

Developer Magic™: ProDev WorkShop and MegaDev Overview

Document Number 007-2582-001

CONTRIBUTORS

Written and illustrated by John C. Stearns

Production by Laura Cooper

Engineering contributions by Lia Adams, Jim Ambras, Trevor Bechtel, Wes Embry, Alan Foster, Christine Hanna, David Henke, Marty Itzkowitz, Mahadevan Iyer, Lisa Kvarda, Stuart Liroff, Song Liang, Allan McNaughton, Michey Mehta, Sudhir Mohan, Ashok Mouli, Anil Pal, Andrew Palay, Tom Quiggle, Kim Rachmeler, Jack Repenning, Paul Sanville, Ravi Shankar, John Templeton, Michele Chambers Turner, Shankar Unni, Mike Yang, Jun Yu, and Doug Young.

© Copyright 1995 Silicon Graphics, Inc.— All Rights Reserved

This document contains proprietary and confidential information of Silicon Graphics, Inc. The contents of this document may not be disclosed to third parties, copied, or duplicated in any form, in whole or in part, without the prior written permission of Silicon Graphics, Inc.

RESTRICTED RIGHTS LEGEND

Use, duplication, or disclosure of the technical data contained in this document by the Government is subject to restrictions as set forth in subdivision (c) (1) (ii) of the Rights in Technical Data and Computer Software clause at DFARS 52.227-7013 and/or in similar or successor clauses in the FAR, or in the DOD or NASA FAR Supplement. Unpublished rights reserved under the Copyright Laws of the United States. Contractor/manufacturer is Silicon Graphics, Inc., 2011 N. Shoreline Blvd., Mountain View, CA 94043-1389.

Silicon Graphics is a registered trademark, and Developer Magic, IRIX IM, Indy, POWER Onyx, IRIS ViewKit, Power Fortran Accelerator, OpenGL, Open Inventor, ShowCase, IRIS Inventor, and Graphics Library are trademarks of Silicon Graphics, Inc. ClearCase is a trademark of Atria Software, Inc. UNIX is a registered trademark in the United States and other countries, licensed exclusively through X/Open Company, Ltd. X Window System is a trademark of the Massachusetts Institute of Technology. OSF/Motif is a registered trademark of the Open Software Foundation. PostScript is a registered trademark of Adobe Systems, Inc.

Developer Magic™: ProDev WorkShop and MegaDev Overview
Document Number 007-2582-001

Contents

List of Figures v

List of Tables vii

ProDev WorkShop and MegaDev Overview 1

Using the ProDev WorkShop Debugger 3

Debugger User Model 4

Where to Find Debugger Information 10

Navigating Through Code With the Static Analyzer and C++ Browser 12

Static Analyzer User Model 13

Where to Find Static Analyzer Information 16

C++ Browser User Model 16

Pinpointing Performance Problems With the Performance Analyzer 19

Performance Analyzer User Model 19

Where to Find Performance Analyzer Information 26

Determining the Thoroughness of Test Coverage With Tester 27

Tester User Model 27

Where to Find Tester Information 30

Recompiling Within the ProDev WorkShop Environment With Build Manager 30

Making Quick Changes With Fix and Continue 31

Fix and Continue User Model 31

Where to Find Fix and Continue Information 34

Debugging IRIX IM Programs 34

Features of the IRIX IM Analyzer 35

Where to Find IRIX IM Analyzer Information 39

Building Application Interfaces With RapidApp 40

RapidApp User Model 41

Where to Find RapidApp Information 43

List of Figures

Figure 1	Major Areas of the Main View Window	4
Figure 2	Typical Debugger Views Accessible at a Breakpoint	6
Figure 3	Array Visualizer	8
Figure 4	Machine-Level Debugger Views	9
Figure 5	Main Static Analyzer Window	12
Figure 6	Queries Available from the Static Analyzer Queries Submenus	15
Figure 7	C++ Browser Class View Window with Query Menus	17
Figure 8	Performance Panel With Task Menu Displayed	20
Figure 9	Performance Analyzer Main Window	22
Figure 10	Usage View (Graphs) Window: Lower Graphs	23
Figure 11	Malloc Error View	24
Figure 12	Major Areas of the Main Tester Window	29
Figure 13	Using Fix+Continue	32
Figure 14	The IRIX IM Analyzer Window	35
Figure 15	IRIX IM Analyzer Widget Tree Examiner	36
Figure 16	IRIX IM Analyzer Trace Examiner	38
Figure 17	RapidApp Window Displaying Container Palette	40
Figure 18	Creating a Widget	42

List of Tables

Table 1	Where to Find Debugger Information in the <i>CASEVision/WorkShop User's Guide</i> 10
Table 2	Where to Find Static Analyzer Information in the <i>CASEVision/WorkShop User's Guide</i> 16
Table 3	Where to Find C++ Browser Information in the <i>CASEVision/WorkShop MegaDev User's Guide</i> 18
Table 4	Performance Analyzer Views and Data 25
Table 5	Where to Find Performance Analyzer Information in the <i>CASEVision/WorkShop User's Guide</i> 26
Table 6	Tester Command Line Interface Summary 28
Table 7	Where to Find Tester Information in the <i>CASEVision/WorkShop User's Guide</i> 30
Table 8	Where to Find Fix and Continue Information in the <i>CASEVision/WorkShop MegaDev User's Guide</i> 34
Table 9	Where to Find IRIX IM Analyzer information in the <i>Developer Magic: IRIX IM Analyzer User's Guide</i> 39
Table 10	Where to Find RapidApp Information in the <i>Developer Magic: RapidApp User's Guide</i> 43

ProDev WorkShop and MegaDev Overview

Welcome to ProDev WorkShop and MegaDev, two major components in the Developer Magic™ software development environment. ProDev WorkShop contains the core software development tools; MegaDev contains advanced features for the development of C and C++ applications. These powerful, highly visual tools help you understand your program's structure and operation so that you can diagnose very difficult, traditionally time-consuming problems in a short amount of time. With them, you can develop applications for the entire Silicon Graphics® product line, from Indy™ to POWER Onyx™ workstations.

Note: In the past, the software development environment was called CASEVision™; that name has been replaced by Developer Magic. In addition to ProDev WorkShop and MegaDev, the Developer Magic environment includes ProMPF—a special module for multi-process Fortran programming—and IDO (IRIX Development Option)—the base compiler and libraries. Some of the documentation may still use the CASEVision name; those documents will be updated soon.

The ProDev WorkShop core tools provide:

- **Comprehensive control over the debugging process**—You can set simple breakpoints with the click of a mouse button or define complex conditions for your traps. ProDev WorkShop's fast data watch points with kernel support are especially adept at tracking memory corruption problems.
- **Visual debugging environment for examining data in your active program**—ProDev WorkShop provides convenient, graphical views of variables, expressions, large arrays, and data structures. If you prefer a tty-style interface, you can always dump values directly using WorkShop's Debugger command line.
- **Powerful static analysis for understanding your program**—You can view the structure of your program and relationships such as call trees, function lists, class hierarchies, and file dependencies. And you can get this information whether or not the program can be compiled.

- **The ability to collect performance and coverage information during test runs**—ProDev WorkShop’s Performance Analyzer lets you see where your program spends its time and pinpoint performance bugs, including those due to memory problems. The Tester tool shows you which source lines and basic blocks are covered in your tests.
- **Convenient recompiling from within the ProDev WorkShop environment**—WorkShop’s standard build tools let you view file dependencies and compiler requirements and fix compile errors conveniently.

The MegaDev tools provide:

- **Quick recompiles for simple changes**—The Fix and Continue tool lets you make simple changes without having to go through a major recompile and relinking, dramatically reducing the number of edit-compile-debug cycles.
- **Ability to analyze structures and relationships in C++ code**—The C++ Browser provides global graphical and textual views of interclass relationships, including inheritance, containment, and interactions within a set of classes.
- **Specialized debugging for IRIX IM™ applications**—MegaDev’s IRIX IM Analyzer lets you solve the special problems in IRIX IM application development. You can look at object data, set breakpoints at the object or X protocol level, trace X and widget events, and tune performance.
- **Rapid application development**—The RapidApp tool lets you create graphical interfaces for C++ applications quickly and easily. RapidApp lets you build graphical interfaces by dragging and dropping interface elements (based on IRIX IM widgets and IRIS ViewKit™-style components) onto a template window.

This overview gives you a broad exposure to the ProDev WorkShop and MegaDev tools as well as pointers to the documentation for getting detailed information. The overview is organized as follows:

- “Using the ProDev WorkShop Debugger”
- “Navigating Through Code With the Static Analyzer and C++ Browser”
- “Pinpointing Performance Problems With the Performance Analyzer”

- “Determining the Thoroughness of Test Coverage With Tester”
- “Recompiling Within the ProDev WorkShop Environment With Build Manager”
- “Making Quick Changes With Fix and Continue”
- “Debugging X/Motif Programs”
- “Building Application Interfaces With RapidApp”

In addition to the ProDev WorkShop and MegaDev tools, you can separately purchase:

- **Developer Magic Pro MPF**—a visual code parallelization tool used with the Power Fortran Accelerator™ to help balance parallel loops in Fortran applications
- **Developer Magic ClearCase™**—a toolset for version control, configuration management, and process control for software organizations
Note: If you use ClearCase, SCCS, or RCS, you can check source files directly into or out of ProDev WorkShop and MegaDev.
- **Developer Magic Tracker**—an application builder for creating change control and change tracking systems. It can be integrated with ClearCase.

Using the ProDev WorkShop Debugger

The Debugger is a UNIX® source-level debugging tool that provides special windows (views) for displaying program data and execution state as the program executes. The Debugger lets you set various types of breakpoints and watch points where you can conveniently view data such as variables, expressions, structures, large arrays, call stacks, and machine-level values. The WorkShop Debugger goes far beyond the capabilities of *dbx*. It includes fast data watchpoints and other types of traps; graphical views for displaying local variables, source-level expressions, array variables, and data structures; and debugging at the machine level.

Debugger User Model

All WorkShop activities can be accessed from the Main View window, which is illustrated in Figure 1.

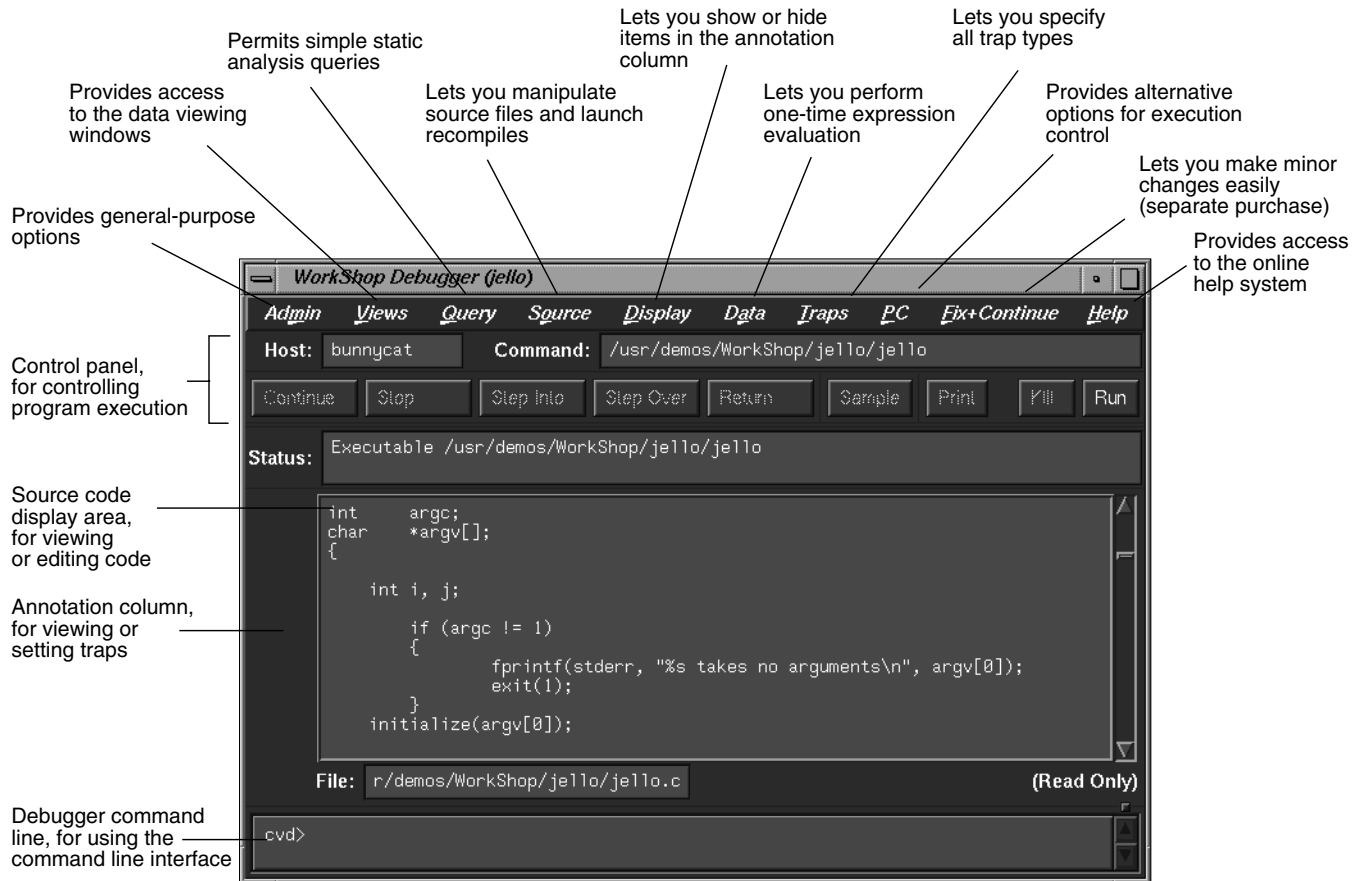


Figure 1 Major Areas of the Main View Window

The basic model for using the Debugger is to:

1. Invoke the Debugger by typing:

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

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

The argument *executable* is the name of the executable file for the process you want to run. It is optional; you can invoke the Debugger first and specify the executable later.

