



Prism™ 7.0 Software User's Guide

Sun Microsystems, Inc.
4150 Network Circle
Santa Clara, CA 95054 U.S.A.
650-960-1300

Part No. 817-0088-10
February 2003, Revision A

Send comments about this document to: docfeedback@sun.com

Copyright 2003 Sun Microsystems, Inc., 4150 Network Circle, Santa Clara, California 95054, U.S.A. All rights reserved.

Sun Microsystems, Inc. has intellectual property rights relating to technology embodied in the product that is described in this document. In particular, and without limitation, these intellectual property rights may include one or more of the U.S. patents listed at <http://www.sun.com/patents> and one or more additional patents or pending patent applications in the U.S. and in other countries.

This document and the product to which it pertains are distributed under licenses restricting their use, copying, distribution, and decompilation. No part of the product or of this document may be reproduced in any form by any means without prior written authorization of Sun and its licensors, if any.

Third-party software, including font technology, is copyrighted and licensed from Sun suppliers.

Parts of the product may be derived from Berkeley BSD systems, licensed from the University of California. UNIX is a registered trademark in the U.S. and in other countries, exclusively licensed through X/Open Company, Ltd.

Sun, Sun Microsystems, the Sun logo, AnswerBook2, docs.sun.com, Sun HPC ClusterTools, Prism, Forte, Sun Performance Library, RSM, UltraSPARC, and Solaris are trademarks or registered trademarks of Sun Microsystems, Inc. in the U.S. and in other countries.

All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. in the U.S. and in other countries. Products bearing SPARC trademarks are based upon an architecture developed by Sun Microsystems, Inc.

The OPEN LOOK and Sun™ Graphical User Interface was developed by Sun Microsystems, Inc. for its users and licensees. Sun acknowledges the pioneering efforts of Xerox in researching and developing the concept of visual or graphical user interfaces for the computer industry. Sun holds a non-exclusive license from Xerox to the Xerox Graphical User Interface, which license also covers Sun's licensees who implement OPEN LOOK GUIs and otherwise comply with Sun's written license agreements.

Use, duplication, or disclosure by the U.S. Government is subject to restrictions set forth in the Sun Microsystems, Inc. license agreements and as provided in DFARS 227.7202-1(a) and 227.7202-3(a) (1995), DFARS 252.227-7013(c)(1)(ii) (Oct. 1998), FAR 12.212(a) (1995), FAR 52.227-19, or FAR 52.227-14 (ALT III), as applicable.

DOCUMENTATION IS PROVIDED "AS IS" AND ALL EXPRESS OR IMPLIED CONDITIONS, REPRESENTATIONS AND WARRANTIES, INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT, ARE DISCLAIMED, EXCEPT TO THE EXTENT THAT SUCH DISCLAIMERS ARE HELD TO BE LEGALLY INVALID.

Copyright 2003 Sun Microsystems, Inc., 4150 Network Circle, Santa Clara, California 95054, Etats-Unis. Tous droits réservés.

Sun Microsystems, Inc. a les droits de propriété intellectuels relatants à la technologie incorporée dans le produit qui est décrit dans ce document. En particulier, et sans la limitation, ces droits de propriété intellectuels peuvent inclure un ou plus des brevets américains énumérés à <http://www.sun.com/patents> et un ou les brevets plus supplémentaires ou les applications de brevet en attente dans les Etats-Unis et dans les autres pays.

Ce produit ou document est protégé par un copyright et distribué avec des licences qui en restreignent l'utilisation, la copie, la distribution, et la décompilation. Aucune partie de ce produit ou document ne peut être reproduite sous aucune forme, par quelque moyen que ce soit, sans l'autorisation préalable et écrite de Sun et de ses bailleurs de licence, s'il y en a.

Le logiciel détenu par des tiers, et qui comprend la technologie relative aux polices de caractères, est protégé par un copyright et licencié par des fournisseurs de Sun.

Des parties de ce produit pourront être dérivées des systèmes Berkeley BSD licenciés par l'Université de Californie. UNIX est une marque déposée aux Etats-Unis et dans d'autres pays et licenciée exclusivement par X/Open Company, Ltd.

Sun, Sun Microsystems, le logo Sun, AnswerBook2, docs.sun.com, Sun HPC ClusterTools, Prism, Forte, Sun Performance Library, RSM, UltraSPARC, et Solaris sont des marques de fabrique ou des marques déposées de Sun Microsystems, Inc. aux Etats-Unis et dans d'autres pays.

Toutes les marques SPARC sont utilisées sous licence et sont des marques de fabrique ou des marques déposées de SPARC International, Inc. aux Etats-Unis et dans d'autres pays. Les produits portant les marques SPARC sont basés sur une architecture développée par Sun Microsystems, Inc.

L'interface d'utilisation graphique OPEN LOOK et Sun™ a été développée par Sun Microsystems, Inc. pour ses utilisateurs et licenciés. Sun reconnaît les efforts de pionniers de Xerox pour la recherche et le développement du concept des interfaces d'utilisation visuelle ou graphique pour l'industrie de l'informatique. Sun détient une licence non exclusive de Xerox sur l'interface d'utilisation graphique Xerox, cette licence couvrant également les licences de Sun qui mettent en place l'interface d'utilisation graphique OPEN LOOK et qui en outre se conforment aux licences écrites de Sun.

LA DOCUMENTATION EST FOURNIE "EN L'ÉTAT" ET TOUTES AUTRES CONDITIONS, DECLARATIONS ET GARANTIES EXPRESSES OU TACITES SONT FORMELLEMENT EXCLUES, DANS LA MESURE AUTORISÉE PAR LA LOI APPLICABLE, Y COMPRIS NOTAMMENT TOUTE GARANTIE IMPLICITE RELATIVE A LA QUALITE MARCHANDE, A L'APTITUDE A UNE UTILISATION PARTICULIERE OU A L'ABSENCE DE CONTREFAÇON.



Contents

Preface xxi

1. Prism Environment 1

Overview 1

 The Prism Environment's Operating Modes 2

 MP Mode Summary 2

Look and Feel of the Prism Programming Environment 3

Loading and Executing Programs 5

Debugging 5

Visualizing Data 6

Editing and Compiling 6

Obtaining Online Help and Documentation 6

Customizing the Prism Programming Environment 7

2. Using the Prism Environment 9

Before Entering the Prism Environment 9

 Supported Languages and Compilers 10

 Compiling and Linking Your Program 10

 Combining Compiler Options 10

 Setting Up Your Working Environment 10

DISPLAY Variable	11
Starting the Prism Environment	11
Starting in MP Program Mode	11
Specifying Multiple Processes	12
Specifying One Process per Available CPU	13
Specifying a Threaded View of Programs	13
Specifying a Nonthreaded View of Programs	14
Specifying a Threaded View for a Single-Process Program	15
Running MP Prism in the Background	15
Attaching to a Running MP Job	15
Starting in Serial Programming Mode	16
Attaching to a Running Serial Program	16
After Starting the Prism Environment	17
Loading a Multiprocess Program From Within the Prism Environment	17
After Loading a Program	17
The Prism Prompt	18
Enabling Support for Spawned MPI Processes	18
Specifying Which Spawned Executables to Debug	19
Attaching to a Job or Process	22
Associating a Program or Process With a Core File	22
Specifying X Toolkit Options	22
Specifying Input and Output Files	23
Specifying Runtime Environment Options	23
Passing Command-Line Options to Secondary Sessions	23
Executing Commands Within the Prism Environment	24
Using a Mouse	24
Using Keyboard Alternatives to the Mouse	25
Issuing Commands	26

Using the Menu Bar	27
Keyboard Accelerators	27
Using the Source Window	28
Moving Through the Source Code	28
Selecting Text	29
Splitting the Source Window	29
Using the Line-Number Region	31
Using the Command Window	32
Using the Command Line	33
Using the History Region	33
Redirecting Output	35
Logging Commands and Output	36
▼ To Specify the Name of a Log File.	36
▼ To Append the Log to an Existing File	36
▼ To Turn Off Logging	36
Rerunning a Prism Session That Was Saved to a Log File	37
Writing Expressions in the Prism Environment	37
How the Prism Environment Chooses the Correct Variable or Procedure	37
Using Qualified Names	38
Using Fortran Intrinsic Functions in Expressions	39
Using C Arrays in Expressions	40
Using Array-Section Syntax in C Arrays	41
Hints for Detecting NaNs and Infinities	41
Using Fortran 90 Generic Procedures	42
Issuing Solaris Commands	44
▼ To Issue Solaris Commands From Within the Prism Environment	44
Changing the Current Working Directory	45
Setting and Displaying Environment Variables	45

