/sbin/cvconfig [rsc sccs ptools clearcase]</code> You must have root permissions to run <code>cvconfig</code> . Selecting any option displays a shell in which you can access the configuration management tool. The selections in the submenu are:
Versioning submenu	Contains the following options: <ul style="list-style-type: none"> • <code>CheckIn</code> — Saves the source file and checks it into the database as a new version. • <code>CheckOut</code> — Recalls the source file from the tool's database if you have the proper authority, locks it, and makes it editable. • <code>UncheckOut</code> — Cancels the checkout, with no changes registered.
Close	Dismisses the Source View window.

A.2.4 Process Meter

The Process Meter window is brought up by choosing the following from the Main View menu bar:

```
Views
  ->Process Meter
```

The Process Meter monitors resource usage of a running process without saving the data. Figure 62, page 214, shows the Process Meter in its default configuration (with only the User Time and Sys Time fields active).

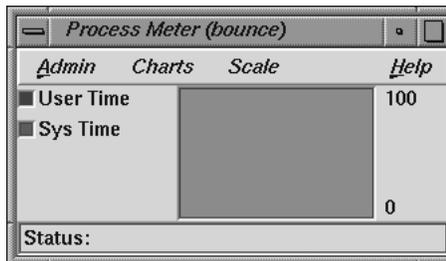


Figure 62. Process Meter

The Process Meter contains its own menu bar that contains the Admin, Charts, Scale, and Help menus. The Admin menu is the same as that described in Section A.3.1.1, page 217. The Help menu is the same as that described in Section A.1.10, page 205. The other menus are described in the following sections.

A.2.4.1 Charts Menu

The Charts menu contains a set of toggles that allow you to choose which charts are displayed in the Process Meter window. You can display as many charts simultaneously as you wish. The following choices are available:

- User/Sys Time (the default)
- Major/Minor Faults
- Context Switches
- Bytes Read/Written
- Read/Write Sys Calls
- Other Sys Calls
- Total Sys Calls
- Signals
- Process Size

A.2.4.2 Scale Menu

The Scale menu allows you to set the time scale for the processes displayed in the Process Meter window. A menu allows you to choose a time scale from 2 seconds to 10 minutes.

A.2.5 Controlling Multiple Processes

The Process menu allows you to control processes and threads.



Figure 63. Process Menu

The Process menu has the following options:

Change focus to
this entry...

Opens a dialog that allows you to switch the process or thread currently focused on in the Main View window to the process or thread selected in the Multiprocess View window. Selecting a call stack entry changes the Main View window's focus to that process or thread and positions the cvmain window at the offset of the selected call stack.

Create a new
window...

Brings up a new Main View window for the selected process or thread.

Goto...

Opens a dialog box that allows you to enter the name of a thread on which focus should be switched. This is useful when multiple threads, all at the same location, are collapsed into a

single line. While `Change focus to this entry...` always takes you to the first thread, `Goto...` allows you to jump to any thread.

`Add...`

Opens a dialog in which you can select from a list of process ids. Selecting a process id (PID) in the dialog and pressing OK will cause the process to be attached and added into the `Multiprocess View` window's list of processes.

`Remove`

After you select a process/thread (by highlighting it), click on `Remove` to remove it from the list of processes in the `Multiprocess View` window.

Note: A process in a `sproc` share group cannot be removed from the process group.

A.3 Ada-specific Windows

This section discusses the `Task View` and `Exception View` windows that are specific to Ada code.

A.3.1 Task View

Select the following from the `Main View` window menu bar to call up the `Task View` window:

```
Views
  ->Task View
```

The `Task View` window is an Ada-specific view that provides you with task and call stack information. If you do not have Ada installed on your system, the `Task View` menu option of the `Views` menu is grayed out.

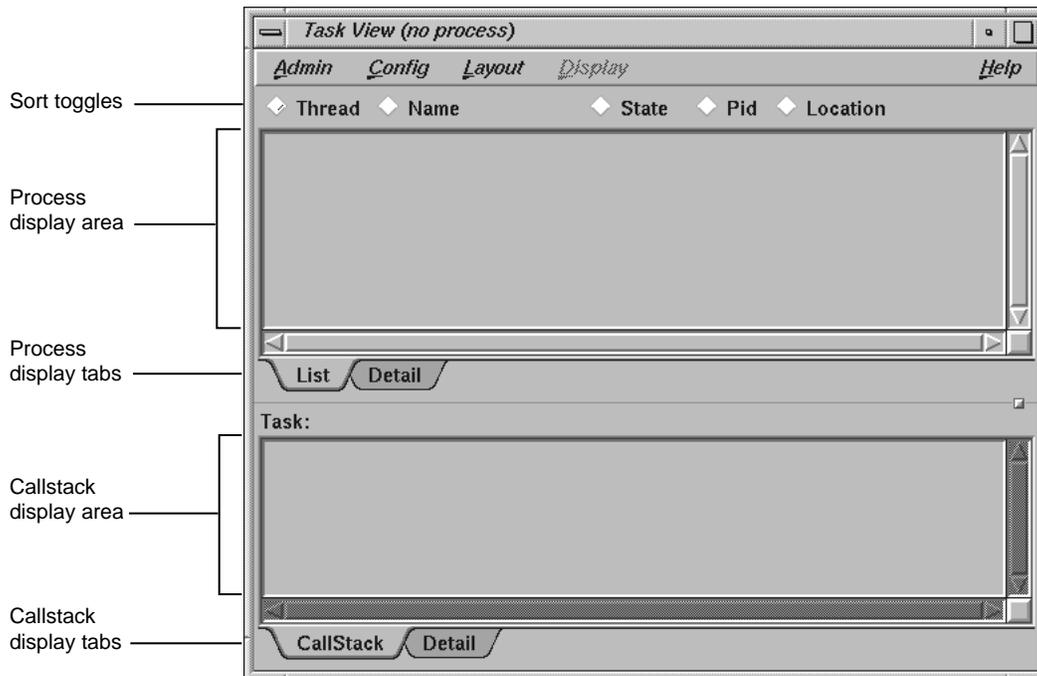


Figure 64. Task View Window

The Task View menu bar contains the Admin, Config, Layout, Display, and Help menus. The Help menu is the same as that described in Section A.1.10, page 205. Other menus are described in the following sections.

A.3.1.1 Admin Menu

The Admin menu contains the following options:

Active	This toggle activates the current window in a set of cloned windows.
Clone	Creates a clone of the current window. This action is not supported in the current release, and the option is grayed out.

Save As Text...	Launches the Save Text dialog. This dialog allows you to save your current session as text in a file you designate.
Close	Closes the current window.

A.3.1.2 Config Menu

The Config menu contains the following item:

Preferences...	Launches the Task View Preference dialog that allows you to set maximum call stack depth shown in Task View. Default depth is 32 frames.
----------------	------------------------------------------------------------------------------------------------------------------------------------------

A.3.1.3 Layout Menu

The Layout menu contains the following toggles:

Task List	Causes only the CallStack View to be shown.
Single Task	Causes only the Process Display to be shown.

A.3.1.4 Display Menu

The Display menu is divided into the Task List Format and Callstack Format sections. The Task List Format toggle buttons control which buttons appear in the toggle sort list, as well as what information is displayed in the Process Display area. The Callstack Format toggles control the amount of information to be displayed in the Callstack Display area of the Task View window.

The Task View Display menu contains the following toggles:

Thread/task	Displays thread/task number. This toggle is active by default.
Status	Displays process status. This toggle is active by default.
PID	Displays PID number.
Location	Displays routine name and location in the current source file.
Arg Values	Allows you to set the argument values in the Callback Display. This toggle is active by default.

Arg Names	Allows you to set the argument names in the Callback Display. This toggle is active by default.
Arg Types	Allows you to display argument types in the Callback Display.
PC	Allows you to set the program counter (PC) in Task View.

In addition to menus, Task View also contains the following items from which you can select to vary the display:

Sort toggles	Allow you to sort the process list by Thread, Name, State, Pid, or Location, depending on which of the buttons is active. Default selection is Thread.
Process display tabs	Allow you to view a list of tasks or details of the currently running (highlighted) task.
Callstack display tabs	Allow you to view all call stack information or call stack details of the currently selected process.

A.3.2 Exception View

Select the following from the Main View window menu bar to display the Exception View window:

```
Views
  ->Exception View
```

Exception View is an Ada-specific view that allows you to set traps on exceptions and control exception handling. This view works only if the Ada compiler is installed. By default, this view displays only the following predefined Ada exceptions:

- Constraint errors
- Program errors
- Storage errors
- Tasking errors

In addition, a single breakpoint is set on any unhandled exception.
Figure 65 shows a typical Exception View window.

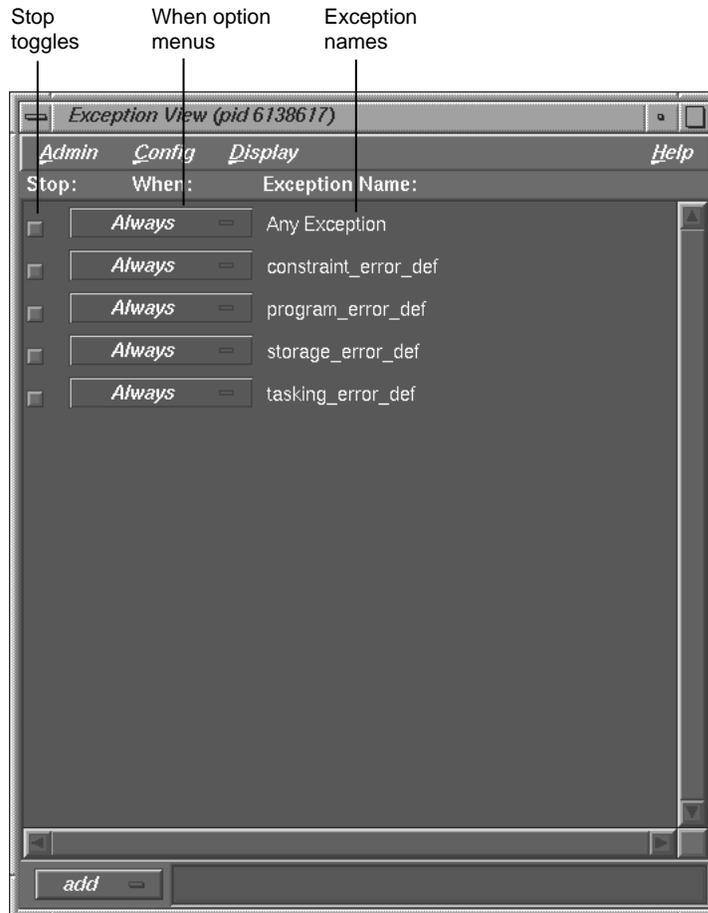


Figure 65. Exception View

The Admin menu has the following options:

- | | |
|--------|----------------------------------------------------------|
| Active | Activates the current window in a set of cloned windows. |
| Clone | Creates a clone of the current window. |

Save As Text... Launches the Save Text dialog. This dialog allows you to save your current session as text in a file you designate.

Close Closes the current window.

The Config menu has the following options:

Load Exceptions... Opens the Load User Defined Exceptions dialog that allows you to add additional exceptions to the predefined Ada exceptions.

Save Exceptions... Opens the Save User Defined Exceptions dialog that allows you to save any user-defined exceptions to the predefined Ada exceptions.

The Display menu has the following options:

Delete All Deletes all exception traps.

Clear All Traps Clears all exception traps. Clearing traps is not the same as deleting traps. Clearing only temporarily affects traps while deleting removes them permanently.

Reset All Buttons Resets all button actions.

The Stop boxes toggle on and off to indicate whether a trap is active.

The When control menus allow you to determine when an exception trap fires. The following choices are available:

Always Stops any time the exception is raised.

WhenOthers Stops when caught by a when others handler rather than an explicit handler or when unhandled.

Unhandled Stops when the exception is unhandled.

In the un-labeled text field at the bottom right of the window you can enter a single, fully qualified Ada exception name or a single, fully qualified Ada unit name. Depending on whether the `add`, `remove`, or `find` mode is active; pressing `Enter` will cause one of the following actions to occur:

- `add` mode:
 - Single exception: Adds single exception to the exception list
 - Library unit name: Adds all exceptions found in that library unit name to the exception list
- `remove` mode:
 - Single exception: Removes single exception from the exception list
 - Library unit name: removes all exceptions found in that library unit name from the exception list
- `set` mode:
- `clear` mode:
- `find` mode:
 - Single exception: positions top of the exception list to single exception
 - Library unit name: positions top of the exception list to the first exception found in given library unit name

A.4 X/Motif Analyzer Windows

The X/Motif Analyzer provides specific debugging support for X/Motif applications. There are various examiners for different X/Motif objects, such as widgets and X Window System graphics context, that might be difficult or impossible to inspect using ordinary debugger functionality. See Chapter 11, page 159 for a comprehensive discussion and tutorial regarding the X/Motif Analyzer.

To access the X/Motif Analyzer window, pull down the `Views` menu and select `X/Motif Analyzer` (see Figure 66, page 223).

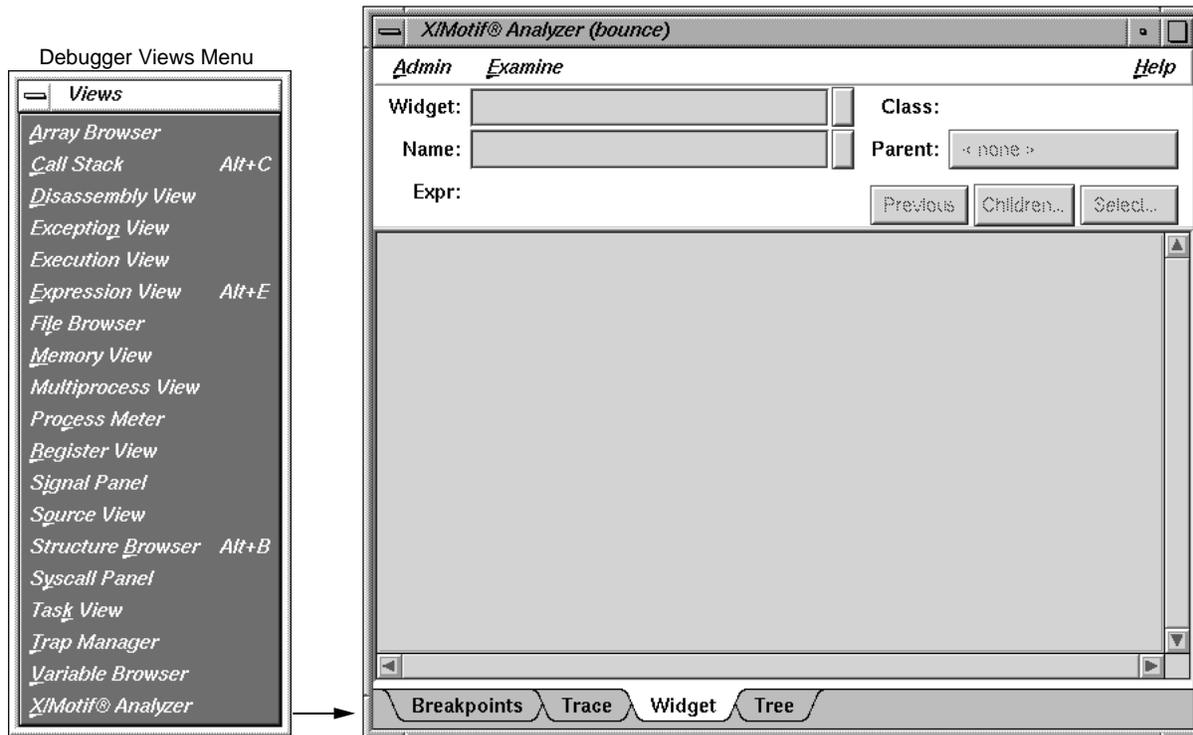


Figure 66. Launching the X/Motif Analyzer Window

A.4.1 Global Objects

Through the X/Motif Analyzer is made up of several different examiner windows, a number of objects, such as the Admin menu, Examine menu, Help menu and several text bars, remain constant throughout window changes. The following examiners are available and discussed in the sections indicated:

- Breakpoints Examiner, Section A.4.2, page 226
- Trace Examiner, Section A.4.3, page 240
- Widget Examiner, Section A.4.4, page 242
- Tree Examiner, Section A.4.5, page 243
- Callback Examiner, Section A.4.6, page 246
- Window Examiner, Section A.4.7, page 247

- Event Examiner, Section A.4.8, page 248
- Graphics Context Examiner, Section A.4.9, page 248
- Pixmap Examiner, Section A.4.10, page 249
- Widget Class Examiner, Section A.4.11, page 250

A.4.1.1 Admin Menu

The Admin menu offers the following menu selections:

Active	Activates the current window in a set of cloned windows. In the current release, this toggle is always active.
Clone	Creates a clone of the current window. This action is not supported in the current release and the option is grayed out.
Save As Text...	Launches the <code>Save Text</code> dialog. This dialog allows you to save your current session as text in a file you designate. Note: This selection is not available for examiners that are graphical displays, such as the Breakpoints Examiner, the Tree Examiner, and the Pixmap Examiner.
Close	Closes the current window.

A.4.1.2 Examine Menu

The Examine menu offers the following options:

Selection	Selects the currently highlighted object for examination. Note: You must first highlight the name of an object before you select this option.
Widget	Uses the current selection as input to the widget examiner, then opens that examiner (see Section A.4.4, page 242, for information).
Widget Tree	Switches the window view to the widget tree examiner (see Section A.4.5, page 243, for information).

Widget Class	Switches the window view to the widget class examiner (see Section A.4.11, page 250, for information).
Window	Switches the window view to the window examiner (see Section A.4.7, page 247, for information).
X Event	Switches the window view to the X Event examiner (see Section A.4.8, page 248, for information).
X Graphics Context	Switches the window view to the X graphics context examiner (see Section A.4.9, page 248, for information).
X Pixmap	Switches the window view to the X pixmap examiner (see Section A.4.10, page 249, for information).

A.4.1.3 Examiner Tabs

In addition to access through the Examine menu, each examiner can be accessed through a tab at the bottom of each view (see Figure 67).



Figure 67. Examiner Tabs

When first launched, the X/Motif Analyzer has the following four tabs from left-to-right:

- Breakpoints
- Trace
- Widget
- Tree

As you select other examiners through the `Examine` menu, new tabs are added for the new examiners.

To delete a tab:

1. Select the tab you want to delete.
2. Right-click and select `Remove Examiner` from the pop-up menu.

The selected tab will disappear.

Note: You can not remove the first four tabs.

The X/Motif Analyzer will also bring up new examiner windows whenever they are needed (see Chapter 11, page 159 for more information).

Click on the collapsed tabs to the right to display them.

A.4.1.4 Return Button

Both the `Widget` and `Name` text fields have return buttons (see Figure 68, page 228) just to the right. Clicking these buttons causes the X/Motif Analyzer to respond exactly as if you had pressed `Return` on your keyboard.

Note: Only the `Breakpoint Examiner` and `Widget Examiner` have a `Return` button.

A.4.2 Breakpoints Examiner

The `Breakpoints` examiner is not really an examiner, but a control area where you can set widget-level breakpoints. The `breakpoints` examiner is divided into three areas (see Figure 68, page 228):

- The widget specification area that contains the same information as that in the `Widget` examiner. You can select a widget address, name, or class in this area, as well as move to the widgets parents or children, or select a widget in the application. In cases where the breakpoint type does not apply to widgets (for example, input-handler breakpoints), this area is blank.
- The parameter specification area, the contents of which vary according to the type of breakpoint you are setting. For example, for callback breakpoints, this area contains the callback name and client data; for event-handler breakpoints, it contains the event type and the client data, and so on.

- The breakpoint area, which contains the breakpoint name, a search field, and the Add, Modify, Delete, and Step To buttons. Since the Search text field and the four buttons appearing in the Breakpoints area function the same way no matter which Breakpoint Type is selected, descriptions for these items are included here, and will not be described in each of the eight remaining subsections. These are:

Search text field	Allows you to perform a text search through your breakpoints.
Add button	Allows you to add a new breakpoint.
Modify button	Allows you to change the selected breakpoint's settings.
Delete button	Deletes the selected breakpoint.
Step To button	Allows you to step to the next condition. Step To creates a temporary breakpoint, resumes the process, and waits until the process stops. This temporary breakpoint acts exactly like an ordinary breakpoint, save that the Step To button automatically resumes the process and puts up a busy cursor until the condition becomes true.

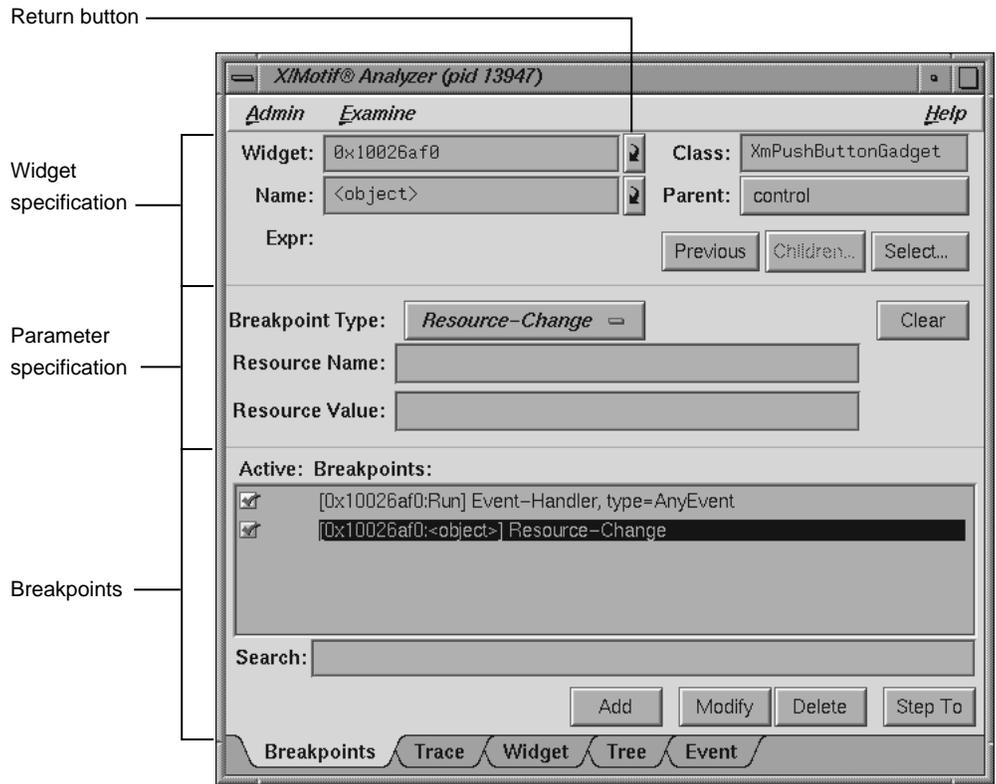


Figure 68. Breakpoints Examiner Display in the X/Motif Analyzer Window

The control area has eight different breakpoint types that it can examine. These types are set through the Breakpoint Type options. The following Breakpoint Type options are available:

- | | |
|---------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Callback | Widget callback installed by XtAddCallback. Parameters include callback name and client_data XtPointer value. See Section A.4.2.1, page 229, for more information. |
| Event-Handler | Widget event handler installed by XtAddEventHandler. Parameters include X event type and client_data XtPointer value. See Section A.4.2.2, page 231, for more information. |

Resource-Change	Resource change caused by <code>XtSetValues</code> or <code>XtVaSetValues</code> . Parameters include resource name and resource value, both strings. See Section A.4.2.3, page 233, for more information.
Timeout-Procedure	Timeout callback installed by <code>XtAppAddTimeOut</code> . Parameters include <code>client_data XtPointer</code> value. See Section A.4.2.4, page 234, for more information.
Input-Handler	Input callback installed by <code>XtAppAddInput</code> . Parameters include <code>client_data XtPointer</code> value. See Section A.4.2.5, page 235, for more information.
State-Change	Various widget state changes (for example, managed or realized). Parameters include widget state. See Section A.4.2.6, page 236, for more information.
X-Event	X event received by target application. Parameters include X event type. See Section A.4.2.7, page 237, for more information.
X-Request	X request received by target application. Parameters include X request type. See Section A.4.2.8, page 239, for more information.

A.4.2.1 Callback Breakpoints Examiner

When the `Callback` option of the `Breakpoint Type` option button in the `Breakpoints Examiner` is selected, the `Callback Breakpoints Examiner` is displayed.

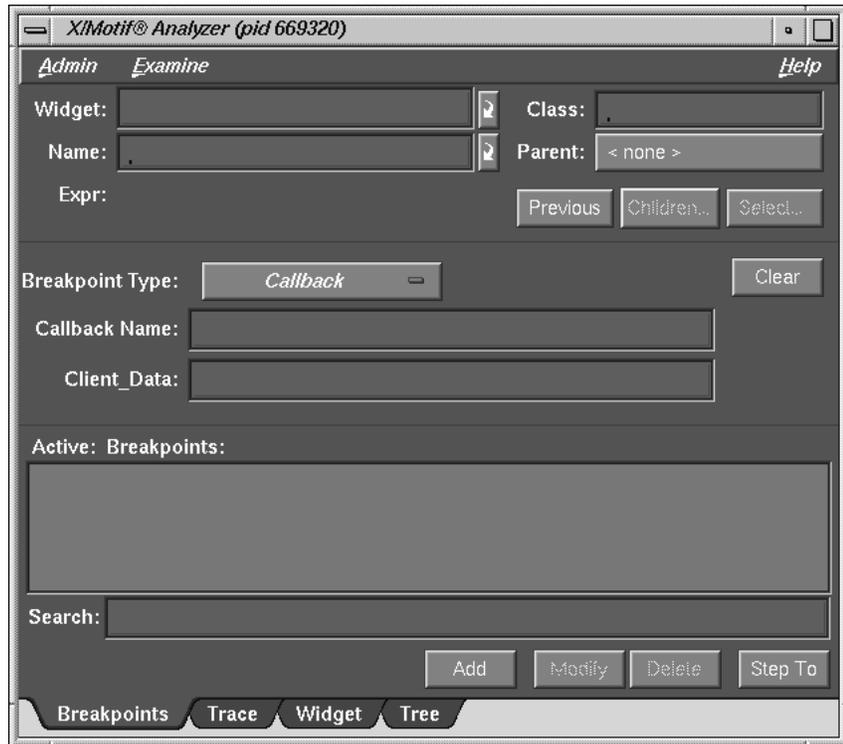


Figure 69. Callback Breakpoints Examiner

The Callback Breakpoints examiner contains the following items:

Widget text field	Allows you to choose a widget to examine by entering the widget address.
Name text field	Allows you to choose a widget to examine by entering the widget name.
Class text field	Allows you to choose a widget to examine by entering the widget's class. Leave the field blank or enter Any to select all widgets.
Parent text field	Allows you to move to the parent of the currently selected widget.
Previous button	Moves you to the previously selected widget.

Children... button	Shows you the widget's children (it is grayed out if the selected widget cannot have children).
Select... button	Allows you to select the widget in the target process.
Breakpoint Type option button	Allows you to select the type of breakpoint you wish to set. In this section, <code>Callback</code> is selected.
Clear button	Clears all the current breakpoint selections and text fields.
Callback Name text field	Allows you to set the Name of the callback for the breakpoint.
Client_Data text field	Allows you to pass and get back pointer values for <code>Client_Data</code> .

A.4.2.2 Event-Handler Breakpoints Examiner

When the `Event-Handler` option of the `Breakpoint Type` option button in the Breakpoint Examiner is selected, the examiner appears as shown in Figure 70, page 232.

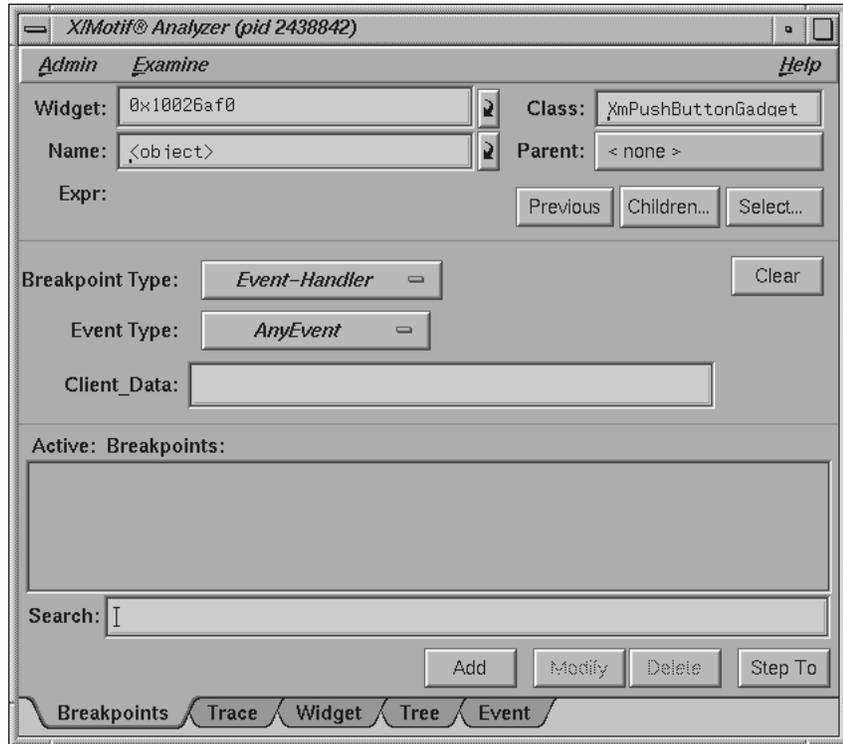


Figure 70. Event-Handler Breakpoints Examiner

The Event-Handler Breakpoints examiner contains the following items:

Widget text field	Allows you to choose a widget to examine by entering the widget address.
Name text field	Allows you to choose a widget to examine by entering the widget name.
Class text field	Allows you to choose a widget to examine by entering the widget's class. Leave the field blank or enter Any to select all widgets.
Parent text field	Allows you to move the parent of the currently selected widget.
Previous button	Moves you to the previously selected widget.

Children... button	Shows you the widget's children (it is grayed out if the selected widget cannot have children).
Select... button	Allows you to select the widget in the target process.
Breakpoint Type option button	Allows you to select the type of breakpoint you wish to set. In this section, Event-Handler is selected.
Clear button	Clears all the current breakpoint selections and text fields.
Event Type option button	Allows you to set the event type for a given breakpoint.
Client_Data text field	Allows you to pass and get back pointer values for the Client_Data.

A.4.2.3 Resource-Change Breakpoints Examiner

When the Resource-Change option of the Breakpoint Type option button in the Breakpoint Examiner is selected, the examiner appears as shown in Figure 68, page 228.