The *corefile* option lets you invoke the Debugger and specify a core file (with its executable) to try to determine why a program crashed.

The `-host` option lets you specify a remote host on which the target executable will be run; the Debugger runs locally. This option is useful if:

- you don't want the Debugger windows to interfere with the application you are debugging.
 - you are supporting an application remotely.
 - you don't want to use the Debugger on the target system for another reason.
2. Set stop traps, that is, breakpoints, in the source code.

Simple traps are set by clicking the left mouse button in the annotation column to the left of the source code display or by using the Traps menu. More complex traps, including watch points, can be set and managed from the Trap Manager, Signal Panel, and Syscall Panel, which can be accessed from the Views menu. You can also set traps by typing them at the Debugger command line in Main View. You can stop a process at any time by clicking the *Stop* button in the Main View control area.
 3. Start the program by clicking the *Run* button in Main View.
 4. When the process stops at a breakpoint or other stopping point of interest, you can examine the data in the Debugger view windows (accessed from the Views menu).

You can display view windows at any time; they update automatically each time the program stops. Figure 2 shows four typical Debugger views and indicates how you access them from the Views menu.

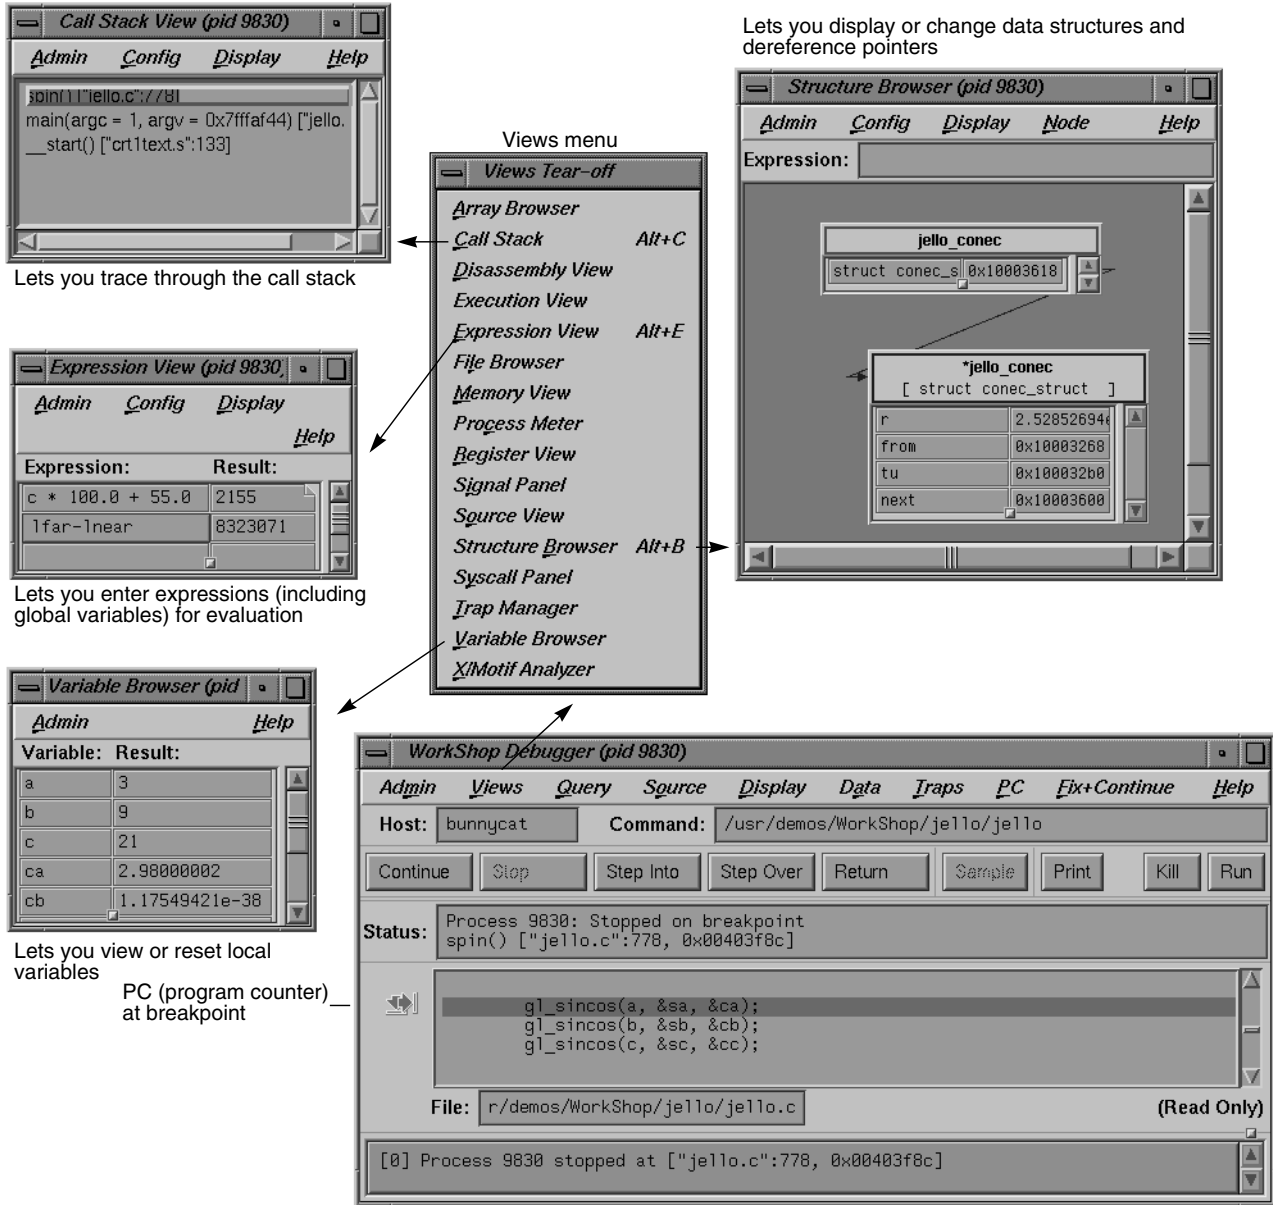


Figure 2 Typical Debugger Views Accessible at a Breakpoint

Figure 3 shows the Array Visualizer, a powerful view for examining data in arrays of up to 100 x 100 elements. You can look for problem areas in a 3D rendering of the array, click on the area of interest, and view the numerical values in a spreadsheet format. In Figure 3, the hue option has been set so that the values appear in a color spectrum from blue (lowest) to red (highest) with out-of-range anomalies appearing in gray. Note the high point coming out of the 3D image; it demonstrates how anomalies in large arrays stand out.

If you need to debug your program at the machine level, you can use Register View, Disassembly View, and Memory View, as shown in Figure 4. These are accessed from the Views menu in the Debugger Main View as well.

5. Use the control options in Main View to continue execution (see Figure 1).

From any breakpoint, you have these options:

- The *Continue* button runs the program until the next breakpoint.
- The “Continue To” selection in the PC menu proceeds to a specified source line. Placing the cursor in a line specifies it.
- The “Jump To” selection in the PC menu goes to a specified line (by the cursor), skipping over any intermediate code.
- The *Step Into* button continues execution by one step or a number specified by holding down the right mouse button over the *Step Into* button and selecting the number from the dialog box. The process then continues the specified number of source lines and enters any called functions.
- The *Step Over* button similarly proceeds a specified number of lines but executes intermediate functions without stepping into them.
- The *Return* button executes the remaining instructions in a function and stops on return from that function.

6. Check out the source code that needs to be fixed.

If you find a bug and are using an integrated source control program such as ClearCase, RCS, or SCCS, you can check out the source code from Main View (or Source View, an alternate editing window).

Choose “Check Out” from the Versioning submenu in the Source menu.

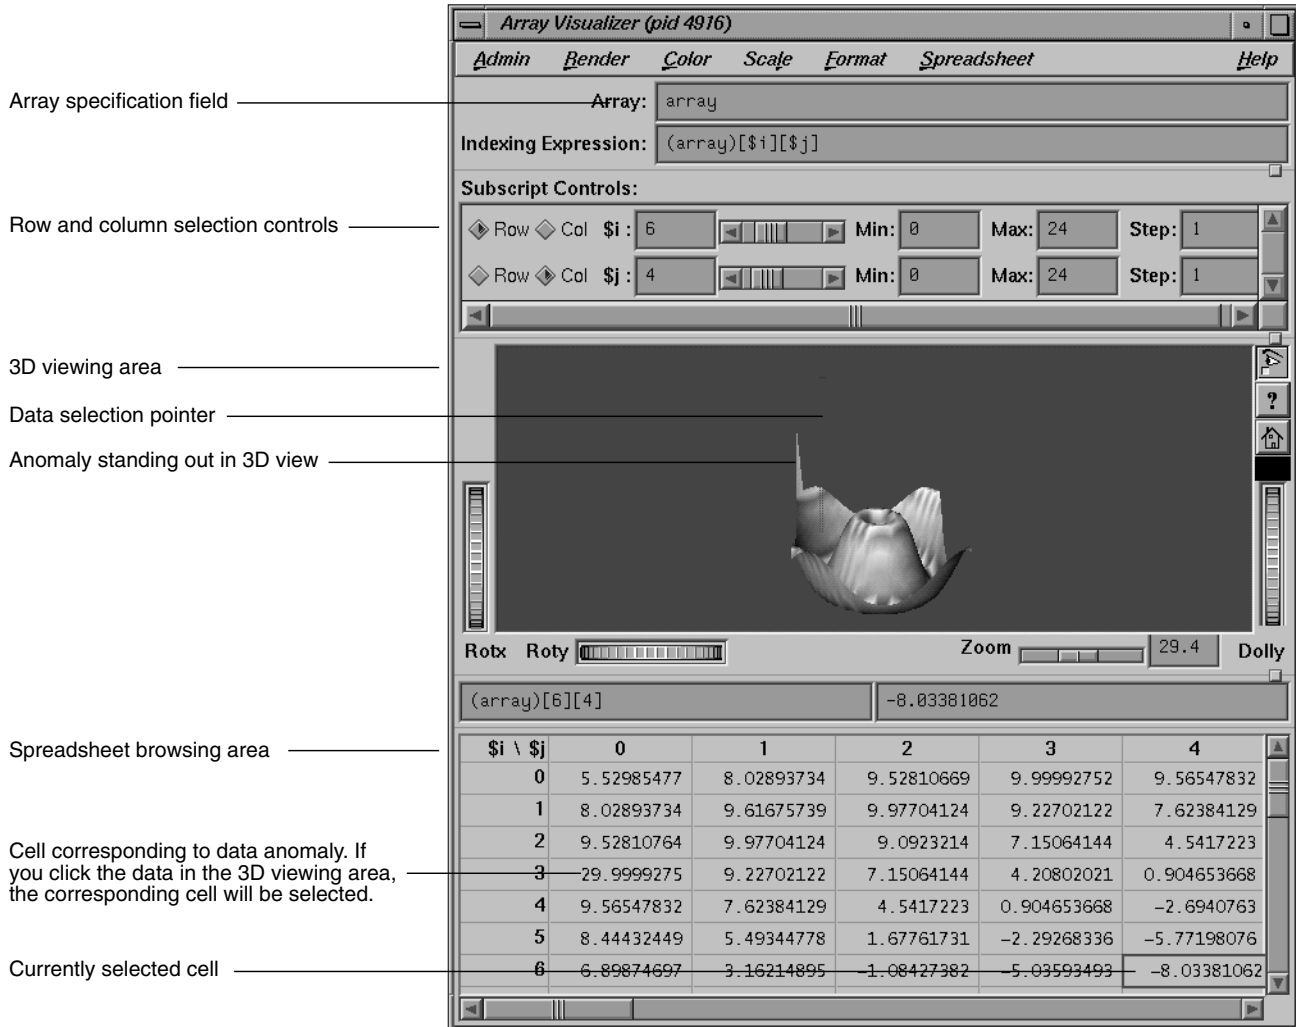
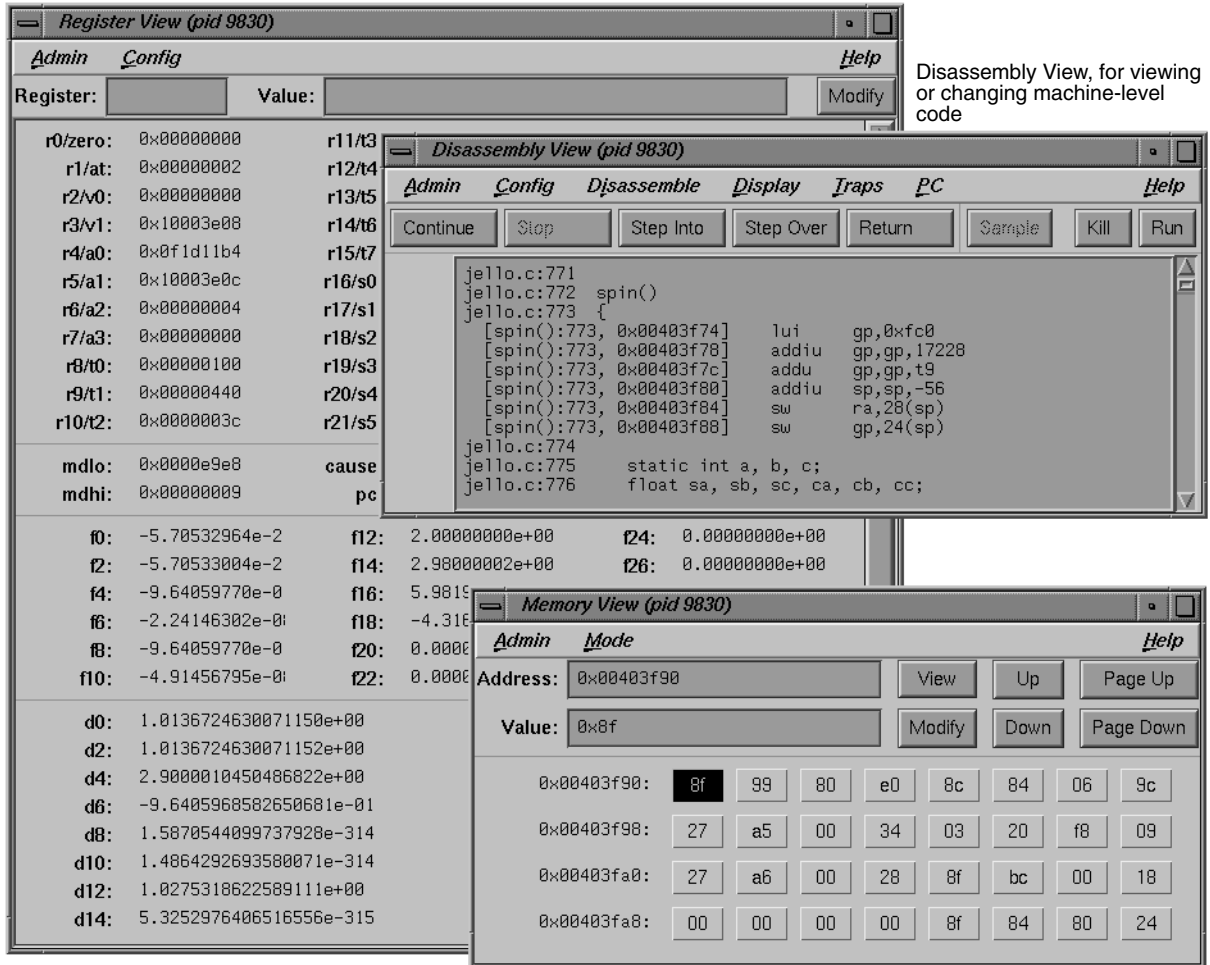