Killing Processes Within the Prism Environment	45
▼ To Kill a Process or Job Running Within the Prism Environment	45
▼ To Kill a Spawned Prism Session	45
Leaving the Prism Environment	46
▼ To Exit a Single-Job Prism Session	46
▼ To Quit a Spawned Prism Session	47
3. Loading and Executing a Program	49
Loading a Program	49
▼ To Load a Program From the Menu Bar	50
▼ To Load a Program From the Command Window	51
What Happens When You Load a Program	51
Associating a Core File With a Loaded Program	52
▼ To Associate a Core File With a Loaded Program	52
▼ To Examine the Core File of a Local Process	52
Attaching to a Running Message-Passing Process	53
▼ To Attach to a Running Message-Passing Process	53
▼ To Attach to Multiple Jobs When Starting Prism	54
Detaching From a Running Process	54
Executing a Program in the Prism Environment	55
▼ To Run a Program	55
Program I/O	56
Status Messages	56
Stepping and Continuing Through a Program	57
Waiting for and Interrupting Processes	57
Execution Pointer	59
Rerunning Spawned Prism Sessions	60
Controlling Programs With the Commands-Only Interface	60
Using Psets in the Prism Environment	61

Using the Psets Navigator	62
Using the Psets Window	62
Predefined Psets	65
Defining Psets	66
Naming Psets	68
Combining Named Psets and Pset Expressions	71
Viewing Pset Contents From the Psets Window	71
Viewing Pset Contents From the Command Line	73
Deleting Psets	75
The Current Pset	75
Current Pset and Predefined Psets	76
Current Pset and Variable Psets	77
The Current Process	78
Scope in the Prism Environment	79
The <code>cycle</code> Pset	80
Cycle Visualizer Window	81
Hiding Threads From Psets	81
Using Psets in Commands	82
▼ To Use a Pset Qualifier	83
Using Unbounded Psets in Commands	83
Using Snapshots of Unbounded Psets in Commands	85
Referring to Nonexistent Thread Identifiers	87
Using the Prism Environment With Sun MPI Client/Server Programs	87
Choosing the Current File and Function	88
▼ To Change the Current File	88
▼ To Change the Current Function or Procedure	89
Creating a Directory List for Source Files	90
▼ To Add a Directory to the Search Path	90

4. Debugging a Program	93
Overview of Events	93
Using the Event Table	95
Description of the Event Table	95
Adding an Event	97
Deleting an Existing Event	98
Editing an Existing Event	99
Disabling and Enabling Events	99
Saving Events	100
Events Taking Pset Qualifiers	101
Events and Dynamic Psets	102
Events and Variable Psets	102
Actions in Events	103
Events and Deleted Psets	104
Setting Breakpoints	104
Using the Line-Number Region	105
What Happens in a Split Source Window	106
Using the Event Table and the Events Menu	106
Setting a Breakpoint Using Commands	107
Examples of the <code>stop</code> Command	108
Tracing Program Execution	109
Using the Event Table and Events Menu	110
Using the Command Window	111
Displaying and Moving Through the Call Stack	112
▼ To Display the Call Stack	112
Moving Through the Call Stack	113
Displaying the Where Graph	113
Panning and Zooming in the Where Graph	115

Combining Debug and Optimization Options	120
Interpreting Interaction Between an Optimized Program and the Prism Environment	120
Accessing Variables in Optimized Routines	121
Debugging Spawned Sun MPI Processes	122
Debugging Spawned Sessions Using the Commands-Only Interface	123
Prism Commands With Special Functions in Spawned Sessions	123
Error Conditions Arising From Spawned Sessions	124
Examining the Contents of Memory and Registers	125
▼ To Display Memory	126
▼ To Display the Contents of Registers	127
5. Visualizing Data	129
Overview of Data Visualization	129
Printing and Displaying	129
Visualization Methods	130
Data Visualization Limits	131
Choosing the Data to Visualize	131
▼ To Print or Display a Variable or Expression at the Current Program Location	131
▼ To Print or Display From the Source Window	132
▼ To Print or Display From the Events Menu	133
▼ To Print or Display From the Event Table	134
▼ To Print or Display From the Command Window	134
▼ To Print or Display the Contents of a Register	135
▼ To Set the Context	136
▼ To Specify the Radix	137
Working With Visualizers	137
Using the Data Navigator in a Visualizer	139
Using the Display Window in a Visualizer	139

Using the Options Menu	140
Updating and Closing the Visualizer	152
Saving, Restoring, and Comparing Visualizers	153
▼ To Save the Values of a Variable	153
▼ To Restore the Data	154
▼ To Compare the Data	155
Visualizing Structures	157
Expanding Pointers	158
More About Pointers in Structures	160
Augmenting the Information Available for Display	161
Printing the Type of a Variable	163
▼ To Print the Type of a Variable From the Menu Bar	163
▼ To Print the Type of a Variable From the Source Window	164
▼ To Print the Type of a Variable From the Command Window	164
What Is Displayed	164
Changing the Radix of Data	165
Printing Pointers as Array Sections	166
▼ To Print an Array by Section	166
▼ To View a Pointer as a One-Dimensional Array	167
▼ To Dereference an Array of Pointers	167
▼ To Cast Pointers	168
Visualizing Multiple Processes	168
▼ To Find Out the Value and Process Number for an Element	170
▼ To Open a <code>Cycle</code> Visualizer Window	171
Visualizing MPI Message Queues	171
▼ To Launch the MPI Queue Visualizer	172
▼ To Select the Queue to Visualize	172
▼ To Zoom Through Levels of Message Detail	172

▼	To Control the Values of Message Labels	176
▼	To Sort Messages	176
▼	To Display Message Fields	177
	Interpreting Message Dialog Fields	178
	Displaying Communicator Data	179
	Displaying MPI Message Queue Text	182
	Displaying and Visualizing Sun S3L Arrays	183
▼	To Display the Data Type of an Array Handle	185
▼	To Create an S3L Parallel Array	185
▼	To Display and Visualize Sun S3L Parallel Arrays	185
▼	To Visualize the Layouts of S3L Parallel Arrays	188
▼	To Print or Display an S3L Array Using the <code>layout</code> Intrinsic	189
6.	Editing and Compiling Programs	191
	Editing Source Code	191
	▼	To Start the Default Editor on the Current Source File 191
	Using the <code>make</code> Utility	192
	Creating the Makefile	192
	Using the Makefile	192
7.	Getting Help	195
	Prism Online Help Systems	195
	▼	To Get Help in the Prism Environment 195
	Using the Browser-Based Help System	196
	Choosing Selections From the Help Menu	196
	Getting Help on Using the Mouse	196
	Obtaining Help From the Command Window	197
	Obtaining Online Documentation	197
	Viewing Manual Pages	197

8. Customizing the Prism Programming Environment	199
Initializing the Prism Environment	199
Customizing MP Prism Mode	200
Using the Tear-Off Region	201
Adding Menu Selections to the Tear-Off Region	201
Adding Prism Commands to the Tear-Off Region	202
Creating Aliases for Commands and Variables	203
▼ To Create an Alias for a Prism Command	203
▼ To Remove an Alias	203
▼ To Set Up an Alternative Name for a Variable or Expression	204
Changing Prism Resource Defaults	204
▼ To Launch the Prism Customize Utility	204
Changing a Resource Setting	205
Resource Descriptions	206
Where the Prism Environment Stores Your Changes	208
Changing Prism Environment Defaults	209
Adding Prism Resources to the X Resource Database	211
Specifying the Editor and Its Placement	212
Specifying the Window for Error Messages	212
Changing the Text Fonts	213
Changing Colors	213
Changing Keyboard Translations	216
Changing Keyboard Translations in Text Widgets	216
Changing General Motif Keyboard Translations	217
Changing Xterm Use With I/O	218
Changing the Way the Prism Environment Signals an Error	218
Changing the make Utility to Use	219
Changing How the Prism Environment Treats Stale Data in Visualizers	219

Specifying a Different Browser for Displaying Help 219

Changing the Way the Prism Environment Handles Fortran 90 Generic
Procedures 220

9. Troubleshooting 223

Avoid Using the `-xs` Compiler Option 223

Keep `.o` Files after Compilation 223

Expect a Pause After Issuing the First `run` Command 224

Monitor Your Use of Color Resources 224

Expect Only Stopped Processes to Be Displayed in the Where Graph 224

Use Only the MP Mode of the Prism Environment to Load MPI Programs 225

Verify That `/opt/SUNWlslf/bin` Is in Your PATH 225

MP Prism Environment Freezes on Startup 226

A. Commands-Only Mode of the Prism Environment 227

B. C++ and Fortran 90 Support 231

C. The Scalar Mode of the Prism Environment 243

▼ To Launch the Prism Environment in Scalar Mode 243

▼ To Attach To a Running Process From Within the Prism Environment 245

Figures

FIGURE 1-1	The Prism Programming Environment's Main Window	4
FIGURE 2-1	Prism Sessions Created by Calls to <code>MPI_Comm_spawn_multiple</code>	21
FIGURE 2-2	Pop-up Menu in Source Window	29
FIGURE 2-3	Split Source Window	30
FIGURE 2-4	Line-Number Region	31
FIGURE 2-5	Command Window With History Region	32
FIGURE 2-6	Generic Procedure Dialog Box	43
FIGURE 2-7	Subprocess Warning	46
FIGURE 3-1	Open Program Filter	50
FIGURE 3-2	Run (args) Dialog Box	55
FIGURE 3-3	Pset Navigator Controls	62
FIGURE 3-4	Psets Window (Nonthreaded)	63
FIGURE 3-5	Psets Window (Threaded)	64
FIGURE 3-6	File Window	89
FIGURE 3-7	Use Dialog Box	91
FIGURE 4-1	Event Table	96
FIGURE 4-2	Pset Field in the Prism Environment's Event Table	101
FIGURE 4-3	Stop <loc> Dialog Box	106
FIGURE 4-4	Where Window	112
FIGURE 4-5	Where Graph	114