The Resource-Change Breakpoints examiner contains the following items:

Widget text field	Allows you to choose a widget to examine by entering the widget address.
Name text field	Allows you to choose a widget to examine by entering the widget name.
Class text field	Allows you to choose a widget to examine by entering the widget's class. Leave the field blank or enter Any to select all widgets.
Parent text field	Allows you to move the parent of the currently selected widget.
Previous button	Moves you to the previously selected widget.
Children... button	Shows you the widget's children (it is grayed out if the selected widget cannot have children).
Select... button	Allows you to select the widget in the target process.

Breakpoint Type option button	Allows you to select the type of breakpoint you wish to set. In this section, Resource-Change is selected.
Clear button	Clears all the current breakpoint selections and text fields.
Resource Name text field	Allows you to set the resource name for the breakpoint.
Resource Value text field	Allows you to set the resource value for the breakpoint.

A.4.2.4 Timeout-Procedure Breakpoints Examiner

When the Timeout Procedure option of the Breakpoint Type option button in the Breakpoint Examiner is selected, the examiner appears as shown in Figure 71.

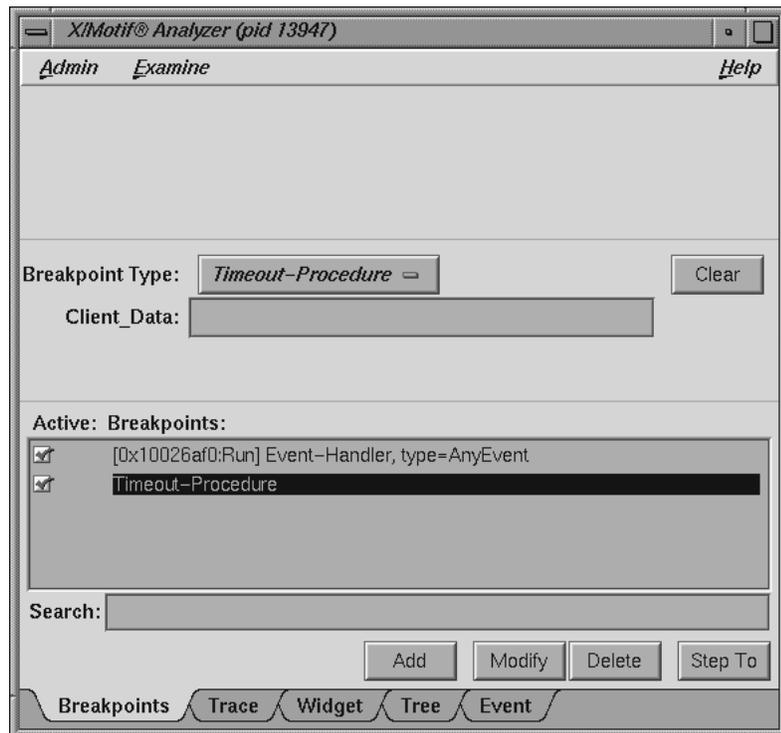


Figure 71. Timeout-Procedure Breakpoints Examiner

The Timeout-Procedure Breakpoints examiner contains the following items:

Breakpoint Type option button	Allows you to select the type of breakpoint you wish to set. In this section, Timeout-Procedure is selected.
Clear button	Clears all the current breakpoint selections and text fields.
Client_Data text field	Allows you to pass in and get back pointer values for the Client_Data.

A.4.2.5 Input-Handler Breakpoints Examiner

When the Input-Handler option of the Breakpoint Type option button in the Breakpoint Examiner is selected, the examiner appears as shown in Figure 72.

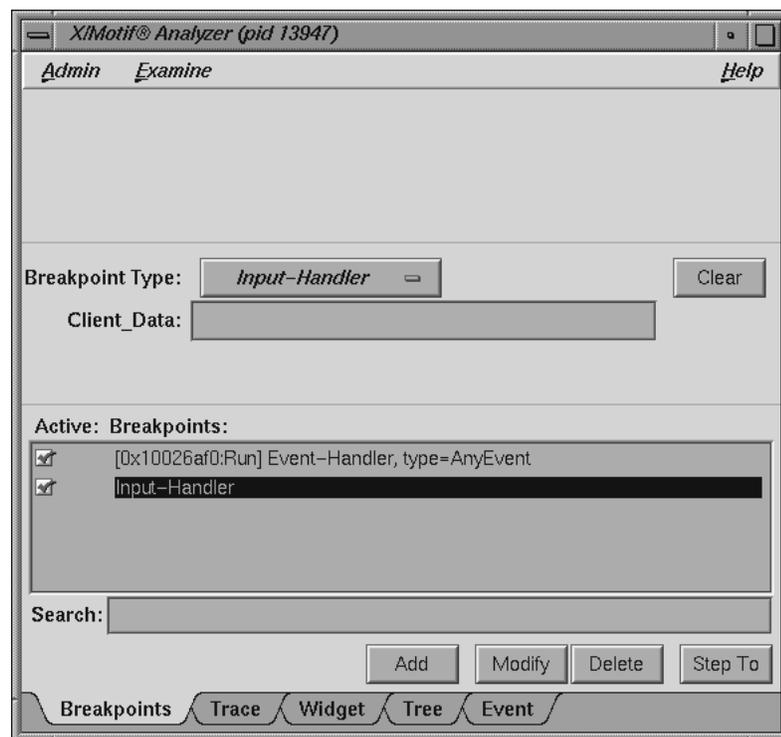


Figure 72. Input-Handler Breakpoints Examiner

The Input-Handler Breakpoints examiner contains the following items:

Breakpoint Type option button	Allows you to select the type of breakpoint you wish to set. In this section, Input-Handler is selected.
Clear button	Clears all the current breakpoint selections and text fields.
Client_Data text field	Allows you to pass in and get back pointer values for the Client_Data.

A.4.2.6 State-Change Breakpoints Examiner

When the State-Change option of the Breakpoint Type option button in the Breakpoint Examiner is selected, the examiner appears as shown in Figure 73.

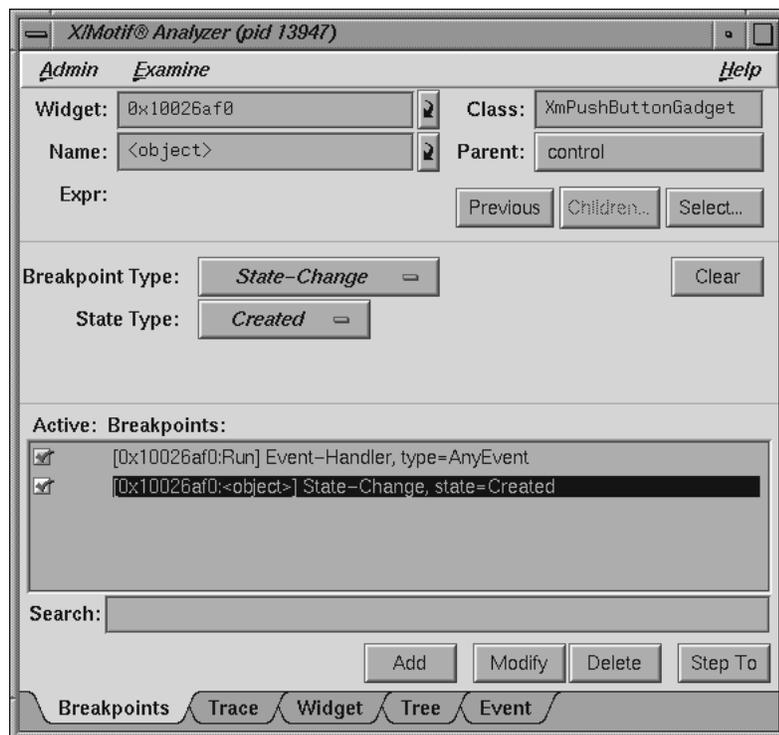


Figure 73. State-Change Breakpoints Examiner

The State-Change Breakpoints examiner contains the following items:

Widget text field	Allows you to choose a widget to examine by entering the widget address.
Name text field	Allows you to choose a widget to examine by entering the widget name.
Class text field	Allows you to choose a widget to examine by entering the widget's class. Leave the field blank or enter Any to select all widgets.
Parent button	Allows you to move the parent of the currently selected widget.
Previous button	Moves you to the previously selected widget.
Children... button	Shows you the widget's children (it is grayed out if the selected widget cannot have children).
Select... button	Allows you to select the widget in the target process.
Breakpoint Type option button	Allows you to select the type of breakpoint you wish to set. In this section, State-Change is selected.
Clear button	Clears all the current breakpoint selections and text fields.
State Type option button	Allows you to set the state change type for a given breakpoint.

A.4.2.7 X-Event Breakpoints Examiner

When you select the X-Event option of the Breakpoint Type option button in the Breakpoint Examiner, the examiner appears as shown in Figure 74, page 238.

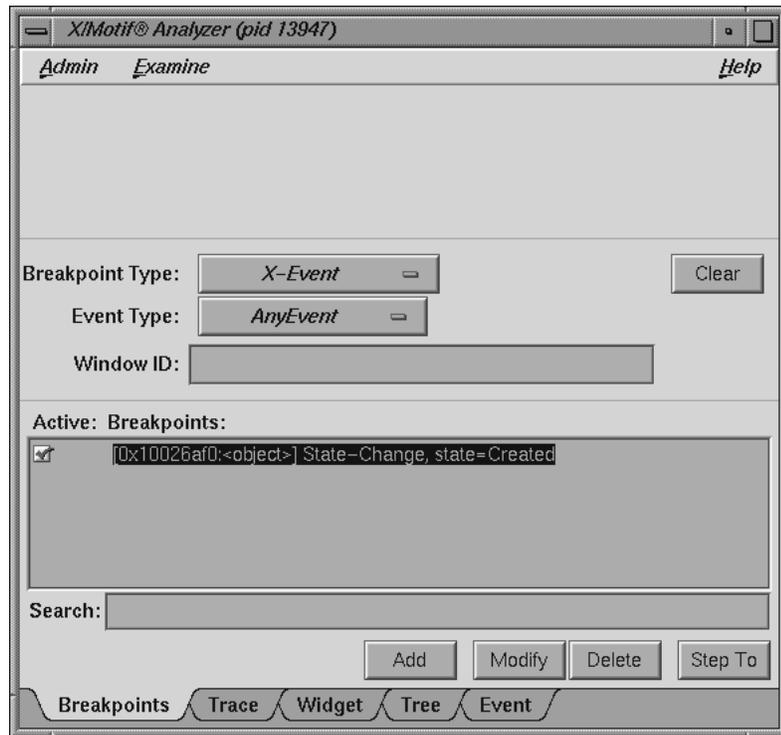


Figure 74. X-Event Breakpoints Examiner

The X-Event Breakpoints examiner contains the following items:

Breakpoint Type button	Allows you to select the type of breakpoint you wish to set. In this section, X-Event is selected.
Clear button	Clears all the current breakpoint selections and text fields.
Event Type button	Allows you to set the event type for a given breakpoint.
Window ID text field	Allows you to set the Window ID value for the breakpoint.

A.4.2.8 X-Request Breakpoints Examiner

When the X-Request option of the Breakpoint Type option button in the Breakpoint Examiner is selected, the examiner appears as shown in Figure 75, page 239.

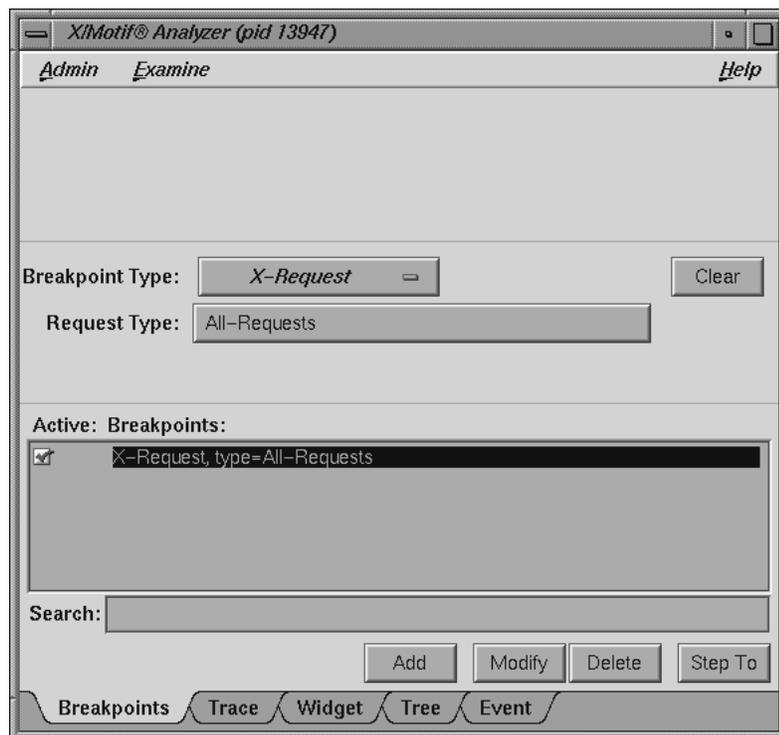


Figure 75. X-Request Breakpoints Examiner

The X-Request Breakpoints examiner contains the following items:

Breakpoint Type option button	Allows you to select the type of breakpoint you wish to set. In this section, X-Request is selected.
Clear button	Clears all the current breakpoint selections and text fields.

Request Type button Launches the Request Type Selection dialog (see Figure 76). This dialog allows you to select the type of X-Request used for your breakpoint. The information displayed is in outline form; selecting a given item selects all its subitems. For example, if you select Window-Category, CreateWindow, ChangeWindowAttributes, GetWindowAttributes, and so on are also selected.

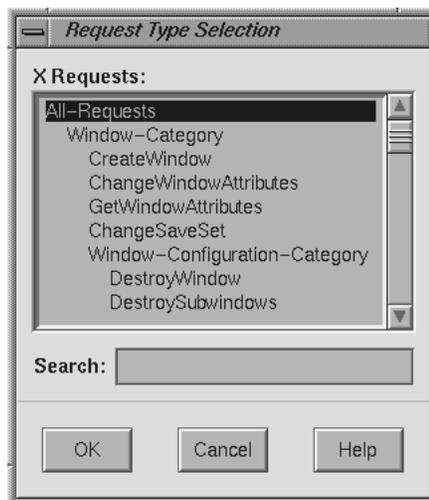


Figure 76. Request Type Selection Dialog

A.4.3 Trace Examiner

The Trace examiner (see Figure 77, page 241) is a control area where you can trace the execution of your application and collect the following data:

- X Server Events
- X Server Requests
- Widget Event Dispatch Information

- Widget Resource Changes (through XtSetValues)
- Widget State Changes (`create`, `destroy`, `manage`, `realize`, and `unmanage`)
- Xt Callbacks (`widget`, `event handler`, `work proc`, `timeout`, `input`, and `signal`)

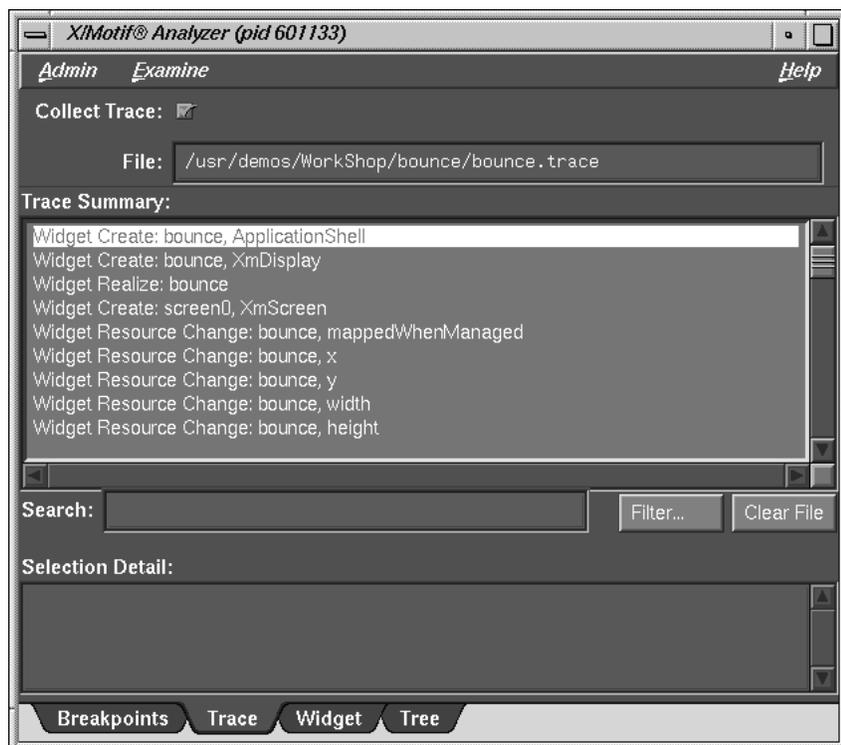


Figure 77. Trace Examiner

The Trace examiner contains the following items:

- | | |
|----------------------|------------------------------------------------------------------------------------------------------------------------------|
| Collect Trace toggle | Allows you to turn the tracing on and off. |
| File text field | Allows you to select the file name for the trace. If no file is selected, a default <i>filename</i> for the trace is chosen. |

Search text field	Allows you to perform an incremental, textual search for the trace list.
Filter... button	Launches a dialog that allows you to select the trace entry types you want displayed in the list.
Clear File button	Erases the trace file. Any subsequent trace information goes to the beginning of the file.

A.4.4 Widget Examiner

The Widget examiner (see Figure 78) displays the internal Xt widget structure, as well as the Xt inheritance implementation using nested C constructs.

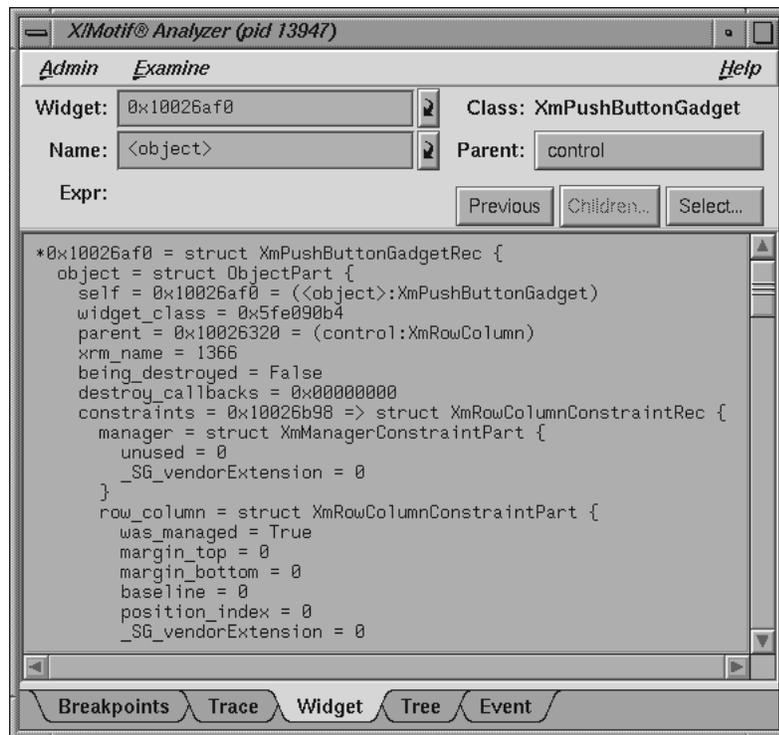


Figure 78. Widget Examiner

The Widget examiner contains the following items:

Widget text field	Allows you to choose a widget to examine by entering the widget address.
Name text field	Allows you to choose a widget to examine by entering the widget name.
Parent button	Allows you to move the parent of the currently selected widget.
Previous button	Moves you to the previously selected widget.
Children... button	Shows you the widget's children. (It is grayed out if the selected widget does not have children.)
Select... button	Allows you to select the widget in the target process.

A.4.5 Tree Examiner

The Tree examiner (see Figure 79, page 244) displays the widget hierarchy.

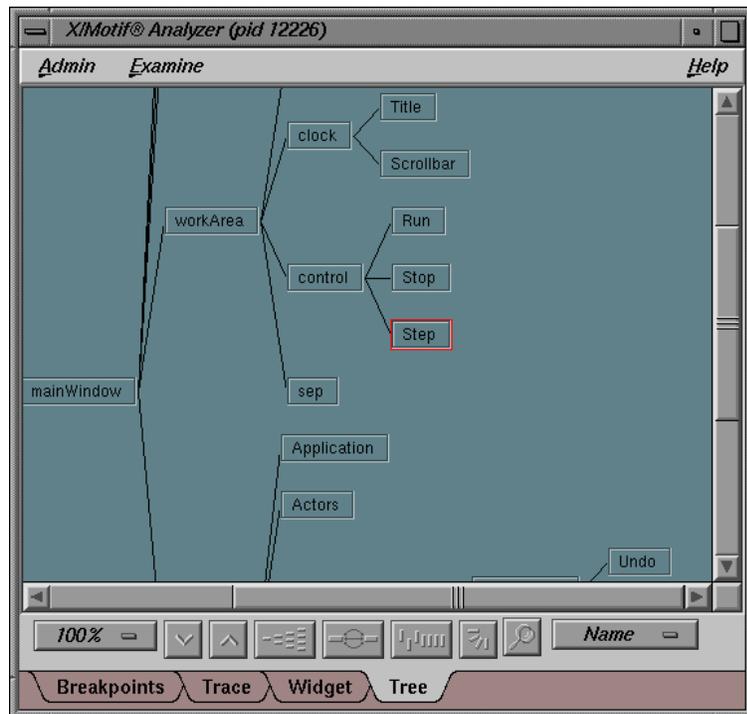


Figure 79. Tree Examiner

You may double-click a node to view that widget in the Widget examiner.

If the Tree examiner is currently selected, it will not automatically fetch the current widget tree each time the process stops. To force retrieval of the widget tree, select another examiner and then go back to the Tree examiner. Or, click on the Tree tab.

The graphical buttons across the bottom of the Tree Examiner window from left-to-right have the following functions.

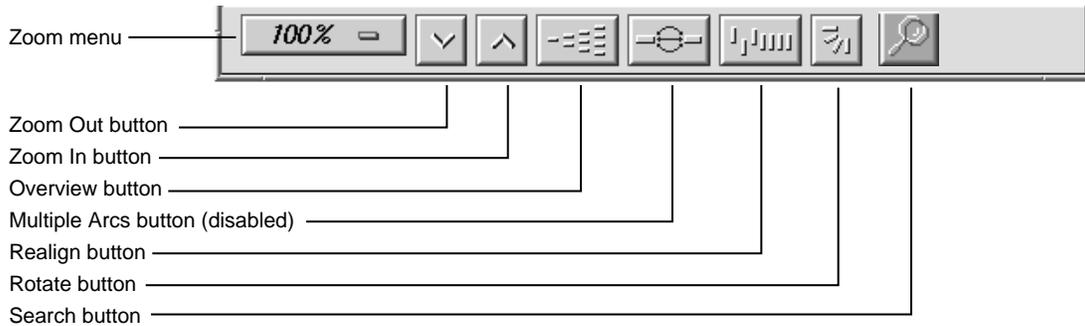


Figure 80. Tree Examiner Window Graphical Buttons

Zoom percentage indicator	Click on this indicator to bring up a submenu from which you can select a different percentage.
Zoom Out button	Decreases the zoom percentage
Zoom In button	Increases the zoom percentage
Overview button	Shows the entire tree structure in another window along with an indication of which portion of the tree is currently being displayed in the full-sized Tree Examiner window.
Multiple Arcs button	This button shows/hides multiple connections between nodes on the graph. For example, if main calls foo several times, you will see a line (arc) for all calls made.
Realign button	This button resets the graph back to its original configuration (before you began working with it). Also, hidden nodes will reappear.
Rotate button	Change the format of the window from up-down to left-right, or vice-versa.
Search button	Search for text in the tree using a Widget Hierarchy Search Text window.
Widget View Type Option button	Click on this indicator to change the kind of information being displayed in the nodes of the tree. This will bring up a submenu which allows you to choose Name, Class, ID (widget address) or Window (indicate <gadget> or address).

A.4.6 Callback Examiner

The Callback examiner (see Figure 81) automatically appears when the process is stopped somewhere in a callback. It first displays the callstack frame. Then it displays information about the widget in the callback. Finally, it displays the proper callback structure contained in the `call_data` argument to the callback procedure, based on the widget type and the callback name.

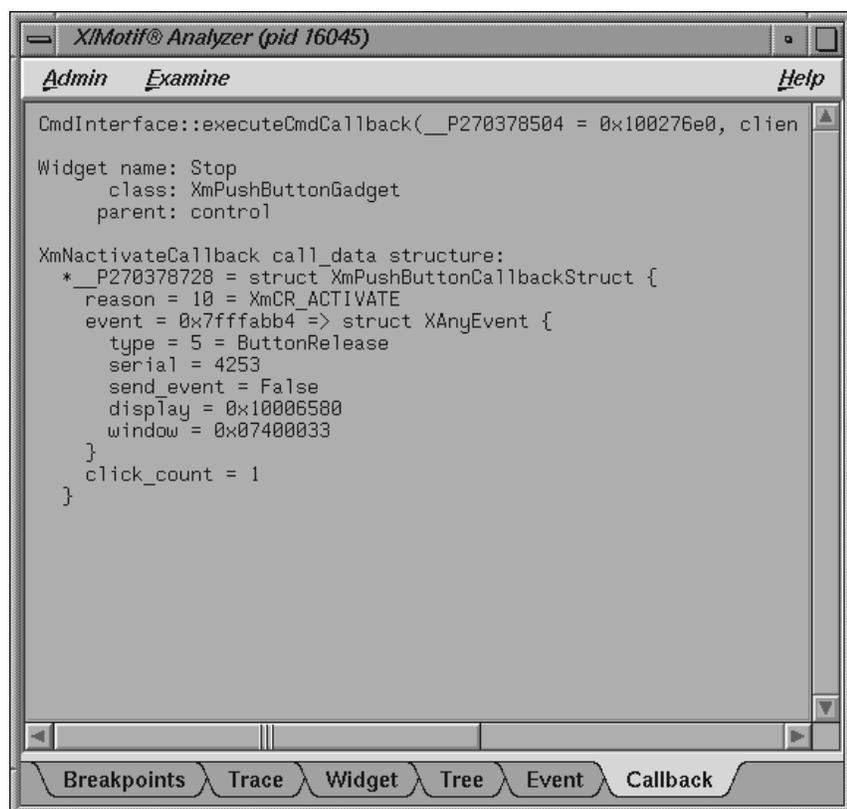


Figure 81. Callback Examiner

A.4.7 Window Examiner

The Window examiner (see Figure 82) displays window attributes for an X window and the parent and children window IDs. These attributes are returned by `XGetWindowAttributes`, with decoding of the visual structure, enums, and masks.

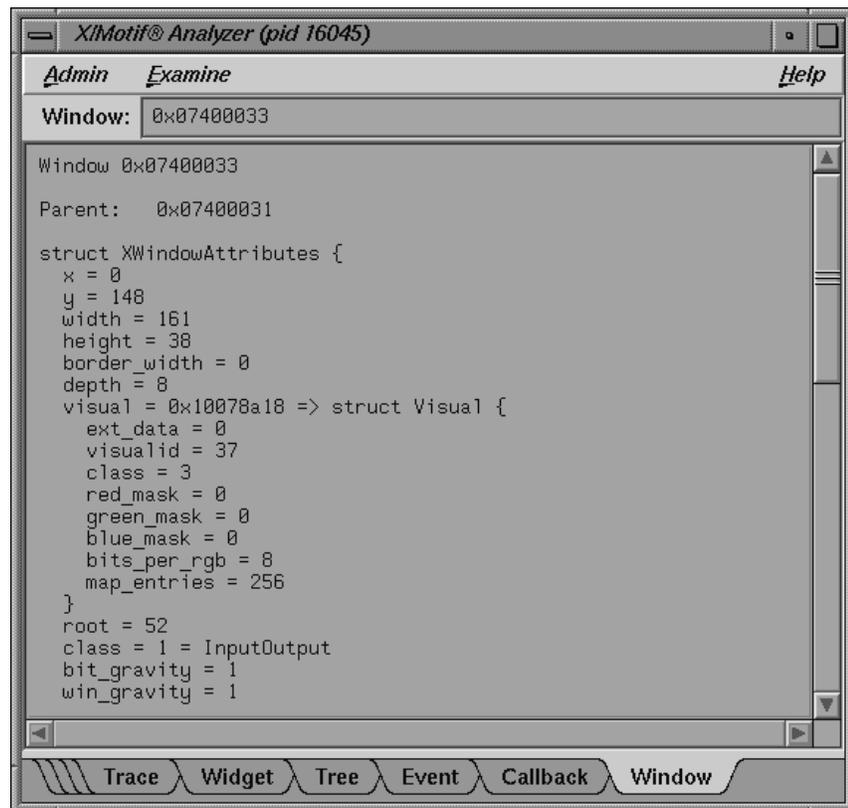


Figure 82. Window Examiner

The Window examiner contains the window text field that displays the address of the window that is being examined. You may change to a different window by entering a new address and pressing Enter.

A.4.8 Event Examiner

The Event examiner (see Figure 83) displays the event structure for an XEvent pointer. The proper XEvent union member is used, and enums and masks are decoded.

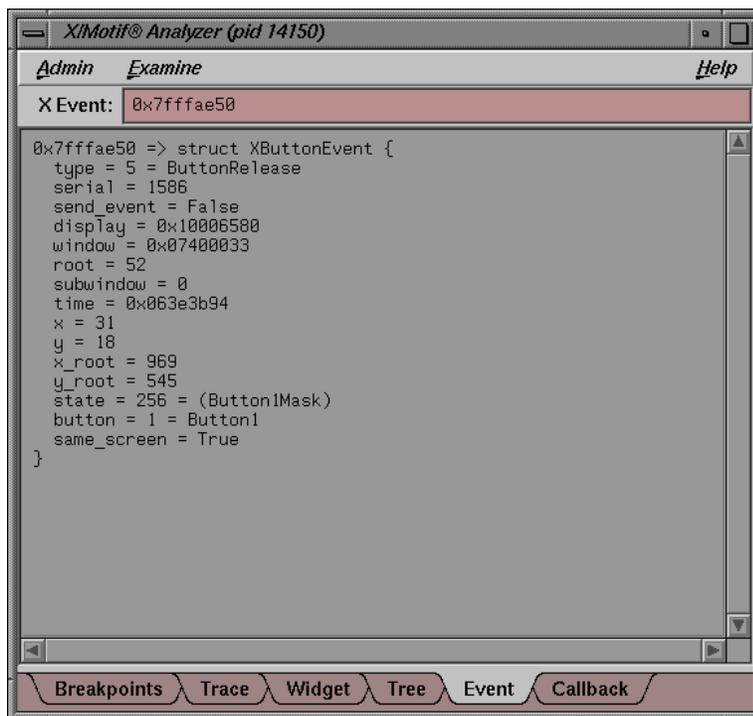


Figure 83. Event Examiner

The Event examiner contains the X Event text field, which displays the address of the X event that is being examined. You may change to a different X event by entering a new address and pressing Enter.