Figure 3 Array Visualizer

Register View, for viewing or changing the contents of registers



Disassembly View, for viewing or changing machine-level code

Memory View, for viewing or changing the contents of memory addresses

Figure 4 Machine-Level Debugger Views

7. Fix any problems in your code using the source code display area in Main View, Source View, or the editor of your choice.

Both Main View and Source View let you do simple editing and annotate the code with trap indicators. Source View also lets you display test data from the Performance Analyzer and Tester in the annotation column. If you prefer to view source code in a text editor other than Source View, add the line

```
*editorCommand: editor
```

to your *.Xdefaults* file, where *editor* is the command for the editor you wish to use.

8. Recompile using Build Manager.

Build Manager has two windows: Build View and Build Analyzer. Build View lets you compile, view compile error lists, and access the offending code in Source View or an editor of your choice. Build Analyzer lets you view build dependencies and recompilation requirements, and access source files. Build View uses the UNIX *make* facility as its default build software. Although Build Analyzer determines dependencies using *make*, you can substitute the build software of your choice, any *make* that runs on Silicon Graphics platforms.

Where to Find Debugger Information

To find out more about the Debugger, refer to Table 1.

Table 1 Where to Find Debugger Information in the *CASEVision/WorkShop User's Guide*

Topic	See ...
General Debugger information	Chapter 1, "Getting Started with the WorkShop Debugger"
Debugger tutorial	Chapter 2, "A Short Debugger Tutorial"
Debugger interaction with source files	Chapter 3, "Debugger: Managing Source Files"
Managing windows while performing multiple tasks	Chapter 4, "Debugger: CASEVision Project Session Management"
Comprehensive trap information	Chapter 5, "Debugger: Setting Traps in WorkShop"

Table 1 (continued) Where to Find Debugger Information in the *CASEVision/WorkShop User's Guide*

Topic	See ...
Controlling execution in a process (stepping, jumping, etc.)	Chapter 6, "Debugger: Controlling Process Execution"
Examining Debugger data in general at the source level	Chapter 7, "Examining Debugger Data"
Tracing through the call stack	"Tracing through Call Stack View" on page 98
Entering expressions to be evaluated at stopping points	"Evaluating Expressions" on page 101
Viewing or changing the values of variables	"Examining Variable Values" on page 109
Examining data in arrays using the 3D or spreadsheet format	"Examining Array Variables" on page 111
Determining the data structures of variables	"Looking at Data Structures" on page 126
Using the Debugger command line	"Examining Data at the Command Line" on page 137
Examining debugger data at the machine level	Chapter 8, "Machine-level Debugging"
Using the debugger to trap memory allocation problems	Chapter 9, "Debugger: Detecting Heap Corruption"
Debugging multiprocess programs	Chapter 10, "Multiple Process Debugging"

Navigating Through Code With the Static Analyzer and C++ Browser

The ProDev WorkShop Static Analyzer is a source code analysis and navigation tool for analyzing source code written in C, C++, or Fortran. (The C++ Browser has additional features for C++ and is described in "C++ Browser User Model" on page 16.) The Static Analyzer shows you the code's structure (graphically or in text format) including information on function calls, definitions of variables, file dependencies, macro locations, class hierarchies, file dependencies, and many other structural details for understanding your code. In addition, you can make specific queries, such as showing everywhere a function is used. You can even analyze programs that don't compile, a particularly nice feature for those porting code.

The Static Analyzer works by reading through source code files that you specify and creating a database of functions, macros, variables, files, and (for C++) classes and methods and their relationships. The main Static Analyzer window with a typical call graph is illustrated in Figure 5.

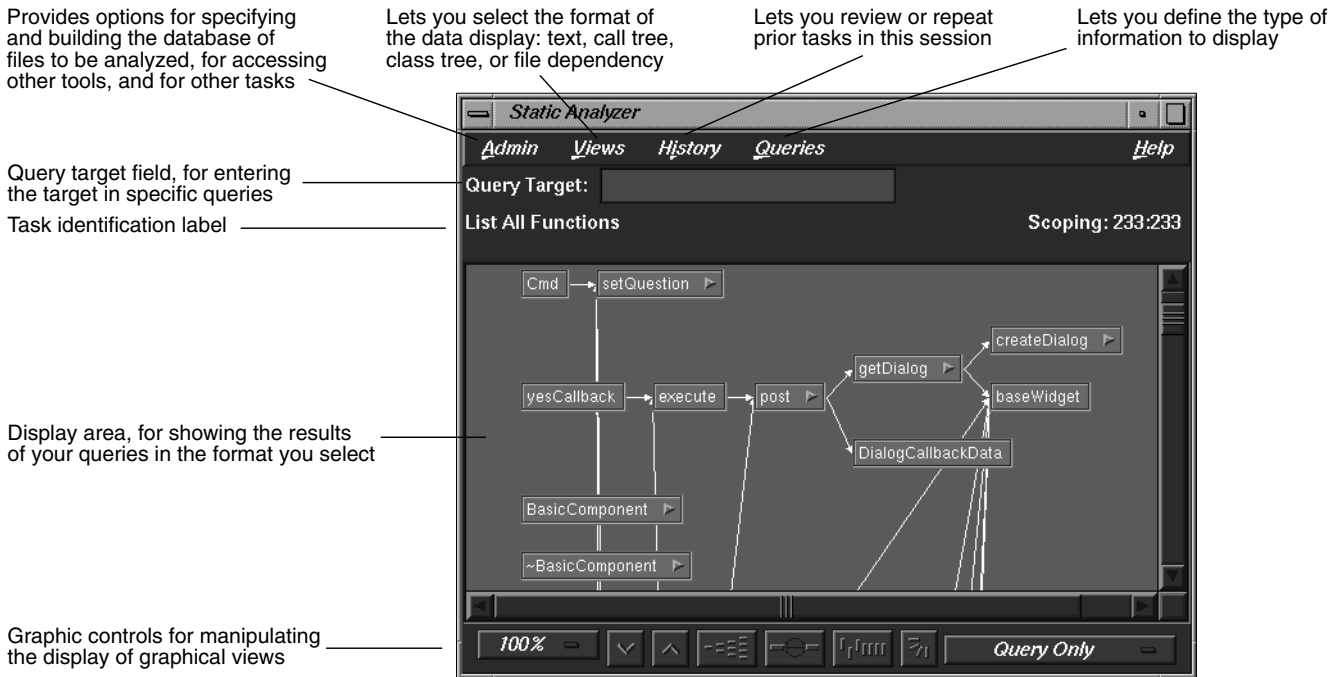


Figure 5 Main Static Analyzer Window

Static Analyzer User Model

Follow these steps for using the Static Analyzer:

1. Invoke the Static Analyzer, either by typing `cvstatic` or by selecting “Static Analyzer” from the Launch submenu in any ProDev WorkShop Admin menu (preferably from the directory where your source is located).
2. Decide which files are to be analyzed.

You designate which files are to be analyzed in a special file called a *fileset*. A fileset is a regular ASCII file with a format of one entry per line, each line separated from the next by a carriage return. The entries can be regular expressions, filenames, or included directories preceded by the designator `-I`.

To specify a fileset, you can

- create the fileset manually using a text editor
- use the Fileset Editor, which is accessed from the Admin menu in the Static Analyzer window
- let the Static Analyzer create the fileset automatically at startup by defaulting to the files in the current directory that match the expression `*.[cCfF]`
- let the Static Analyzer create the fileset automatically at startup from the command line by typing `cvstatic` with the `-executable` flag and designating an executable
- use the compiler to create a fileset (and database) by adding the `-sa, <dbdirectory>` option to your makefile

Note: Many programs are so big that a query covering the entire scope is useless due to its size and complexity. There are two ways to keep the scope of your analysis at a manageable size: (1) Limit the number of files to be analyzed or (2) avoid queries that begin with “List All ...”.

3. Decide how you are going to build the database.

Before you can specify a fileset, you need to decide how you are going to build the database. You can choose to create the database in *scanner* mode (the default), which is fast but not sensitive to any specific programming language, or in *parser* mode, which uses the compiler and is slower but more thorough. Use scanner mode for large programs or

for programs that do not compile. Scanner mode is particularly suited to porting situations. Parser mode is better when you have code that compiles and you need to determine language-specific relationships, particularly in Fortran and C++.

4. Build the database.
5. Perform your queries.

Queries are selected from the Queries menu in the Static Analyzer. They fall into 10 categories, as shown in Figure 6. Remember that the “List All ...” queries can produce overwhelming results for large programs.

6. View (and save) the results.

The Static Analyzer has four ways of presenting data, which are selected from the Views menu:

- “Text View” displays query results in a text format. In addition to listing the queried items, it indicates the source filename and line number, and includes the actual source line.
- “Call Tree View” applies to function queries. It presents the data in a graphical format with *nodes* (rectangles) representing functions and *arcs* (arrows) representing calls to functions.
- “Class Tree View” applies to C++ class queries. It presents a class inheritance tree with nodes representing classes and arcs representing parent-child class relationships.
- “File Dependency View” applies to file queries. It presents a graph, with nodes representing files and arcs representing include relationships.

If you want to save a query in a graphical view, you can save a PostScript® version by selecting “Save Query...” from the Admin menu and print it out at your leisure.