FIGURE 4-6	Where Graph, Zoomed Out One Level	116
FIGURE 4-7	Where Graph, Zoomed Out to the Maximum	117
FIGURE 4-8	Where Graph, Zoomed In	118
FIGURE 4-9	Where Graph of a Threaded Program, Zoomed in to Show Thread Stripes	119
FIGURE 5-1	Print Dialog Box	132
FIGURE 5-2	Print Dialog Box	133
FIGURE 5-3	Visualizer for a Three-Dimensional Array	138
FIGURE 5-4	Options Menu in a Visualizer	141
FIGURE 5-5	Histogram Visualizer	142
FIGURE 5-6	Dither Visualizer	143
FIGURE 5-7	Threshold Visualizer	144
FIGURE 5-8	One-Dimensional Graph Visualizer	145
FIGURE 5-9	Surface Visualizer	146
FIGURE 5-10	Vector Visualizer	147
FIGURE 5-11	Visualization Parameters Dialog Box	148
FIGURE 5-12	Threshold Visualizer With a Ruler	150
FIGURE 5-13	Statistics for a Visualizer	151
FIGURE 5-14	Set Context Dialog Box	152
FIGURE 5-15	Saving a Visualizer's Data to a File	154
FIGURE 5-16	Diff With Dialog Box	156
FIGURE 5-17	Structure Visualizer	157
FIGURE 5-18	Structure Visualizer, With One Pointer Expanded	158
FIGURE 5-19	Zooming Out in a Structure Visualizer	160
FIGURE 5-20	Visualizer in the Prism Environment (Threshold Representation)	169
FIGURE 5-21	Queue Visualizer at Zoom Level One	173
FIGURE 5-22	Queue Visualizer at Zoom Level Two	174
FIGURE 5-23	Queue Visualizer at Zoom Level Three	175
FIGURE 5-24	Queue Visualizer at Zoom Level Four	175
FIGURE 5-25	Message Dialog Box	178
FIGURE 5-26	Communicator Dialog Box	180

- FIGURE 5-27 Data Type Dialog Box 181
- FIGURE 6-1 The make Window 193
- FIGURE 7-1 xman Window 198
- FIGURE 8-1 Tear-Off Region 201
- FIGURE 8-2 Tear-Off Region Dialog Box 201
- FIGURE 8-3 Customize Window 205

Tables

TABLE 2-1	Passing Command Line Options to Secondary Sessions	24
TABLE 2-2	General Keyboard Alternatives to Mouse Control	25
TABLE 2-3	Text-Entry Keyboard Alternatives	26
TABLE 2-4	Keyboard Accelerators for Main Menu Selections	27
TABLE 2-5	Prism Identifier Syntax	38
TABLE 3-1	Status Messages	56
TABLE 3-2	Examples of Pset Composition	71
TABLE 4-1	Error Messages Related to Debugging of Spawned Processes	124
TABLE 4-2	Memory Address Formats	126
TABLE 4-3	UltraSPARC Registers	127
TABLE 5-1	Row Sort Criteria	176
TABLE 5-2	Column Sort Criteria	177
TABLE 5-3	Message Dialog Box Fields	178
TABLE 5-4	Communicator Dialog Box Fields	180
TABLE 5-5	Data Type Dialog Box Fields	181
TABLE 8-1	Sample Visualizer Colors	208
TABLE 8-2	Prism Resources	209

Preface

The *Prism User's Guide* explains how to use the Prism™ environment to develop, execute, debug, and visualize data in serial and parallel programs.

These instructions are intended for application programmers developing serial or parallel programs that are to run on a Sun HPC ClusterTools™ system. It is assumed you know the basics of developing and debugging programs, as well as the basics of the system on which you will be using the Prism environment. Some familiarity with the Solaris™ debugger `dbx` is helpful but not required. Prism is based on the X and OSF/Motif standards. Familiarity with these standards is also helpful but not required.

How This Book Is Organized

Chapter provides an introduction to the Prism environment.

Chapter contains instructions for using Prism features of broadest interest.

Chapter explains how to load and run programs in the Prism environment.

Chapter provides instructions for debugging programs in the Prism environment.

Chapter explains how to use the Prism environment's visualization capabilities to examine data in useful ways.

Chapter provides instructions for editing and compiling programs within the Prism environment.

Chapter explains how to get help from the Prism environment's online help system.

Chapter explains how to customize certain aspects of the Prism programming environment.

Chapter contains tips for recognizing and correcting problems.

Appendix provides instructions for using the commands-only interface to the Prism environment.

Appendix describes the Prism environment's support for debugging C++ and F90 programs.

Appendix discusses the Prism environment's scalar mode.

Using UNIX Commands

This document may not contain information on basic UNIX® commands and procedures such as shutting down the system, booting the system, and configuring devices.

See one or both of the following for this information:

- AnswerBook2™ online documentation for the Solaris operating environment
- Other software documentation that you received with your system

Typographic Conventions

Typeface*	Meaning	Examples
AaBbCc123	The names of commands, files, and directories; on-screen computer output	Edit your <code>.login</code> file. Use <code>ls -a</code> to list all files. % You have mail.
AaBbCc123	What you type, when contrasted with on-screen computer output	% su Password:
<i>AaBbCc123</i>	Book titles, new words or terms, words to be emphasized. Replace command-line variables with real names or values.	Read Chapter 6 in the <i>User's Guide</i> . These are called <i>class</i> options. You <i>must</i> be superuser to do this. To delete a file, type <code>rm filename</code> .

* The settings on your browser might differ from these settings.

Shell Prompts

Shell	Prompt
C shell	<i>machine-name%</i>
C shell superuser	<i>machine-name#</i>
Bourne shell and Korn shell	\$
Bourne shell and Korn shell superuser	#

Related Documentation

Application	Title	Part Number
Sun HPC ClusterTools Documentation	<i>Read Me First: Guide to Sun HPC ClusterTools Documentation</i>	817-0096-10
Sun HPC ClusterTools Software	<i>Sun HPC ClusterTools 5 Product Notes</i>	817-0081-10
	<i>Sun HPC ClusterTools 5 Installation Guide</i>	817-0082-10
	<i>Sun HPC ClusterTools 5 Performance Guide</i>	817-0090-10
	<i>Sun HPC ClusterTools 5 Administrator's Guide</i>	817-0083-10
	<i>Sun HPC ClusterTools 5 User's Guide</i>	817-0084-05
Sun MPI Programming	<i>Sun MPI 6.0 Programming and Reference Guide</i>	817-0085-10
Sun S3L	<i>Sun S3L 4.0 Programming Guide</i>	817-0086-10
	<i>Sun S3L 4.0 Reference Manual</i>	817-0087-10
Prism™ graphical programming environment	<i>Prism 7.0 Software Reference Manual</i>	817-0089-10

Accessing Sun Documentation

You can view, print, or purchase a broad selection of Sun documentation, including localized versions, at:

<http://www.sun.com/documentation>

Sun Welcomes Your Comments

Sun is interested in improving its documentation and welcomes your comments and suggestions. You can email your comments to Sun at:

docfeedback@sun.com

Please include the part number (817-0088-10) of your document in the subject line of your email. Table listing other documents that are related to this book or product.

Prism Environment

The Prism programming environment is an integrated graphical environment within which you can develop, execute, and debug programs. It provides an easy-to-use, flexible, and comprehensive set of tools for performing all aspects of serial and message-passing programming. Prism software operates on terminals or workstations running the Solaris operating environment under either the OpenWindows™ environment or the Sun Common Desktop Environment (CDE). The Prism environment also provides a commands-only option, which enables you to operate on any terminal, but without the graphical interface.