A.4.9 Graphics Context Examiner

The Graphics Context examiner (see Figure 84, page 249) displays the X graphics context attributes that are cached by Xlib in the form of an XGCValues structure. Enums and masks are decoded.

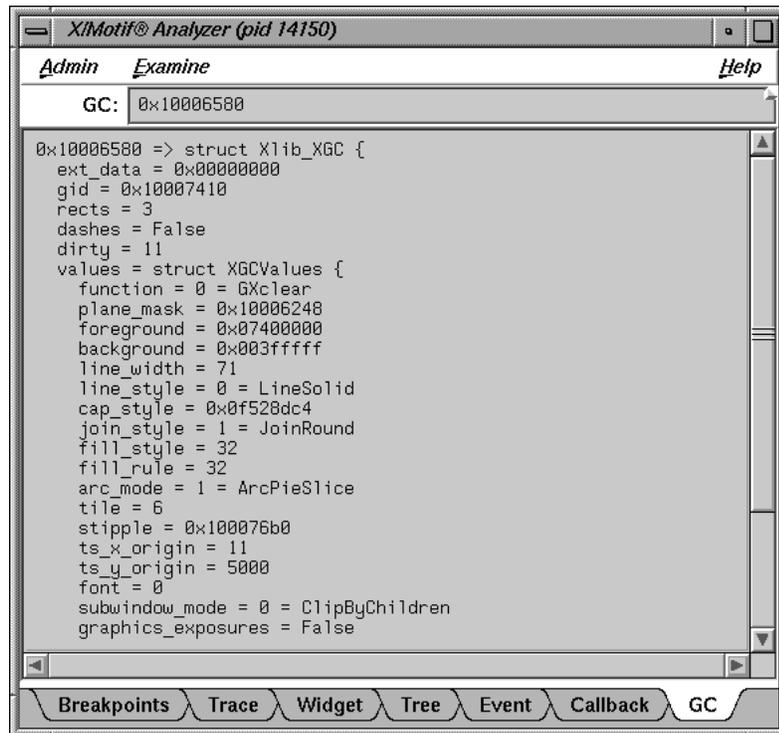


Figure 84. Graphics Context Examiner

The Graphics Context examiner contains the GC text field that displays the address of the graphics context that is being examined. You may change to a different context by entering a new address and pressing Enter.

A.4.10 Pixmap Examiner

The Pixmap examiner (see Figure 85, page 250) displays basic attributes of an X pixmap, like size and depth. It also attempts to provide an ASCII display of small pixmaps, using the units digit of the pixel values.

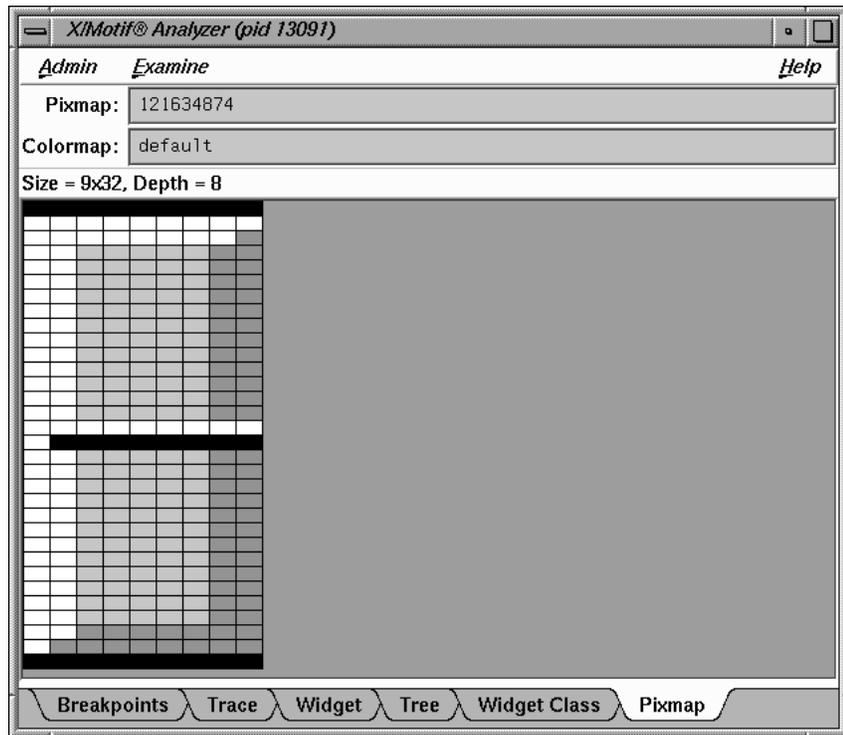


Figure 85. Pixmap Examiner

The Pixmap examiner displays the contents of an X pixmap. To specify an X pixmap identifier, enter a numeric expression in the top text field of the window. Then, use `default` as the colormap identifier to specify the default X colormap for your screen. In the pixmap display, left-click on a pixel to see the pixel value, position, and red-green-blue intensities.

A.4.11 Widget Class Examiner

The Widget Class examiner (see Figure 86, page 251) displays the X_t widget class structure, as well as the X_t inheritance implementation using nested C constructs.

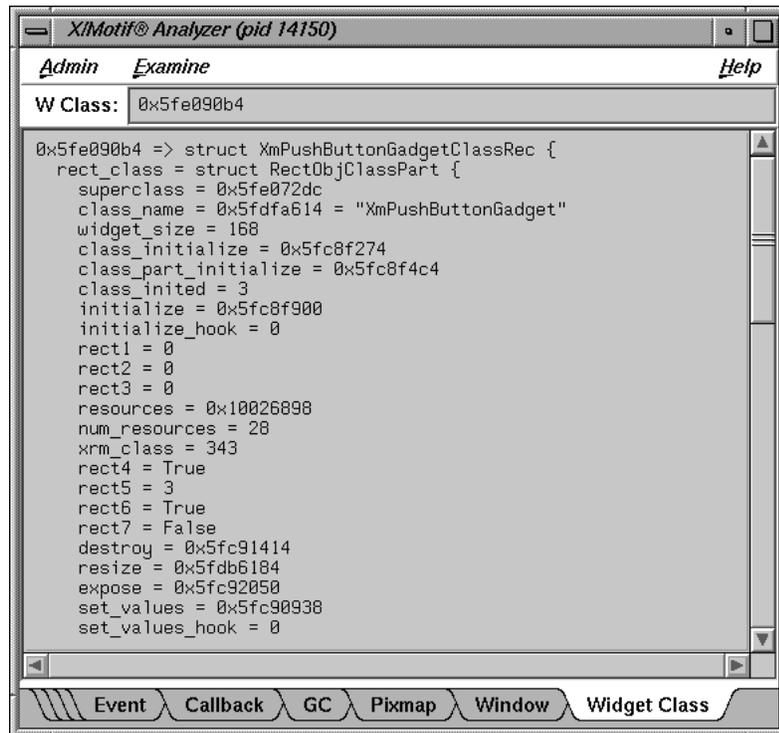


Figure 86. Widget Class Examiner

The Widget Class examiner contains the `W Class` text field, which displays the address of the widget class that is being examined. You may change to a different widget class by entering a new address and pressing `Enter`.

A.5 Trap Manager Windows

In addition to setting traps by using the command line, the `View` menu of the Main View window provides you with three views specific to trap management:

- Trap Manager
- Signal Panel
- Syscall Panel

Call up the Trap Manager window from the Main View window menu bar as follows:

Views
->Trap Manager

A.5.1 Trap Manager

The Trap Manager allows you to set, edit, and manage traps (used in both the Debugger and Performance Analyzer). The X window is shown in Figure 87.

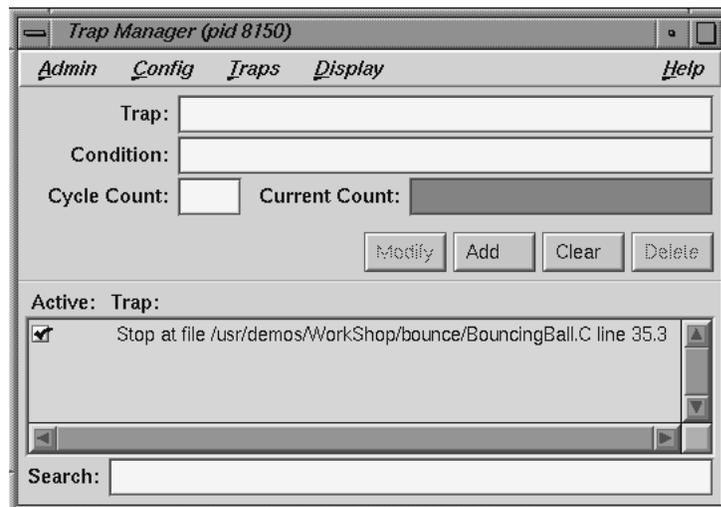


Figure 87. Trap Manager Window

The Trap Manager window contains the following items (besides the menu bar, which is discussed below):

- | | |
|------------------------|-------------------------------------|
| Trap text field | Contains a description of the trap. |
| Condition text field | Contains the condition of the trap. |
| Cycle Count text field | Displays the current cycle count. |

Current Count text field	Displays the current trap count.
Full button	Allows you to toggle between display of full and partial path names.
Modify button	Allows you to change the selected breakpoint's settings.
Add button	Allows you to add a new breakpoint.
Clear button	Clears all the current breakpoint selections and text fields.
Delete button	Deletes the selected breakpoint.
Active label	If selected with a check mark, the trap is enabled.
Trap display area	Contains a description of each trap, and a toggle to indicate whether or not the trap is active.
Search text field	Allows you to perform an incremental textual search for the trap list.

The Trap Manager window has a menu bar which contains the Admin, Config, Traps, Display, and Help menus. The Admin menu is the same as that described in Section A.3.1.1, page 217. The Help menu is the same as that described in Section A.1.10, page 205. The other menus are described in the following sections.

A.5.2 Config Menu

The Config menu contains the following items:

Load Traps...	Brings up the File dialog allowing you to load the traps from a file.
Save Traps...	Brings up the File dialog allowing you to save the current traps to a file.

A.5.3 Traps Menu

The Traps menu has options that allow you to set traps under a number of conditions. The following conditions are available:

At Source Line	Highlight a line in the Main View window's source pane before selecting this option to set a breakpoint at the selected line.
Entry Function	Highlight a function name in the Main View window's source pane before selecting this option to set a breakpoint at the entry to the function.
Exit Function	Highlight a function name in the Main View window's source pane before selection this option to set a breakpoint at the exit from the function.
Stop Trap Default	Causes a trap created to be a "stop" trap. Toggles with the Sample Trap Default.
Sample Trap Default	Causes a trap created to be a "sample" trap. Toggles with the Stop Trap Default.
Group Trap Default	ON/OFF toggle to cause a trap created to have the "pgrp" option.
Stop All Default	ON/OFF toggle to cause a trap created to have the All option.

A.5.4 Display Menu

The Display menu contains the following item:

Delete All	Deletes all traps from the trap list.
------------	---------------------------------------

A.5.5 Signal Panel

The Signal Panel displays the signals that can occur. You can specify which signals trigger traps and which are to be ignored. The Signal Panel is shown in Figure 88, page 255.



Figure 88. Signal Panel

The Signal Panel contains an Admin menu (described in Section A.3.1.1, page 217) and a Help menu (described in Section A.1.10, page 205). Each signal trigger trap in the display has a toggle associated with it. In addition, the panel has a Search text field.

Note: When debugging IRIX 6.5 pthreads, the Signal Panel is inaccessible if more than one thread is active.

A.5.6 Syscall Panel

The Syscall Panel allows you to set traps at the entry to or exit from system calls. The Syscall Panel is shown in Figure 89, page 256.

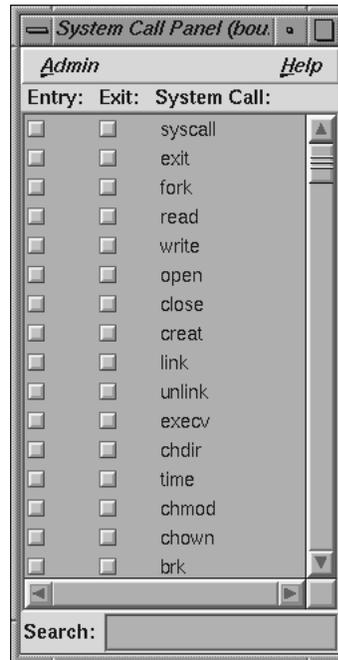


Figure 89. Syscall Panel

The Syscall Panel contains an Admin menu (described in Section A.3.1.1, page 217) and a Help menu (described in Section A.1.10, page 205). Each system call in the display has two toggles associated with it: one to set a trap on entry, one to set a trap on exit. In addition, the panel has a Search text field.

A.6 Data Examination Windows

There are several windows that are used primarily to examine your program's data:

- Array Browser Window, Section A.6.1, page 257
- Call Stack View Window, Section A.6.2, page 271
- Expression View Window, Section A.6.3, page 273

- File Browser Window, Section A.6.4, page 276
- Structure Browser Window, Section A.6.5, page 277
- Variable Browser Window, Section A.6.6, page 288

A.6.1 Array Browser Window

To examine numeric, pointer, or character string data in an array variable, select `Array Browser` from the `Views` menu at a point in the process where the variable is present. The `Array Browser` allows you to view elements in a multi-dimensional array (up to 100 x 100 elements), presented in a spreadsheet and graphically, if desired.

Indexing
Expression

For Fortran, the expression may be an array or a dummy array variable name. If the last dimension of the array is unspecified (*), a subscript value of 1 is assumed initially.

For C and C++, the entry may be an array or a pointer. If pointers are used, the expression is treated as though it were a single element, in which case you need to use the subscript controls to see more than the first element.

The expression used to view an element in the array. It is filled in automatically when you specify an array to view.

The expression supplied is language-specific. It represents the indexing expression used in the language to access a particular element. The subscripts are specified by special indexing variables (\$i, \$j, \$k, and so forth) that can be manipulated in the subscript controls area.

The Subscript Controls area allows you to control which elements in the variable are displayed and allows you to shift the current element. The number of dimensions in the array governs the number of controls that are displayed. A close-up view of the subscript controls area appears in Figure 91.

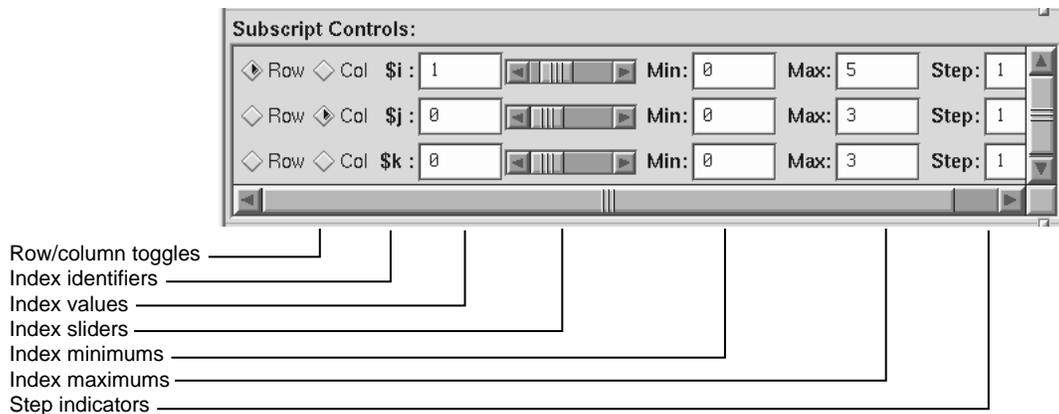


Figure 91. Subscript Controls Area in the Array Browser

The Subscript Controls area provides the following features:

Row/column toggles	Control whether an index variable represents rows or columns (or neither) in the spreadsheet area. You are not limited by the number of dimensions of an array, but you can only view a two-dimensional orthogonal slice of the array at a time.
Index identifiers	Indicate to which subscript the controls in the row refer.
Index values	Show the value of the subscript for the element currently in the focus cell. You can enter a different value if you wish.
Index sliders	Let you move the focus cell along the particular dimension.
Index minimums	Identify the beginning visible element in that particular dimension.
Index maximums	Identify the last visible element in that particular dimension. If you have an unspecified array, you can use this field to specify the last element in the vector to be displayed in the spreadsheet.
Step indicators	Specifies the increment between adjacent elements in the dimension to be displayed. A value of 1 displays consecutive data. Specifying some n greater than 1 allows you to display every n element in a vector.
Control area scroll bars	Allows you to expose hidden portions of the subscript control area if your window is not large enough for viewing all of the controls.

The spreadsheet area is where numeric data is displayed. It can show two dimensions at a time (indicated in the upper left corner of the matrix). The column indexes run along the top of the matrix and the row indexes are displayed along the left column. The spreadsheet area has scroll bars for viewing data elements not currently visible in the viewing area. Figure 92, page 261, shows a close-up of the spreadsheet area.

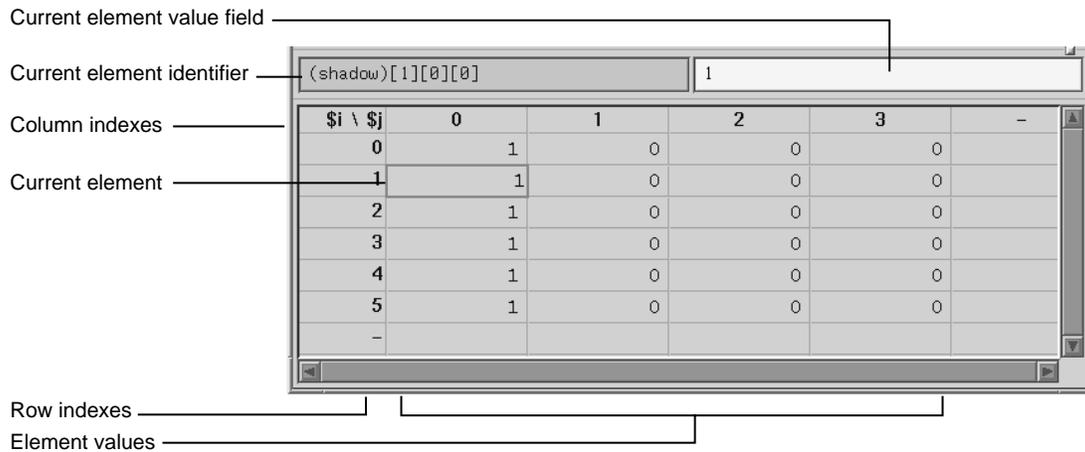


Figure 92. Array Browser Spreadsheet Area

The current element is highlighted by a colored rectangle in the spreadsheet area. Its corresponding expression is shown in the current element identifier field, and the value is shown in the current element value field.

A.6.1.1 Spreadsheet Menu

The Spreadsheet menu allows you to change the appearance of data in the spreadsheet area. It provides these selections:

- Column Width... Allows you to specify the width of the spreadsheet cells in terms of characters. For instance, a value of 12 indicates that 12 characters, including punctuation and digits are viewable.
- Wrapped Display Allows you to display a single dimension of an array wrapped around the entire spreadsheet area. The index value for an element is determined by adding the appropriate row index and column index values.

Figure 93, page 262, shows an example of a wrapped array. There is only one index \$i. The current cell is element 4 in the array (by adding 3 and +1).

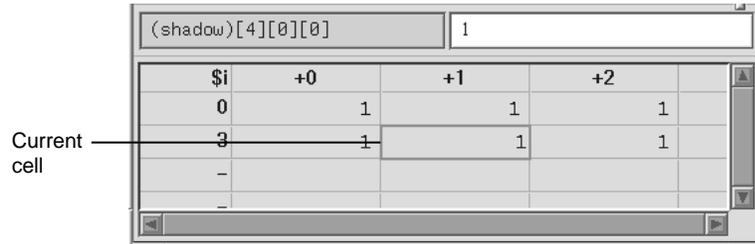


Figure 93. Example of Wrapped Array

A.6.1.2 Format Menu

The **Format** menu displays a separate menu that you allows you to display the elements in the following formats:

- | | |
|----------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Default toggle | Toggles the default format. |
| Value submenu | Contains the following display toggles for formatted values: <ul style="list-style-type: none"> • Decimal • Unsigned • Octal • Hex • Float • Exponential • Char • Wide Character • String |
| Type | Allows listing by data type. |
| Bit Size | Allows listing by bit size. |

The graphical display area presents array data in a three-dimensional graph in one of the following formats:

- Surface (polyhedron)
- Bar chart
- Points
- Multiple lines (array vectors)

A.6.1.3 Render Menu

Note: The Render tear-off menu is available only if you are have a Silicon Graphics workstation with Open Inventor installed on it.

You select the graphical display mode through the Render menu. The Render menu has the following options:

Surface	Exhibits the data as a solid using the data values as vertices in a polyhedron.
Bar Chart	Presents the data values as 3-D bar charts.
Points	Simply plots the data values in 3-D space.
Multi Line	Plots and connects the data values in each row.
None	Allows you to display with no graphical display, in effect turning off graphical display mode.

A.6.1.4 Color Menu

Note: The Color tear-off menu is available only if you are have a Silicon Graphics workstation with Open Inventor installed on it.

If the Color menu is grayed out when the Array Browser window first opens, select the Surface option of the Render menu. The Color menu provides the following options:

Monotone Ramp	Displays the data values in a single tone, with lower numbers being darker and higher values lighter in tone.
Hue Ramp	Displays the data values in a spectrum of colors ranging from blue (lowest values) through green, yellow, orange, and red (highest values).

Exception

Allows you to flag certain conditions by color, usually for the purpose of spotting bad data. When you select Exception, the controls shown in Figure 94, page 264 appear in the window.

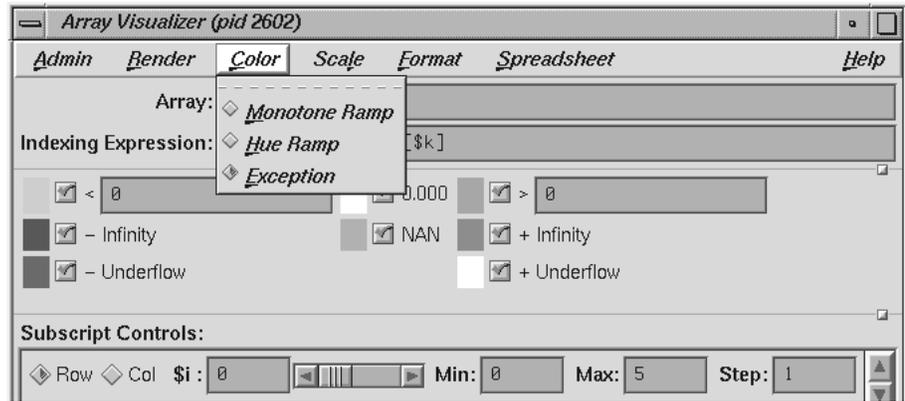


Figure 94. Color Exception Portion of Array Browser Window

Thus, you can highlight data values less than or greater than specified values, values of plus or minus infinity, values of plus or minus underflow, zero values, and NaN (not a number) values.

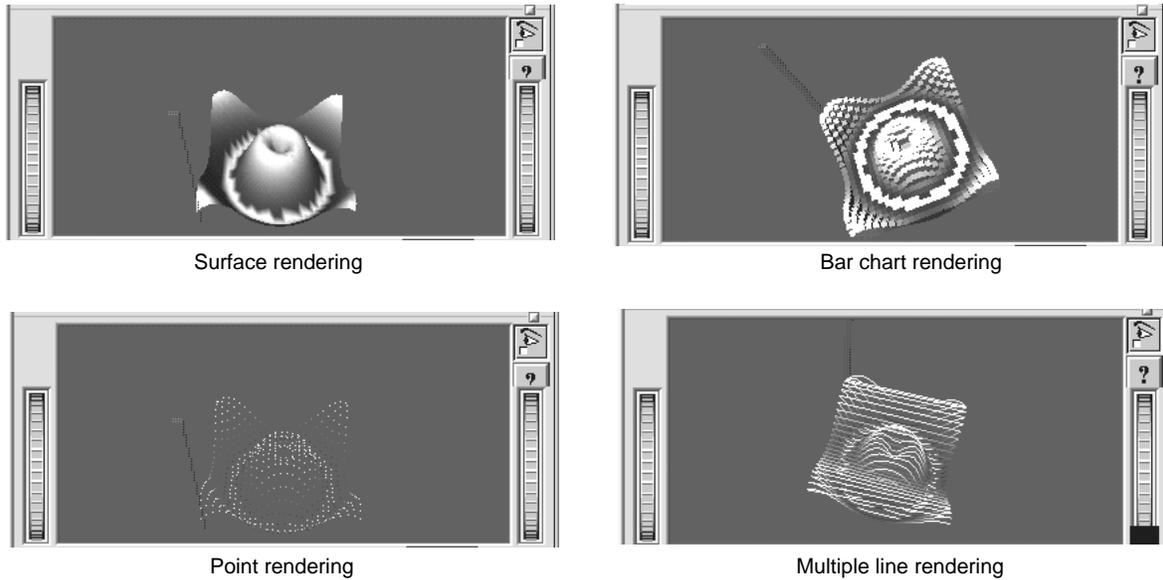


Figure 95. Array Browser Graphic Modes

A.6.1.5 Scale Menu

Note: The *Scale* tear-off menu is available only if you have a Silicon Graphics workstation with Open Inventor installed on it.

If the *Scale* menu is grayed out when the *Array Browser* window first opens, select the *Surface* option of the *Render* menu.

The *Scale* menu provides options for changing the ratio of the *z*-dimension, which represents the value of the element. The number on the left represents the value of the *x* and *y*-dimensions (which are always the same as each other). The number on the right is the *z*-dimension.

Manipulating the *z*-dimension affects the ease of spotting differences in values. If your data is scattered over a narrow range of values, you may wish to heighten the graph by selecting 10:1 as your scale; this exaggerates the values in the *z*-dimension. If your data is in a wide range, selecting 1:2 or 1:10 as the scale will minimize the differences, flattening the graph.

A.6.1.6 Examiner Viewer Controls

Note: The Examiner Viewer is available only if you have a Silicon Graphics workstation with Open Inventor installed on it.

The graphical display uses controls and menus from Examiner Viewer. Examiner Viewer is based on a camera metaphor and borrows terms from the film industry, such as *zoom* and *dolly*, in naming its controls. The graphical display area of the window is shown in Figure 96, page 267, with its main controls and menus. Note that the buttons on the upper right side of the graphical display area may not be visible if the area is too small; you can expose them by moving either the upper or lower sash to enlarge the display area. (The lines between the window panes include a small box to the right. Click-drag this box to change the size of the panes.)

Examiner Viewer provides these controls for viewing the graph. The right side buttons provide the following actions:

view mode Toggles between a view-only mode (closed eye) and manipulation mode (open eye).

In view-only mode, the cursor appears as an arrow and the graph cannot be moved. Clicking on a portion of the graph selects the corresponding array element in the spreadsheet.

In manipulation mode, the cursor appears as a hand and you can move the graph. Dragging the graph with the left mouse button down moves the graph in any direction as if it were in a trackball; a quick movement spins the graph. Dragging the graph with the left mouse button and the `Ctrl` key rolls (rotates) the graph in the plane of the screen. Dragging the graph with the middle mouse button moves it without changing the viewing angle.

If you drag the graph with both the left and middle mouse buttons down, the graph will appear to move into or out of the window (this is the same as the *dolly* thumbwheel, which is described in this section).

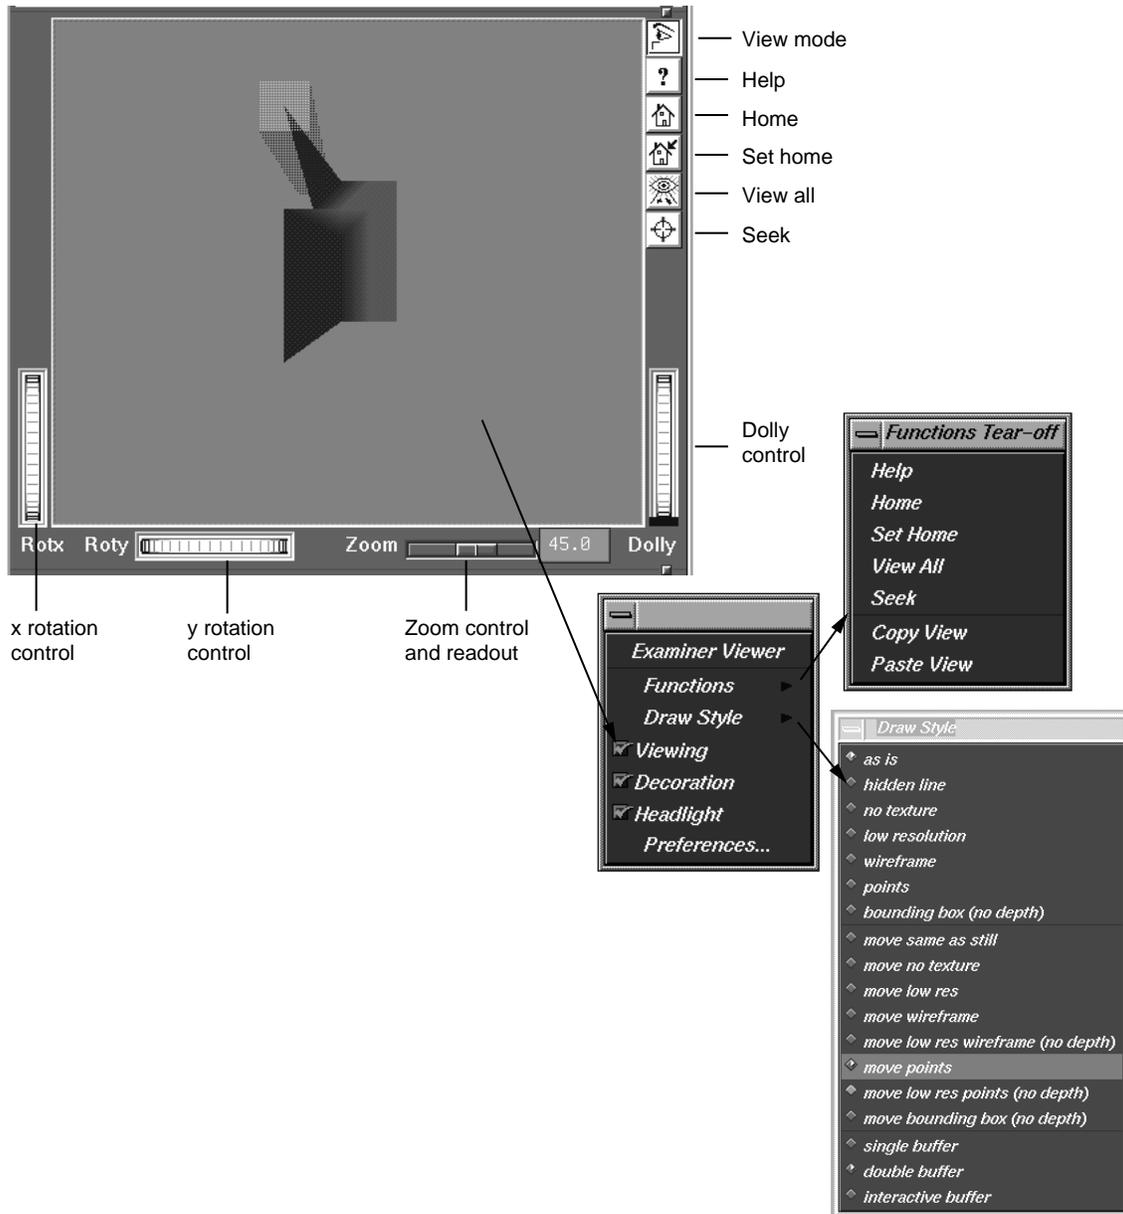


Figure 96. Examiner Viewer with Controls and Menus

help	Runs a special help system containing Inventor Viewer information.
home	Repositions the graph in its original viewing position.
set home	Changes the home (original viewing) position for subsequent use of the home button.
view all	Repositions the display area so that the entire graph is visible.
seek	Provides a special cursor that allows you to reposition the graph in the center of the display area or allows you to center the view on a point you select with the cursor. See <i>Seek to point <or object></i> in the Preferences dialog box.