7. Access the source code.

Double-clicking any node in a graph or item in Text View brings up the Source View window containing the corresponding source code. Double-clicking any arc (arrow) displays Source View with the corresponding call site or file inclusion.

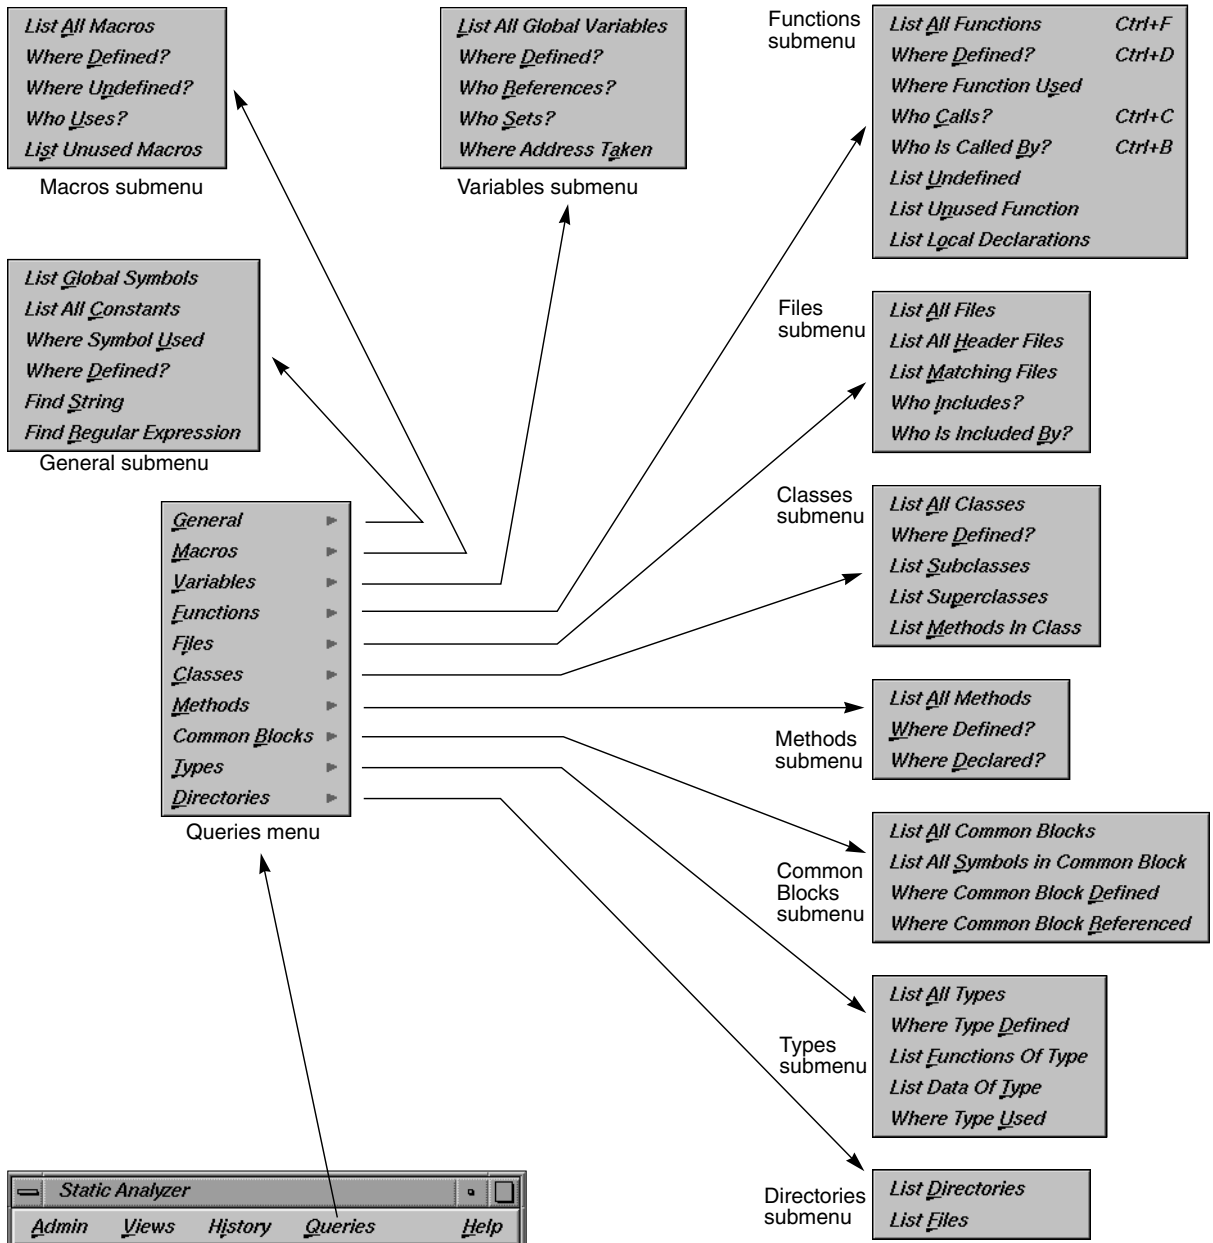


Figure 6 Queries Available from the Static Analyzer Queries Submenus

Where to Find Static Analyzer Information

To find out more about the Static Analyzer, refer to Table 2.

Table 2 Where to Find Static Analyzer Information in the *CASEVision/WorkShop User's Guide*

Topic	See ...
General Static Analyzer description	Chapter 12, "Introduction to the WorkShop Static Analyzer"
Static Analyzer tutorial	Chapter 13, "A Sample Session with the Static Analyzer"
Specifying a fileset	"Fileset Specifications" on page 238
Building a database using scanner mode	"Scanner Mode" on page 247
Building a database using parser mode	"Parser Mode" on page 248
Performing queries	Chapter 15, "Static Analyzer: Queries"
Static Analyzer viewing formats	Chapter 16, "Static Analyzer: Views"
Strategies for analyzing large programs	Chapter 17, "Static Analyzer: Working on Large Programming Projects"

C++ Browser User Model

The C++ Browser user model is similar to the Static Analyzer user model. After you build the database (which must be done in parser mode), you access the C++ Browser by selecting "C++ Browser" from the Static Analyzer Admin menu.

The C++ Browser lets you display different sets of information about classes, class members, and interclass relationships through these three views:

- **Class View**— displays member and class information in an expandable, hierarchical outline format with the members of the current class in the left pane and related classes on the right (see Figure 7). Clicking the diamond-shaped icons next to the headings in the list hides or displays the associated information.

Like the Static Analyzer, you have numerous queries available through the Query menu. In addition, if you select an item in either of the Class View lists and hold down the right mouse button, you can access the Queries menu specific to that type of item, that is, methods, data members, classes, and so on.

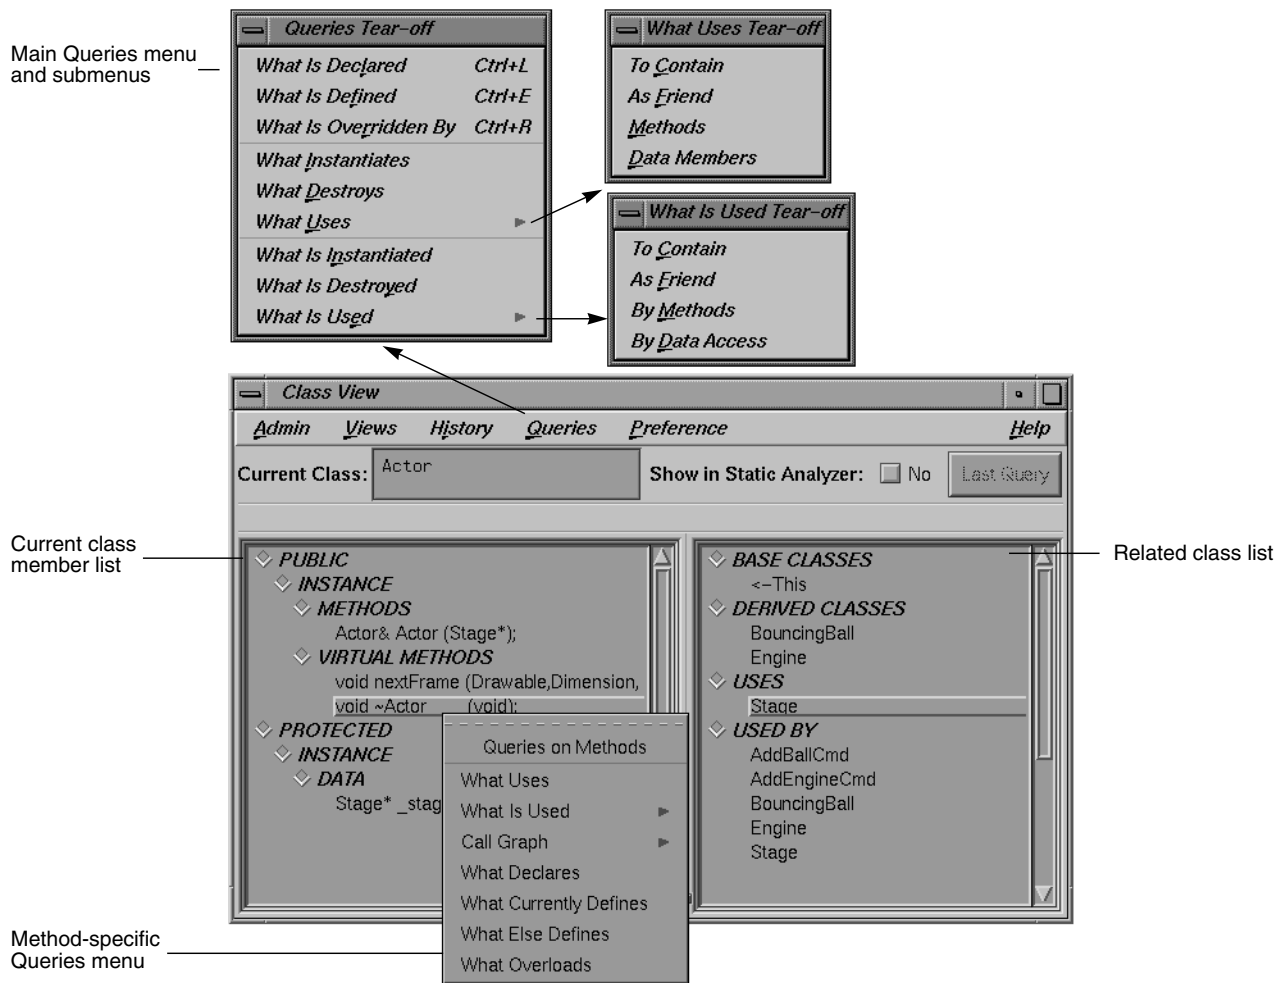


Figure 7 C++ Browser Class View Window with Query Menus

You can also create man pages describing classes by selecting “Generate man pages...” from the Class View Admin menu. You simply specify an individual class or all of them and the C++ Browser fills in the man page template for you.

- **Class Graph**—displays the class hierarchy for the current class in the Class View window with nodes as classes and arcs as parent-child relationships. Class Graph can show four types of relationships: inheritance, containment, interaction, and friends. You can display all classes, limit the scope of classes derived from the current class, or get a butterfly view showing the immediate base and derived classes of the current class.
- **Call Graph**—displays the calling relationships of methods and virtual methods selected from Class View with options for customizing the display of the graph.

To find out more about the C++ Browser, refer to Table 3.

Table 3 Where to Find C++ Browser Information in the *CASEVision/WorkShop MegaDev User's Guide*

Topic	See ...
General C++ Browser description	Chapter 5, “Getting Started with the C++ Browser”
C++ Browser tutorial	Chapter 6, “Using the C++ Browser: A Sample Session”
Class View	“Using the Class View Outline Lists” on page 83 and “Class View Window” on page 117
Viewing graph data	“Class Graph and Call Graph Displays” on page 147
Call Graph	“Using the Call Graph Window” on page 107 and “Call Graph Window” on page 156
Class Graph	“Using the Class Graph Window” on page 101 and “Class Graph Window” on page 154

Pinpointing Performance Problems With the Performance Analyzer

The ProDev WorkShop Performance Analyzer helps you understand how your program performs so that you can correct any problems. In performance analysis, you run experiments to capture performance data and see how long each phase or part of your program takes to run. You can then determine if the performance of the phase is slowed down by the CPU, I/O activity, memory, or a bug and attempt to speed it up.

A menu of predefined tasks is provided to help you set up your experiments. With the Performance Analyzer views, you can conveniently analyze the data. These views show CPU utilization and process resource usage (such as context switches, page faults, and working set size), I/O activity, and memory usage (to capture such problems as memory leaks, bad allocations, and heap corruption).

The Performance Analyzer has three general techniques for collecting performance data:

- Counting—It can count the exact number of times each function and/or basic block has been executed. This requires *instrumenting* the program, that is, inserting code into the executable to collect counts.
- Profiling—It can periodically examine and record the program's PC (program counter), call stack, and resource consumption.
- Tracing—It can trace events that affect performance, such as *reads* and *writes*, system calls, page faults, floating point exceptions, and *mallocs*, *reallocs*, and *frees*.