This chapter is organized as follows:

- “Overview” on page 1
- “Look and Feel of the Prism Programming Environment” on page 3
- “Loading and Executing Programs” on page 5
- “Debugging” on page 5
- “Visualizing Data” on page 6
- “Editing and Compiling” on page 6
- “Obtaining Online Help and Documentation” on page 6
- “Customizing the Prism Programming Environment” on page 7

Overview

When you want to debug or analyze a program that already exists, you can simply load the executable code into the Prism environment and immediately begin using the data visualization, performance analysis, and debugging tools built into the Prism environment. If the program does not yet exist, you can start the Prism environment, call up an editor and UNIX shell from within the Prism environment, and then perform all the usual tasks associated with developing a new program. That is, while in the Prism environment, you can:

- Invoke an editor of your choice and write and edit the source code.

- Compile the program, linking in libraries of interest.
- Run the program.

Once you have produced executable code, you can begin the program analysis and debugging tasks, while staying in the same environment in which you developed the program.

The Prism Environment's Operating Modes

The Prism environment has two modes of operation, *MP* and *scalar*. Their respective roles can be summarized as follows:

- MP mode is used for debugging multiprocess programs or multiple threads in threaded serial programs.
- Scalar mode is used for debugging nonthreaded serial programs or a single thread in a threaded serial program.

The Prism environment is intended primarily for use in developing and debugging message-passing programs and multithreaded serial programs. Consequently, the chief focus of this manual is on using the Prism environment in MP mode.

MP Mode Summary

The MP mode provides a superset of Prism features that make it possible to access and control separate processes and threads. Since this capability is at the heart of debugging message-passing and multithreaded code, the Prism environment *must* be operating in MP mode when debugging programs of these kinds.

When operating MP mode:

- The Prism environment creates a separate debug process, called a *node Prism*, for each MP process and uses these Prism processes to collect information about the MP processes. If an MP process contains multiple threads, the node Prism associated with that MP process will also be responsible for debugging those threads.
- The logic used to control multiple threads is encapsulated within the node Prism processes and is therefore active only in MP mode. Consequently, MP mode is required to debug multithreaded programs, even if the program is not itself multiprocessing.
- Each Prism debug process runs on the same node as the MP process to which it is attached.
- If an MPI program spawns MP processes through calls to `MPI_Comm_spawn()` or `MPI_Comm_spawn_multiple()`, the Prism environment creates a new debug session to support the newly spawned processes. The initial Prism session is

referred to as the primary session and any debug sessions it creates are called secondary debug sessions. Each secondary debug session creates its own set of debug processes, one for each MP process in the spawned message-passing job.

- If calls to `MPI_Comm_spawn_multiple()` specify multiple executables, the Prism environment will create a separate secondary Prism session for each spawned executable. These secondary Prism sessions will include separate debug (node Prism) processes for every MP process in the spawned executables.

Prism's MP mode supports the grouping of program processes into logical sets, called *psets* (for process sets). When defining a pset, you can specify one or more attributes that distinguish certain MP processes from all other MP processes in the program. The resulting pset can then be used to restrict the effects of Prism operations to those processes within the defined pset. For example, a pset could be defined as all processes in which the value of a variable exceeds some threshold.

The Prism environment also provides several predefined process sets (psets)—that is, psets that have characteristics that are likely to be of interest in most MP mode debugging sessions. For example, three of the predefined psets are:

- `all` (all processes in the program)
- `running` (all processes that are currently running)
- `error` (all processes that have encountered an error)

Look and Feel of the Prism Programming Environment

FIGURE 1-1 shows the main window of the Prism environment with a program loaded. It is within this window that you debug and analyze your program. You can operate with a mouse, use keyboard equivalents of mouse actions, or issue keyboard commands.

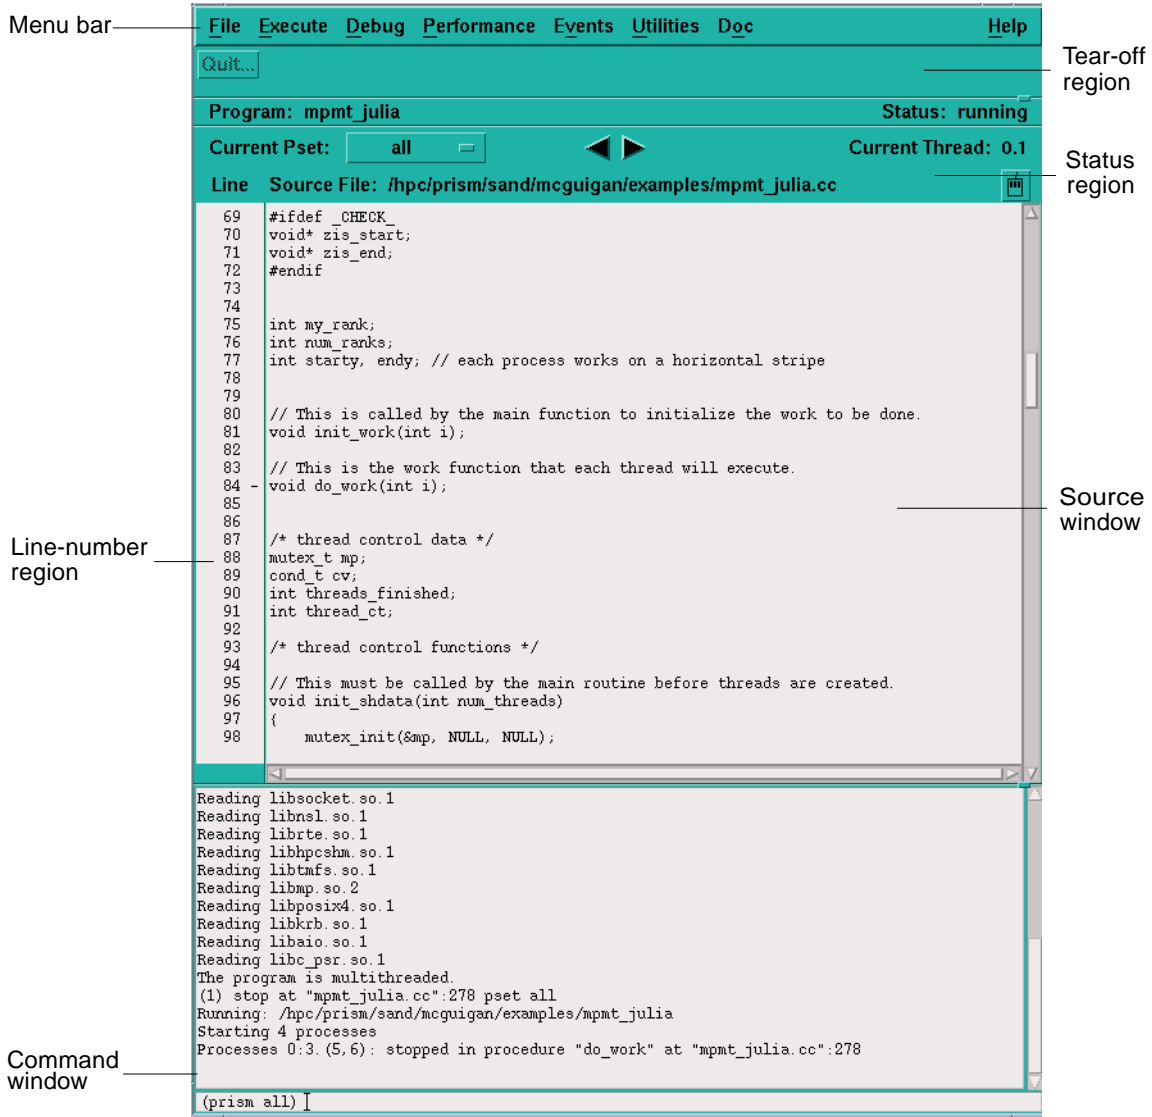


FIGURE 1-1 The Prism Programming Environment's Main Window

Clicking on items in the *menu bar* displays pull-down menus that provide access to most of the Prism environment's functionality.

You can add frequently used menu items and commands to the *tear-off region*, below the menu bar, to make them more accessible.

The *status region* displays the program's name and messages about the program's status.

The *source window* displays the source code for the executable program. You can scroll through this source code and display any of the source files used to compile the program. When a program stops execution, the source window updates to show the code currently being executed. You can select variables or expressions in the source code and print their values or obtain other information about them.

The *line-number region* is associated with the source window. You can click to the right of a line number in this region to set a breakpoint at that line.

The *command window* at the bottom of the main Prism environment window displays messages and output from the Prism environment. You can type commands in the command window rather than use the graphical interface menus.

General aspects of using these areas are discussed in Chapter 2.

Loading and Executing Programs

You can load an executable program into the Prism environment when you start it up, or any time afterward. You can also attach to a program that is already running. Once a program is loaded, you can run the program or step through it. You can interrupt execution at any time.

See “Loading and Executing a Program” on page 49 for a discussion of these topics.

Debugging

The Prism environment enables you to perform standard debugging operations such as setting breakpoints and tracepoints and displaying and moving through the call stack. Chapter 4 discusses these topics.