The following controls let you move the graphic display:

x rotation thumbwheel	Rotates the graph around its x-axis.
y rotation thumbwheel	Rotates the graph around its y-axis.
dolly thumbwheel	Changes the size of the graph and adjusts the angles to maintain perspective. The dolly control simulates moving the viewing camera back and forth with respect to the graph.

A.6.1.7 Examiner Viewer Menu

You access the Examiner Viewer menu by holding down the right mouse button in the graphical display area. The Examiner Viewer menu provides the following options (see Figure 96, page 267 for illustration):

Functions	Displays a submenu with the selections Help, Home, Set Home, View All, and Seek, which are the same as the right mouse button controls described in the previous section, and the Copy View and Paste View selections. These operate like standard copy and paste editing commands, enabling you to transfer graphs.
-----------	----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------

Draw Style	<p>Displays a submenu that controls how the graph is displayed. The top group of options, from <code>as is</code> to <code>bounding box (no depth)</code> control how the graph is displayed when it is static. These override any Render menu selections.</p> <p>The middle (<code>move...</code>) group of options control how the graph is displayed while in motion.</p> <p>The last three option, <code>single</code>, <code>double</code>, and <code>interactive</code>, refer to buffering techniques used in moving the graph. These affect the smoothness of the movement.</p>
Viewing	<p>The same as the <code>view mode</code> button described in the previous section. When it is off, you can select points from the graph to display in the spreadsheet but cannot move the graph. When on, it allows you to manipulate the graph.</p>
Decoration	<p>Displays the right side buttons when it is on and hides them when it is off.</p>
Headlight	<p>Controls the shadow effect on the graph. When it is on, the light appears to come from the camera.</p>
Preferences	<p>Causes the <code>Examiner Viewer Preferences Sheet</code> dialog to display.</p>



Figure 97. Examiner Viewer Preference Sheet Dialog

The Examiner Viewer Preference Sheet dialog provides the following options:

Seek animation time	Allows you to specify the time it takes for the graph to be repositioned after you change the seek point. Set to 0 for instant seek. See also Seek to point <or object>.
Seek to point <or object>	Seek to point uses the picked point and surface normal to align the camera. Seek to object uses only the object center to align the camera.
Seek distance wheel	This wheel controls how close to the camera the object will appear. This distance can be either an absolute distance or a percentage of the distance to the picked point.

Camera Zoom slider	This slider allows you to set (in degrees) the camera height angle (only perspective camera). If you wish, you may set this value in the field immediately to the right of the slider.
Zoom slider ranges from	Range is 1.0 to 140.0 by default, but you can reset this if you wish.
Auto clipping planes	Centers the graph in your view if enabled. If disabled, it allows you to move the graph out of view at either end of the z-axis. This is useful if you wish to focus on data above or below a set value.
Stereo Viewing	This may be turned on to see the scene in stereo (special glasses required). The offset between the left and right eye can be specified using the camera rotation thumbwheel.
Enable spin automation	When selected, you can cause the camera to continue spinning. To animate, left-click and drag in the direction of your desired spin, then release while spinning. The graph spins as if by a trackball. To stop, left-click anywhere.
Show point of rotation axes	Displays a set of three axes. You can move the graph around the x and y axes using the thumbwheel controls described in the previous section. When this option is on, you can set the size of the axes in pixels.

A.6.2 Call Stack View Window

The Call Stack View (Figure 98, page 272) window displays call stack entries when a process has stopped.

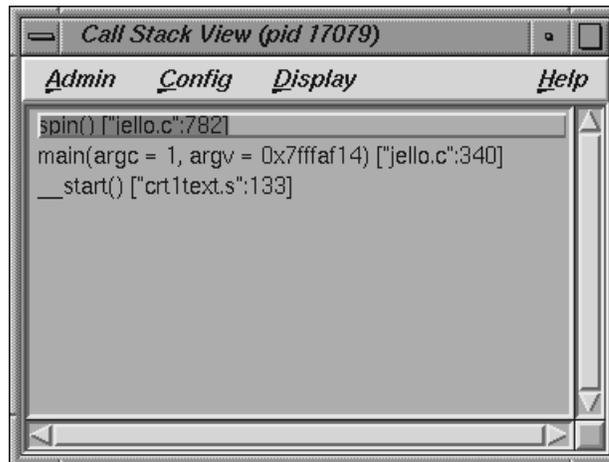


Figure 98. Call Stack View

The source display in the Main View window has two special annotations that are synchronized with the Call Stack View:

- The location of the current program state is indicated by a large arrow representing the PC. The source line to which the arrow points is highlighted (usually in green).
- The location of the call to the item selected in the Call Stack View window is indicated by a small arrow representing the current context. The source line becomes highlighted (usually in blue-green).

The Call Stack View contains its own menu bar, which contains the *Admin*, *Config*, *Display*, and *Help* menus. The *Admin* menu is the same as that described in Section A.3.1.1, page 217. The *Help* menu is the same as that described in Section A.1.10, page 205. The other menus are described in the following sections.

A.6.2.1 Config Menu

The Config menu contains the following option:

Preferences...	Launches the Call Stack View Preferences dialog that allows you the option of setting the maximum depth of the Call Stack View.
----------------	---------------------------------------------------------------------------------------------------------------------------------

A.6.2.2 Display Menu

The Display menu contains the following toggles which change what is displayed in each entry of the Call Stack View in addition to the subroutine/function name:

Arg Values	Allows you to display argument values. Default is on.
Arg Names	Allows you to display argument names. Default is on.
Arg Types	Allows you to display argument types. Default is off.
Location	Allows you to display function location. Default is on.
PC	Allows you to display the program counter (PC). Default is off.

A.6.3 Expression View Window

The Expression View window is shown in Figure 99, page 274. Expression View displays a collection of expressions that are evaluated each time the process stops or the context changes.

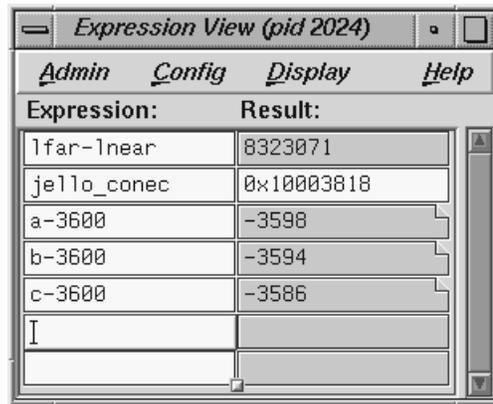


Figure 99. Expression View

In addition to the items on the menu bar, the window has two pop-up menus: the Language menu and the Format menu. The Admin menu is the same as that described in Section A.3.1.1, page 217. The Help menu is the same as that described in Section A.1.10, page 205. The other menus are described in the following sections.

A.6.3.1 Config Menu

The Config menu contains the following options:

- | | |
|---------------------|---------------------------------------------------------------------------------------------------------------------|
| Load Expressions... | Launches the Load Expressions dialog box that allows you to choose source file from which to load your expressions. |
| Save Expressions... | Launches the Save Expressions dialog box that allows you to choose a file to which you can save your expressions. |

A.6.3.2 Display Menu

The Display menu contains the following option:

Clear All Clears all fields in the view.

A.6.3.3 Language Pop-up Menu

The Language pop-up menu contains three buttons that allow you to select one of three languages for evaluation: C, C++, or Fortran. The Language pop-up is invoked by holding down the right mouse button while the cursor is in the Expression column.

A.6.3.4 Format Pop-up Menu

The Format pop-up menu is displayed by holding down the right mouse button in the Result column.

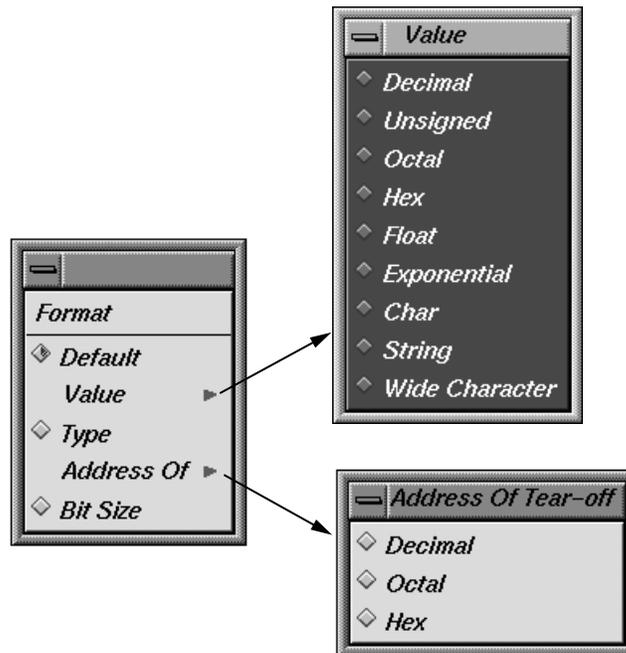


Figure 100. Expression View Format Pop-up with Submenus

The Format popup contains the following options:

Default	Sets the format to the default values.
Value	Displays a submenu from which you can select a value of Decimal, Unsigned, Octal, Hex, Exponential, Float, Char, String, or Wide Character type.
Type	Displays a submenu from which you can select a type of Decimal, Octal, or Hex.
Bit size	Sets the format to Bit Size.

A.6.4 File Browser Window

The File Browser window displays a list of source files used by the current executable. Double-click on a file in the list to load it into the source display area in the Main View or Source View windows. Some files may be listed due to subroutines/functions being resolved from system libraries. If you select such a file, you may get the following message:

```
Unable to find file <xxx.c>
```

This is because the source for system library routines may not be stored on your system.

The Search field allows you to find files in the list.

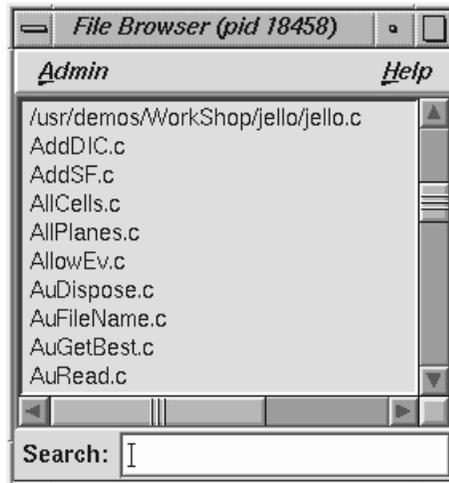


Figure 101. File Browser Window

The File Browser contains an Admin menu (described in Section A.3.1.1, page 217) and a Help menu (described in Section A.1.10, page 205). In addition, the browser has a Search field.

A.6.5 Structure Browser Window

The Structure Browser window allows you to examine data structures and the relationships of the data within them. It displays complex data structures as separate graphical objects, using arrows to indicate relationships. A sample Structure Browser is shown in Figure 102, page 278, with the Config, Display, Node, and Format menus displayed.

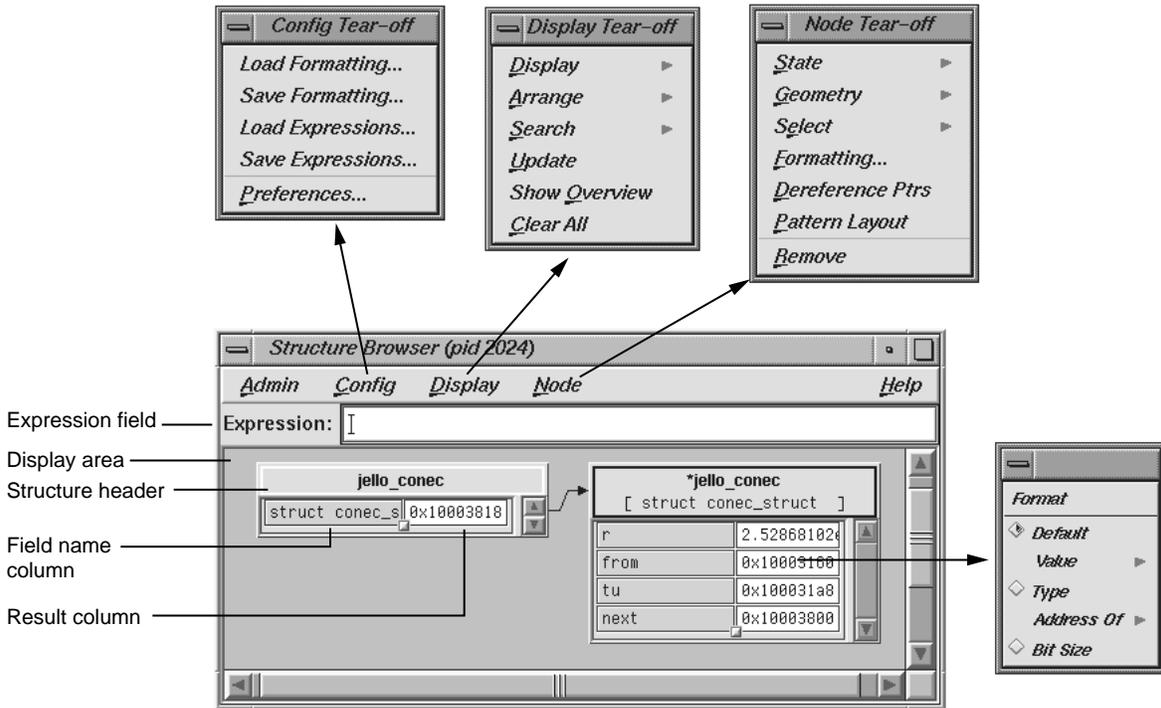


Figure 102. Structure Browser with Menus Displayed

The structure name is entered in the Expression field. It then appears as an object or set of objects in the display area in the lower portion of the window. Each structure has a header identifying the structure, color coded by data type. Below the header are two columns: the left displays the field name and the right displays the field's value. If a displayed structure exceeds the size of the Structure Browser window, scroll bars appear.

While the Admin menu contains options for selecting the active structure, cloning structures, saving them as text, and closing the Structure Browser window; the four additional menus contain options that allow you to change the way data is displayed. The following data-display menus are available:

- Config** Provides options for saving and reusing type-specific formats and expressions. You can also set preferences regarding how objects of a given type are to be displayed.

Display	Provides display options for all objects in the display area.
Node	Provides options for selected objects in the display area only.
Format	Allows you to change or reformat a specific value in the result column. To access this pop-up menu, hold down the right mouse button while the cursor is over the result column.

A.6.5.1 Using the Structure Browser Overview Window to Navigate

WorkShop provides the Structure Browser Overview window (from the Show Overview option in the Display menu) as another way to navigate around the display.

This window is a reduced-scale view of the requested structures. The structures are represented by solid rectangles color-coded by data type. The relative position of the currently visible area is represented by a transparent rectangle. This rectangle can be dragged (using the left mouse button) to change the display of the Structure Browser. Clicking the left mouse button in an area of this window repositions the currently visible area.

A.6.5.2 Entering Expressions

The Structure Browser accepts any valid expression. If the result type is simple, a structure displays showing the type and value. If the result type is a pointer, it is automatically de-referenced until a non-pointer type is reached. If the result type is a structure or union, an object is displayed showing the structures' fields and their values. After the expression is entered, the Expression field clears. The Structure Browser can display unrelated structures at the same time; you simply enter new structures by using the Expression field.

The Expression field is also used to enter strings used in searches.

A.6.5.3 Working in the Structure Browser Display Area

Within the display area, you select objects by clicking in the node headers. Shift-clicks add the selected object to the current selection. You can drag selected objects using the middle mouse button.

Clicking the right button while the cursor is in the right column of an object displays the `Format` menu, which is used to change the display. You can set a default format or request that results be displayed by value, type, address, or size in bits.

Holding down the right button in the header of an object brings up the `Node` pop-up menu, which is the same as the `Node` menu in the menu bar. It is used to change the way selected objects are displayed. When you left-click in the header of an object, it turns on the resizer, which allows you to change the size of the object. You will see a small square (handle) at the upper-right and lower-left corners. Left-click-drag the handle to resize the object. Middle-click-drag the handle to move the object.

Arrows show relationships among the displayed structures. If a member field is not visible in a structure, its arrow tail is displayed at the top or bottom of the scrolling area for fields. Otherwise, its tail is adjacent to its field.

Double-clicking a value field (right column) for a pointer changes the display so that the data structure it points to is displayed.

Double-clicking a member field (left column) puts the full expression for that member in the `Expression` field.

A.6.5.4 Structure Browser Display Menu

The `Display` menu controls the way structures appear in the display area. The `Display` menu provides the following options:

`Display`

Determines contents of the display. The `Display` option has the following two options:

- `Expression` — Displays the structure of the expression entered in the `Expression` field.
- `Selection` — Displays the structure based on the text you have selected in the source code pane in the Main View window.

`Arrange`

Rearranges the currently selected nodes.

`Arrange` has the following two options: (See Figure 103, page 281.)

- `Tree` — Arranges nodes into a tree-type formation, that is, the hierarchy descends from left to right and child structures are shown as branches to the right of the parent.

- **Linked List** — Arranges nodes into a linked list formation, that is, horizontally.

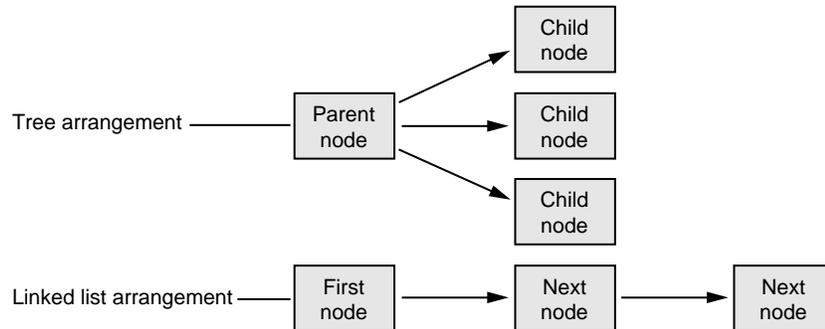


Figure 103. Tree and Linked List Arrangements of Structures

Search

Allows you to select structures in the display area that contain the *string* specified in the Expression field. Search has the following four options:

- **Name** — Selects structures whose names contain the specified string.
- **Type** — Selects structures whose types contain the specified string.
- **Field Name** — Selects structures that have a field whose name contains the specified string.
- **Value** — Selects structures that have a field value containing the specified string.

Update

Explicitly updates the displayed structures. This happens automatically in the current Structure Browser when the process stops. This can be used in an inactive Structure Browser to update it. It can also be used to update the display after changes have been made in other Debugger views.

- | | |
|---------------|--------------------------------------------------|
| Show Overview | Brings up the Structure Browser Overview window. |
| Clear All | Clears all structures from the display area. |

A.6.5.5 Node Menu

The Node menu gives you options that apply to currently selected objects.

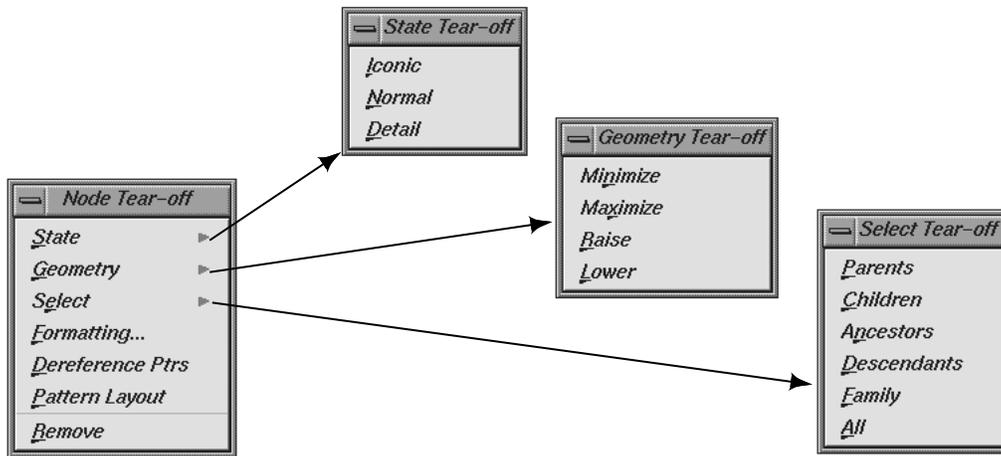


Figure 104. Node Menu

Descriptions for these menu items are as follows:

State

Controls the display of nodes. There are three options:

- **Iconic** — Displays the node header only.
- **Normal** — Uses the default chart display but hides those fields selected to be invisible.
- **Detail** — Uses the default chart display and shows all fields.

Geometry

Manages graphical objects in the display area. There are four options:

- `Minimize` — Sets the vertical size of an object to the default minimum number of fields. The initial default is four fields but can be changed through either the `Formatting` selection from the `Node` menu or the `Preferences...` selection from the `Config` menu.
- `Maximize` — Displays the object as large vertically as necessary to fit all of the fields.
- `Raise` — Raises the selected object(s) to the top of the display.
- `Lower` — Lowers the selected object(s) to the bottom of the display.

Select

Allows you to select objects in various ways. There are six options:

- `Parents` — Selects all objects that have pointers pointing to a selected object.
- `Children` — Selects all objects pointed to by any fields in a selected object.
- `Ancestors` — Selects all objects pointed to a selected object or pointing to an object that has a descendant pointing to a selected object.
- `Descendant` — Selects all objects pointed to by any fields in a selected object or pointed to by any children of a selected object.
- `Family` — Selects all ancestors and descendants of a selected object.
- `All` — Selects all objects.

Formatting	Brings up the type formatting dialog for this type.
Dereference Ptrs	Dereferences any pointers in selected objects.
Pattern Layout	Displays selected structures that are connected by pointers to position related structures in the same way.
Remove	Removes selected object from the display.

A.6.5.6 Formatting Fields

Each field in a data structure has certain display characteristics. These can be specified for all objects in the Structure Browser Preferences dialog box or for type-specific objects only in the Structure Browser Type Formatting dialog box. To display the Structure Browser Preferences dialog box, select Preferences... from the Config menu (see Figure 105).

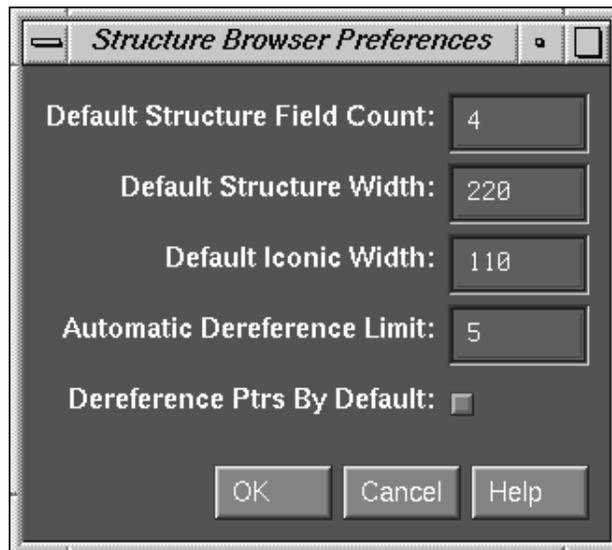


Figure 105. Structure Browser Preferences Dialog

The dialog has the following fields:

Default Structure Field Count	Sets the number of fields to be displayed initially.
Default Structure Width	The width in pixels of the object.
Default Iconic Width	The width in pixels of the object when it is in iconic form.
Automatic Dereference Limit	Limits the number of structures that are automatically dereferenced.
Dereference Ptrs By Default	Toggles automatic dereferencing on and off.

To bring up the Structure Browser Type Formatting dialog box, select the set of structures under consideration and select Node Formatting from the Node menu (see Figure 106, page 286).

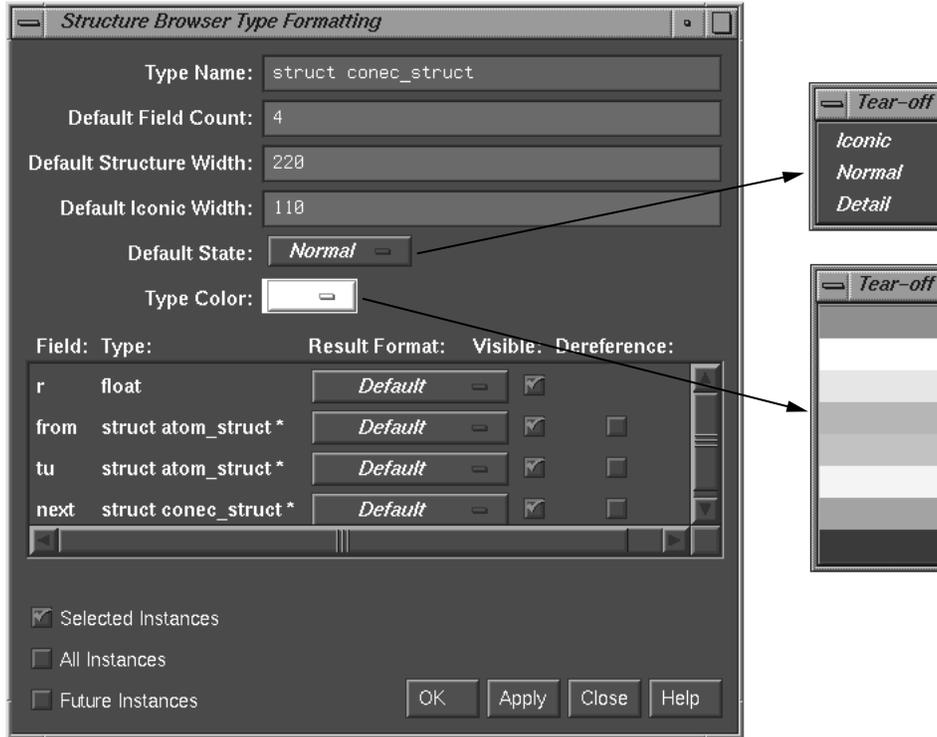


Figure 106. Structure Browser Type Formatting Dialog

The dialog box has the following fields:

Type Name	Displays the current data type.
Default Field Count	Lists number of fields to be displayed initially for objects of that type.
Default Structure Width	Displays the width in pixels of the object.
Default Iconic Width	Displays the width in pixels of the object when it is in iconic form.
Default State	Brings up a pop-up menu that allows you to specify whether structures are first displayed as icons (Iconic), with the minimum number of fields displayed (Normal) or with all fields displayed (Detail).

Type Color

Provides a submenu for color coding. It allows you to select a color for the header and overview rectangles for objects of a given type.

For structure and union types, the list box shows all the fields with their types. For each field, you can change the result format to one of the following types:

- Default
- Decimal
- Unsigned
- Octal
- Hex
- Float
- Exponential
- Char
- String
- Wide Character
- Type
- Dec addr
- Oct addr
- Hex addr
- Bit Size

You can also specify whether a field is visible in normal state, and if it is a pointer field, whether it should be automatically dereferenced.

Once you specify the format for this type, you can apply it to any combination of the following through the toggle buttons in the bottom left portion of the window:

- Selected instances
- All existing instances
- Any future instances of this type

A.6.6 Variable Browser Window

The Variable Browser window allows you to view and change the values of local variables and arguments at a specific point in a process. (Global variables can be viewed or changed using Expression View or the Evaluate Expression selection from the Data menu for one-shot evaluations.) In addition to providing values, the Variable Browser is useful for getting a quick list of the local variables in a scope without having to search for their names. A sample Variable Browser window with the Language and Format menus displayed is shown in Figure 107, page 289.

Typically, you inspect variable values at the following points:

- At a breakpoint
- At a frame in a call stack
- As you step through a process

A useful technique is to set a trap at the entry to a function and inspect the values of the variables there. Some variables may be in an uninitialized state at that point. You can then step through the function and make sure that no uninitialized variables are used inadvertently.

Note: In programs compiled by using the `-n32 -g` compiler options, you cannot click on or print an unused variable as you could when using the `-o32 -g` options.

A.6.6.1 Entering Variable Values

The Variable Browser allows you to change the values of variables in the window. You simply enter the new value in the result column and press Enter. Thus, you can force new values into the process and see their effect.

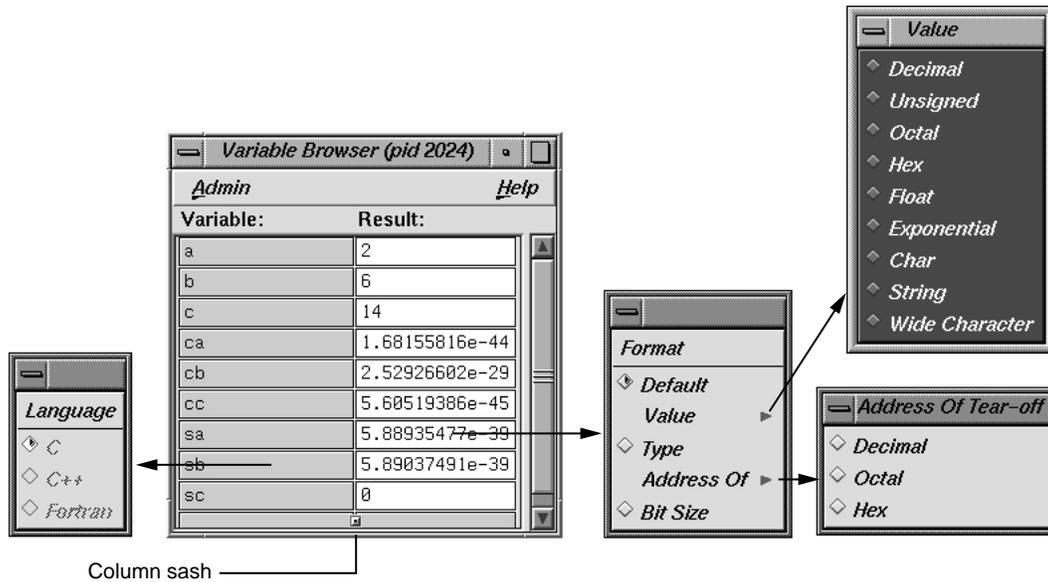


Figure 107. Variable Browser with Menus Displayed

A.6.6.2 Changing Variable Column Widths

By using the sash between the columns, you can adjust the relative widths of the Variable and Result columns (see Figure 107, page 289). For example, you may wish to adjust for short variable names and long result values.