Performance Analyzer User Model

1. Set up a general experiment to determine areas for improvement in your program.

To set up a performance experiment, select "Performance Task..." from the Admin menu in the Debugger Main View or type `cvsped` at the command line. The Performance Panel window is displayed; it provides a task menu from which you select predefined experiment tasks (see Figure 8). At this point, you probably haven't formed a

hypothesis yet about where the performance problems lie. If this is the case, select the “Determine bottlenecks, identify phases” task. This is useful for determining the general problem areas within the program.

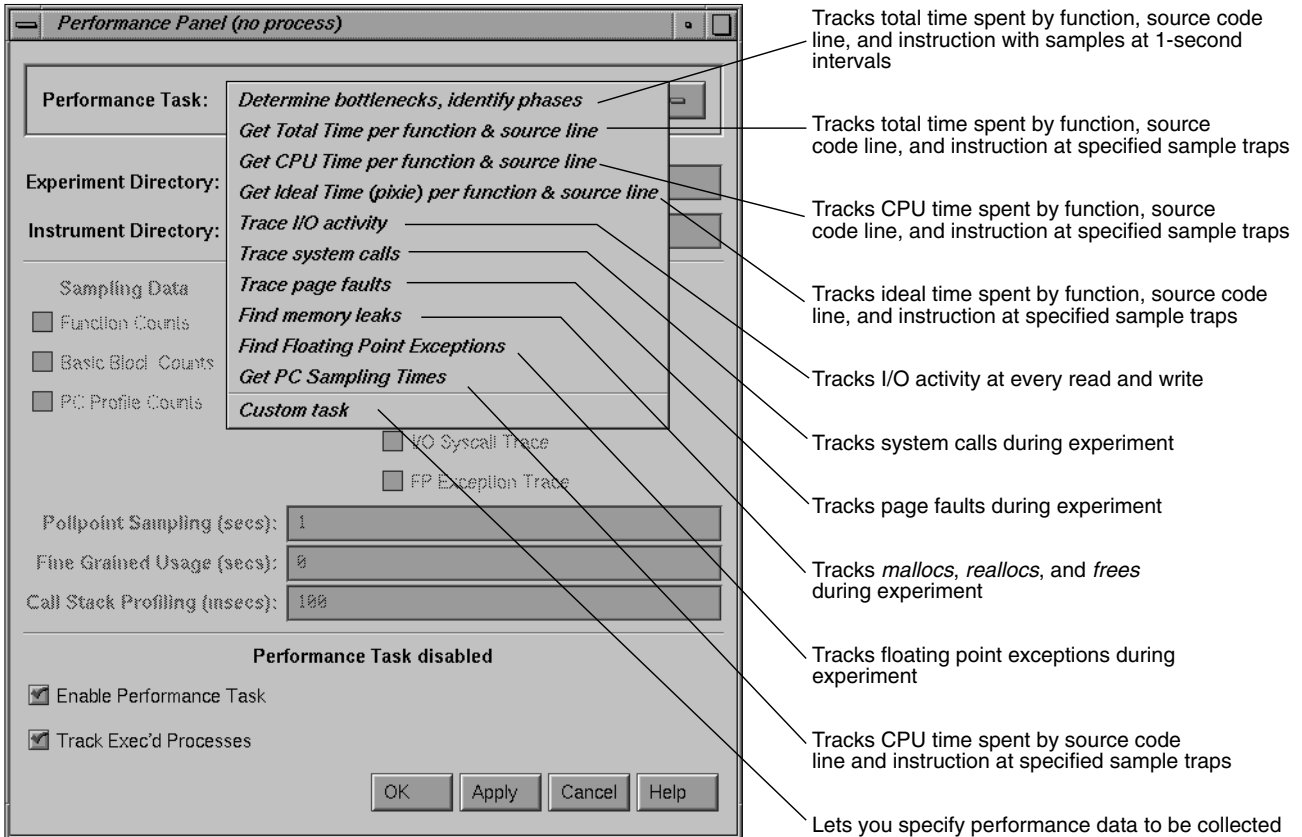


Figure 8 Performance Panel With Task Menu Displayed

2. Start the program by clicking the *Run* button in Main View.

This runs the experiment and collects the performance data, which is written to a directory *test0000* (or a name of your choice); *test0000* is the identification for your experiment.

3. Analyze the results in the Performance Analyzer window and the Usage View (Graphs) window.

After the experiment has finished, you can display the results in the Performance Analyzer window by selecting “Performance Analyzer” from the Launch submenu in any ProDev WorkShop Admin menu or by typing `cvperf -exp experimentname`. The results from a typical performance analysis experiment appear in Figure 9, the main Performance Analyzer window, and Figure 10, which shows a subset of the graphs in the Usage Views (Graphs) window. You should be able to determine where the phases of execution occur so that you can set *sample traps* between them. Sample traps collect performance data at specified times and events in the experiment.

Function list area displays functions with their performance data from the experiment, time spent in the function including its called functions, and time spent in the function excluding calls

Usage chart area indicates general resource usage during the experiment

Time line area indicates where samples were taken. The calipers let you limit the scope of the analysis to a portion of the experiment. Double-clicking a sample point displays the call stack that occurred there.

Caliper and sample point selector controls

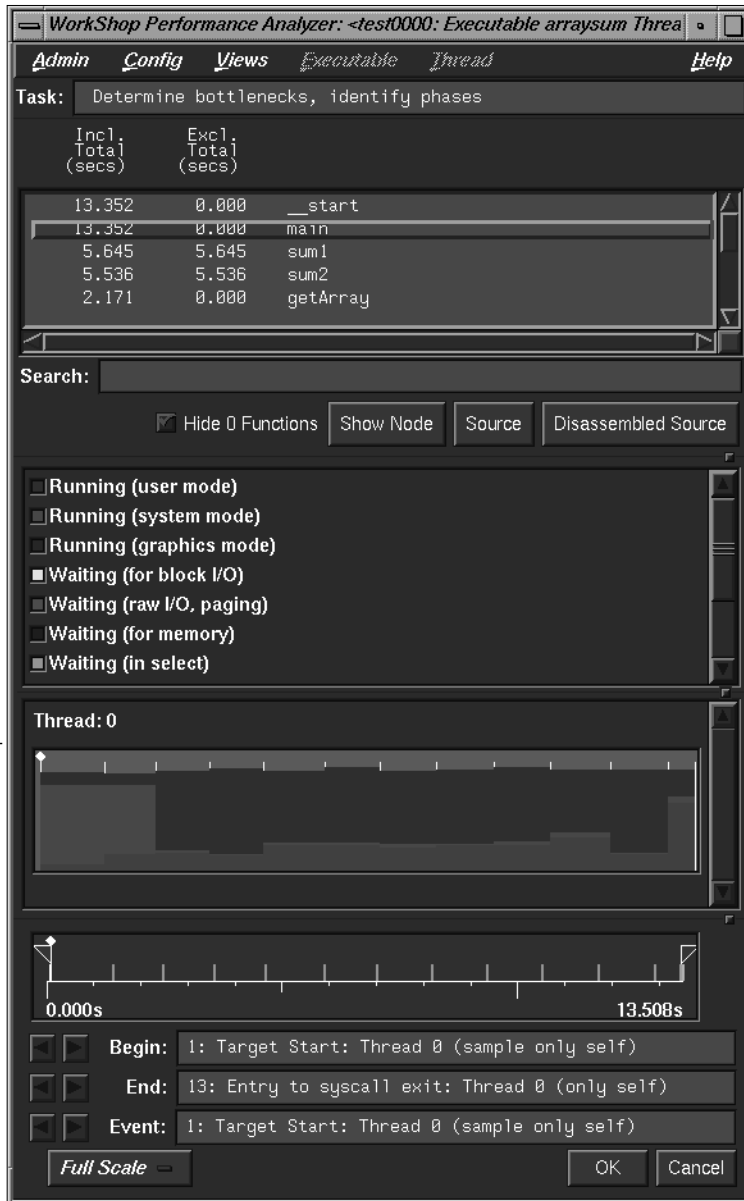


Figure 9 Performance Analyzer Main Window

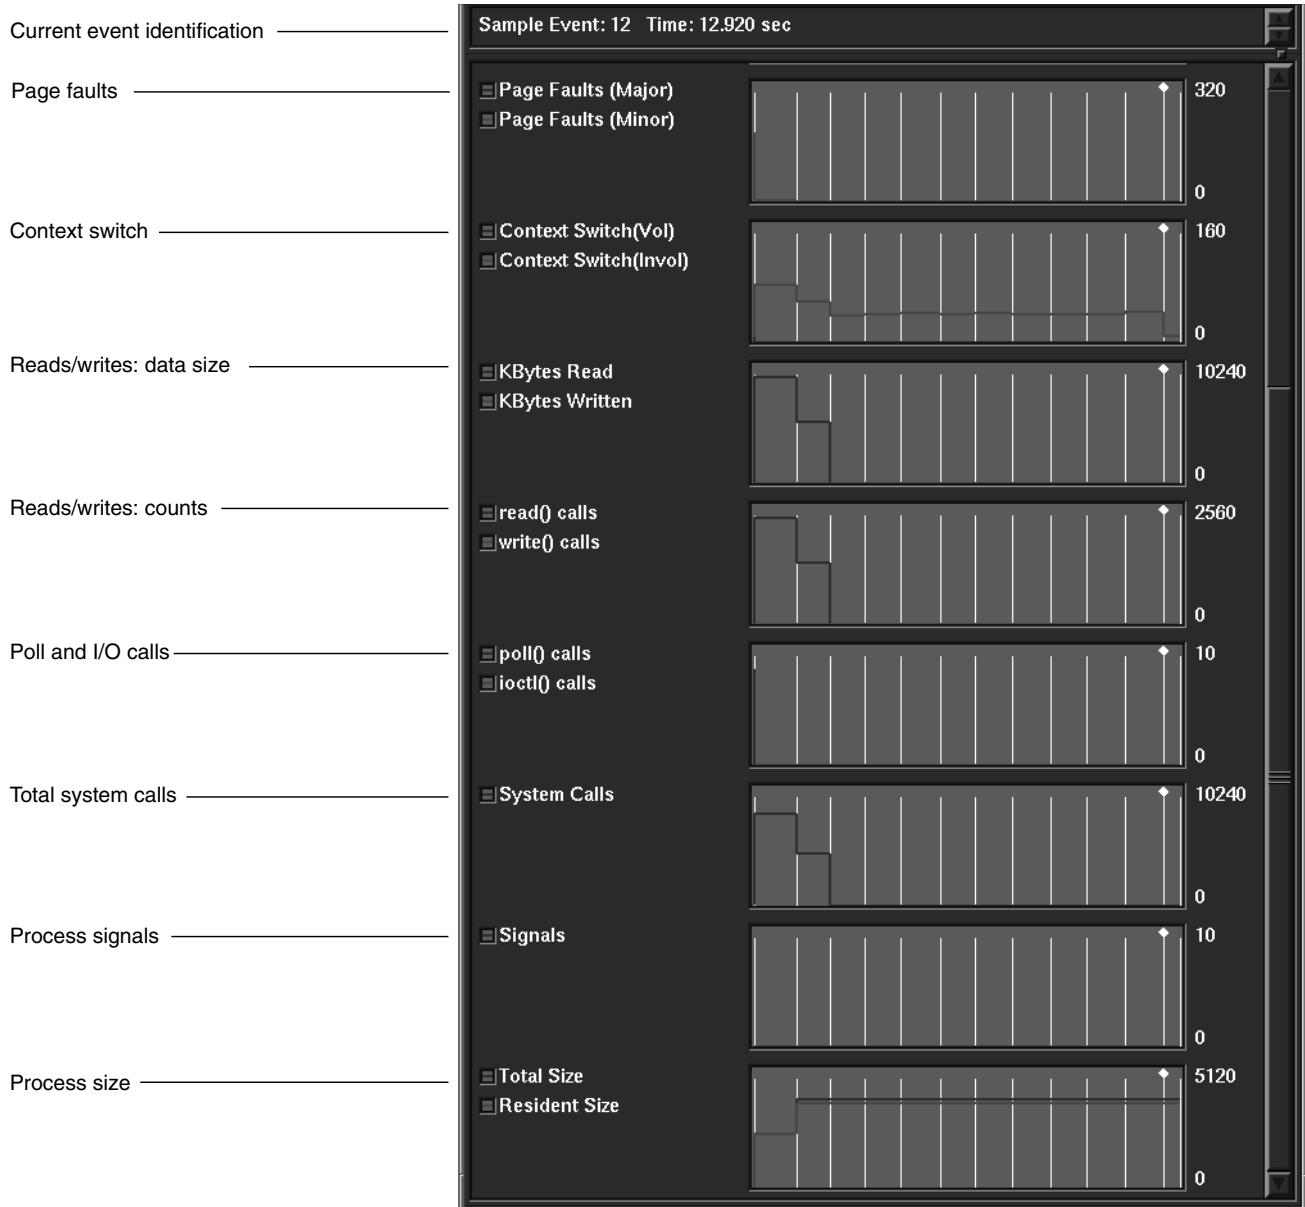


Figure 10 Usage View (Graphs) Window: Lower Graphs

4. Set sample traps at the start and end of each phase.

Setting sample traps between phases isolates the data to be analyzed on a phase-by-phase basis. Sample traps are set by selecting “Sample”, “Sample at Function Entry”, or “Sample at Function Exit” from the Set Trap submenu in the Traps menu in the Debugger Main View or through the Traps Manager.