Visualizing Data

It is often important to obtain a visual representation of the data elements that make up an array or parallel variable. In the Prism environment, you can create *visualizers* that provide standard representations of variables or expressions. For example,

- In the *text* representation, the data are shown as numbers or characters.
- In the *colormap* representation, each data element is mapped to a color, based on a range of values and a color map that you specify. (This representation is available only on color workstations.)
- In the *threshold* representation, each data element is mapped to either black or white, based on a cutoff value that you can specify.

A *data navigator* lets you manipulate the display window relative to the data being visualized. Options are available that let you update a visualizer or save a snapshot of it.

See Chapter 5 for additional discussion, covering aspects of visualization specific to the MP Prism environment.

Editing and Compiling

You can call up the editor of your choice within the Prism environment to edit source code (or any other text files). If you change your source code and want to recompile, the Prism environment also provides an interface to the UNIX `make` utility. Editing and compiling are described in more detail in Chapter 6.

Obtaining Online Help and Documentation

The Prism environment features a comprehensive online help system. Help is available for each menu, window, and dialog box in the Prism programming environment.

Online help and documentation are described in more detail in Chapter 7.

Customizing the Prism Programming Environment

You can change many aspects of the way the Prism environment operates. You can create customized command buttons in the tearoff region of the main Prism window, create aliases for commands and variables, and change the Prism environment default settings. These customizations are discussed in Chapter 8.

Using the Prism Environment

This chapter describes, at an introductory level, various aspects of using the Prism environment. Succeeding chapters provide more detailed instruction on performing specific operations within the Prism environment.

The best way to learn how to use the Prism environment is to try it out for yourself as you read this chapter. The chapter is organized into the following sections:

- “Before Entering the Prism Environment” on page 9
- “Starting the Prism Environment” on page 11
- “Executing Commands Within the Prism Environment” on page 24
- “Using the Menu Bar” on page 27
- “Using the Source Window” on page 28
- “Using the Line-Number Region” on page 31
- “Using the Command Window” on page 32
- “Writing Expressions in the Prism Environment” on page 37
- “Using Fortran 90 Generic Procedures” on page 42
- “Issuing Solaris Commands” on page 44
- “Killing Processes Within the Prism Environment” on page 45
- “Leaving the Prism Environment” on page 46

Before Entering the Prism Environment

This section describes the programming conditions under which you can make use of the Prism environment’s features.

Supported Languages and Compilers

You can work on Sun Fortran (77, 90), C, and C++ programs within the Prism environment. However, support for debugging Fortran 90 and C++ programs is limited. These limitations are described in Appendix B.

The Prism environment supports the following compilers from Forte Developer 6 update 2 and the Sun ONE Studio 7, Compiler Collection:

- Fortran 77
- Fortran 90
- C
- C++

Compiling and Linking Your Program

To use the Prism environment's debugging features, compile and link each program module with the `-g` compiler option to produce the necessary debugging information.

Note – The `-g` option overrides certain optimizations. For example, in C++ the `-g` option turns on debugging and turns off inlining of functions. The `-g0` (zero) option, on the other hand, turns on debugging, but does not affect inlining of functions. You cannot debug inline functions with this option. For Fortran 77 codes, the `-g` option conflicts with the `-auto-inlining` and `-depend` options.

Combining Compiler Options

If you compile programs with both the debugging option `-g` and an optimization option, such as `-xO[1, 2, 3, 4, 5]`, the combined options change the behavior of several Prism commands. This topic is covered in “Combining Debug and Optimization Options” on page 120.

Setting Up Your Working Environment

To enter the Prism environment, you must be logged in to a terminal or workstation that is running OpenWindows or the Common Desktop Environment (CDE).

DISPLAY Variable

Make certain that your `DISPLAY` environment variable is set for the terminal or workstation from which you are running OpenWindows or CDE. For example, if your workstation is named `valhalla`, you can issue this command (C shell example):

```
% setenv DISPLAY valhalla:0
```

Starting the Prism Environment

You can launch the Prism programming environment in four different ways:

- “Starting in MP Program Mode” on page 11
- “Attaching to a Running MP Job” on page 15
- “Starting in Serial Programming Mode” on page 16
- “Attaching to a Running Serial Program” on page 16

Starting in MP Program Mode

When you start the Prism environment in MP program mode, you are essentially starting it along with a parallel processing environment. That parallel processing environment uses the number of processes, resource manager, and any other characteristics specified by the `mprun(1)` command. The syntax is:

```
mprun mprun-options prism prism-options
```

For example:

```
% mprun -np 4 -x lsf prism
```

This example starts a multiprocessing environment with four processes under the control of the LSF resource manager, a host prism process, and four node prism processes, one for each of the four processes reserved by `mprun(1)`. Valid *prism-options* for starting in MP program mode are:

```
prism-options --> [ -C | -CX ]  
[ Xoption ]  
[ < infile | > outfile ]  
[ -install | -threads | -nothreads ]  
[ program ]
```

If you include the name of a program on the `prism` command line, the Prism environment automatically loads the program when it completes the start process.

If you don't specify a program on the command line, the Prism environment starts without loading an MPI program. You can, however, load a program for debugging from within the Prism environment. See "Loading a Multiprocess Program From Within the Prism Environment" on page 17 for more information.

You can specify other arguments with the program name to control various aspects of how the Prism environment deals with the program once it is loaded. Examples of this are provided in the following subsections.

Note – These descriptions all apply to launching and using the Prism environment in MP mode with its graphical user interface (GUI) enabled. See Appendix A for instructions on launching and using Prism's commands-only mode. See Appendix C for instructions on launching and using Prism's scalar mode.

The number of processes you specify with the *mprun*-options have special meaning to the Prism environment. The following subsections illustrate the effect, including the special cases of `-np 0` and `-np 1`:

- "Specifying Multiple Processes" on page 12
- "Specifying One Process per Available CPU" on page 13
- "Specifying a Nonthreaded View of Programs" on page 14
- "Specifying a Threaded View for a Single-Process Program" on page 15

Specifying Multiple Processes

When you want the Prism environment to debug a message-passing program with multiple processes distributed across multiple CPUs and nodes, simply use the `mprun` command's `-np` option to specify the number of processes that you want to run. When the program starts, the resource manager determines how the processes should be mapped to the available resources.

For example, the following command starts the Prism environment, loads the program `a.out` and, when the program is executed, causes eight instances of the program to be run under the control of the LSF resource manager.

```
% mprun -np 8 -x lsf prism
```

The Prism environment creates a separate debug process for each MP process and uses these Prism processes to collect information about the MP processes and, if threads are present, about the threads.

It also creates a single administrative process that communicates with the Prism debug processes and provides the interface to the user. This main Prism process is referred to as *Host Prism* or the *host process*.

Each Prism debug process runs on the same node as the MP process to which it is attached.

If a program spawns additional MP processes during execution, the Prism environment creates a new debug session to support the newly spawned MP processes. This secondary debug session creates a set of secondary debug processes, one per spawned MP process. An MPI program can dynamically spawn MP processes via calls to the library functions `MPI_Comm_spawn()` or `MPI_Comm_spawn_multiple()`. See “Enabling Support for Spawned MPI Processes” on page 18 for information about debugging spawned MP processes.

Specifying One Process per Available CPU

If you want to have one MP process running on each available CPU in a cluster, specify 0 for *numprocs* in the arguments to `mprun`.

```
% mprun -np 0 -x lsf prism a.out
```

Note – You can also specify a single process to be run per node, regardless of how many CPUs the nodes contain. To do this, however, you need to pass additional runtime environment values to the program, which is done either through the arguments to `mprun` or the arguments to the resource manager under which the program runs. This topic is discussed, with examples, in “Specifying Runtime Environment Options” on page 23.

Specifying a Threaded View of Programs

To debug multiple threads, the Prism environment must be in MP mode with its multithread support features enabled.

The Prism environment automatically adopts a multithreaded view of a loaded program if that program has been linked to the `libmpi_mt` library or the `libthread` library.

However, when loading a threaded program that has *not* been linked to the `libmpi_mt` library, use the `-threads` argument to direct the Prism environment to view the program as threaded. Use of this option is illustrated below.

▼ To Specify a Threaded View of a Program

- **Type:**

```
% mprun mprun-options [ rm-spec ] prism -threads program
```

In the following example, assume that `a.out` contains threads but was compiled without a link to `libmpi_mt`.

```
% mprun -np 4 -x lsf prism -threads a.out
```

The program `a.out` is executed as four processes and the Prism environment creates a separate debug process to handle each MP process. For each MPI process that is multithreaded, the corresponding debug process gathers debug information about each thread as well.

Note – If you specify the `-threads` option for a program that does not include threads, the Prism environment ignores the option.

Specifying a Nonthreaded View of Programs

If you load a program that has been linked to the `libmpi_mt` library, but want to view only the main thread, specify the `-nothreads` option. This disables the Prism environment's support for multiple threads so that it views only the main stream of execution in that program.