A.6.6.3 Viewing Variable Changes

The Debugger views that are involved with variables (that is, the Variable Browser and Expression View) have indicators that show when the variable has changed since the last breakpoint. If you click the indicator, you can view the previous value. The variable change indicators for a Variable Browser window are shown in Figure 108, page 290.

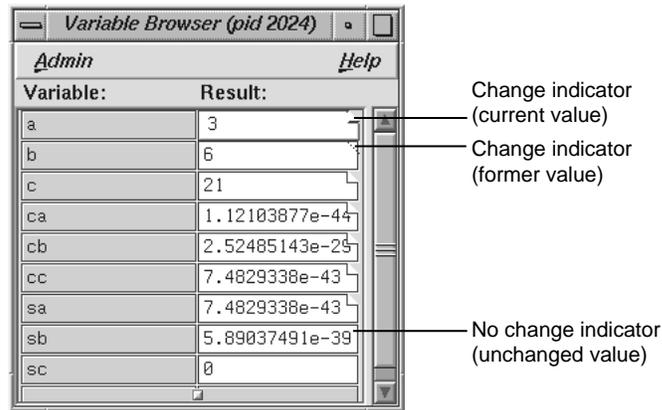


Figure 108. Typical Variable Change Indicators

A.7 Machine-level Debugging Windows

The Debugger offers three views useful in debugging at the machine level: the Disassembly View, Register View, and Memory View.

A.7.1 The Disassembly View Window

The Disassembly View window allows you to look at machine-level code rather than source-level code. A typical Disassembly View window is brought by selecting the following from the Main View window menu bar:

```
Views
->Assembly View
```

The window is shown in Figure 109, page 291, with the Disassemble menu displayed.

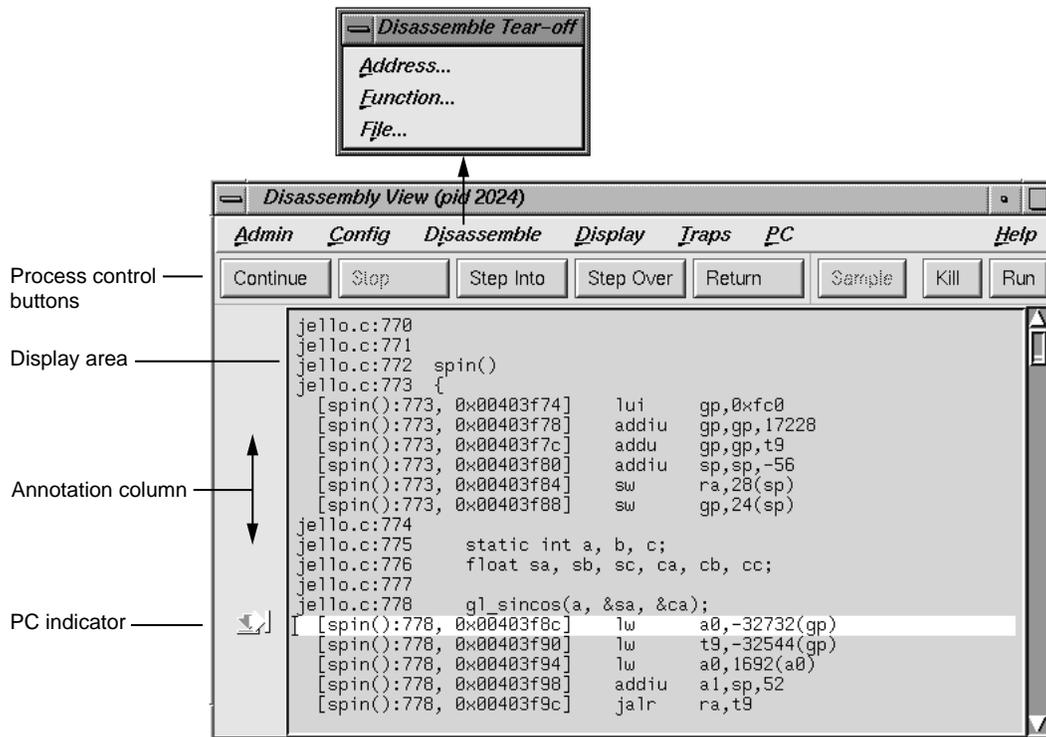


Figure 109. The Disassembly View Window with the Disassembly Menu Displayed

A.7.1.1 Similarities with Main View Window

At the top of the window are the same process control buttons as those in the Main View window. They behave the same way except for *Step Into* and *Step Over*, which do machine-level instruction stepping instead of source-level. Remember that you can select the number of steps by holding down the right mouse button over the *Step Into* and *Step Over* buttons.

The *Admin*, *Display*, *Traps*, and *PC* items on the menu bar basically work same as their counterparts in the Main View window. The exception is that the *PC* submenu selections *Continue To* and *Jump To* are based on machine-level instructions rather than source-level steps.

The *Disassemble* menu is discussed in Section A.7.1.2, page 292. The *Preferences...* selection from the *Config* menu is discussed in Section A.7.1.3, page 294.

You can set traps either by using the Traps menu or by clicking in the annotation column to the left of the source display area that contains the disassembled code.

A.7.1.2 The Disassemble Menu

The Disassemble menu allows you to display disassembled code. It contains the following items:

Address... Allows you to disassemble a specified number of lines, starting from a specified source line address (see Figure 110).

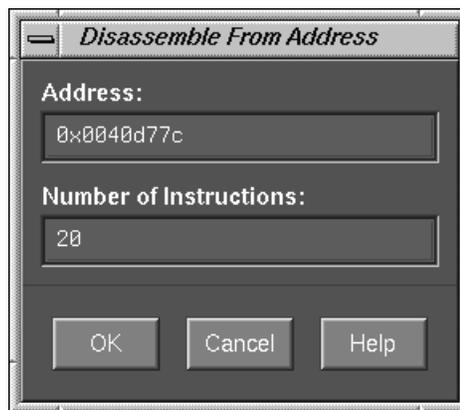


Figure 110. The Disassemble From Address Dialog

Function... Allows you to disassemble a specified number of lines, starting from the beginning address of a specified function name (see Figure 111, page 293).

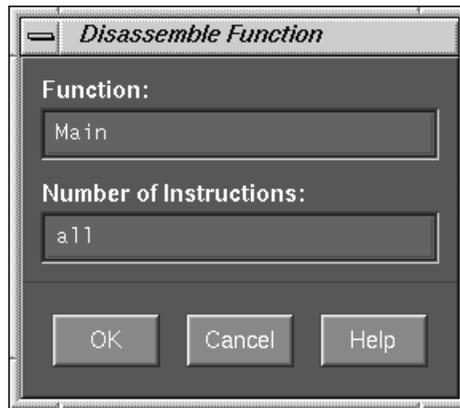


Figure 111. The Disassemble Function Dialog

File...

Allows you to disassemble a specified number of lines, starting from the address corresponding to a specified line number in a specified file (refer to Figure 112, page 294). If you have a current selection in the Main View window or the Source View window, its file and cursor position are used as the default file name and line number, respectively.

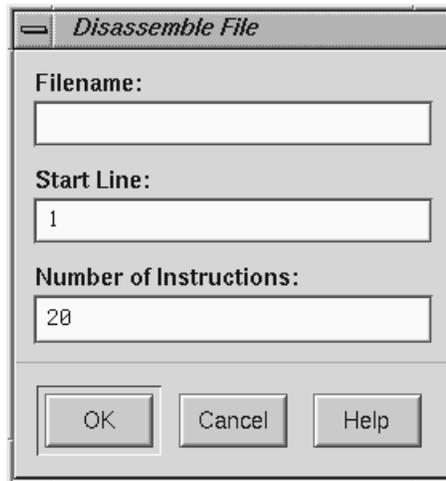


Figure 112. The Disassemble File Dialog

A.7.1.3 The Config Menu Preferences Dialog

Selecting Preferences . . . from the Config menu brings up the Disassembly View Preferences dialog box (shown in Figure 113, page 295) so that you can select the display preferences you desire.

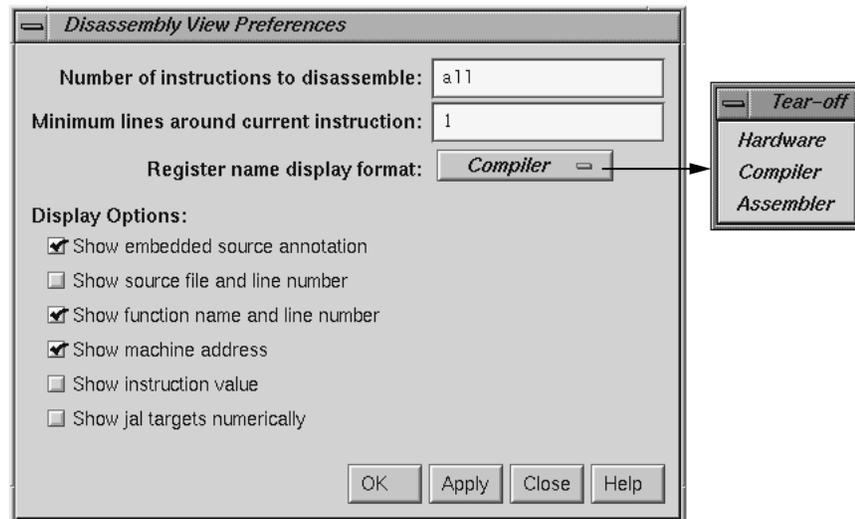


Figure 113. The Disassembly View Preferences Dialog with Display Format Menu

The dialog box provides you with the following options:

Number of instructions to disassemble	Controls the default number of disassembly lines shown when the process stops. This number appears in the dialog boxes selected from the Disassemble menu (see Figure 110, page 292, Figure 111, page 293, and Figure 112, page 294). The default is all instructions, indicating that the entire function will be disassembled.
Minimum lines around current instruction	Controls the display of the disassembled code, enabling you to view at least the specified number of instructions before and after the current instruction.
Register name display format	Controls how register names are displayed. The available modes are Hardware, Compiler, and Assembler.
Show embedded source annotation	When ON, displays source and disassembly statements interleaved. When OFF, displays disassembly statements only.

Show source file and line number	Displays the filename and file position along with each machine instruction.
Show function name and line number	Displays the function name and file position along with each machine instruction.
Show machine address	Displays the memory address of each machine instruction.
Show instruction value	Displays the instruction word along with each machine instruction.
Show jal targets numerically	Controls whether the target address of a jal instruction is displayed as a hex address or symbolic label.

A.7.2 The Register View Window

Register View window allows you to examine and modify register values. You bring it up by selecting Register View from the Views menu in the Main View window. Figure 114, page 297, shows a typical Register View window that has been resized to show all available registers.

The Register View window displays each register with its current value. A question mark (?) displayed immediately before a register value signifies that the value is suspect; it may not be valid for the current frame. This can occur if a register is not saved across a function call. A colored marker indicates that a register value has changed since the last time the process stopped.

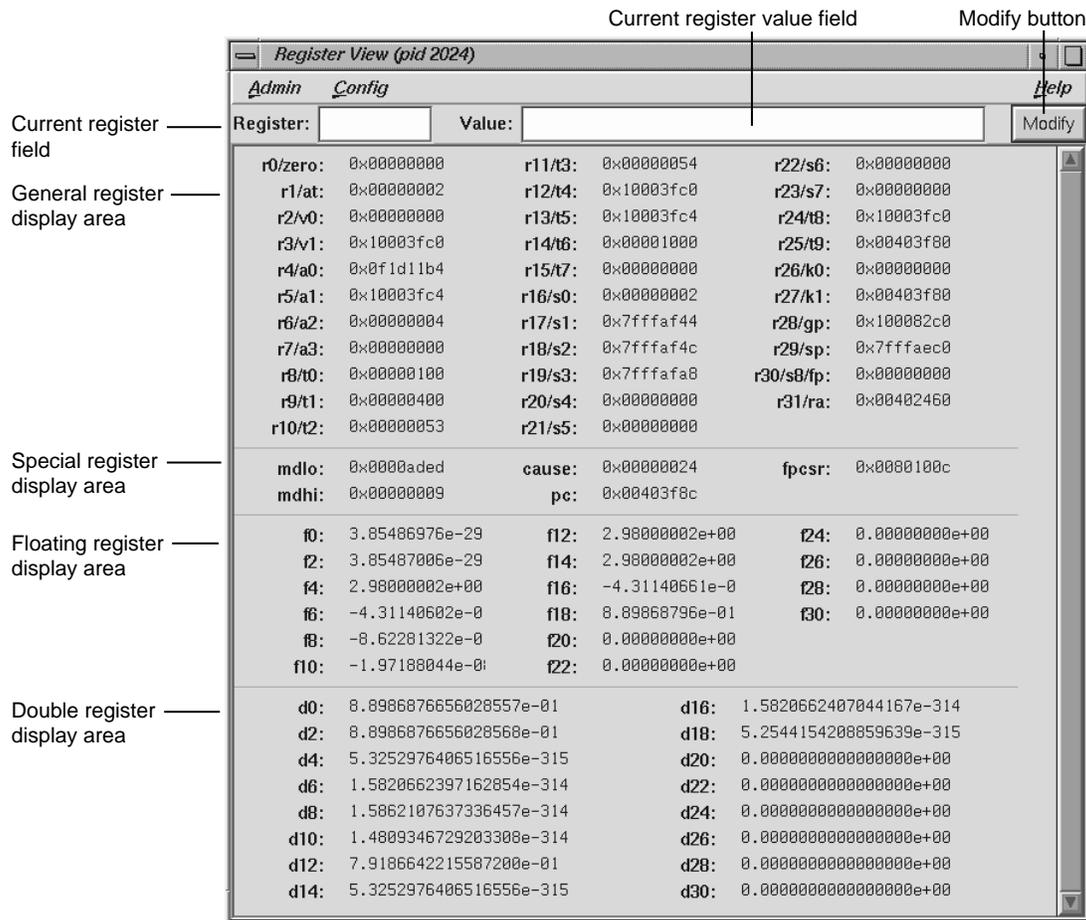


Figure 114. The Register View Window

A.7.2.1 The Register View Window

The major features of the Register View window are the following:

- Current register field** Identifies the currently selected register. You can switch to a different register by entering its name (either by hardware name or by alias) in this field and pressing Enter. You can also switch registers by clicking on the new register in the display area.

Current register value field	Shows the contents of the selected register. You can assign a new value to a register by entering either a literal or an expression into the Value field. You then click on the <code>Modify</code> button to change the value or press <code>Enter</code> .
Register display area	Shows the registers organized into four groups: general, special, floating, and double. Note that the general registers are identified by both their hardware and software names. For systems with 32-bit processors (O2s, for example), double precision registers represent a pair of floating-point registers. For systems with 64-bit processors (Origin2000s, for example), float registers are not displayed at all. Floating-point calculations are done in double precision registers.

Note: The special registers `p0`, `p1`, and `p2` do not appear in Figure 114, page 297. These are used for instrumentation and appear only when instrumentation has taken place.

A.7.2.2 Changing the Register View Display

The `Preferences...` selection in the `Config` menu allows you to change the Register View display. It brings up the Register View Preferences dialog box (see Figure 115, page 299).

The register display toggle buttons let you specify which types of registers are to be displayed by default.

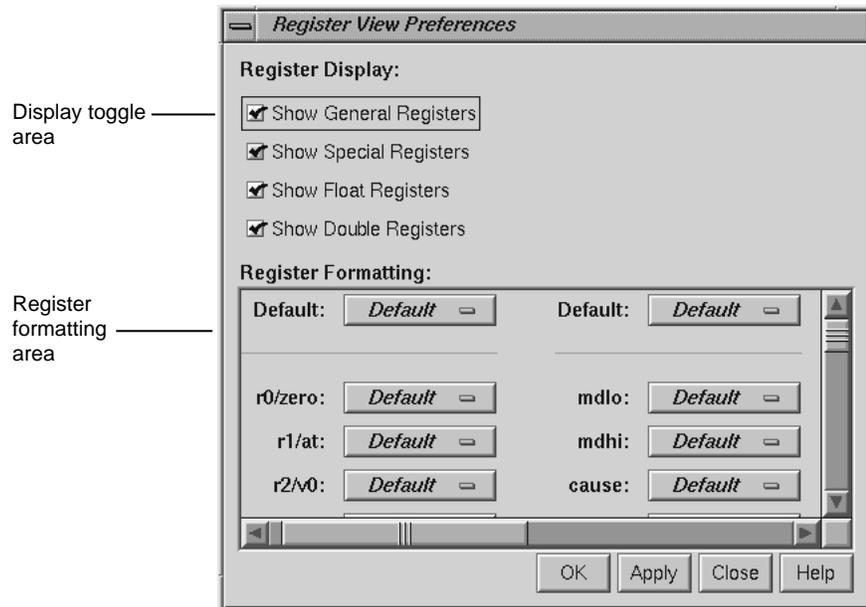


Figure 115. The Register View Preferences Dialog

The register formatting area allows you to select formats for any of the registers.

The default fields in the top row let you change defaults for the four major types, which are set as follows:

- General registers — hexadecimal
- Special registers — hexadecimal
- Float registers — floating point
- Double registers — floating point

The rows in the register formatting area let you change the modes for the individual registers.

A.7.3 The Memory View Window

The Memory View window allows you to examine and modify memory. A typical Memory View window appears in Figure 116, page 300.

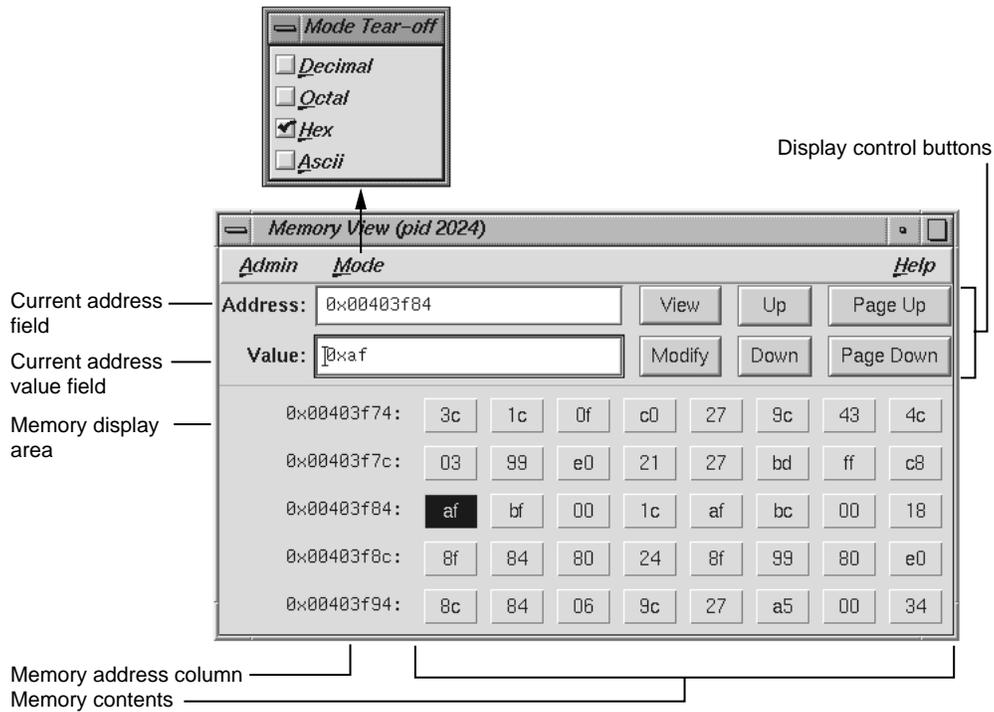


Figure 116. The Memory View Window with the Mode Submenu Displayed

A.7.3.1 Viewing a Portion of Memory

To view a portion of memory, enter the beginning memory location in the Address field. You can enter the literal value or an expression that evaluates to an integer address. These address specifications must be in the language of the current process as indicated by the call stack frame. The syntax of this expression depends upon the language of the program. For example, you can enter `0x7fff4000+4` as the memory address when stopped in a C function or enter `$7fff4000+4` as the equivalent for a Fortran routine. Press Enter while the cursor is in the field or click the View button to display the contents of that location and the subsequent locations in the display area. This also displays the contents of the first address in the Value field where it can be modified.

The memory display area shows the contents of individual byte addresses. The column at the left of the display shows the first address in the row. The contents at that address are shown immediately to its right, followed by the

contents of the next seven byte locations. If you enlarge the `Memory View` window, you can see additional rows of memory.

A.7.3.2 Changing the Contents of a Memory Location

To change the contents of a memory location, you select the address to be changed, either by direct entry or by clicking on the byte value in the display area. You can enter a single value or a sequence of hex byte values separated by spaces (for example, `00 3a 07 b2`) in the `Value` field. You can also enter a quoted string to change a consecutive range of values to the ASCII values of that string. Pressing `Enter` while the cursor is in the `Value` field or clicking the `Modify` button substitutes the new value(s) starting at the specified location.

A.7.3.3 Changing the Memory Display Format

The `Mode` menu allows you to change the format of the value field and byte locations to either decimal, octal, hex, or ASCII.

A.7.3.4 Moving around the `Memory View` Display Area

The four control buttons at the upper right of the window help you move around the display area. These buttons are:

<code>Up</code>	Moves displayed bytes up a single row.
<code>Down</code>	Moves displayed bytes down a single row.
<code>Page Up</code>	Moves displayed bytes upward by as many rows as are currently displayed.
<code>Page Down</code>	Moves displayed bytes downward by as many rows as are currently displayed.

A.8 Fix+Continue Windows

The `Fix+Continue` utility interacts with several Debugger windows. The `Main View` and `Source View` windows access the `Fix+Continue` utility from the menu bars. The standard output results of running redefined code are displayed in the `Execution View` window (refer to Section A.8.4, page 309).

Special line numbers (in decimal notation) applied to redefined functions appear in several views.

Note: Fix+Continue functionality within the debugger is limited to C/C++ programs compiled with the `-o32` compiler option.

The following windows, devoted entirely to Fix+Continue functions, can be brought up as follows from the Main View window menu bar:

```
Fix+Continue
  ->View
    ->[Pulldown menu...]
```

- Status window
- Message window
- Build Environment window

This section describes Fix+Continue menu selections and windows.

The Fix+Continue menu is available from the Main View menu bar. The menu selections operate on the selected code or the file shown in the Source View window. The Fix+Continue menu is also available from Source View window and from the Fix+Continue Status window.

A.8.1 Fix+Continue Status Window

The Fix+Continue Status window (see Figure 117, page 303) provides you with a summary of the modifications you have made during your session. It also allows you quick access to your modifications and somewhat expanded menu options.

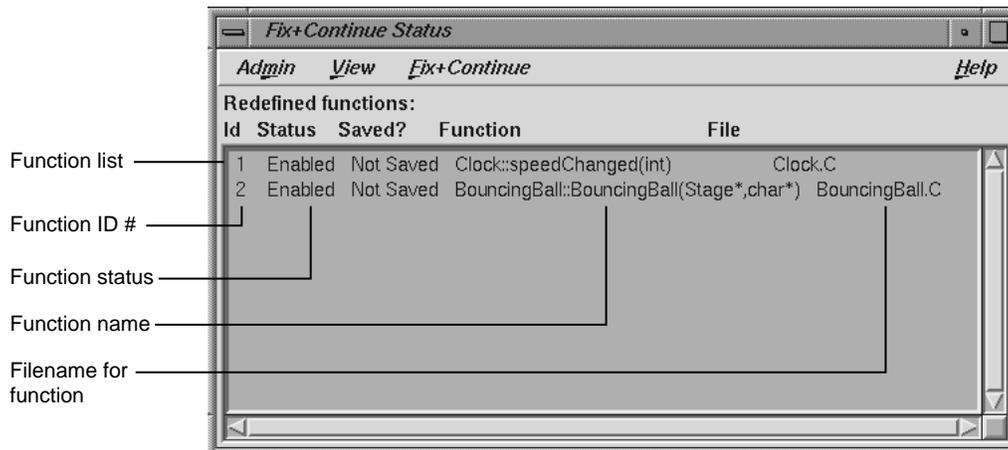


Figure 117. Fix+Continue Status Window

The function ID number, status, name, and file name are displayed in the window. Double-clicking a line in the window brings up the corresponding source in the Debugger main window.

Status window menus and submenus are described in Figure 118, page 304.

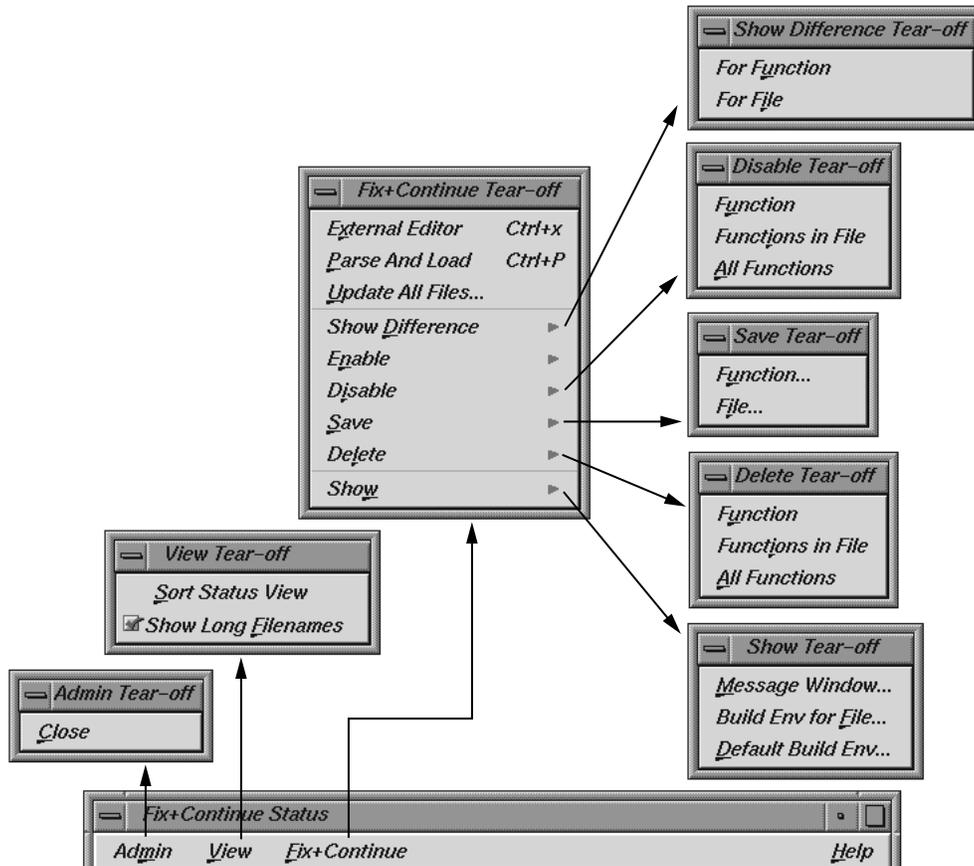


Figure 118. Fix+Continue Status Window Menus

A.8.1.1 Admin Menu

The Admin menu contains an option for closing the window.

Close Closes the status window.

A.8.1.2 View Menu

The View menu contains options for sorting information in the window and displaying file names.

Sort Status View Sorts the information in the status view according to the field currently selected.

Show Long FileNames	Toggles among absolute (long) path names or base names.
------------------------	---------------------------------------------------------

A.8.1.3 Fix+Continue Menu

The Fix+Continue menu available from the Fix+Continue Status window is somewhat different from that available through the Debugger Main View. It contains a number of options and submenus that are described below. These options and submenus are active on the item you select in the source pane of the Main View window. Before using this menu, you should select an item by clicking on it. The following options and submenus are available:

External Editor	Allows you to edit with an external editor such as vi, rather than the Debugger's default editor.
Parse And Load	Parses your modified code and loads it for execution. You can execute the modified code by clicking on the Run or Continue command buttons in the Main View window.
Update All Files...	Launches the Save File+Fixes As... dialog that allows you to update the current session while saving all the modifications to the appropriate files.
Show Difference submenu	Allows you to show the difference between the original source and your modified code. You can show the difference in the code in one of the two following ways: <ul style="list-style-type: none"> • For Function — Shows the differences for the current function only. • For File — Shows the differences for the entire file that contains the current function.
Enable submenu	Allows you to enable the changes in your modified code in one of the following ways: <ul style="list-style-type: none"> • Function — Enables the changes in the current function. • Functions in File — Enables the changes to the current function in its own file. • All Functions — Enables the changes to all functions in the modified code.

Disable submenu	Has the same menu choices as the Enable submenu, but disables rather than enables.
Save submenu	<p>Allows you to save your code changes to a file. You can save the changes in one of the following ways:</p> <ul style="list-style-type: none">• <code>Function...</code> — Launches the File dialog, allowing you to save only the current function to a file.• <code>File...</code> — Launches the Save File+Fixes As... pop-up window allowing you to save the entire file that contains the current function.
Delete submenu	Has the same menu choices as the Save submenu, but deletes rather than saves.
Show submenu	<p>Allows you to launch any of the following windows:</p> <ul style="list-style-type: none">• <code>Message Window</code> — Launches a <code>Fix+Continue Error Messages</code> window for the selected item. See Section A.8.2, page 307, for more details.• <code>Build Env for File</code> — Launches a <code>Fix+Continue Build Environment</code> window for the file shown in the Source View window. See Section A.8.3, page 308, for more details on the <code>Fix+Continue Build Environment</code> window.• <code>Default Build Env</code> — Launches the <code>Fix+Continue Build Environment</code> window to show the options that are to be used in cases where they could not be obtained from the target. See Section A.8.3, page 308, for details on the <code>Fix+Continue Build Environment</code> window.

A.8.2 Fix+Continue Error Messages Window

The Fix+Continue Error Messages window contains a list of errors and other system messages that pertain to your source modifications, parses, and attempts to run your modified source.