5. Select your next experiment from the Task Menu in the Performance Panel and run it by clicking the *Run* button in the Main View window.

You need to form a hypothesis about the performance problem and select an appropriate task (see Figure 8) for your next experiment. There are trade-offs in selecting tasks—experiments can collect huge amounts of data and may perturb the results in some cases. The strategies for selecting tasks are discussed in detail in Chapter 20, “Setting Up Performance Analysis Experiments,” in the *CASEVision/WorkShop User’s Guide*.

6. Analyze the results using the Performance Analyzer main window, its views, or Source View with performance data annotations displayed.

A typical Performance Analyzer view, Malloc Error View, is shown in Figure 11. The Performance Analyzer provides results in the windows listed in Table 4.

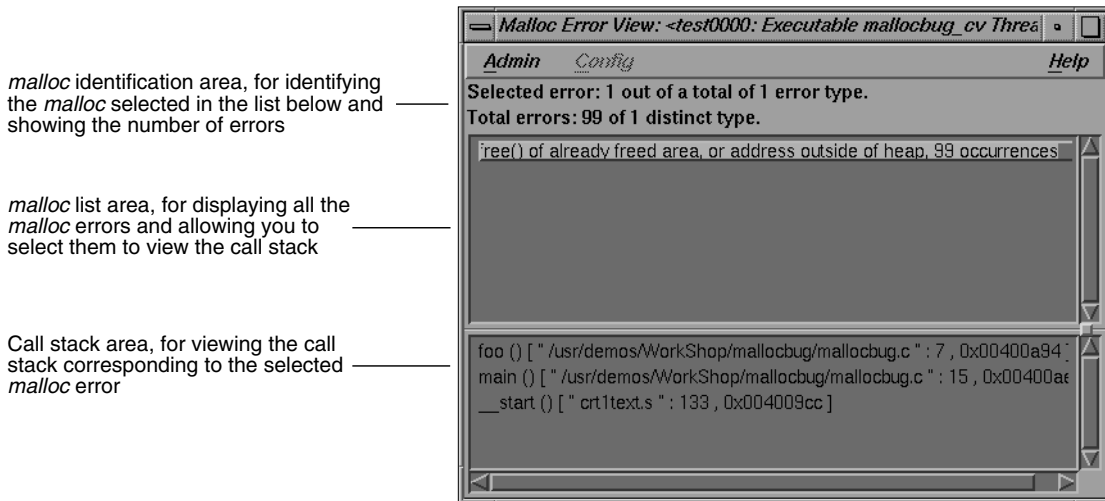


Figure 11 Malloc Error View

Table 4 Performance Analyzer Views and Data

Performance Analyzer Window	Data Provided
Performance Analyzer main window	Function list with performance data, usage chart showing general resource usage over time, and time line for setting scope on data
Call Stack View	Call stack recorded when selected event occurred
Usage View (Graphs)	Specific resource usage over time, shown as graphs
Usage View (Numerical)	Specific resource usage for selected (by caliper) time interval, shown as numerical values
Call Graph View	A graph showing functions that were called during the time interval, annotated by the performance data collected
I/O View	A graph showing I/O activity over time during the time interval
Malloc View	A list of all <i>mallocs</i> , their sizes and number of occurrences, and, if selected, their corresponding call stack within the selected time interval
Malloc Error View	A list of <i>malloc</i> errors, their number of occurrences, and if selected, their corresponding call stack within the time interval
Leak View	A list of specific leaks, their sizes and number of occurrences, and if selected, their corresponding call stack within the time interval
Heap View	A generalized view of heap memory within the time interval
Source View	The ProDev WorkShop text editor window showing source code annotated by performance data collected
Working Set View	The instruction coverage of dynamic shared objects (DSOs) that make up the executable, showing instructions, functions, and pages that were not used within the time interval
Cord Analyzer	The Cord Analyzer is not actually part of the Performance Analyzer, but it works with data from Performance Analyzer experiments. It lets you try out different ordering of functions to see the effect on performance.

Where to Find Performance Analyzer Information

To find out more about the Performance Analyzer, refer to Table 5.

Table 5 Where to Find Performance Analyzer Information in the *CASEVision/WorkShop User's Guide*

Topic	See ...
General Performance Analyzer information	Chapter 18, "Introduction to the Performance Analyzer"
Performance analysis theory	"Sources of Performance Problems" on page 312
General Performance Analyzer tutorial	Chapter 19, "Performance Analyzer Tutorial"
Memory leak tutorial	"Memory Experiment Tutorial" on page 415
Setting up performance analysis experiments including task selection	Chapter 20, "Setting Up Performance Analysis Experiments" for details and "The Performance Panel" on page 362 for a summary
Setting sample traps	"Setting Traps in Main View and Source View" on page 73 and "Setting Traps in Trap Manager" on page 76
Performance Analyzer main window	"The Performance Analyzer Main Window" on page 374
Usage View (Graphs) window	"Usage View (Graphs)" on page 387
Watching an experiment without collecting data in the Process Meter	"Process Meter" on page 393
Usage View (Numerical) window	"Usage View (Numerical)" on page 393
Tracing I/O calls using the I/O View window	"I/O View" on page 396
Call Graph View window	"Call Graph View" on page 397
Finding memory problems	"Analyzing Memory Problems" on page 405
Specifying performance annotations for Source View and Call Graph View	"Config Menu" on page 382
Call Stack View window	"Call Stack" on page 418
Improving working set behavior	"Analyzing Working Sets" on page 419

Determining the Thoroughness of Test Coverage With Tester

Tester is a software quality assurance toolset for measuring dynamic coverage over a set of tests. It tracks the execution of functions, individual source lines, arcs, blocks, and branches.

Tester User Model

This section describes the user model for designing a single test. After you have your instrumentation file and test directories set up, you can automate your testing and create larger test sets. Tester has both a command line interface (see Table 6) and a graphical user interface (see Figure 12).

1. Plan your test.
2. Create (or reuse) an instrumentation file.

The instrumentation file defines the coverage data you wish to collect in this test.
3. Apply the instrument file to the target executable(s).

This creates a special executable for testing purposes that collects data as it runs.
4. Create a test directory to collect the data files.
5. Run the instrumented version of the executable to collect the coverage data.
6. Analyze the results.

Tester produces a wide variety of column-based reports. Most are available in both interfaces: command line and graphical. The reports can show source and assembly line coverage; coverage of functions; arc coverage, that is, coverage of function calls; call graphs indicating caller and callee functions and their counts; basic block counts; count information for assembly language branches; summaries of overall coverage; and argument tracing.

Table 6 Tester Command Line Interface Summary

Command Category	Command Name	Description
general	cvcov cattest	Describes the test details for a test, test set, or test group
	cvcov lsinstr	Displays the instrumentation information for a particular test
	cvcov lstest	Lists the test directories in the current working directory
	cvcov mktest	Creates a test directory
	cvcov rmttest	Removes tests and test sets
	cvcov runinstr	Adds code to the target executable to enable you to capture coverage data, according to the criteria you specify
	cvcov runttest	Runs a test or a set of tests
coverage analysis	cvcov lssum	Shows the overall coverage based on the user-defined weighted average over function, line, block, branch, and arc coverage
	cvcov lsfun	Lists coverage information for the specified functions in the program that was tested
	cvcov lsblock	Displays a list of blocks for one or more functions and the count information associated with each block
	cvcov lsbranch	Lists coverage information for branches in the program, including the line number at which the branch occurs
	cvcov lsarc	Shows arc coverage, that is, the number of arcs taken out of the total possible arcs
	cvcov lscall	Lists the call graph for the executable with counts for each function
	cvcov lsline	Lists the coverage for native source lines
	cvcov lssource	Displays the source annotated with line counts
	cvcov lstrace	Shows the argument tracing information
	cvcov diff	Shows the difference in coverage for different versions of the same program
test set	cvcov mktset	Makes a test set
	cvcov addtest	Adds a test or test set to a test set or test group
	cvcov deltest	Removes a test or test set from a test set or test group

Table 6 (continued) Tester Command Line Interface Summary

Command Category	Command Name	Description
	cvcov optimize	Selects the minimum set of tests that give the same coverage or meet the given coverage criteria as the given set
test group	cvcov mktgroup	Creates a test group that can contain other tests or test groups; targets are either the target libraries or DSOs

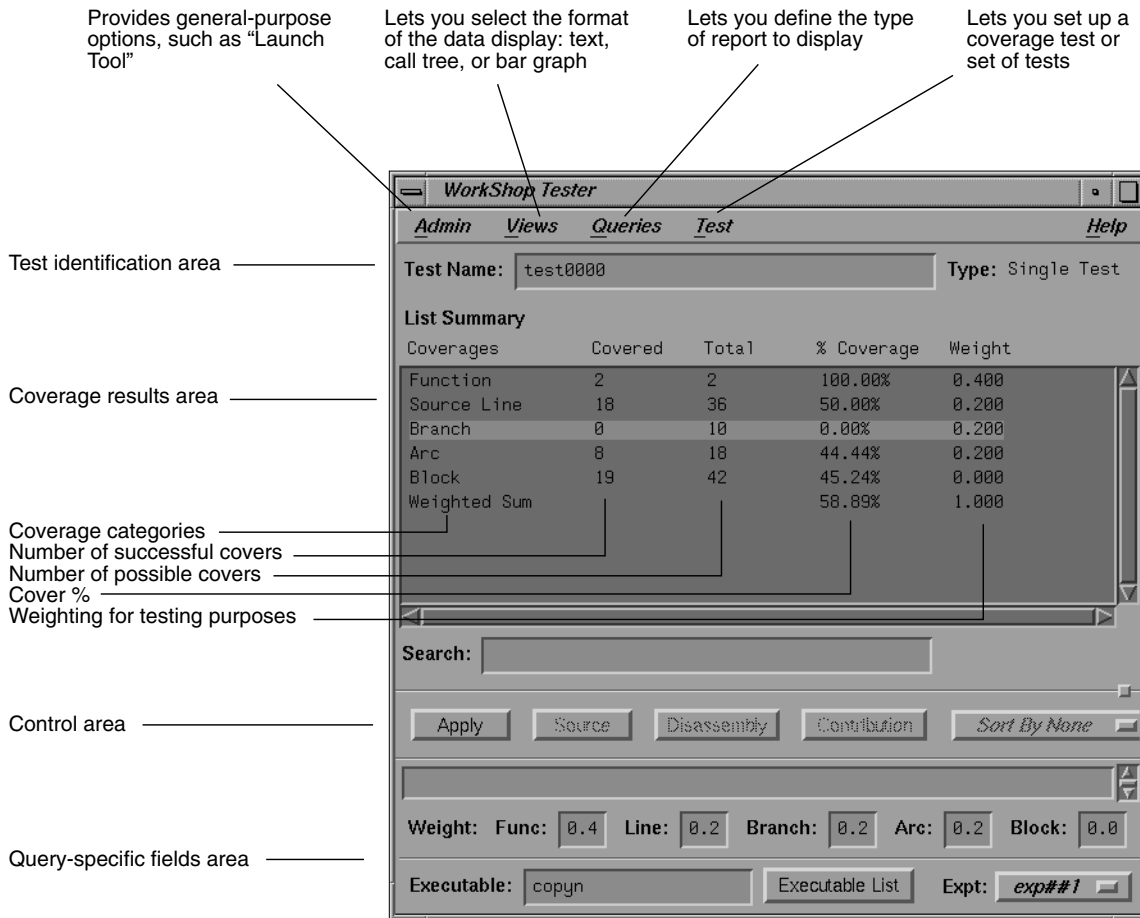


Figure 12 Major Areas of the Main Tester Window

Where to Finder Tester Information

To find out more information about Tester, refer to Table 7.

Table 7 Where to Find Tester Information in the *CASEVision/WorkShop User's Guide*

Topic	See ...
General Tester information	"Tester Overview" on page 435
Automated testing	"Automated Testing" on page 446
Command line interface tutorial	Chapter 23, "Tester Command Line Interface Tutorial"
Graphical user interface tutorial	Chapter 25, "Tester Graphical User Interface Tutorial"
Command line interface details	Chapter 24, "Tester Command Line Reference"
Graphical user interface details	Chapter 26, "Tester Graphical User Interface Reference"

Recompiling Within the ProDev WorkShop Environment With Build Manager

The Build Manager lets you view file dependencies and compiler requirements, fix compile errors conveniently, and compile software without leaving the WorkShop environment. It provides two views:

- Build View—for compiling, viewing compile error lists, and accessing the code containing the errors in Source View (the ProDev WorkShop editor) or an editor of your choice.
- Build Analyzer—for viewing build dependencies and recompilation requirements and accessing source files.