▼ To Specify a Nonthreaded View of a Program

- **Type:**

```
% mprun mprun-options [ rm-spec ] prism -nothreads program
```

In the following example, assume that `a.out` was compiled with a link to `libmpi_mt`.

```
% mprun -np 4 -x lsf prism -nothreads a.out
```

This starts the Prism environment in MP mode with multithread support disabled. The program `a.out` will be executed as four processes, but the Prism debug processes will be aware only of the primary thread in each MPI process.

Note – When the Prism environment opens a nonthreaded view of a program that uses threads, it issues a warning that thread debugging has been disabled. The Prism environment issues this warning for all programs linked with `libmpi`.

Specifying a Threaded View for a Single-Process Program

The MP mode is required for debugging multiple threads, even in a serial (single-process) program. This means you must specify the `-np` option of the `mprun` command to enable the MP mode, but with a `numprocs` value of 1.

▼ To Specify a Threaded View of a Single-Process Program

- **Type:**

```
% mprun -np 1 [ rm-spec ] prism -threads program
```

In the following example, the Prism environment starts in MP mode with support for debugging multiple threads enabled. The program `a.out` is executed as a single process.

```
% mprun -np 1 -x lsf -threads a.out
```

Running MP Prism in the Background

If you attempt to run the message-passing Prism environment in the background without specifying the `-n` flag to `mprun`, the Prism process hangs.

Once hung, you can either kill the Prism process manually from another terminal window using the `mpkill` command or put the Prism process in the foreground in your shell with the `fg` command, thereby releasing it from the freeze.

To run an MP Prism process in the background, add the `-n` flag to the `mprun` command line.

For example,

```
% mprun -n -np 4 prism &
```

Attaching to a Running MP Job

You can attach the Prism environment to an MP program that is already running. One instance of the Prism environment is created for each process in the program, plus one for the host. All you need to provide is the job name or job names to which you want the Prism environment attached. The syntax for attaching Prism to a running parallel program is:

```
% prism [ program | -jid | -jid-list ] prism-options
```

Valid *prism-options* for attaching to a running MP program are:

```

prism-options --> [ program-name | - jid |- jid-list ]
[ -C | -CX ]
[ xoption ]
[ < infile | > outfile ]
[ - install | -threads | -nothreads ]

```

The Prism environment loads the jobs you specify into the Prism environment. The processes are interrupted, and you can then work with the program in the Prism environment as you normally would. For more information, see “Attaching to a Running Message-Passing Process” on page 53.

Starting in Serial Programming Mode

You can start the Prism environment for use with serial (scalar) programs, by using this syntax:

```
% prism program-name prism-options
```

Valid *prism-options* for starting in serial program mode are:

```

prism-options --> [ -C | -CX ]
[ xoption ]
[ < infile | > outfile ]
[ - install | -threads | -nothreads ]

```

Attaching to a Running Serial Program

You can attach the Prism environment to a running serial program by using this syntax:

```
% prism [ program | - pid ] prism-options
```

Valid *prism-options* for starting in serial program mode are:

```

prism-options --> [ -C | -CX ]
[ xoption | core-file ]
[ < infile | > outfile ]
[ - install | -threads | -nothreads ]

```

When attaching to a running serial program in this manner, the Prism environment must be started on the same node on which the process is running.

After Starting the Prism Environment

There are several things you should know after starting the Prism environment:

- “Loading a Multiprocess Program From Within the Prism Environment” on page 17
- “After Starting the Prism Environment” on page 17
- “The Prism Prompt” on page 18
- “Enabling Support for Spawned MPI Processes” on page 18
- “Specifying Which Spawned Executables to Debug” on page 19
- “Attaching to a Job or Process” on page 22
- “Associating a Program or Process With a Core File” on page 22
- “Specifying X Toolkit Options” on page 22
- “Specifying Input and Output Files” on page 23
- “Specifying Runtime Environment Options” on page 23
- “Passing Command-Line Options to Secondary Sessions” on page 23

Loading a Multiprocess Program From Within the Prism Environment

After the Prism environment is started, you can load programs in either of the following ways:

- Use the Open selection from the File menu on the menu bar. This procedure is explained in “To Load a Program From the Menu Bar” on page 50.
- Specify the `load` command, together with the name of the program, in the Prism command window. This procedure is explained in “To Load a Program From the Command Window” on page 51.

After Loading a Program

Once a program is successfully loaded, the following conditions are in effect:

- The program’s name appears in the Program field in the main window.
- The source file containing the program’s main function appears in the source window.
- If you loaded the program using the Open menu option, the Open dialog box disappears.
- The status region displays the message `not started`.

You can now issue commands to execute and debug the program.

Note – If the program’s source file can’t be found—either when trying to load at Prism startup or when trying to load from within the Prism environment—the Prism command window displays a warning message. You can instruct the Prism environment to search for the file in other directories with the Use selection on the File menu. See “Creating a Directory List for Source Files” on page 90 for details.

The Prism Prompt

In the MP mode of the Prism environment, the Prism prompt includes the current pset, such as `(prism all)`.

Enabling Support for Spawned MPI Processes

If the program to be debugged includes calls to either `MPI_Comm_spawn()` or `MPI_Comm_spawn_multiple()`, you need to set the `follow_spawn` Prism environment variable to `on` before the program reaches the first spawn function call. Otherwise, the Prism environment cannot create a new debug session to handle the spawned processes. You can set this environment variable by entering the following in the Prism command window.

```
(prism all) set follow_spawn = on
```

When `follow_spawn` is set to `on` and a spawn function call occurs, the Prism environment creates one or more secondary debug sessions, depending on the following circumstances:

- A call to `MPI_Comm_spawn()` spawns processes for a single executable with a single set of arguments. Consequently, the Prism environment needs to create only one secondary debug session to manage the processes associated with the spawned executable.
- A call to `MPI_Comm_spawn_multiple()` can spawn processes for multiple executables or for a single executable with multiple sets of arguments. For these function calls, the Prism environment creates a separate secondary debug session for each distinct spawned executable or set of arguments. Each of the secondary Prism debug sessions includes enough secondary Prism processes to attach one to each spawned process in that debug session.

Like the primary Prism processes, each secondary Prism process is located on the same node as the spawned MPI process to which it is attached.

Some Prism commands that are issued in a primary Prism debug session affects secondary Prism sessions. Others do not. For information about which commands cross primary/secondary Prism session boundaries, see “Prism Commands With Special Functions in Spawned Sessions” on page 123.

Also, secondary Prism sessions may not acquire every Prism option that is specified in the primary Prism session that spawned them. For a list of Prism options and their acquisition by secondary sessions, see “Passing Command-Line Options to Secondary Sessions” on page 23.

To disable Prism support for debugging dynamically spawned processes, set the `follow_spawn` environment variable to `off`.

Note – If an X terminal window is created for a secondary Prism session, it identifies the job it belongs to in its title bar. The job’s identity is in the format “*aout:jid*”, where *aout* is the program name and *jid* is the ID of the job. See the *Sun HPC ClusterTools User’s Guide* for information about job IDs.

Specifying Which Spawned Executables to Debug

Since the Prism environment creates a separate debugging session for each executable spawned by a call to `MPI_Comm_spawn_multiple()`, aggressive use of this function could lead to hundreds of Prism debug windows being opened. This could seriously complicate the task of tracking all the processes being debugged.

You can avoid this problem by limiting the creation of secondary debug sessions to a subset of the spawned executables. To do this, set the `debug_spawn_aout` variable, listing the names of the spawned executables to be debugged. The syntax for setting the `debug_spawn_aout` variable is

```
(prism all) set debug_spawn_aout = "aout_list"
```

For example, if a call to `MPI_Comm_spawn_multiple()` in the program `alpha` spawns the executables `alpha`, `beta`, and `gamma`, you can limit the creation of secondary debug sessions to `beta` and `gamma` by entering the following in the Prism command line:

```
(prism all) set debug_spawn_aout = "beta gamma"
```

The primary Prism session continues to debug `alpha` processes and secondary Prism sessions debug `beta` and `gamma`. The executable `alpha` is still spawned, but the Prism primary debug session does not create a secondary debug session for it.

If you expect to use the same *about_list* in multiple successive debug sessions, you can store and reuse it automatically by storing it in the optional Prism initialization file `.prisminit`. When this file exists, the Prism environment automatically executes the commands contained in the file when it starts up.

For example, if you add the following lines to the `.prisminit` file:

```
set follow_spawn=on
set debug_spawn_about="beta gamma"
```

the next time the Prism environment starts up, it will have spawn debugging enabled, but restricted to the spawned executables `beta` and `gamma`.