You can highlight the source line where the error occurred by double-clicking the appropriate line in the window. The Fix+Continue Error Messages window contains the following buttons:

Clear	Clears all the parsing errors and warnings.
Next	Puts a check mark on the next unchecked error warning entry in the parse messages. It displays the corresponding file and line in the Source view, highlighting it according to the type of error or warning. The Next option does not work after all the entries in the messages are ticked.
Rescan	Erases all the ticks, so that you can rescan all the error warnings from the beginning.

A.8.2.1 Admin Menu

The Admin menu allows you to perform either of the following operations:

Clear All	Clears all messages in the window.
Close	Closes the window.

A.8.2.2 View Menu

The View menu allows you to set any of the following toggles:

Show Warnings	Causes compile warnings to be displayed in the parse errors list.
Append Parse Messages	Causes parse messages to be appended to the parse errors list.
Append Load Messages	Causes load messages to be appended to the load errors list.

A.8.3 Fix+Continue Build Environment Window

This section describes the Fix+Continue Build Environment window (see Figure 119, page 308). The Fix+Continue Build Environment window provides you with the build information for your source code in your current environment. It displays the command that was used to build your executable and the name of the file that contains the function that you currently have selected.

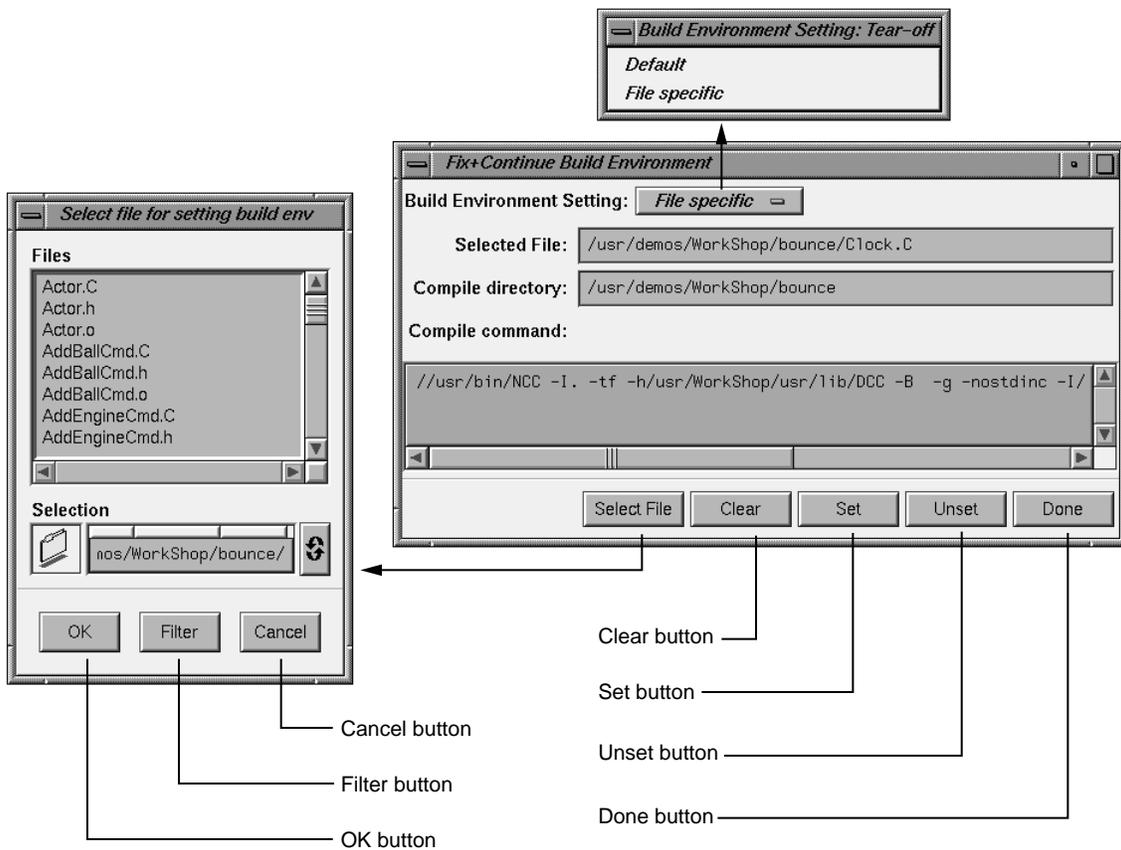


Figure 119. Fix+Continue Build Environment Window

The compiler and associated flags that were used to compile the file are normally gathered from the target. You can use this window to make any changes to these flags.

The window allows you to select your build environment setting through the `Build Environment Setting` toggle that contains the following two options:

<code>Default</code>	Sets the build environment to default that is displayed in the window.
<code>File Specific</code>	Sets the build environment to that of the file that contains the currently selected function. You can change the file by clicking the <code>Select File</code> button, which launches the <code>File</code> dialog.

The `Fix+Continue Build Environment` window also contains the following buttons:

<code>Select File</code>	Launches the <code>File</code> dialog and allows you to select a file from which to set the build environment.
<code>Clear</code>	Clears the window.
<code>Set</code>	Sets the build environment to what is displayed in the window.
<code>Unset</code>	Unsets the build environment.
<code>Done</code>	Dismisses the window.

A.8.4 Changes to Debugger Views

When you use `Fix+Continue`, views change to show redefined functions or stopped lines containing redefined functions.

A.8.4.1 Main View

All `Fix+Continue` actions are available through the `Fix+Continue` menu on the `Main View` window. See Figure 120, page 310, for details.

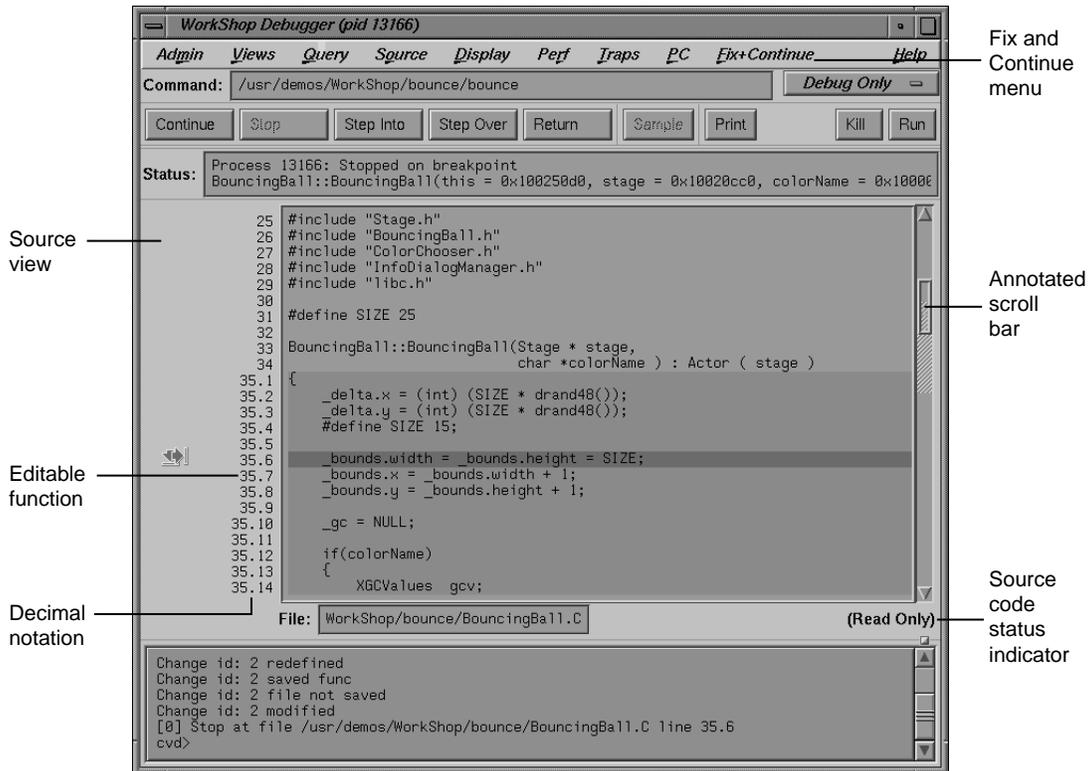


Figure 120. Debugger Main View Window

You can select commands from the Fix+Continue menu or enter them at the Debugger command line. The source code status is Read Only. Color coding shows the differences between editable code, enabled redefinitions, disabled definitions, and breakpoints. Line numbers in redefined functions have decimal notation that is used for every reference to the line number. The integer portion of the decimal is the same as the first line of the function. This ensures that compiled source code line numbers remain unchanged.

A.8.4.2 Command Line Interface

The Debugger command line interface accepts Fix+Continue commands and reports status involving redefined functions or files. Figure 121, page 311, shows a function successfully redefined using the command line. Change id 1 was previously redefined and assigned the number 1.

Specify function with
Change id 1

```

cvd> redefine 1
"/d2/people/cgeary/interp/fermat.c":18.1> {
"/d2/people/cgeary/interp/fermat.c":18.2> int i;
"/d2/people/cgeary/interp/fermat.c":18.3> int result = 1;
"/d2/people/cgeary/interp/fermat.c":18.4> for (i = 0; i < n; i++)
"/d2/people/cgeary/interp/fermat.c":18.5> result *= a;
"/d2/people/cgeary/interp/fermat.c":18.6> return result;
"/d2/people/cgeary/interp/fermat.c":18.7> }
"/d2/people/cgeary/interp/fermat.c":18.8> .
Change id: 1 redefined
      1      enabled /d2/people/cgeary/interp/fermat.c      power
Change id: 1 activated
cvd>

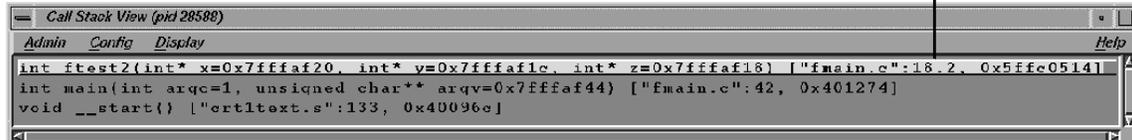
```

Figure 121. Command Line Interface with Redefined Function

A.8.4.3 Call Stack View

The Call Stack View recognizes redefined functions. It uses the decimal notation for line numbers, as shown in Figure 122.

Decimal notation for line number



```

Call Stack View (pid 28588)
Admin Config Display Help
int ftest2(int* x=0x7ffffaf20, int* y=0x7ffffaf1c, int* z=0x7ffffaf18) ["fmain.c":18.2, 0x5ffc0514]
int main(int argc=1, unsigned char** argv=0x7ffffaf44) ["fmain.c":42, 0x401274]
void __start() ["ertltext.s":133, 0x400960]

```

Figure 122. Call Stack View

A.8.4.4 Trap Manager

The Trap Manager recognizes redefined functions. It uses the decimal notation for line numbers, as shown in Figure 123, page 312.

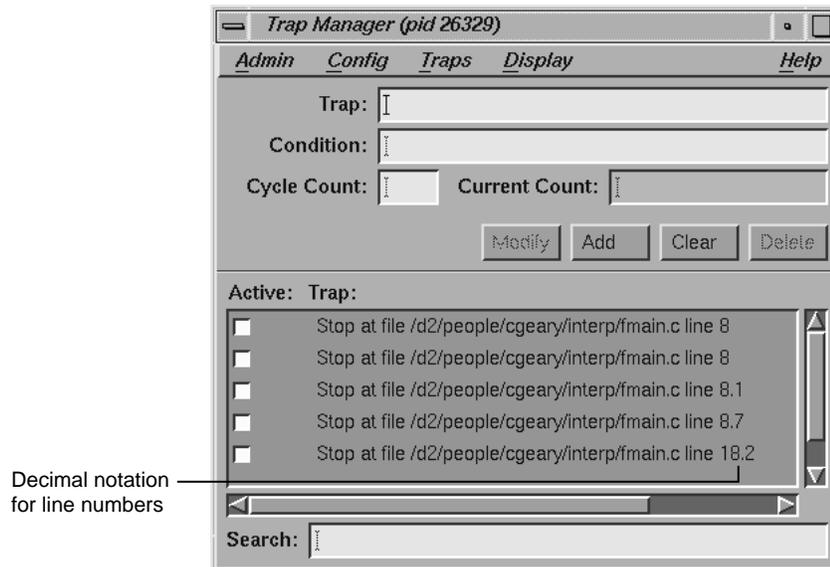


Figure 123. Trap Manager Window with Redefined Function

A.9 Debugger Command Line

You can debug programs by entering dbx-style commands at the `cvd>` prompt found at the bottom of the Main View window (see Figure 57, page 184). For more information, refer to the *dbx User's Guide*.

A.9.1 Syntax for dbx-style Commands

The syntax for the debugging commands is as follows:

```
add_source { "filename":line_number }
```

(For C and C++ only, when compiled with `-o32`.) Prompts you to add source code lines (for example, `add_source "fmain.c":15.2`). *line_number* must be within the body of a function. Entering a period (.) specifies the end of your input. The source lines you provide are added after the specified line. This command returns an ID existing or new, depending on whether the function affected has already been changed or not. The resulting new definition of the function is executed on its

entry next time. See also `delete_source` and `replace_source`.

`alias` [*shortform command*]

Lists all aliases without arguments. With arguments, it assigns *command* to *shortform*.

`assign` *expression1=expression2*

Assigns *expression2* to *expression1*.

`attach` *pid*

Attaches to specified process ID (*pid*).

`call` *function_name* [*argument, ...*]

Executes the specified function with any arguments supplied.

`catch` [*signal_name* | *all*]

With no arguments, lists signals to be trapped. If a signal is specified, it's added to the list. If *all* is specified, it traps all signals.

`clearcalls`

Cancels pending function calls.

`cont` *in function_name*

Continues execution from the current line to the entry to the specified function.

`cont` *to line_number*

Continues execution from the current line until the specified line, if there is not an intervening breakpoint to stop execution.

`continue` [*all*]

Continues executing a program, or all processes, after a breakpoint. You can use both `c` and `cont` as aliases for `continue`.

`continue` [*signal*]

Sends specified signal and continues executing a program after a breakpoint.

`corefile` [*filename*]

With no arguments, reports whether data referencing commands reference a core file. If so, displays the current core file. With *filename* provided, specifies core file to be debugged.

`delete` {all | [*trap_number* [, *trap_number*, ...]]}

Deletes all traps. The `all` option deletes all breakpoints.

`delete_changes` {*func_spec* | -all | {-file *filename*}}

(For C and C++ only when compiled with `-o32`.) Deletes changes corresponding to the selected functions (for example, `delete_changes getNumbers -file fmain.c`). Once IDs are deleted, you will not be able to use the IDs again because the IDs associated with the selected functions are released. The default is `-all`. See also `save_changes`.

`delete_source` { "*filename*":*line_number*[,*line_number*]}

(For C and C++ only, when compiled with `-o32`.) Deletes the given line(s) if *line_number* or *line_number* (range) is within the body of a function. An example is:

```
delete_source "fmain.c":8.6,8.7
```

This command returns an ID existing or new, depending on whether the function affected has already been changed or not. The resulting new definition of the function is executed on its entry next time.

`delete trap_number [,trap_number , ...]`

Deletes the specified breakpoint numbered *trap_number*, as obtained from the list generated when the `status` command is issued.

`detach`

Detaches from the current process.

`disable all`

Deactivates all traps.

`disable_changes {func_spec | -all | {-file filename}}`

(For C and C++ only, when compiled with `-o32`.) Disables specified changes for selected functions (for example, `disable_changes getNumbers -file fmain.c`. Nothing happens if the selected function is already disabled. The compiled definition of the function is executed on its next entry. You can invoke this command when the process is stopped or on a running process when a function entry breakpoint is set.

`disable trap_number [,trap_number , ...]`

Deactivates the specified breakpoint numbered *trap_number*, as obtained from the list generated when the `status` command is issued.

`display [expression , ...]`

With *expression*, adds *expression* to the list of expressions displayed whenever the process stops. With no arguments, lists all expressions. See `undisplay` to delete an expression.

`down [number]`

Moves down the specified number of frames in the call stack. `down` moves away from the direction of the caller.

`dump`

Prints local variable values.

`enable all`

Reactivates all inactive traps.

`enable_changes {func_spec | -all | {-file filename}}`

(For C and C++ only when compiled with `-o32`.) Enables specified changes for selected functions (for example, `enable_changes getNumbers -file fmain.c`. Nothing happens if the selected function is already enabled. The latest accepted definition of the function is redefined on its next entry. You can invoke this command when the process is stopped or on a running process when a function entry breakpoint is set.

`enable trap_number [,trap_number, ...]`

Reactivates the specified breakpoint numbered `trap_number`, as obtained from the list generated when the `status` command is issued.

`expression/[count] [format] or expression,[count] /[format]`

Prints the contents of the memory address specified by `expression`, according to the specified format. `count` represents the number of formatted items. The following format options are available:

<code>d</code>	Prints a short word in decimal.
<code>D</code>	Prints a long word in decimal.
<code>o</code>	Prints a short word in octal.
<code>O</code>	Prints a long word in octal.
<code>x</code>	Prints a short word in hexadecimal.
<code>X</code>	Prints a long word in hexadecimal.
<code>b</code>	Prints a byte in octal.
<code>c</code>	Prints a byte as a character.
<code>s</code>	Prints a string of characters that ends in a null byte.
<code>f</code>	Prints a single-precision real number.
<code>g</code>	Prints a double-precision real number.

`file` [*filename*]

Displays the name of the current or specified file (*filename*). If a file is specified, it becomes the current file.

`func` [*func_name*]

Moves to the source code corresponding to the specified frame in the call stack or to the function in the executable if not on the stack.

`givenfile` [*filename*]

With no arguments, displays name of current object file. With *filename*, specifies object file to be debugged.

`goto` *linenumber*

Skips over lines going directly to the specified line number in the current routine. Unlike `dbx(1)`, `cvd(1)` does not begin execution at the specified line.

`ignore` [*signal_name* | `all`]

With no arguments, lists those signals not to be trapped. If a signal is specified, this command removes it from the list of signals to be trapped. If `all` is specified, ignores all signals.

`kill` [*pid*][`all`]

Kills the specified process currently controlled by the Debugger or kills all processes.

`list` [*from-line#*[`{:line_count}` | `{,to-line#}`]] | [*function_name*]

Lists source lines beginning at *from-line#*. If no additional argument is specified, the default for *line_count* is 10. If *line_count* is specified, a total of *line_count* lines are listed. If *function_name* is specified, the lines from the given function are listed.

`list_changes [func_spec | -all | {-file filename}]`

Lists one or more lines using the following syntax:

`change_id isEnabled filename function_spec`

For example:

```
4 enabled foo.c foo
8 disabled A.c++ A::bingo
```

The default is `list_changes -all`.

`next [int]`

Steps over the specified number of source lines. This command does not step into procedures. The default is one line.

`nexti [int]`

Steps over the specified number of machine instructions. This command does not step into procedures. The default is one instruction.

`print expression [, expression, ...]`

Prints the value of the specified expression(s). If the expression is a character pointer or array, both the string and address print. You can use `p` as an alias.

`printd expression [, expression, ...]`

Prints the value of the specified expression(s) in decimal format. You can use `pd` as an alias.

`printf string [, expression1 [, expression2, ...]]`

Prints the value(s) of the specified expression(s) in the format specified by the *string* string. The `printf` command supports all formats of the IRIX `printf` command except `%s`. For a list of formats, see the `printf(1)` man page.

`printo expression [,expression , ...]`

Prints the value of the specified expression(s) in octal format. You can use `po` as an alias.

`printregs`

Prints the contents of the registers.

`printx expression [,expression , ...]`

Prints the value of the specified expression(s) in hexadecimal format. You can use `px` as an alias.

`pwd`

Displays the current directory.

`quit`

Exits the debugging session.

`redefine func_spec[-edit |{ -read filename[line_number, line_number]}]`

(For C and C++ only.) Specifies a new body for a function. The new definition is checked, and errors (if any) are printed. The new function body is redefined on the next function entry. Breakpoints (if set) on the old definition are put on the new definition based on their relative line number position from the beginning of the function definition. (Note that some breakpoints may not make it to the new definition.) You can invoke this command when the process is stopped or on a running process when a function entry breakpoint is set. There are three ways to provide a new definition:

- `-edit` pops up an editor of your choice containing the current definition of the function. The specification of the new definition is complete when you exit the editor. You may not leave the editor open.
- `-read` takes the contents of the file specified (within the line numbers if given) as the new function definition.

- No option allows you to type in replacement code from the next line. A period in the first column on a fresh line terminates the definition. For example:

```
redefine getNums
"/usr/fmain.c":8.1> {
"/usr/fmain.c":8.2> printf(``In getNums.\n``);
"/usr/fmain.c":8.3> }
"/usr/fmain.c":8.4> .
```

```
replace_source { "filename":line_number[ ,line_number]}
```

(For C and C++ only when compiled with `-o32`.) Prompts you to type in replacement source if *line_number* or *line_number* (range) is within the body of a function. The source lines you provide replace the specified line(s). An example is:

```
replace_source "fmain.c":12. This command returns an existing or new id depending on whether the function affected has already been changed or not. The resulting new definition of the function is executed on its entry next time. See also add_source and delete_source.
```

```
rerun
```

Runs the program again using the same arguments.

```
return
```

Continues executing the current procedure and returns to the next sequential line in the calling function.

```
run [[all] | [argument_list]]
```

Runs the program (s). If an *argument_list* is specified, it is used as the arguments to be supplied to the program.

```
runtime_check func_spec [-options key [key, ...]]
```

(For C and C++ only when compiled with `-o32`.) Enables all run-time checking options by default. If `-options` is specified, then run-time checking is restricted to the *keys*. The result of the runtime checks are printed the next time the specified function (*func_spec*) is entered. You can invoke this command on a stopped or a running process.

```
save_changes {func_spec| {-file filename}} [-[w|a]] filename_to_save
```

(For C and C++ only when compiled with `-o32`.) Saves (enabled or disabled) function redefinitions or an entire file to another file (*filename_to_save*). The following example shows how to save a function definition:

```
save_changes getNumbers getNumbersFunc
```

If you specify the `-file` option, then before saving to *filename_to_save*, all function changes are applied to the compiled source of the file (with the condition that the file has had only its functions redefined, and has not been edited since the last build). An example of saving an entire file is the following:

```
save_changes -file fmain.c fmain.c
```

The `-w` option replaces the *filename_to_save*. The `-a` option appends to the *filename_to_save*. An example of adding a function to a file is the following:

```
save_changes -file fmain.c -a fmain.c
```

See also `delete_changes`.

```
setbuildenv ["filename"] compiler-flag-list
```

(For C and C++ only, when compiled with `-o32`.) Overrides default build environment flags (compiler options). Without *filename*, the flags are passed along with `-c -g` flags to the compiler for any function in any file except those set separately with `setbuildenv`. An example is the following:

```
setbuildenv -DnameA -Idir
```

If *filename* is given, this command sets separate flags specifically for that file. For example, consider the following:

```
setbuildenv "fermat.c" -DnameB -Ianotherdir
```

See also `unsetbuildenv`.

```
sh [shell_command]
```

Calls a shell if no arguments; otherwise, executes the specified shell command.

`showbuildenv ["filename"]`

(For C and C++ only, when compiled with `-o32`.) Lists all the build environment flags set. `showbuildenv` with a *filename* lists any build environment specifications that have been set separately with `setbuildenv "filename"`.

`show_changes [func_spec | -all | {-file filename}]`

(For C and C++ only, when compiled with `-o32`.) Prints the code of all enabled redefinitions of the specified function(s). The default is `show_changes -all`. See also `enable_changes` and `disable_changes`.

`show_diff {func_spec | {-file filename}}`

(For C and C++ only, when compiled with `-o32`.) Launches a `xdiff` comparing the compiled source and its latest redefinition for the specified function. If `-file filename` is specified, `xdiff` shows the difference between the compiled file and the file with all redefinitions applied to the compiled source of the file (with the condition that the file has had only its functions redefined, and has not been edited since the last build).

`showthread [full] [thread] {number | all}`

Shows brief status information about threads. If `full` is specified, prints full status information. You can request status information for a specific thread by number or you can request information for all threads. The `thread` qualifier does not affect command output.

`source filename`

Executes commands in the specified file.

`status`

Displays a list of currently set breakpoints and traces.

`step [int]`

Steps the specified number of source lines. This command steps into procedures. The default is one line.

`stepi [int]`

Steps the specified number of machine instructions. This command steps into procedures. The default is one instruction.

`stop {all | pgrp} in [filename]`

If set to `all`, stops all members of the specified process group whenever the trap is encountered. If set to `pgrp`, all members of the process group will apply the trap.

`stop at ["filename":] line_number [if expression]`

Traps at the specified line in the specified file. If the `if` option is used, the trap fires only if *expression* is true. If specified, the file name must be enclosed in double-quotation marks. As an example, to specify a stop at line 5 in `myfile.c`, the syntax is:

```
stop at "myfile.c":5
```

`stop exception {all | item}`

Stops on all C++ exceptions or exceptions that throw the base type `item`. Do not include complex expressions using operators such as `*` and `&` in your type specification for an exception breakpoint.

`stop exception {all | [item [, item]]}`

Stops on all C++ exceptions or exceptions that either have no handler or are caught by an unexpected handler. If you specify `item`, stops on exceptions that throw the base type `item`. Do not include complex expressions using operators such as `*` and `&` in your type specification for an exception breakpoint.

`stop in [filename:] function_name [if expression]`

Traps at the entry to the specified function. If the `if` option is used, then the trap fires only if *expression* is true. If the *filename* is given, the function is assumed to be in that file's scope. If specified, the *filename* must be enclosed in double-quotation

marks. As an example, to specify a stop in function `func1` in `myfile.f` only if `n` is 20, the syntax is:

```
stop in "myfile.f":func1 if n .eq. 20
```

```
syscall {catch | ignore} {call | return} {sys_call_name | all}
```

The `catch` option adds a system call to the list of system calls to be trapped. The `ignore` option removes a system call from the system call trap list. The `call` option specifies the entry to the system call and `return` signifies the return from the call.

```
trace [variable] at [{"filename [line_number ":" | function_name] \ [if expression]}
```

Traces the specified `variable`. You can specify a file and/or test condition. You can also specify a line number or a function where the trace is to take place.

```
unalias aliasname
```

Cancels the alias specified as *aliasname*.

```
undisplay [displaynumber, ...]
```

Stops display of expression with specified *displaynumber* when the process stops. Removes the expression from the display list.

```
unsetbuildenv ["filename"]
```

(For C and C++ only, when compiled with `-o32`.) Disregards the default build environment flags if specified earlier. For all functions in files that don't have an overriding build environment, `unsetbuildenv` passes only the `-c` and `-g` flags.

If *filename* is given, this command disregards the build environment flags specified for the file earlier. Further redefinition of the functions in the file use the default build environment flags, if set. See also `setbuildenv`.

`up` [*number*]

Moves up the specified number of frames in the call stack. `up` moves in the direction of the caller.

`use` [*path*]

Uses the specified path to search for source files.

`watch` *identifier* [`write` | `read`]

Causes program to stop when the *identifier* is written or read, depending on whether `write` or `read` is specified. If neither is specified, the default is `write`.

`whatis` *identifier*

Displays all the qualifications of the specified variable.

`when at` ["*filename*":] *line_number* {*command*[: *command* ...]}

When your program reaches the specified *line_number*, the commands specified are executed before the program resumes execution.

`when in` ["*filename*":] *function_name* {*command* [: *command* ...]}

When your program enters the specified function, the commands specified are executed before the program resumes execution.

`which` [*identifier*]

Displays the qualification of the specified variable.

`where` [`thread` | *thread-id*]

Performs a stack trace showing the activation levels of a program or, optionally, of the specified thread. You can obtain *thread-ids* from the first column of output of the `showthread` command.

A.10 Blocking Kernel System Calls

The following are the kernel system calls (syscalls) that actually block continued pthreads. There are numerous library routines, such as `printf`, that can use one of these blocking system calls.

It would be impractical here to list all library routines which utilize a blocking syscall. Nevertheless, as a user you should know, for example, that if you call the `printf` library routine it eventually calls `writev()`, a blocking system call, and thus may block continued pthreads.

<code>accept</code>	accept a connection on a socket
<code>close</code>	close a file descriptor
<code>creat</code>	create a new file or rewrite an existing one
<code>dmi</code>	SGI specific. Used to implement the interface defined in X/Open document <i>Systems Management: Data Storage Management (XDSM) API</i>
<code>fcntl</code>	file and descriptor control. Provides for control over open descriptors.
<code>fsync</code>	synchronize a file's in-memory state with that on the physical medium
<code>getmsg / getpmsg</code>	get next message off a stream
<code>ioctl</code>	control device. Performs a variety of control functions on devices and streams.
<code>lockf</code>	record locking on files. Allow sections of a file to be locked.
<code>mq_open</code>	open/create a message queue
<code>msgsnd / msgrcv</code>	message send and message receive
<code>msync</code>	synchronize memory with physical storage
<code>nanosleep</code>	high resolution sleep
<code>open</code>	open for reading and writing
<code>pause</code>	suspend process until signal is received
<code>poll</code>	input/output multiplexing
<code>putmsg / putpmsg</code>	send a message on a stream
<code>read / readv / pread</code>	read from a file

recv / recvfrom / recvmsg	receive a message from a socket
select	synchronous I/O multiplexing
semget / semctl / semop	semaphore handling
send / sendto / sendmsg	send a message from a socket
sginap	times sleep and processor yield function
write / writev / pwrite	write on a file

Using the Build Manager [B]

WorkShop lets you compile software without leaving the WorkShop environment. Thus, you can look for problems using the WorkShop analysis tools (Static Analyzer, Debugger, and Performance Analyzer), make changes to the source, suspend your testing, and run a compile. WorkShop provides two tools to help you compile:

- *Build View*—for compiling, viewing compile error lists, and accessing the code containing the errors in *Source View* (the WorkShop editor) or an editor of your choice. *Build View* helps you find files containing compile errors so that you can quickly fix them, recompile, and resume testing.
- *Build Analyzer*—for viewing build dependencies and recompilation requirements and accessing source files.

Build View uses the UNIX `make(1)` facility as its default build software. Although `cvmake` can be set up to run any program instead of `make` (for example, `gnumake`), `cvbuild` will only parse and display standard makefiles (in particular, it does not understand `gnu make` constructs).

B.1 Build View Window

You can access the *Build View* window from the WorkShop analysis tools, from the command line (by typing `cvmake`), or from the *Build Analyzer* (see next section).

To access *Build View* from WorkShop, select *Recompile* from the *Source* menu in the *Main View* window in the *Debugger* or from the *File* menu in *Source View* (for more information on the *Main View* and *Source View* windows, refer to Chapter 1, page 1). Selecting *Recompile* detaches the current executable from the WorkShop analysis tools and displays *Build View*. You can edit the *Directory* and *Target(s)* fields as needed and click *Build* to compile. If the source compiles successfully, the new executable is reattached when you reenter the WorkShop analysis tools.