For more information on Build Manager, see Chapter 11, "Using the Build Manager," in the *CASEVision/WorkShop User's Guide*.

Making Quick Changes With Fix and Continue

Fix and Continue is part of the Developer Magic MegaDev module. The Fix and Continue feature lets you make minor changes to your code from within WorkShop without having to recompile and link the entire system. You issue Fix and Continue commands in the Debugger Main View window, either by selecting them from the Fix+Continue menu or typing them in directly in the Debugger command line area.

With Fix and Continue, you can edit a function, parse the new function, and continue execution of the program being debugged. Fix and Continue enables you to speed up your development cycle significantly. For example, a program that takes 5 minutes to rebuild through a conventional compile might take 45 seconds using Fix and Continue.

Fix and Continue lets you:

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

Figure 13 shows you the WorkShop Main View during a Fix and Continue session and explains how to use the Fix and Continue menu.

Fix and Continue User Model

1. Invoke the Debugger as you normally would by typing:

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

See "Debugger User Model" on page 4.

2. Navigate to the function to be changed.

You can get to the function numerous ways, by selecting "Search..." from the Source menu, typing **func** *functionname* at the Debugger command line, or simply scrolling to the location. If you did not use **func** *functionname*, you need to place the cursor inside the function.

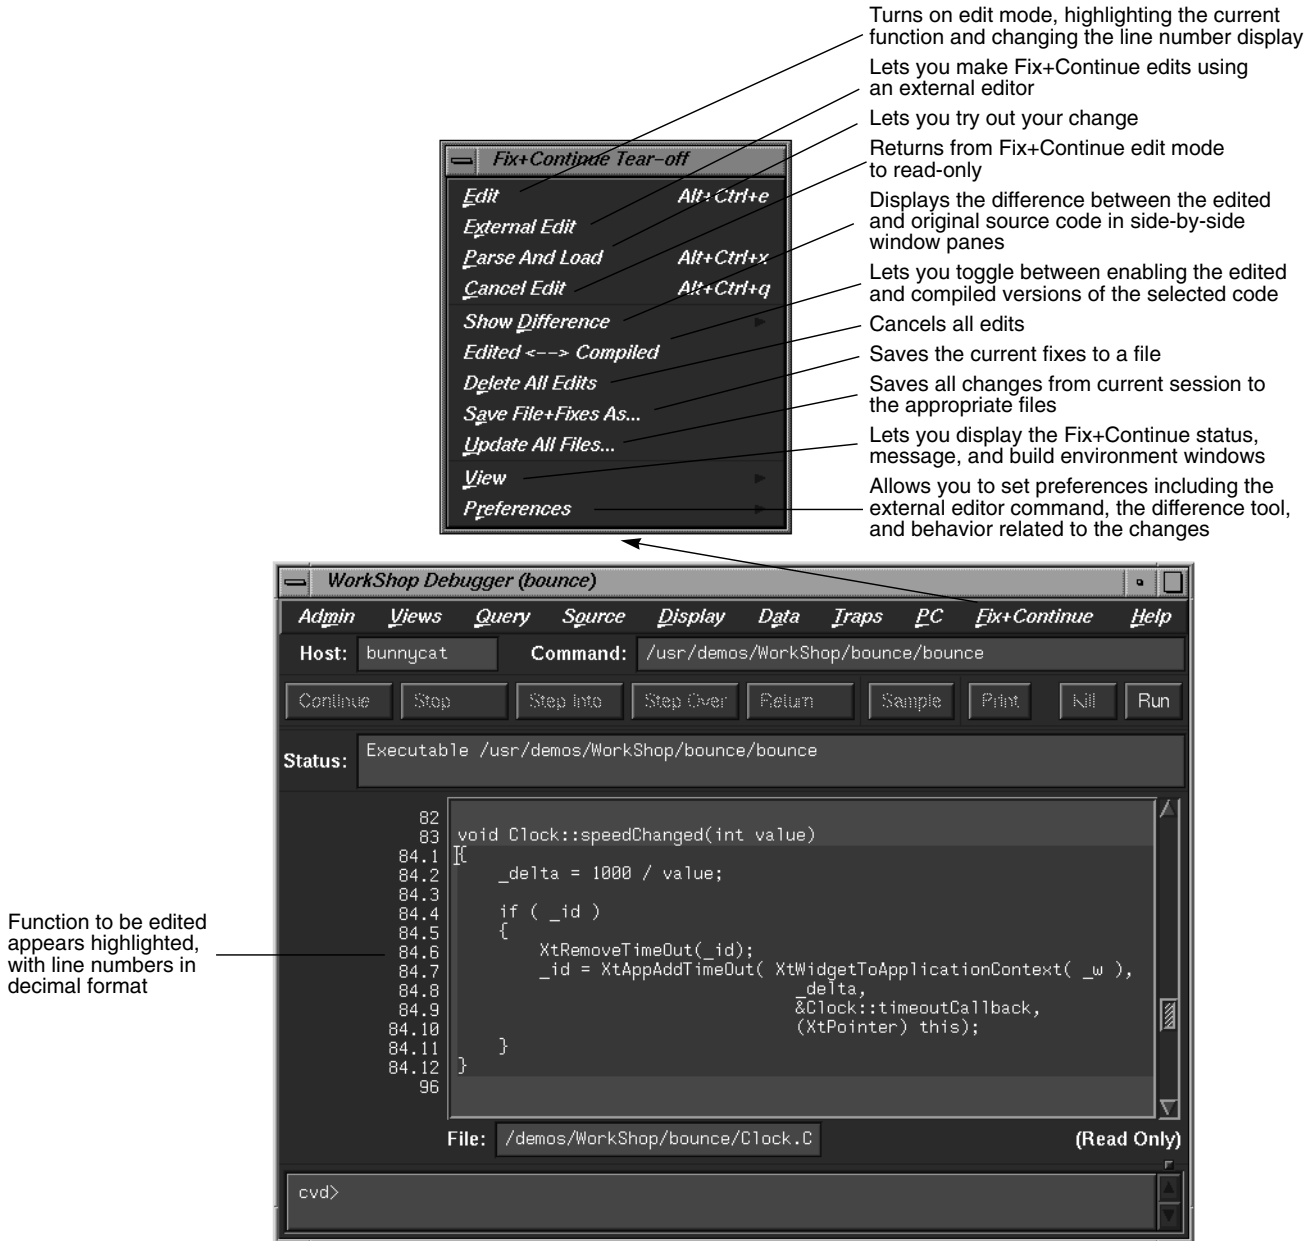


Figure 13 Using Fix+Continue

3. Select "Edit" from the Fix+Continue menu.

This turns on edit mode, highlighting the function source code. If line numbers are displayed, those in the selected function appear with a two-part number separated by a decimal point. The left part represents the starting line number of the function in the source file before selecting "Edit". The right part rennumbers the source within the function to make it easier to keep track of added new lines.

4. Make your changes to the source code.

You can do this directly in Main View or you can use a preferred editor by selecting "External Edit" from the Fix+Continue menu.

5. Try out your changes.

Selecting "Parse And Load" adds your changes to the executable you are debugging. The changed function will get executed the next time it is invoked. If you stopped in the edited function, the Debugger will let you continue from the corresponding line in the new function, barring certain restrictions.

6. If the changes are satisfactory, save them for later compiling.

"Save File+Fixes As ..." saves the fixes in your current file. "Update All Files..." saves all fixes in your current session.

At any point, you can make comparisons with your old code. "Show Difference" displays the old and new source code in a side-by-side format. "Edited<-->Compiled" lets you toggle between the old and new executables making it easy to verify or demonstrate your bug fix.

Where to Find Fix and Continue Information

To find out more information about Fix and Continue, refer to Table 8.

Table 8 Where to Find Fix and Continue Information in the *CASEVision/WorkShop MegaDev User's Guide*


Topic	See ...
General information	Chapter 2, "Getting Started with Fix and Continue"
Fix and Continue tutorial	Chapter 3, "Using Fix and Continue: A Sample Session"
Detailed command information	Chapter 4, "Fix and Continue Reference"

Debugging IRIX IM Programs


The IRIX IM Analyzer provides special debugging support for IRIX IM applications and is available from the WorkShop Views menu. The IRIX IM Analyzer operates in a number of modes (referred to as *examiners*) for examining different types of IRIX IM objects. The IRIX IM Analyzer provides information unavailable through conventional debuggers. It also lets you set widget-level breakpoints and collect X event history.

When you first invoke the IRIX IM Analyzer, it comes up in its Widget Examiner mode. You can switch to the other examiners through the Examiner menu or by clicking the tabs at the bottom of the window (See Figure 14).

Examiner menu lets you select different types of data for examination.



Data display area shows data appropriate to the type of examiner.



Examiner tabs provide a quick way to select examiners

Figure 14 The IRIX IM Analyzer Window

Features of the IRIX IM Analyzer

The IRIX IM Analyzer provides the following types of examiners:

- Widget examiner—identifies a widget's ID, name, class, and parent, and displays its definitions.
- Widget tree examiner—displays the widget hierarchy (see Figure 15). The widgets can be displayed by name, class, or ID by selecting from

the widget display menu. Double-clicking a widget node switches to the widget examiner and displays the data for the selected widget.

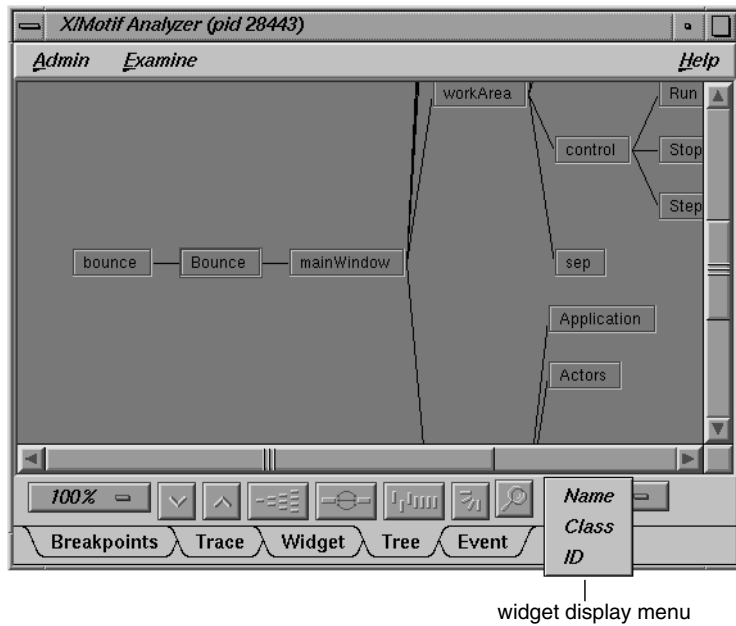


Figure 15 IRIX IM Analyzer Widget Tree Examiner

- Breakpoints examiner—lets you set breakpoints at the widget and widget class level. You can set breakpoints at
 - callback functions
 - widget events
 - resource changes caused
 - timeout callback functions
 - input callback functions
 - widget state changes
 - X events
 - X requests

- Trace examiner—lets you trace the execution of your application and collect the following types of data:
 - X Server Events
 - X Server Requests
 - widget event dispatch information
 - widget resource changes
 - widget state changes
 - Xt callbacks

Figure 16 is a typical example of the trace examiner. The events appear in a list. Double-clicking an event displays its details.

- Callback examiner—comes up automatically when the process stops in a callback. It displays
 - the callstack frame for the callback function
 - widget information
 - the callback data structure
- Window examiner—identifies the window, its parent and any children, and displays window attribute information
- Event examiner—displays the event structure for a given XEvent pointer
- Graphics context (GC) examiner—displays the X graphics context attributes for a given GC pointer
- Pixmap examiner—displays the basic attributes of an X pixmap, including size and depth, and can provide an ASCII display of small pixmaps, using the units digit of the pixel values
- widget class examiner—displays the widget class attributes for a given widget class pointer

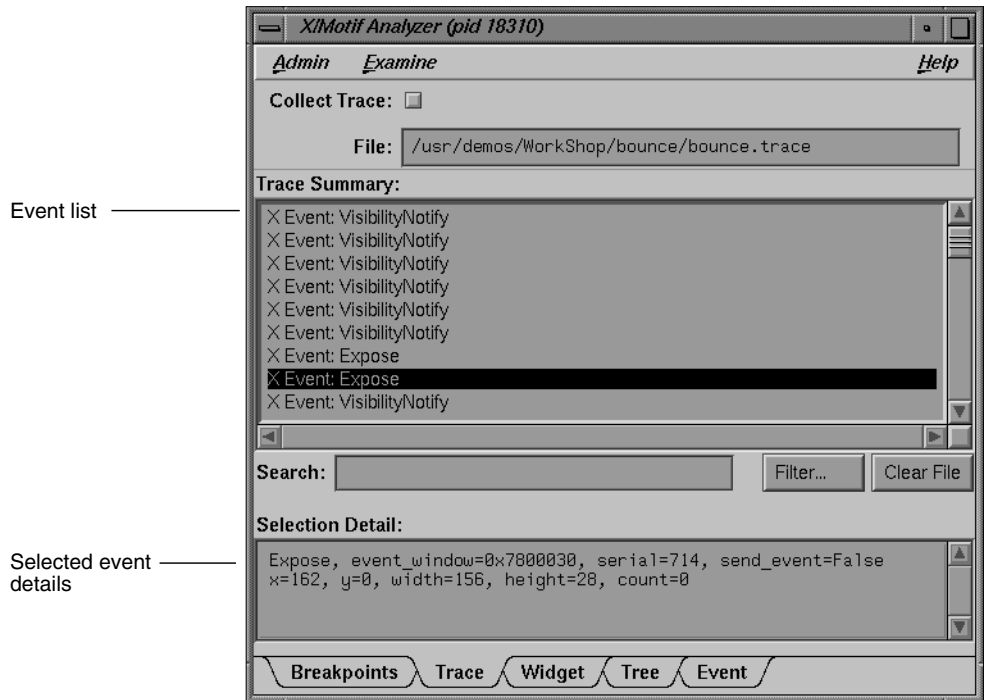


Figure 16 IRIX IM Analyzer Trace Examiner

Where to Find IRIX IM Analyzer Information

To find out more information about the IRIX IM Analyzer, refer to Table 9.