When debugging multiple sets of executables that have been created by calls to `MPI_Comm_spawn_multiple()`, the Prism process numbers might not coincide with the MPI ranks of the processes. In FIGURE 2-1 alpha first spawns beta. Then beta spawns further instances of beta and gamma. The rank of each instance of alpha, beta, and gamma in their respective `MPI_COMM_WORLD` is zero. The processes in each instance number from 0 (zero) to $n-1$. However, in the Prism session debugging beta, there are three instances of beta. Their processes, in the Prism session, number from 0 (zero) to $3n-1$.

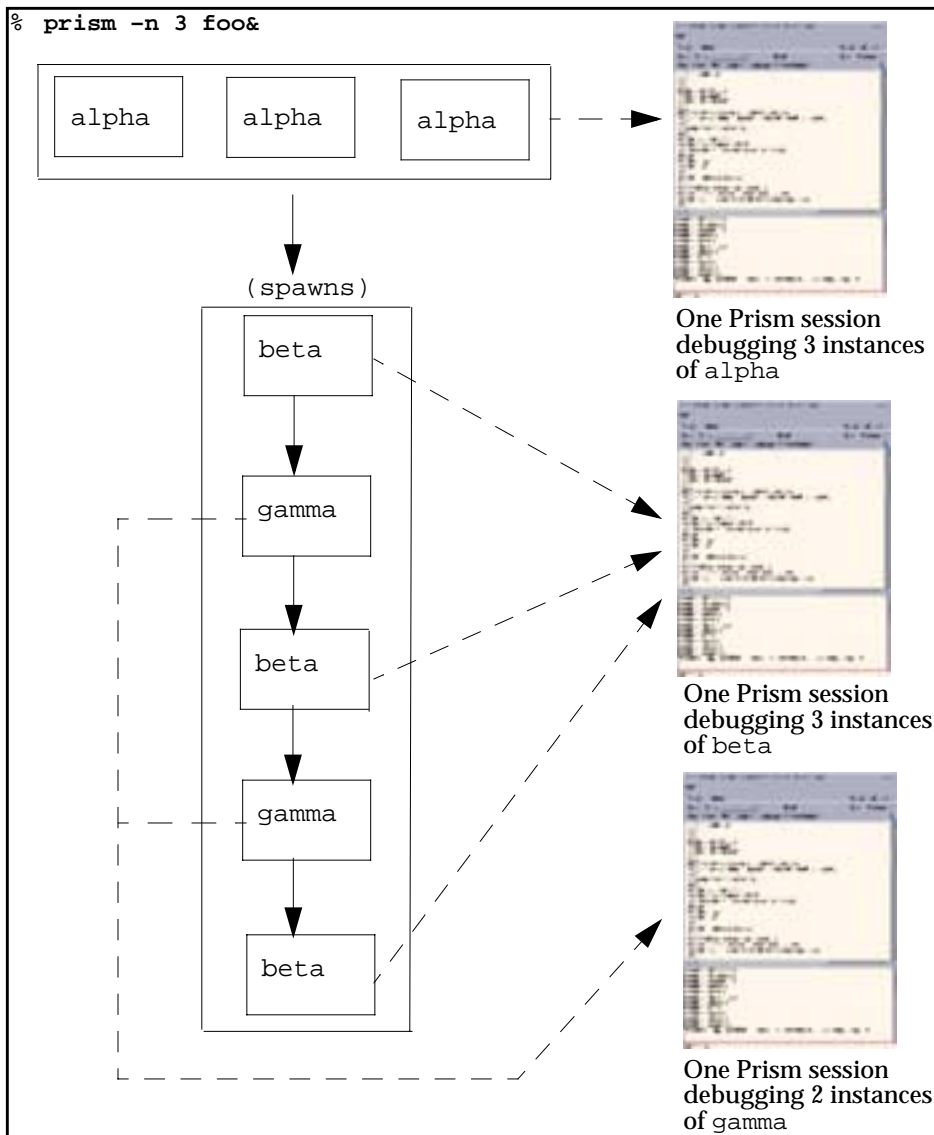


FIGURE 2-1 Prism Sessions Created by Calls to `MPI_Comm_spawn_multiple`

For information about using the Prism environment to debug multiple sessions, see “Debugging Spawned Sun MPI Processes” on page 122.

Attaching to a Job or Process

You can attach to a message-passing job or process that is currently running. If attaching to a process, the Prism environment must run on the same system on which the process is running. If attaching to a job, the Prism environment must be running on the same cluster as the job.

See “Attaching to a Running Message-Passing Process” on page 53 for more information about attaching to and detaching from a running serial process or a message-passing job.

Associating a Program or Process With a Core File

You can associate a core file with a single-process program or a particular process of a multiprocess program.

▼ To Associate a Program or Process With a Core File at Startup

- **Type:**

```
% mprun mprun-options prism program corefile
```

or (if you have launched the Prism environment and loaded a single-process program),

```
(prism all) core corefile
```

See “Associating a Core File With a Loaded Program” on page 52 for more information about core files.

Specifying X Toolkit Options

You can include most standard X toolkit command-line options when you issue the `prism` command; for example, you can use the `-geometry` option to change the size of the main Prism window. See your X documentation for information on these options. Also, note these limitations:

- The `-font`, `-title`, and `-rv` options have no effect.
- The `-bg` option is overridden in part by the setting of the `Prism.textBgColor` resource, which specifies the background color for text in the Prism environment; see “Changing Colors” on page 213.

X toolkit options are ignored if you use `-C` to run the Prism environment with the commands-only interface.

Specifying Input and Output Files

▼ To Specify an Input File

- **Type:**

```
% mprun mprun-options prism < input-file
```

This specifies a file from which the Prism environment is to read and execute commands upon startup.

Note – When the `TERM` environment variable is not set before starting the Prism environment, the Prism environment does not echo input when commands are read from a file.

▼ To Specify an Output File

- **Type:**

```
% mprun mprun-options prism > log-file
```

This specifies a file to which the Prism environment commands and their output are to be logged.

If you have created a `.prisminit` initialization file, the Prism environment automatically executes the commands in the file when it starts up. See “Initializing the Prism Environment” on page 199 for information on `.prisminit`.

Specifying Runtime Environment Options

You can specify runtime environment options either with the `mprun` command options or with the command options of the resource manager. See the *Sun HPC ClusterTools User Guide* for more information.

Passing Command-Line Options to Secondary Sessions

When debugging programs that make calls to `MPI_Comm_spawn()` or `MPI_Comm_spawn_multiple()`, the Prism environment creates special node Prism processes to debug the processes created by the spawn calls. These special node Prism processes are sometimes called *secondary* Prism sessions.

Secondary Prism sessions acquire some, but not all, options that you have set when you launch the primary Prism session. The acquisition status of Prism command-line options is described in TABLE 2-1.

TABLE 2-1 Passing Command Line Options to Secondary Sessions

Command Option Set in Primary Prism Session	Acquired by Secondary Prism Sessions
[-C -CX]	Yes
[-n -np]	No
[Xoption ...]	Yes
[core-file pid jid_list]	No
[< infile]	No
[> outfile]	No
[-install]	Yes
[-threads -nothreads]	Yes
[-]	No

Executing Commands Within the Prism Environment

The following three methods can be used for performing most Prism environment actions:

- “Using a Mouse” on page 24
- “Using Keyboard Alternatives to the Mouse” on page 25
- “Issuing Commands” on page 26

Using a Mouse

You can point and click with a mouse in the Prism environment to choose menu items and to perform actions within windows and dialog boxes.

In any window where you see this mouse icon:

you can left-click on the icon to obtain information about using the mouse in the window.



Note – The Prism environment assumes that you have a standard three-button mouse.

Using Keyboard Alternatives to the Mouse

You can use a keyboard to perform many of the same functions you can perform with a mouse. This section lists the supported keyboard alternatives.

In general, to use a keyboard alternative, the *focus* must be in the screen region where you want the action to take place. The focus is generally indicated by the *location cursor*, which is a heavy line around the region.

General keyboard alternatives to mouse control are listed in TABLE 2-2.

TABLE 2-2 General Keyboard Alternatives to Mouse Control

Key Name	Description
Tab	Use the Tab key to move the location cursor from field to field within a window or dialog box. Each button in a window or box constitutes one field. The location cursor highlights the relevant button when you tab to the field.
Shift-Tab	Use the Shift-Tab keys to perform the same function as Tab, but in the reverse direction.
Return	Use the Return key to choose a highlighted choice in a menu or to perform the action associated with a highlighted button in a window or dialog box.
Arrow keys	Use the up, down, left, and right arrow keys to move within a field. For example, when the location cursor highlights a list, you can use the up and down arrow keys to move through the choices in the list. In some windows that contain text, pressing the Control key along with an up or down arrow key scrolls the text one-half page.
F1	Use the F1 key instead of the Help button to obtain help about a window or dialog box.
F10	Use the F10 key to move the location cursor to the menu bar.