The *Build View* window has three major areas:

- Build Process Control Area, Section B.2, page 330
- Transcript Area, Section B.3, page 331
- Error List Area, Section B.4, page 332

B.2 Build Process Control Area

The build process control area lets you run or stop the build and view the status. See Figure 124.

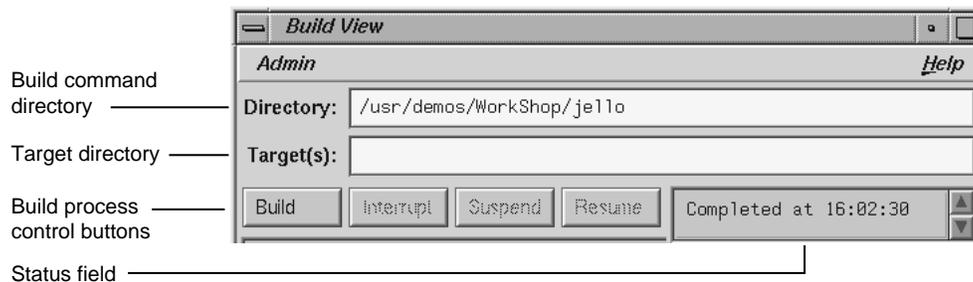


Figure 124. Build Process Control Area in Build View Window

The directory in which the build will run displays in the `Directory:` field at the top of the area. The current directory displays by default. You can specify the build using `make`, `smake`, `pmake`, `clearmake`, or any other builder and any flags or options that the builder understands (see Section B.5.1, page 332, and Section B.5.2, page 333). The target to be built is specified in the `Target(s):` field.

The build process control buttons let you control the build process. The following buttons are available:

<code>Build</code>	Runs (or reruns) a build. If you have modified any files you will be prompted to save the new versions prior to the compile.
<code>Interrupt</code>	Stops a build.
<code>Suspend</code>	Stops a build temporarily.
<code>Resume</code>	Restarts a suspended build.

The status field is to the right of the build process control buttons. It indicates the progress of the build.

B.3 Transcript Area

The transcript area displays the verbatim output from the build. The vertical scroll bar lets you go through the list; the horizontal scroll bar lets you see long messages obscured from view. A sash between the compile transcript area and the error list area lets you adjust the lengths of the lists displayed. See Figure 125.

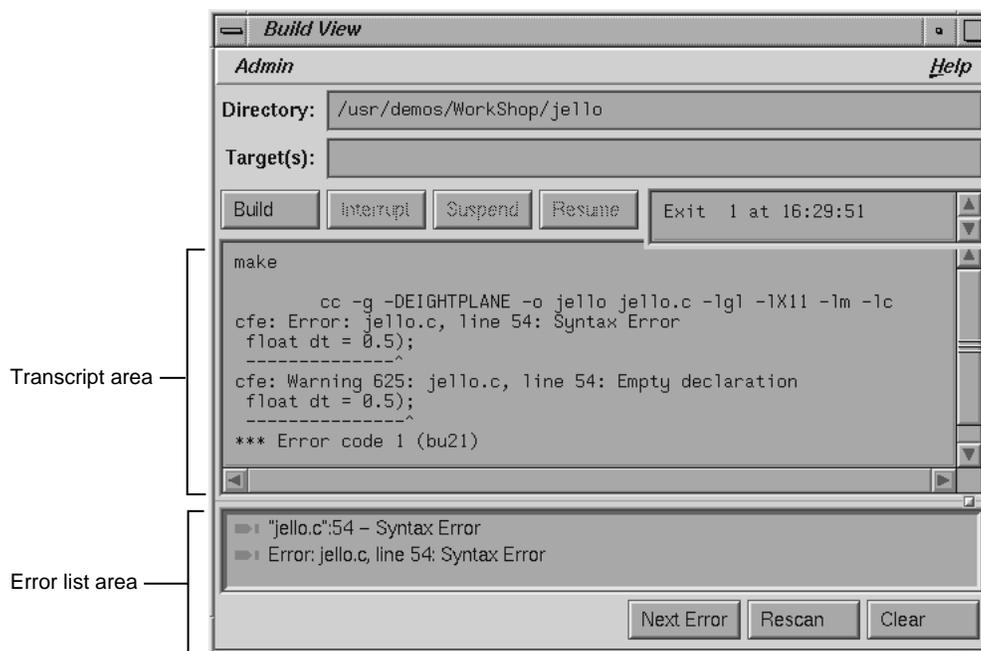


Figure 125. Build View Window with Typical Data

B.4 Error List Area

The error list area consists of the error list display and three control buttons. The following buttons are available:

Next Error	Brings up the default editor scrolled to the next error location. This button is below the error list display.
Rescan	Refreshes the error list display.
Clear	Clears the error list display area.

The error list area displays compile errors (see Figure 125, page 331). The errors are annotated according to their severity level (fatal has a solid icon and the warning icon is hollow). Double-clicking the text portion of an error brings up the default editor scrolled to the error location and displays a check mark to help you keep track of where you are in the error list. Check marks also display when you click the `Next Error` button.

B.5 Build View Admin Menu

The Admin menu in Build View has two selections in addition to the standard `WorkShop` entries:

- `Build View Preferences...`, Section B.5.1, page 332
- `Build Options...`, Section B.5.2, page 333

For information on `Launch Tool`, `Project`, and `Exit` menu selections, refer to Section A.1.1, page 190.

B.5.1 Build View Preferences

The `Preferences...` selection brings up the dialog box shown in Figure 126, page 333. The options are:

Maker Program field	Lets you enter the program you use to build your executable.
Macro Settings field	Lets you enter build macros, such as <code>CFLAGS=-g</code> .
Makefile field	Lets you enter the name of a makefile if you do not wish to use the default.

Discard Duplicate
Errors button

Eliminates subsequent duplicates of errors in the
error list area.

Show Warnings
button

Toggles the option to display warnings in the list.

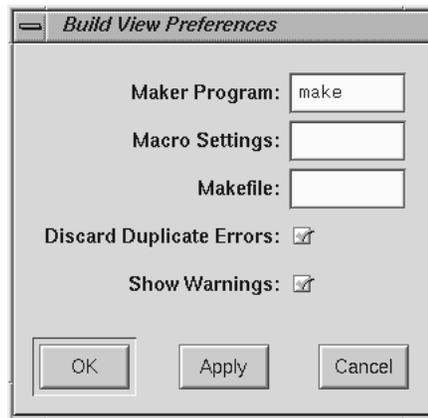


Figure 126. Build View Preferences Dialog

B.5.2 Build Options

The Build Options Dialog lets you add the options shown in Figure 127, page 334, to your make command.

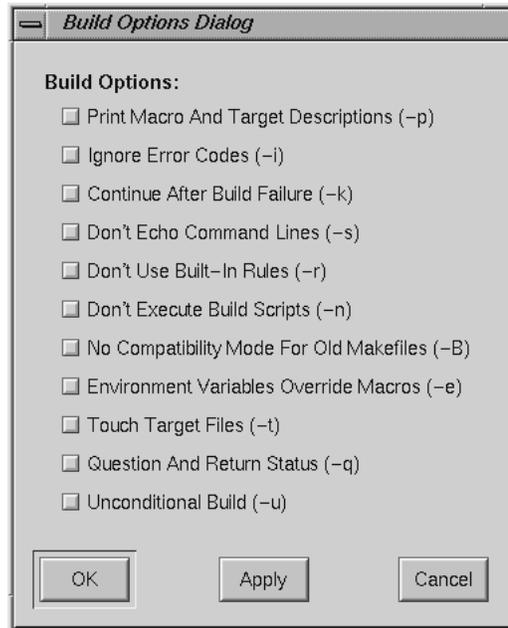


Figure 127. Build Options Dialog

B.5.3 Using Build View

The steps in running a compile using Build View are as follows:

1. Bring up the Build View window.
2. Edit the `Target(s):` and `Directory:` fields as required.
3. Specify your preference regarding duplicate errors and warnings using the Admin menu (optional).
4. Click `Build` to start the build. All compile information displays in the transcript area. Errors are grouped in a list below.
5. Click `Interrupt` to terminate or `Suspend` for a temporary stop, if you want to stop the build. The `Resume` button restarts a suspended build.
6. Double-click an error to bring up your preferred editor with the appropriate source code. A check mark indicates that an error has been accessed.

Note: The default editor is determined by the `editorCommand` resource in the `app-defaults` file. The value of this resource defaults to `wsh -c vi +%d`, which means run `vi` in a `wsh` window and scroll to the current line. If the editor lets you specify a starting line, enter `%d` in the resource to indicate the new line number.

7. Click `Build` to restart the build.

B.6 Build Analyzer Window

The `Build Analyzer` window displays a graph indicating the source files and derived files in the build, and their dependency relationships and current status. Source files refers to input files, such as code modules, documentation, data files, and resources. Derived files refers to output files, such as compiled code. You request builds in `Build Analyzer` by either:

- Double-clicking a derived module
- Making a selection from the `Build` menu

You access `Build Analyzer` from `WorkShop` by selecting `Launch Tool` from the `Admin` menu in `Main View`. Outside of `WorkShop`, you can access `Build Analyzer` by typing `cvbuild` at the command line. A typical `Build Analyzer` window appears in Figure 128, page 336, with the menus displayed.

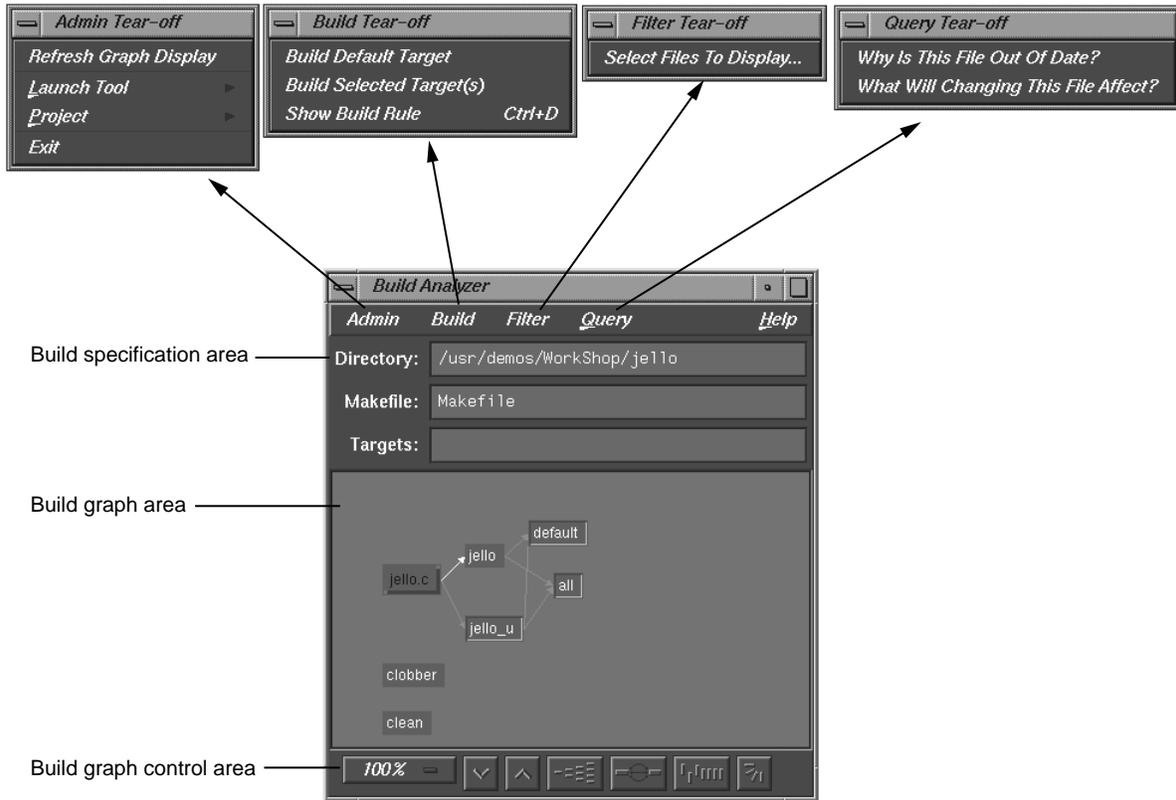


Figure 128. Build Analyzer Window

B.7 Build Specification Area

The three fields in the build specification area identify the working directory, makefile script, and target file(s) for compilation. You can edit the Directory:, Makefile:, and Target(s): fields directly. The Target(s): field also lets you specify a search string for locating a file in the build graph.

B.8 Build Graph Area

The build graph area displays the specified source and derived files and their dependency relationships. Files are depicted as rectangles; dependency relationships are shown as arrows, with the supplying file at the base of the arrow and the dependent file at the head. The colors used to depict the files

depends on your color scheme. *Build Analyzer* differentiates the two types of files by depicting one with light characters on a dark background and the other with dark text on a light background. If you double-click a source file icon, an editor is brought up for that file. Double-clicking a derived file starts a build and displays *Build View*.

In addition to dependency relationships, *Build Analyzer* indicates the status of the files and relationships as follows:

- Source file availability status: *normal* or *checked out*
 - *Normal* means that the source file is read-only and needs to be made writable to be edited. Normal files appear as light rectangles with black text.
 - *Checked out* means that you have a writable version of this file available and can thus edit it. A checked out file appears in a different color (from normal files) with a shadow.
- Derived file compile status: *current* or *obsolete*
 - When applied to a derived file, the term *current* means that none of the files on which the derived file depends have been edited since the derived file was created. Current derived files appear as dark rectangles with white text.
 - *Obsolete* means that one or more of the source files have been modified since the derived file was created. Obsolete files appear in the same color as current derived files but with a colored outline.
- Dependency relationship: *current* or *obsolete*
 - *Current* means that the derived file is up to date with the source files. Note that a relationship can be current even if both files are obsolete. This happens when a file on which both files are dependent has been modified. Current arcs are black.
 - *Obsolete* means that the source file has changed and the derived file has not been updated accordingly. Obsolete arcs appear as colored arrows.

Some typical build graph icons are shown in Figure 129, page 338.

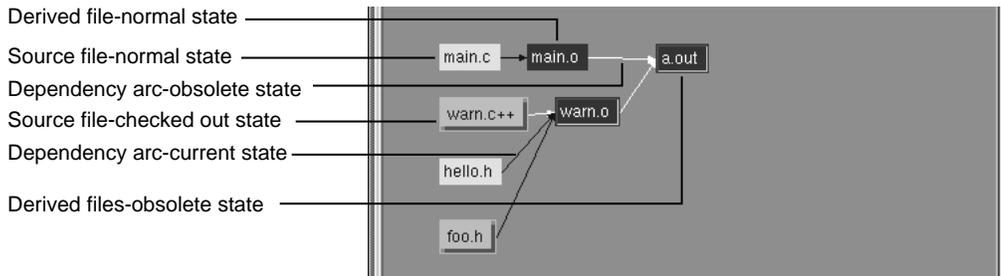


Figure 129. Build Graph Icons

The main.c and hello.h source files are in their normal state. The source files warn.c++ and foo.h are checked out and thus appear highlighted and with dropped shadows. The derived file main.o is current, since it has not changed since the last compile. The black dependency arcs indicate that the source and derived files at either end are current with each other. When an arc is highlighted, it indicates that the source has changed since the last compile. The derived files warn.o and a.out are obsolete because warn.c++ has changed.

B.9 Build Graph Control Area

The build graph control area contains a row of graph control buttons similar to the ones in the WorkShop Static Analyzer and the Call Graph View in the Performance Analyzer. The Overview button is particularly useful in the Build Analyzer because it helps you quickly find obsolete files where a lot of dependencies are involved.

The build graph control area is shown in Figure 130, page 338.

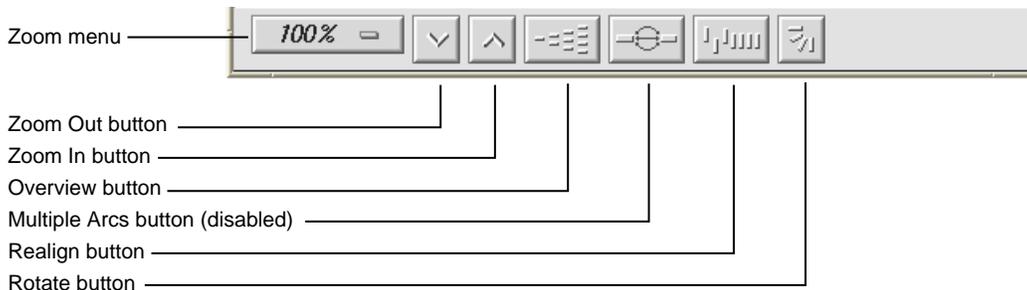


Figure 130. Build Graph Control Area

B.9.1 Build Analyzer Overview Window

Since build graphs can get quite complicated, an overview mode (similar to those in Static Analyzer and Profiling View) is supplied that lets you view the entire graph at a reduced scale. To display the overview window, you click the overview icon (see Figure 130, page 338).

Figure 131, page 339, shows a typical Build Analyzer Overview window with the resulting graph. The window has a movable viewport that lets you select the portion of the build graph displayed in Build Analyzer. Source files that have changed and derived files needing recompilation are highlighted for easy detection. In this particular color scheme, the Build Analyzer Overview window displays normal source files in turquoise, checked out source files in pink, current derived files in dark blue, and obsolete derived files in yellow. Arcs appear only in black in this window.

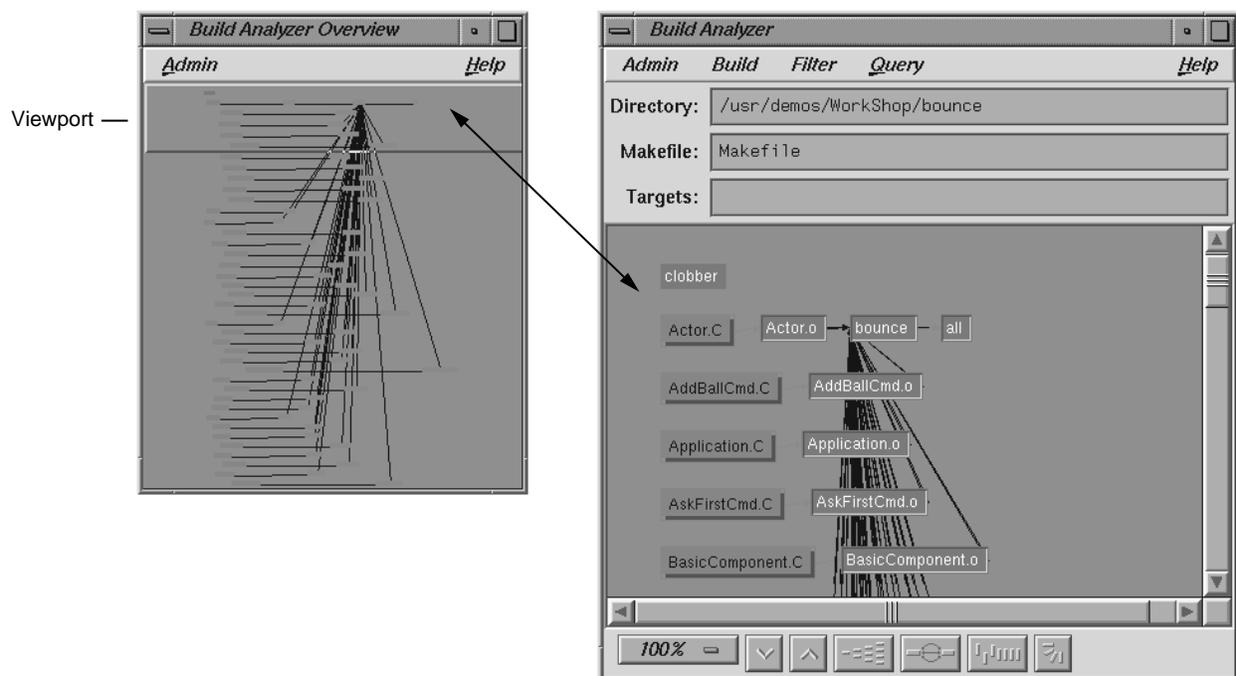


Figure 131. Build Analyzer Overview Window with Build Analyzer Graph

B.9.2 Build Analyzer Menus

The Build Analyzer window contains the following menus:

- Admin
- Build
- Filter
- Query

B.9.2.1 Admin Menu

The Admin menu provides one selection Refresh Graph Display in addition to the standard WorkShop selections.

Refresh Graph Display	Refreshes the window.
Launch Tool	Lets you execute the WorkShop tools. For more information, see the Section A.1.1, page 190.
Project	Lets you control the WorkShop tools operating on the same executable as a group. For more information, see Section A.1.1, page 190.

B.9.2.2 Build Menu

The selections in the Build menu let you perform builds as follows:

Build Default Target	Performs a make with no arguments.
Build Selected Target(s)	Performs the build(s) as entered in the Target(s): field.
Show Build Rule	Displays a dialog box showing the makefile line for the selected node.

B.9.2.3 Filter Menu

The Filter menu has only one selection:

Select files to show in graph	Opens the File Filter dialog box that lets you enter a regular expression to filter files displayed in the build graph.
-------------------------------	-------------------------------------------------------------------------------------------------------------------------

The upper list area lets you specify files to be excluded from the build graph. The lower list is for specifying files to appear in the graph.

B.9.2.4 Query Menu

The Query menu lets you request information about the build graph. The following selections are available:

Why Is This File Out Of Date?	Identifies the source files requiring this file to be recompiled. This query only applies to derived files.
What Will Changing This File Affect?	Shows all derived files dependent on this source file.

A

- accept syscall that blocks continued pthreads, 326
- access to freed memory, 125
- access to uninitialized memory, 125
- action
 - traps term, 70
- Active selection in Admin menu
 - described in jello tutorial, 59
- active toggle, 217, 224
- Ada windows description
 - Task View, 216
- Ada-specific windows description, 216
- add_source filename, dbx-style command, 312
- adding a breakpoint
 - in X/Motif analyzer tutorial, 167
- Address..., selection in disassemble menu, 292
- Admin menu, 304
 - general description, 190
 - Library search path..., 190
- admin menu
 - active toggle, 224
 - clone, 224
- alias, dbx-style command, 313
- all trap debugger command option, 77
- AllowPendingTraps
 - .Xdefaults variable, 71
- arguments, command line, 87
- Arrange, selection in structure browser display menu, 280
- Array Browser
 - subscript controls, 66
 - viewing variables with, 25, 97
 - Views menu option, 4
- Array browser
 - general description, 258
- Array Browser selection in Views menu
 - in jello tutorial, 64
- Array Browser, selection in views menu, 192

- Array field in array browser, 258
- array subscripts, 259
- array variables, 4, 258
- assign, dbx-style command, 313
- assigning values to variables, 102
- At Source Line, Traps menu option, 254
- Attach to forked processes, Multiprocess View
 - preferences option, 209
- attach, dbx-style command, 313
- auto-fork application, 135
- Automatic dereference limit, field in structure browser preferences box, 285

B

- basic debugger usage, 11
- blocking anomalies
 - and pthreads, 155
- boundary overrun, 125
- boundary underrun, 125
- breakpoint, 3
- breakpoint results, viewing, 174
- breakpoint, adding for a widget, 167
- breakpoints tab, 172
- breakpoints, setting (Fix+Continue tutorial), 118
- breakpoints, setting for a class, 169
- Build Analyzer
 - to examine build dependencies, 7
- Build analyzer, 335
- build dependencies
 - examine with Build Analyzer, 7
- Build environment window, 308
- Build manager, 329
- build path, 109
- Build view, 329

C

- C expressions, 103
 - C function calls, 104
 - C tutorial code example, 15
 - C++ exception trap, 72
 - C++ expressions, 104
 - Call stack
 - Views menu option, 4
 - Call Stack selection in Views menu, 58
 - Call Stack View, 93
 - Call stack view, 271
 - Call Stack View Window
 - introduction to, 28
 - Call Stack View with Fix+Continue, 311
 - Call Stack, selection in views menu, 192
 - call, dbx-style command, 313
 - callback breakpoints examiner, 230
 - callback context, viewing, 170
 - callback examiner, 170, 246
 - callstack view
 - in X/Motif analyzer tutorial, 174
 - catch, dbx-style command, 313
 - change id, 109
 - change values
 - in Array Browser, 26
 - changing code from the command line, 116
 - checking for out-of-bounds array accesses, 32
 - checking for uninitialized variable used in calculations, 34
 - classes, examining widget, 168
 - Clear all, selection in structure browser display menu, 282
 - Clear button in Trap Manager
 - in jello tutorial, 57
 - Clear trap selection in Traps menu, 75
 - clearcalls, dbx-style command, 313
 - ClearCase, 7
 - Click for help, selection in help menu, 205
 - Click To Evaluate
 - viewing variables with, 22
 - Click to Evaluate
 - viewing variables with, 97
 - clone current window, 217, 224
 - close current window, 218, 224
 - close syscall that blocks continued pthreads, 326
 - Close, selection in admin menu, 192
 - code
 - redefined vs. compiled, 108
 - code, changing (tutorial), 115
 - code, changing from command line, 116
 - code, comparing original to redefined, 121
 - code, deleting changed, 116
 - code, switching between compiled and redefined
 - in Fix+Continue, 121
- Col button in Array Browser
 - in jello tutorial, 65
 - Column width..., selection in array browser display menu, 261
 - Combine threads at same location,
 - Multiprocess View preferences option, 210
 - Command field in the Main View window, 87
 - command line interface with Fix+Continue, 310
 - command syntaxes for traps, 78
 - comparing code, original vs. redefined, 121
 - comparing function definitions
 - (in Fix+Continue tutorial), 122
 - compiled code
 - distinguished from refined code, 108
 - compiling
 - introductory tips, 12
 - with the malloc library (heap corruption), 126
 - compiling with the malloc library, 126
 - Condition field in trap manager, 82
 - Config menu in structure browser, 278
 - Config menu in trap manager, 76
 - cont in, dbx-style command, 313
 - cont to, dbx-style command, 313
 - continue
 - of single 6.5 POSIX pthread, 153
 - continue signal, dbx-style command, 314
 - Continue button description, 19
 - Continue button in the Main View window, 56, 88

Continue Even If Line Has Changed, toggle in
 Fix+Continue Preferences dialog, 204

Continue To selection in PC menu, 91

Continue to, selection in disassembly view pc
 menu, 292

Continue to, selection in pc menu, 201

continue, dbx-style command, 313

control flow constructs, 160

controlling program execution options, 19

Copy Traps Even on Changed Lines, toggle in
 Fix+Continue Preferences dialog, 204

Copy Traps On Previous Definition, toggle in
 Fix+Continue Preferences dialog, 203

Copy traps to forked processes, Multiprocess
 View preferences option, 209

Copy traps to sproc'd processes
 in multiprocess tutorial, 142

Copy traps to sproc'd processes, Multiprocess
 View preferences option, 209

core file analysis
 for scientific programs, 36

corefile, dbx-style command, 314

corrupt program
 heap corruption tutorial program, 129

creat syscall that blocks continued pthreads, 326

customizing the debugger, 177
 changing X window system resources, 180
 startup file for, 177
 user-defined buttons, 178
 with scripts, 177

cvd
 Main View window (X/Motif analyzer
 tutorial), 162

cvd> command line
 viewing variables with, 21

cvdrc file
 for customizing the debugger, 177

Cycle Count field in Trap Manager
 in jello tutorial, 58

Cycle count field in trap manager, 83

D

data structures, 4

dbx commands, 312

Debugger
 Call Stack View with Fix+Continue, 311
 changes to views with Fix+Continue, 309
 command line interface with
 Fix+Continue, 310
 exiting, 67
 how to start, 47
 Main View window with Fix+Continue, 309
 program execution control, 87
 Trap Manager with Fix+Continue, 311

debugger
 how to customize, 177
 Main View window (X/Motif analyzer
 tutorial), 162

Debugger Command Line, , 312

Debugger command line, 2

Debugger main features, 1

Debugger views, 93, 271

Debugger with Fix+Continue support
 Fix+continue
 debugger support with, 109

debugging
 debugging a multiprocess C program, 139
 debugging a multiprocess fortran program, 146
 fortran multiprocess debugging session, 146
 general fortran debugging hints, 146
 tips and features, 11

debugging with the X/Motif analyzer, 8

Default field count, field in structure browser
 type formatting, 286

Default iconic width, field in structure browser
 preferences box, 285

Default iconic width, field in structure browser
 type formatting box, 286

Default state, field in structure browser type
 formatting box, 286

Default structure field count, field in structure
 browser preferences box, 285

- Default structure width, field in structure
 - browser preferences box, 285
- Default structure width, field in structure
 - browser type formatting box, 286
- delete all, dbx-style command, 314
- delete trap, dbx-style command, 315
- delete_changes, dbx-style command, 314
- delete_source, dbx-style command, 314
- deleting changed code (tutorial), 116
- Dereference ptrs by default, field in structure
 - browser preferences box, 285
- Dereference ptrs, selection in structure browser
 - node menu, 284
- detach, dbx-style command, 315
- Detach, selection in admin menu, 191
- Detail, selection in structure browser
 - submenu, 282
- difference tools
 - in Fix+Continue, 124
- disable all, dbx-style command, 315
- disable, dbx-style command, 315
- disable_changes, dbx-style command, 315
- disabling traps
 - in jello tutorial, 57
- Disassemble File Dialog, 294
- Disassemble Function Dialog, 293
- Disassemble menu in disassembly view, 292
- disassembled code, 5
- Disassembly View
 - Preferences, , 294
 - Views menu option, 5
- Disassembly View, selection in views menu, 192
- display area in structure browser, 280
- DISPLAY environment variable for debugging, 11
- Display menu
 - in Main View window, 197
 - Main View window, 340
- Display menu in structure browser, 279
- Display menu in traps manager, 76
- display, dbx-style command, 315
- Display, selection in structure browser display menu, 280
- divisions by zero

- how to find in scientific programs, 35
- dmi syscall that blocks continued pthreads, 326
- double frees, 125
- down, dbx-style command, 315
- dump, dbx-style command, 315
- DUMPCORE environment variable, 182

E

- edit source code
 - as shown in jello tutorial, 53
- editors
 - fork editor, 6
 - how to access from the Main View window, 6
- editres requests
 - and X/Motif analyzer, 161
- enable all, dbx-style command, 315
- enable trap, dbx-style command, 316
- enable_changes, dbx-style command, 316
- Entry Function, Traps menu option, 254
- environment variables
 - CVDINIT, 177
 - DUMPCORE, 182
 - MALLOC_CLEAR_FREE, 127
 - MALLOC_CLEAR_FREE_PATTERN, 127
 - MALLOC_CLEAR_MALLOC, 127
 - MALLOC_CLEAR_MALLOC_PATTERN, 127
 - MALLOC_FASTCHK, 127
 - MALLOC_MAXMALLOC, 128
 - MALLOC_NO_REUSE, 128
 - MALLOC_TRACING, 128
 - MALLOC_VERBOSE, 128
 - _RLD_LIST, 126
 - setting in Execution View, 92
 - setting to detect heap corruption, 127
- environment variables for debugging, 11
- erroneous frees, 125
- Error messages window, 307
- evaluating expressions, 99
- evaluating expressions in C++, 104
- evaluating expressions in Fortran, 105