Table 9 Where to Find IRIX IM Analyzer information in the *Developer Magic: IRIX IM Analyzer User's Guide*

Topic	See ...
General information	Chapter 1, "Getting Started With the IRIX IM Analyzer"
IRIX IM Analyzer tutorial	Chapter 2, "Using the IRIX IM Analyzer: A Sample Session"
General reference information	Chapter 3, "IRIX IM Analyzer Reference"
Setting breakpoints to capture widget-level information	"Breakpoints Examiner"
Tracing widget-level data through the execution of a program	"Trace Examiner"
Getting information on a specified widget	"Widget Examiner"
Displaying a graph of the widget hierarchy	"Tree Examiner"
Getting information on a specified callback	"Callback Examiner"
Getting information on a specified window	"Window Examiner"
Getting information on a specified X event	"Event Examiner"
Getting information on a specified graphics context	"Graphics Context Examiner"
Getting information on a specified pixmap	"Pixmap Examiner"

Building Application Interfaces With RapidApp

RapidApp is a simple, interactive tool for creating application interfaces. It's integrated with the other WorkShop tools to provide a complete environment for developing object-oriented applications quickly and easily. RapidApp generates C++ code, with interface classes based on the IRIS ViewKit toolkit and IRIX IM. RapidApp also includes predefined interface components that allow you to conveniently use other Developer Magic libraries such as OpenGL™ and Open Inventor™. Applications produced by RapidApp are automatically integrated into the Indigo Magic Desktop environment, letting you take advantage of Silicon Graphics' interface and desktop technology.

Working with RapidApp is similar to using a drawing tool such as Showcase™. A typical RapidApp window is shown in Figure 17. RapidApp lets you create interface elements by clicking icons representing widgets or components in the palette area, positioning them in a template window, and setting their resources in the editing area.

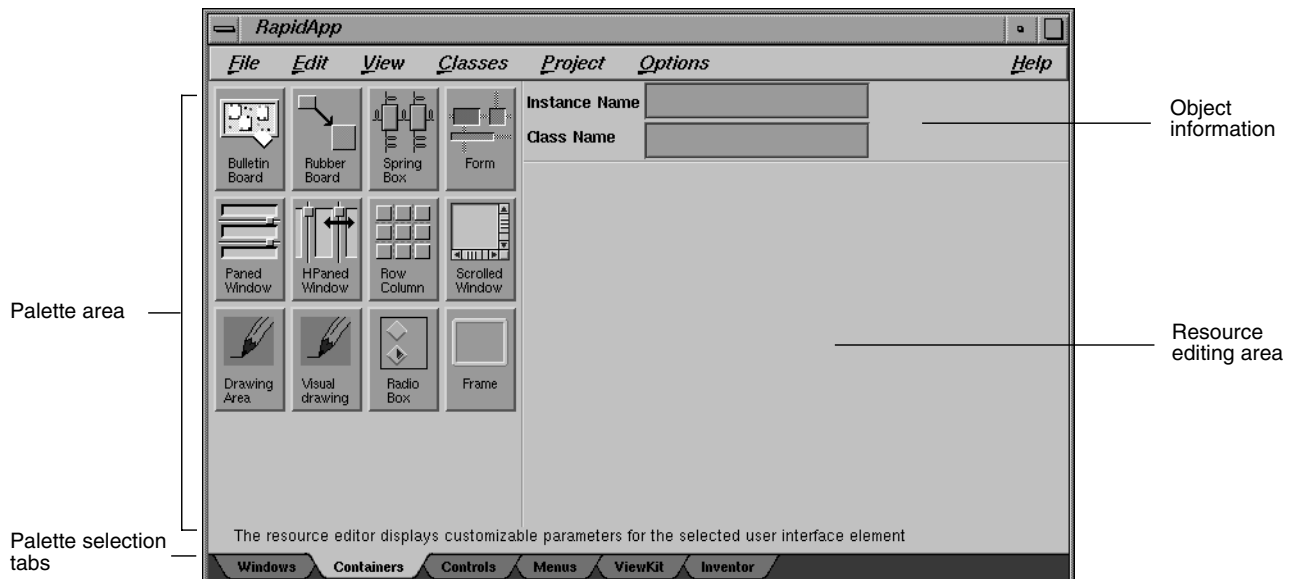


Figure 17 RapidApp Window Displaying Container Palette

RapidApp User Model

RapidApp users should be familiar with IRIX IM, IRIS ViewKit, and C++ programming. Here's the basic user model:

1. Invoke RapidApp by typing `rapidapp` in the directory in which you wish to build your application.

The RapidApp window is displayed as shown in Figure 17. There are six palettes of icon widgets available. The number of palettes and icons available will increase over time as new, useful widgets are developed. The palettes and icons are:

- Container palette—provides container widgets, that is, widgets that can hold other widgets
- Controls palette—provides miscellaneous widgets, typically for controls, fields, and so on.
- Windows palette—provides simple or special-purpose windows and window-oriented controls, such as menu bars and pulldown menus
- ViewKit palette—provides ViewKit components, that is, prepackaged assemblies of widgets from the ViewKit libraries
- Inventor palette—provides viewers, editors, and drawing areas compatible with IRIS Inventor™

The process is then one of selecting containers, populating them with widgets, and assembling them into elements of your user interface.

2. Select a container widget.

A rubber-band box appears, representing the initial default size of the widget. Use the mouse to drag it to a working area on your desktop (or inside another container). After you've positioned the new container widget, you can adjust its size by dragging the corners.

3. Edit the widget's resources.

Customize the widget for your application. RapidApp changes the resource editing area according to the type of widget you are working with. It displays text fields for string resources, radio buttons for Booleans, and menus for resources with multiple values. Figure 18

illustrates the creation of a drawing area container widget. The drawing area icon has been selected from the container palette, the new widget has been placed, and the resource editing area has changed accordingly.

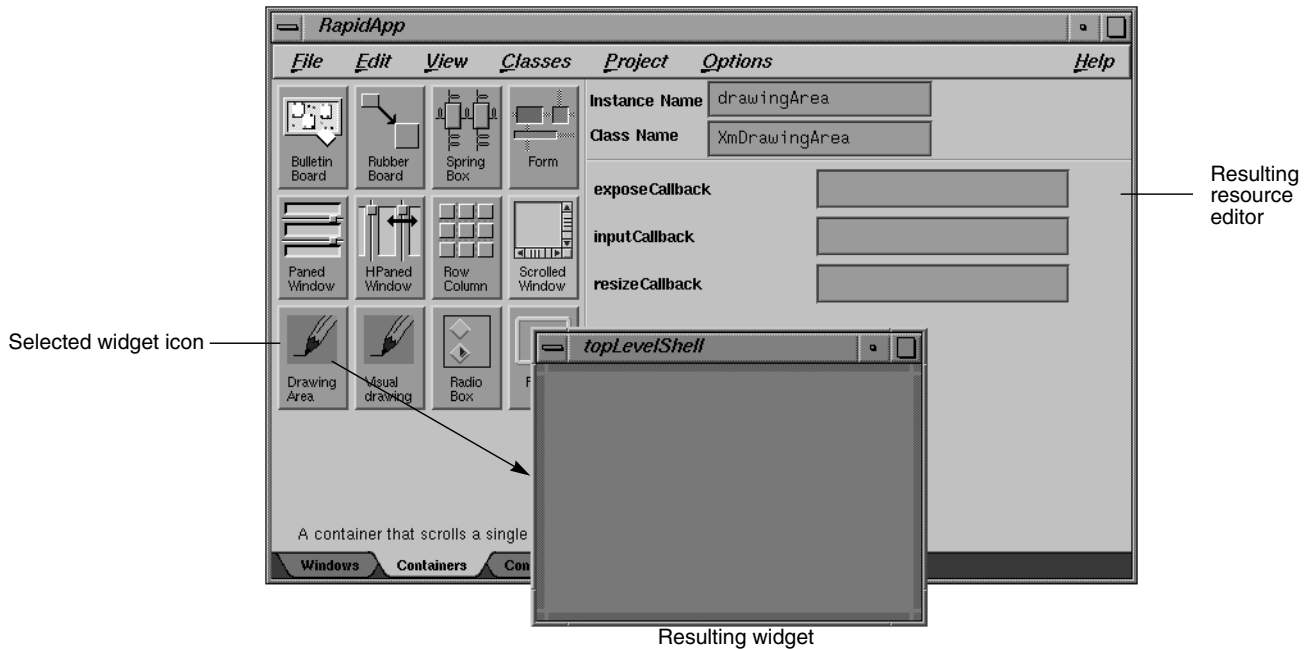


Figure 18 Creating a Widget

4. Select "Play Mode" from the View menu.

This lets you try out the interface design. When you are through trying it out, go back to working on the interface by selecting "Build Mode" from the View menu.

5. Perform any further edits on the widget.
6. Repeat steps 2-5 until your window (or application) is complete.
7. Select "Generate C++" from the Project menu.

This produces the source code (including Makefile) necessary to implement the interface you have designed. It also displays the Builder information window, a shell that displays RapidApp status messages.

8. Select “Edit File ...” from the Project menu to make any necessary adjustments to the source code.

A file selection dialog box displays showing the contents of the directory containing the generated source files. When you choose a file, it will appear in your default editor.

9. Select “Build Application” from the Project menu to compile the new program.

The WorkShop Build View displays and starts a compile going and lets you view any compile errors (see “Recompiling Within the ProDev WorkShop Environment With Build Manager” on page 30).

10. Use the other ProDev WorkShop and MegaDev tools, if necessary, to fix any coding problems.

RapidApp is fully integrated with the rest of the Developer Magic environment so that the full range of tools and libraries are at your disposal for completing your application.

Where to Find RapidApp Information

To find out more information about RapidApp, refer to Table 10.

Table 10 Where to Find RapidApp Information in the *Developer Magic: RapidApp User’s Guide*

Topic	See ...
Understanding the RapidApp window	“The RapidApp Interface” in Chapter 1
Using RapidApp	“Basic Interaction Techniques” in Chapter 1 and “Advanced Techniques and Features” in Chapter 3
General tutorial	“A Simple Example: A Calculator” in Chapter 1
Inventor tutorial	“Example: A Simple Inventor Program” in Chapter 4
OpenGL tutorial	“Example: Using the OpenGL Widget” in Chapter 4

Table 10 **(continued)** Where to Find RapidApp Information in the
Developer Magic: RapidApp User's Guide

Topic	See ...
Windows	"Choosing and Using Windows" in Chapter 3
Containers	"Choosing and Using Containers" in Chapter 3
Generating software code	Chapter 2, "Creating Applications with RapidApp"
Applying the other ProDev tools to RapidApp applications	"Integration with WorkShop for Building and Debugging" in Chapter 2
RapidApp reference information	Appendix C, "Quick Reference"

We'd Like to Hear From You

As a user of Silicon Graphics documentation, your comments are important to us. They help us to better understand your needs and to improve the quality of our documentation.

Any information that you provide will be useful. Here is a list of suggested topics to comment on:

- General impression of the document
- Omission of material that you expected to find
- Technical errors
- Relevance of the material to the job you had to do
- Quality of the printing and binding

Please include the title and part number of the document you are commenting on. The part number for this document is 007-2582-001.

Thank you!

Three Ways to Reach Us



The **postcard** opposite this page has space for your comments. Write your comments on the postage-paid card for your country, then detach and mail it. If your country is not listed, either use the international card and apply the necessary postage or use electronic mail or FAX for your reply.



If **electronic mail** is available to you, write your comments in an e-mail message and mail it to either of these addresses:

- If you are on the Internet, use this address: techpubs@sgi.com
- For UUCP mail, use this address through any backbone site:
[your_site]!sgi!techpubs



You can forward your comments (or annotated copies of manual pages) to Technical Publications at this **FAX** number:

415 965-0964