- event examiner, 248
 - event-handler breakpoints examiner, 231
 - examine menu, 224
 - examiner
 - breakpoint, 167
 - breakpoints, 226
 - callback, 170, 246
 - callback breakpoints, 229
 - event, 248
 - event-handler breakpoints, 231
 - graphics context (GC), 248
 - input-handler breakpoints, 235
 - ixmap, 249
 - resource-change breakpoints, 233
 - state-change breakpoints, 236
 - timeout-procedure breakpoints, 234
 - trace, 240
 - tree, 243
 - tree examiner, 166
 - widget, 242
 - widget class, 250
 - widget examiner, 165
 - window, 171, 247
 - X-event breakpoints, 237
 - examiners
 - overview, in X/Motif analyzer, 159
 - selections, in X/Motif analyzer, 160
 - examining data, 4
 - in jello tutorial, 58
 - examining program data, 93
 - examining widget classes, 168
 - examining widgets
 - in X/Motif analyzer tutorial, 166
 - exception trap, 72
 - Exception View description, 219
 - Exception View, selection in views menu, 192
 - executable
 - run directly from the Main View window, 39
 - execution control buttons, 88
 - in Main View window, 88
 - Execution View, 92
 - Execution View description, 12, 206
 - Execution View, selection in views menu, 192
 - Exit Function, Traps menu option, 254
 - Exit, selection in admin menu, 192
 - exiting the debugger, 67
 - Expression column in Expression View, 99
 - Expression column in expression view, 275
 - expression count, dbx-style command, 316
 - Expression field in Structure Browser
 - in jello tutorial, 63
 - Expression field in structure browser, 279
 - Expression View
 - viewing variables with, 23
 - Expression view, 273
 - Views menu option, 4
 - Expression View selection in Views menu
 - in jello tutorial, 61
 - Expression View Window
 - for evaluating expressions, 99
 - Expression View window
 - viewing variables with, 98
 - Expression View, selection in views menu, 192
 - Expression, selection in structure browser
 - display submenu, 280
 - expressions
 - C++, 104
 - for C operations, 103
 - Fortran, 105
 - External Editor Command, text field in
 - Fix+Continue Preferences dialog, 203
- F**
- fcntl syscall that blocks continued pthreads, 326
 - fibonacci program
 - Fibonacci program used in multiprocessing tutorial, 140
 - File Browser for locating and loading files, 40
 - File Browser, selection in views menu, 193
 - file dbx-style command, 317
 - File Difference Tool, text field in Fix+Continue Preferences dialog, 203
 - File menu, source view, 211

- “File...”, selection in disassemble menu, 293
- files (source)
 - loading, 39
 - managing, 39
- files, comparing source code with xdiff, 124
- files, finding for Fix+Continue, 110
- finding files for Fix+Continue, 110
- Fix+ continue
 - change id, 109
- Fix+Continue
 - Build environment window, 308
 - comparing original to redefined code, 121
 - Error message window, 307
 - GUI, 301
 - keyboard accelerators, 205
 - menu selections and operations, 201
 - Session, 305
 - Show Difference submenu, 202
 - Status window, 302
 - switching between compiled and redefined code, 121
 - tutorial, 111
 - View submenu, 203
- Fix+continue
 - basic cycle description, 107
 - breakpoints, 118
 - build path, 109
 - changing code (tutorial), 115
 - changing code from the command line, 116
 - deleting changed code (tutorial), 116
 - editing a function (tutorial), 113
 - functionality and features, 107
 - interface description, 109
 - introduction and tutorial, 107
 - overview, 8
 - redefining function, 112
 - redefining functions with, 107
 - restrictions, 110
 - sample session (tutorial), 111
 - setting traps for, 118
 - Status window, 121
 - WorkShop integration, 108
- Fix+Continue menu, 305

- Fix+Continue Preferences submenu, 203
- fork application, 135
- fork editor, 6
- Fork Editor, selection in Source menu, 195
- fork processes, 136
- Format menu in Expression View, 99
 - in jello tutorial, 61
- Format menu in expression view, 274, 275
- Format menu in structure browser, 279
- Format menu in variable browser, 288
- Formatting Fields in Structure Browser, , 284
- fortran
 - debugging a multiprocess fortran program, 146
 - fortran multiprocess debugging session, 146
 - general fortran debugging hints, 146
- Fortran 90 tutorial code example, 13
- Fortran expressions, 105
- Fortran function calls, 107
- Fortran variables supported in expressions, 106
- frames, 93, 271
- free call errors
 - heap corruption errors, 126
- free run
 - of single 6.5 POSIX pthread, 153
- fsync syscall that blocks continued pthreads, 326
- func dbx-style command, 317
- function definitions, comparing
 - (in Fix+Continue tutorial), 122
- function, editing, 113
- function, redefining
 - Fix+continue, 112
- Function..., selection in disassemble menu, 292
- functions, identifying, 109

G

- g option for compiling, 12
- Geometry, selection in structure browser node menu, 283
- getmsg syscall that blocks continued pthreads, 326

getpmsg syscall that blocks continued
 pthreads, 326
 getstarted tutorial directory, 13
 getting started with debugger: tutorial, 11
 givenfile dbx-style command, 317
 GLdebug, 191
 GLdebug, selection in admin menu, 191
 Go to line..., selection in Source menu, 195
 goto dbx-style command, 317
 Goto dialog box, 195
 graphics context (GC) examiner, 248
 grelnotes release notes , 9
 Group Trap Default toggle, 75
 Group Trap Default, Traps menu option, 254

H

heap corruption
 access to freed memory, 125
 access to uninitialized memory, 125
 boundary overrun definition, 125
 boundary underrun definition, 125
 compiling with malloc library, 126
 detection, 125
 double frees, 125
 erroneous frees, 125
 how to find heap corruption errors, 126
 trapping errors using the malloc library, 128
 typical problems, 125
 Help menu, 205
 Hide Icons, selection in Display menu, 197
 Hide Line Numbers, selection in display
 menu, 197

I

Iconic, selection in structure browser
 submenu, 282
 Iconify, selection in admin menu, 191
 identifying functions, 109
 ignore dbx-style command, 317

include files
 and Fix+Continue, 110
 index identifiers in array browser, 260
 index maximum specification in array
 browser, 260
 index minimum specification in array
 browser, 260
 index sliders in array browser, 260
 index values in array browser, 260
 Index..., selection in help menu, 205
 Indexing expression field in array browser, 259
 input-handler breakpoints examiner, 235
 Insert source..., selection in Source menu, 195
 integration of MIPSpro WorkShop tools, 5
 interface, command line, 310
 introductory tips and features for debugging, 11
 ioctl syscall that blocks continued pthreads, 326

J

jello program, 48
 Jump To selection in PC menu, 91
 Jump to, selection in disassembly view pc
 menu, 292
 Jump to, selection in pc menu, 201

K

keyboard accelerators in Fix+Continue, 205
 Keys & shortcuts, selection in help menu, 205
 Kill button description, 19
 Kill button in the Main View window, 88
 kill dbx-style command, 317

L

Language menu in Expression View, 99
 in jello tutorial, 61
 Language menu in expression view, 274, 275

- Language menu in variable browser, 288
- Launch, selection in admin menu, 192
- launching debugger in Multiprocess View, 141
- launching the X/Motif analyzer, 8
- launching X/Motif analyzer, 163
- Levels to open, Multiprocess View preferences
 - option, 209
- Library search path dialog box, 190
- Linked list, selection in structure browser
 - display menu, 281
- lint
 - option for debugging scientific programs, 32
- list dbx-style command, 317
- list_changes, dbx-style command, 318
- lmalloc_ss library
 - for finding heap corruption problems, 126
- Load expressions... selection in Expression
 - View—> Config menu, 102
- load files
 - directly into the Main View window, 39
 - through the File Browser Window, 40
 - through the Open dialog box, 41
 - with File Browser, 41
- Load settings..., selection in admin menu, 191
- Load traps... selection in config menu in trap
 - manager, 84
- loading source files, 39
- locally distributed application, 135
- locate files
 - with File Browser, 41
- lockf syscall that blocks continued pthreads, 326

M

- Main View window, 309
 - Command field, 87
 - Continue button, 88
 - control panel, 87
 - Display menu, 197, 340
 - execution control buttons, 88
 - general description, 2
 - Kill button, 88

- PC menu, 91
- Run button, 88
- Sample button, 91
- Status field, 88
- Step Into button, 88
- Step Over button, 90
- Stop button, 88
- Main View window (X/Motif analyzer
 - tutorial), 162
- Make Editable, selection in Source menu, 195
- Make Read Only, selection in Source menu, 195
- malloc call failing
 - heap corruption error, 126
- malloc library
 - compiling with, 126
- MALLOC_CLEAR_FREE_PATTERN, 127
- MALLOC_CLEAR_MALLOC, 127
- MALLOC_CLEAR_MALLOC_PATTERN, 127
- MALLOC_FASTCHK, 127
- MALLOC_FASTCHK environment variable
 - and heap corruption errors, 126
- MALLOC_MAXMALLOC, 128
- MALLOC_NO_REUSE, 128
- MALLOC_TRACING, 128
- MALLOC_VERBOSE, 128
- managing source files, 39
- Maximize, selection in structure browser node
 - submenu, 283
- memalign call with improper alignment
 - heap corruption error, 126
- memory locations, 5
- Memory View
 - Views menu option, 5
- Memory view, 299
- Memory view mode menu, 301
- Memory View, selection in views menu, 193
- menu selections and operations
 - Fix+Continue, 201
- Message window
 - Admin menu, 307
 - buttons, 307
 - View menu, 307

-
- Messages window, 307
 - Minimize, selection in structure browser node submenu, 283
 - Minimum lines around current instruction field in disassembly view preferences box, 295
 - mq_open syscall that blocks continued pthreads, 326
 - msgrcv syscall that blocks continued pthreads, 326
 - msgsnd syscall that blocks continued pthreads, 326
 - msync syscall that blocks continued pthreads, 326
 - multiple process debugging description and introduction, 135
 - multiprocess
 - add and remove processes, 139
 - additional main view windows for, 139
 - debugging a multiprocess fortran program, 146
 - fortran multiprocess debugging session, 146
 - multiprocess traps, 138
 - preferences, 139
 - using trap manager to control trap inheritance (tutorial), 144
 - view control buttons, 138
 - viewing process status, 137
 - multiprocess traps, 77
 - Multiprocess View
 - launching debugger in, 141
 - to control execution, 142
 - Views menu option, 4
 - Multiprocess View description, 206
 - administrative functions, 208
 - control buttons, 208
 - preferences, 209
 - Multiprocess View preferences options
 - Attach to forked processes, 209
 - Combine threads at same location, 210
 - Copy traps to forked processes, 209
 - Copy traps to sproc'd processes, 209
 - Levels to open, 209
 - Resume child after attach on fork, 210
 - Resume child after attach on sproc, 210
 - Resume parent after fork, 210
 - Resume parent after sproc, 210
 - Show Thread Status vs Thread State, 210
 - Show/Hide buttons, 210
 - Show/Hide Header Information, 210
 - Stack Depth, 209
 - Multiprocess View window, 136
 - Multiprocess View, selection in admin menu, 190
 - multiprocessing debugging, 135
- N**
- N... selection in step into menu, 89
 - N... selection in Step Over menu, 89
 - nanosleep syscall that blocks continued pthreads, 326
 - next dbx-style command, 318
 - nexti dbx-style command, 318
 - Node menu in structure browser, 279
 - Node pop-up menu in structure browser, 280
 - Normal, selection in structure node submenu, 282
 - Number of instructions to disassemble field in disassembly view preferences box, 295
- O**
- Open dialog box
 - loading files, 41
 - Open Recent..., selection in Source menu, 194
 - open syscall that blocks continued pthreads, 326
 - Open..., selection in Source menu, 194
 - overflows
 - how to find in scientific programs, 35
 - Overview, selection in help menu, 205
- P**
- path remapping case example, 45
 - path remapping description, 43

- Pattern layout, selection in structure browser
 - node menu, 284
- pause syscall that blocks continued pthreads, 326
- PC, 201
- PC menu, 91, 201
 - Continue To, 91
 - Jump To, 91
- PC menu in disassembly view, 292
- PC menu in main view window, 91
- pending trap
 - definition, 71
- performance analyzer
 - how to access from Main View window, 5
- performance data
 - Sample button, 91
- pgrp trap debugger command option, 77
- pixmap examiner, 249
- poll syscall that blocks continued pthreads, 326
- pollpoint, 4
- pollpoint trap debugger command option, 82
- pread syscall that blocks continued pthreads, 326
- preferences for multiprocesses, 139
- preparing the fileset for X/Motif analyzer
 - tutorial, 162
- print expression dbx-style command, 318
- printf expression dbx-style command, 318
- prinfo expression dbx-style command, 319
- printregs dbx-style command, 319
- printx expression dbx-style command, 319
- process group, 136
- Process menu description, 215
- Process Meter description, 213
 - Charts menu, 214
 - Scale menu, 215
- Process Meter, selection in views menu, 193
- Product information, selection in help menu, 205
- Program counter
 - definition, 59
- program counter, 91, 201
- Program data, 93
- program execution
 - options for controlling, 19

- program execution control, 87
 - Main View control panel, 87
 - PC menu, 91
- program output, tracking, 110
- pthread
 - and blocking anomalies, 155
 - and scheduling anomalies, 153
 - debugging a pthreaded program, 152
 - debugging session, 157
 - differences between 6.4 and 6.5 pthreads, 156
 - how to continue a single pthread, 155
 - pthread debugging hints, 156
 - syscalls which block continued pthreads, 326
 - user-level continue of single 6.5 POSIX pthread, 153
 - viewing pthreaded applications, 138
- putmsg syscall that blocks continued pthreads, 326
- putpmsg syscall that blocks continued pthreads, 326
- pwd dbx-style command, 319
- pwrite syscall that blocks continued pthreads, 327

Q

- quit dbx-style command, 319

R

- Raise, selection in admin menu, 191
- Raise, selection in structure browser node
 - submenu, 283
- RCS, 7
- read syscall that blocks continued pthreads, 326
- Read-Only
 - debugger status, 109
- readv syscall that blocks continued pthreads, 326
- realloc call errors
 - heap corruption error, 126
- Recompile, selection in the Source menu, 195

recompiling code, 7
 recv syscall that blocks continued pthreads, 327
 recvfrom syscall that blocks continued pthreads, 327
 recvmsg syscall that blocks continued pthreads, 327
 redefine dbx-style command, 319
 redefined code
 distinguished from compiled code, 108
 redefining a function in Fix+Continue (tutorial), 112
 redefining functions, 107
 Register name display format field in disassembly view preferences box, 295
 Register View
 Views menu option, 5
 Register view, 296
 Register view formatting, 299
 Register view preferences dialog box, 299
 Register view window, 297
 Register View, selection in views menu, 193
 registers, 5
 release notes, 9
 relnotes, 9
 Remap paths... selection in session menu, 44
 Remove, selection in structure browser node menu, 284
 removing traps with mouse, 73
 replace_source, dbx-style command, 320
 rerun, dbx-style command, 320
 Reset To Factory Defaults, toggle in
 Fix+Continue Preferences dialog, 205
 resource-change breakpoints examiner, 233
 restrictions and limitations of X/Motif analyzer, 161
 Result column in Expression View, 99
 Result column in expression view, 275
 Resume child after attach on fork, Multiprocess View preferences option, 210
 Resume child after attach on sproc, Multiprocess View preferences option, 210
 Resume child after attach on sproc in multiprocess tutorial, 142

Resume parent after fork, Multiprocess View preferences option, 210
 Resume parent after sproc, Multiprocess View preferences option, 210
 Return button description, 20
 Return button in the Main View window
 Main View
 Return button, 90
 return dbx-style command, 320
 row/column toggles in array browser, 260
 RUN button description, 19
 Run button in the Main View window, 88
 run dbx-style command, 320
 runtime_check, dbx-style command, 320

S

Sample At Function Entry, 200
 Sample at Function Entry selection in traps submenu, 74
 Sample at Function Exit selection in traps submenu, 74
 Sample at function exit, selection in traps submenu, 200
 Sample button in the Main View window, 91
 sample session
 setting up for X/Motif analyzer, 162
 sample session setup
 Fix+continue, 111
 sample trap, 72
 sample trap command, 77
 Sample Trap Default, Traps menu option, 254
 sample traps, 3
 save as text, 218, 224
 Save as text..., selection in Source menu, 195
 Save as... selection in Source menu, 194
 Save deactivated code during File Save, toggle in Fix+Continue Preferences dialog, 204
 Save expressions... selection in Expression View —>Config menu, 102
 Save settings..., selection in admin menu, 191

- Save traps... selection in config menu in trap manager, 84
- Save, selection in Source menu, 194
- save_changes, dbx-style command, 321
- saving changes to source file (Fix+Continue tutorial), 118
- scheduling anomalies
 - and pthreads, 153
- scientific programs
 - checking for out-of-bounds array accesses, 32
 - checking for uninitialized variables used in calculations, 34
 - core file analysis for, 36
 - finding divisions by zero and overflows, 35
 - suggestions for debugging serial execution of, 31
 - using lint for debugging, 32
- scripts
 - for customizing the debugger, 177
- Search field in trap manager, 84
- Search window, 27
- "Search", selection in structure browser display menu, 281
- Search..., selection in Source menu, 195
- searching for character strings, 27
- searching in the jello tutorial, 52
- select syscall that blocks continued pthreads, 327
- Select, selection in structure browser node menu, 283
- selection, 224
- selection in structure browser display menu, 280
- semctl syscall that blocks continued pthreads, 327
- semget syscall that blocks continued pthreads, 327
- semop syscall that blocks continued pthreads, 327
- send syscall that blocks continued pthreads, 327
- sendmsg syscall that blocks continued pthreads, 327
- sendto syscall that blocks continued pthreads, 327
- Session submenu, 305
- Set Trap
 - selection in Traps menu, 74
- Set trap, selection in traps menu, 199, 200
- setbuildenv, dbx-style command, 321
- setting traps
 - in jello tutorial, 55
 - introduction, 20
- setting traps at the cvd> command line, 73
- setting traps with the mouse, 73
- setting traps with the Traps menu, 73
- sginap syscall that blocks continued pthreads, 327
- sh dbx-style command, 321
- Show Difference submenu, 202
- Show embedded source annotation field in disassembly view preferences box, 295
- Show Icons, selection in Display menu, 197
- Show instruction value field in disassembly view preferences box, 296
- Show jal target numerically field in disassembly view preferences box, 296
- Show Line Numbers selection in display menu, 197
- Show machine address field in disassembly view preferences box, 296
- Show overview, selection in structure browser display menu, 282
- Show source file and line number field in disassembly view preferences box, 296
- Show Thread Status vs Thread State, Multiprocess View preferences option, 210
- Show/Hide buttons, Multiprocess View preferences option, 210
- Show/Hide Header Information, Multiprocess View preferences option, 210
- show_changes, dbx-style command, 322
- show_diff, dbx-style command, 322
- showbuildenv, dbx-style command, 322
- showthread dbx-style command, 322
- Signal panel, 84
- Signal Panel, selection in views menu, 193
- signal trap debugger command option, 81
- signal traps, 4
- SIGTERM signal, 136
- source annotation column
 - traps, 74

- source code status indicator, 109
- source control
 - using configuration management tools for, 7
- source dbx-style command, 322
- source file
 - saving changes to (Fix+Continue tutorial), 118
- source files
 - comparing code files with xdiff, 124
 - loading, 39
 - managing, 39
- Source view
 - File menu, 211
- Source View description, 210
- Source View, selection in views menu, 193
- spreadsheet area in array browser, 260
- spool processes, 136
- ssmalloc_error, 128
- Stack Depth, Multiprocess View preferences
 - option, 209
- stack frame, 59
- stack frames, 94, 272
- starting the debugger, 47
- starting, program execution, 88
- startup file
 - for customizing the debugger, 177
- State, selection in structure browser node
 - menu, 282
- state-change breakpoints examiner, 236
- static analyzer
 - how to access from the Main View window, 6
- status dbx-style command, 322
- Status field in the Main View window, 87, 88
- Status window, 302
 - Admin menu, 304
 - Fix+continue, 121
 - Fix+Continue menu, 305
 - Fix+Continue Preferences submenu, 203
 - View menu, 304
- status, viewing, 121
- step dbx-style command, 323
- step indicators in array browser, 260
- Step Into button, 12
- Step into button in disassembly view, 291
- Step Into button in the Main View window, 88
- Step Over button, 12
- Step Over button description, 20
- Step over button in disassembly view, 291
- Step Over button in the Main View window, 89
- stepi dbx-style command, 323
- Stop All Default, Traps menu option, 254
- Stop All Defaults toggle, 75
- stop at dbx-style command, 323
- Stop at Function Entry selection in traps
 - submenu, 74
- Stop at function entry, selection in traps
 - submenu, 199
- Stop at Function Exit selection in traps
 - submenu, 74
- Stop at function exit, selection in traps
 - submenu, 200
- Stop button description, 19
- Stop button in the Main View, 88
- stop dbx-style command, 323
- stop exception, dbx-style command, 323
- stop in, dbx-style command, 323
- Stop Into button description, 19
- stop trap, 71
- stop trap command, 77
- Stop Trap Default, Traps menu option, 254
- stop traps, 3
 - in jello tutorial, 55
- stopping, process execution, 88
- Structure Browser
 - viewing variables with (C code only), 97
 - Views menu option, 4
- Structure browser
 - general description, 277
- Structure Browser (C code only)
 - viewing variables with, 27
- Structure browser preferences dialog box, 284
- Structure Browser selection in views menu, 63
- Structure Browser, selection in views menu, 193
- subscript controls in Array Browser, 65
- Subscript Controls panel in Array Browser window, 26

- subscripts
 - array, 259
- Switch executable..., selection in admin menu, 191
- Switch process dialog box, 191
- Switch process..., selection in admin menu, 191
- switching between compiled and redefined code in Fix+Continue, 121
- syscall dbx-style command, 324
- Syscall panel, 84
- Syscall Panel, selection in views menu, 193
- syscall trap debugger command option, 81
- syscalls
 - accept, 326
 - close, 326
 - creat, 326
 - dmi, 326
 - fcntl, 326
 - fsync, 326
 - getmsg, 326
 - getpmsg, 326
 - ioctl, 326
 - lockf, 326
 - mq_open, 326
 - msgrcv, 326
 - msgsnd, 326
 - msync, 326
 - nanosleep, 326
 - open, 326
 - pause, 326
 - poll, 326
 - pread, 326
 - putmsg, 326
 - putpmsg, 326
 - pwrite, 327
 - read, 326
 - readv, 326
 - recv, 327
 - recvfrom, 327
 - recvmsg, 327
 - select, 327
 - semctl, 327
 - semget, 327

- semop, 327
- send, 327
- sendmsg, 327
- sendto, 327
- sginap, 327
- which block continued pthreads, 326
- write, 327
- writv, 327
- system calls
 - traps, 4
 - which block continued pthreads, 326

T

- tab overflow area, 172
- tabs, 173
- tabs, X/Motif Analyzer examiner, 225
- Task View admin menu
 - active toggle, 217
 - clone, 217
 - close, 218
 - save as text, 218
- Task View config menu, 218
- Task View display menu, 218
- Task View layout menu, 218
- Task View window description, 216
- Task View, selection in views menu, 193
- timeout procedure breakpoints examiner, 234
- tools integration, 5
- trace dbx-style command, 324
- trace examiner, 240
- tracking program output, 110
- trap actions, 70
- trap condition, 82
- trap examples, 77
- trap icons, 75
- Trap Manager
 - in jello tutorial, 56
- Trap manager menus, 76
- Trap Manager Window
 - introduction to, 76

- Trap Manager with Fix+Continue, 311
- Trap Manager, selection in views menu, 193
- trap triggers, 70
- trap types, 71
- traps, 3
 - adding a breakpoint for a widget, 167
 - all trap debugger command option, 77
 - C++ exception trap, 72
 - command syntaxes for, 78
 - descriptive overview, 69
 - disabling in jello tutorial, 57
 - enabling and disabling, 84
 - for multiprocesses, 138
 - how to set, 72
 - in jello tutorial, 55
 - introduction, 20
 - multiprocess traps, 77
 - one-time, 91
 - pending trap, 71
 - pgrp trap debugger command option, 77
 - removing with mouse, 73
 - sample trap, 72
 - sample trap command, 77
 - set at signal-name, 81
 - set at time interval, 82
 - set for C++ exception, 82
 - set in filename at line-number, 78
 - set on entry to function, 79
 - set on entry to sys-call-name, 81
 - set on exit from function, 79
 - set on exit from sys-call-name, 81
 - set on instruction-address, 79
 - setting a watchpoint for address and size, 80
 - setting a watchpoint on specified expression, 80
 - setting at the cvd> command line, 73
 - setting conditions, 82
 - setting cycle count, 83
 - setting with mouse, 73
 - setting with the Traps menu, 73, 83
- Signal panel, 84
- stop trap, 71
- stop trap command, 77
- Syscall panel, 84
- trap examples, 78
- triggering, 70
- traps definition, 70
- Trap Manager “Traps” menu options
 - At Source Line, 254
 - Entry Function, 254
 - Exit Function, 254
 - Group Trap Default, 254
 - Sample Trap Default, 254
 - Stop All Default, 254
 - Stop Trap Default, 254
- Traps menu in the Main View window, 74
- Traps menu in trap manager, 76, 83
- traps terminology , , 70
- traps, setting (Fix+Continue tutorial), 118
- tree examiner, 166, 243
- Tree, selection in structure browser display submenu, 281
- trigger
 - traps term, 70
- triggering traps, 70
- true multiprocess program, 135
- tutorials
 - debugging a multiprocess C program, 139
 - Fix+Continue, 111
 - fortran multiprocess debugging session, 146
 - heap corruption tutorial, 129
 - how to load source files, 39
 - introductory: C code, 15
 - introductory: Fortran 90, 13
 - jello program, 47
 - path remapping, 45
 - starting a multiprocess session, 137
 - X/Motif analyzer, 161
- Type color, field in structure browser type formatting box, 287
- Type Formatting Dialog, , 286
- Type name, field in structure browser type formatting box, 286

U

- unalias dbx-style command, 324
- undisplay dbx-style command, 324
- uninitialized variables, 34
- unsetbuildenv, dbx-style command, 324
- up, dbx-style command, 325
- Update, selection in structure browser display menu, 281
- use, dbx-style command, 325
- using the X/Motif analyzer, 159

V

- Variable Browser
 - in jello tutorial, 60
 - View menu option, 4
 - viewing variables with, 23, 98
- Variable browser
 - general description, 288
- Variable Browser selection in Views menu, 60
- Variable Browser, selection in views menu, 193
- variables
 - assigning values to, 102
 - compiled with -n32 -g, 288
 - options for viewing, 20, 96
 - unused, 288
 - viewing at the cvd> command line, 97
 - viewing with Array Browser, 25
 - viewing with Click To Evaluate, 22
 - viewing with Click to Evaluate, 97
 - viewing with cvd> command line, 21
 - viewing with Expression View, 23
 - viewing with Structure Browser (C code only), 27
 - viewing with the Array Browser, 97
 - viewing with the Expressions View window, 98
 - viewing with the Structure Browser (C code only), 97
 - viewing with the Variable Browser, 98
 - viewing with Variable Browser, 23

- versioning
 - and accessing configuration management tools, 7
- Versioning, selection in Source menu, 196
- view changes in debugger, 309
- View menu, 304
- View submenu, 203
- view windows
 - Multiprocess View, 136
- view, call stack, 311
- viewing data, 4
- viewing status (Fix+Continue tutorial), 121
- Views menu in the Main View window, 192

W

- Warn Unfinished Edits Before Continue, toggle
 - in Fix+Continue Preferences dialog, 204
- Warn Unfinished Edits Before Run, toggle in Fix+Continue Preferences dialog, 204
- watch command
 - in jello tutorial, 57
- watch trap debugger command option, 80
- watch, dbx-style command, 325
- Watchpoint
 - definition, 57
- watchpoint, 57
- watchpoints, 4
- whatis dbx-style command, 325
- when at, dbx-style command, 325
- when in, dbx-style command, 325
- where, dbx-style command, 325
- which, dbx-style command, 325
- widget class examiner, 250
- widget class menu item, 225
- widget classes, examining, 168
- widget examiner, 242
- widget hierarchy, 166
- widget item, 224
- widget structure, navigating, 163
- widget tree menu item, 224

widgets, examining, 166
window attributes, viewing, 171
window examiner, 171, 247
WorkShop integration, 108
Wrapped display, selection in array browser
 display menu, 261
write syscall that blocks continued pthreads, 327
writev syscall that blocks continued pthreads, 327

X

X event menu item, 225
X graphics context menu item, 225
X pixmap menu item, 225
X window system resources
 changing to customizing the debugger, 180
X-event breakpoints examiner, 237
X/Motif Analyzer
 global objects description, 223
X/Motif analyzer
 “additional” features, 172
 debugging with, 8
 default view, 165
 inspecting data with, 160
 inspecting the control flow with, 160
 launching, 163
 navigating widget structure, 163
 overview, 159
 restrictions and limitations, 161
 starting, 8
 tracing execution with, 161
X/Motif Analyzer admin menu, 224
 close, 224
 save as text, 224
X/Motif Analyzer breakpoint type option
 button, 228

X/Motif Analyzer breakpoints examiner, 226
 callback, 229
 event-handler, 231
 input-handler, 235
 resource-change, 233
 state-change, 236
 timeout-procedure, 234
 X-event, 237
 X-Request Breakpoints Examiner, 239
X/Motif Analyzer callback examiner, 246
X/Motif Analyzer event examiner, 248
X/Motif Analyzer examine menu
 selection, 224
 widget, 224
 widget class, 225
 widget tree, 224
 X event, 225
 X graphics context, 225
 X pixmap, 225
X/Motif Analyzer examiner tabs, 225
X/Motif Analyzer graphics context (GC)
 examiner, 248
X/Motif Analyzer pixmap examiner, 249
X/Motif Analyzer trace examiner, 240
X/Motif Analyzer tree examiner, 243
 graphical buttons, 245
X/Motif Analyzer widget class examiner, 250
X/Motif Analyzer widget examiner, 242
X/Motif Analyzer window examiner, 247
X/Motif Analyzer window menu item
 examine menu
 window, 225
X/Motif Analyzer windows description, 222
X/Motif Analyzer, selection in views menu, 194
xdiff, 123
xhost setting for debugger, 11