TABLE 2-2 General Keyboard Alternatives to Mouse Control (*Continued*)

Key Name	Description
Meta	Use the Meta key along with the underlined character in the desired menu item to display a menu or dialog box (equivalent to clicking on the item with the mouse). The Meta key has different names on different keyboards; on some it is the Left or Right key.
Control-C	Use the Control-C key combination to interrupt command execution.
Esc	Use the Esc key instead of the Close or Cancel button to close the window or dialog box in which the mouse pointer is currently located.

The keys and key combinations described in TABLE 2-3 work on the command line and in *text-entry boxes*—that is, fields in a dialog box or window where you can enter or edit text.

TABLE 2-3 Text-Entry Keyboard Alternatives

Key Name	Description
Back Space	Deletes the character to the left of the I-beam cursor.
Delete	Same as Backspace.
Control-A	Moves to the beginning of the line.
Control-B	Moves back one character.
Control-D	Deletes the character to the right of the I-beam cursor.
Control-E	Moves to the end of the line.
Control-F	Moves forward one character.
Control-K	Deletes to the end of the line.
Control-U	Deletes to the beginning of the line.

In addition, you can use *keyboard accelerators* to perform actions from the menu bar; see “Keyboard Accelerators for Main Menu Selections” on page 27.

Issuing Commands

You can issue commands in the Prism environment from the command line in the command window. Most commands duplicate functions you can perform from the menu bar. Some functions, however, are only available via commands. See the *Prism*

Reference Manual for complete information about Prism commands. See “Using the Command Window” on page 32 for instructions on entering commands in the command window.

Many commands have the same syntax and perform the same action in both the Prism environment and the Solaris debugger `dbx`. There are differences, however; you should check the reference description of a command before using it.

Using the Menu Bar

The menu bar is the line of titles across the top of the main window of the Prism environment.

Each title is associated with a pull-down menu from which you can perform actions within the Prism environment.

Keyboard Accelerators

A keyboard accelerator is a shortcut that lets you choose a frequently used menu item without displaying its pull-down menu. Keyboard accelerators consist of the Control key plus a function key; you press both at the same time to perform the action. The keyboard accelerator for a keyboard menu selection is displayed next to the name of the selection. If nothing is displayed, there is no accelerator for that selection.

The keyboard accelerators (on a Sun keyboard) are listed in TABLE 2-4.

TABLE 2-4 Keyboard Accelerators for Main Menu Selections

Accelerator	Function
Control-F1	Run
Control-F2	Continue
Control-F3	Interrupt
Control-F4	Step
Control-F5	Next
Control-F6	Where
Control-F7	Up
Control-F8	Down

Using the Source Window

The source window displays the source code for the executable program loaded into the Prism environment. (Chapter 3 describes how to load a program into the Prism environment and how to display the different source files that make up the program.) When you execute the program and execution then stops for any reason, the source window updates to show the code being executed at the stopping place. The Source File field at the top of the source window lists the name of the file displayed in the window.

You can resize the source window by dragging the small resize box at the lower right of the window. If you change the window's size, the new size is saved when you leave the Prism environment.

You cannot edit the source code displayed in the source window. This must be done within an editor, which can be called up from within the Prism environment. See Chapter 6 for instructions on editing programs.

Moving Through the Source Code

You can move through a source file displayed in the source window by using the scroll bar on the right side of the window. You can also use the up and down arrow keys to scroll a line at a time, or press the Control key along with the arrow key to move half a page at a time.

To return to the current execution point, type Control-X in the source window.

▼ To Search for Text in a String or Regular Expression

- **Type:**

`(prism all) /regexp`

or

`(prism all) ?regexp`

The */regexp* command searches forward in the file for the string (or regular expression) that you specify and repositions the file at the first occurrence it finds. The *?regexp* command searches for the string (or regular expression) in the reverse direction.

Selecting Text

You can select text in the source window by dragging over it with the mouse; the text is then highlighted. Or double-click with the mouse pointer on a word to select that word. Left-click anywhere in the source window to deselect text.

Right-click in the source window to display a menu that includes actions to perform on the selected text, as shown in FIGURE 2-2. For example, select Print to display a visualizer containing the value(s) of the selected variable or expression at the current point of execution. See Chapter 5 for a discussion of visualizers and printing. To close the pop-up menu, right-click anywhere else in the main window.

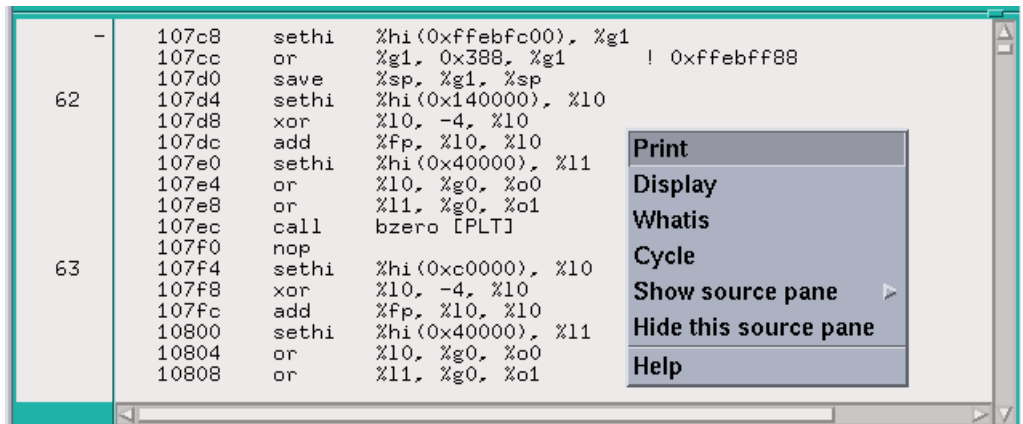


FIGURE 2-2 Pop-up Menu in Source Window

You can display the definition of a function by pressing the Shift key while selecting the name of the function in the source window. This is equivalent to choosing the Func selection from the File menu and selecting the name of the function from the list. When selecting the function name, highlight only the name, not arguments to the function.

Splitting the Source Window

You can split the source window to simultaneously display the source code and assembly code of the loaded program.

▼ To Split the Source Window

1. Load the program of interest.
2. Right-click in the source window to display the pop-up menu.

3. Click on the Show source pane selection in the pop-up menu.

This displays another menu.

4. Choose the Show .s source selection from the menu.

This causes the assembly code for your program to be displayed in the bottom pane of the window, as shown in FIGURE 2-3.

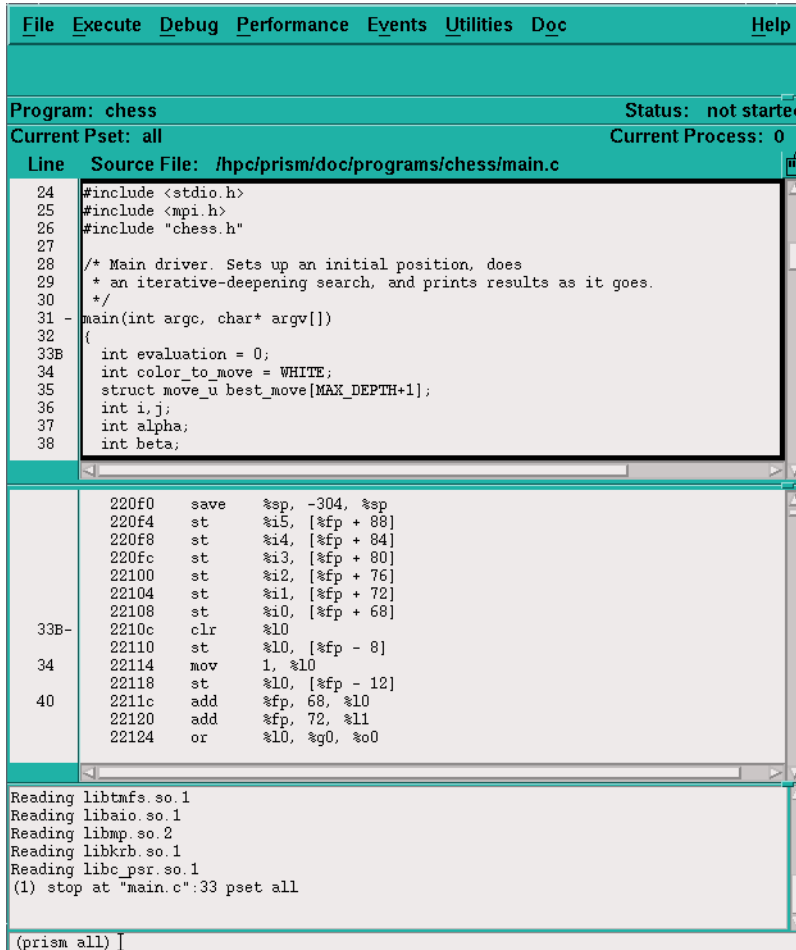


FIGURE 2-3 Split Source Window

▼ To Return to a Single Source Window

1. Right-click in the pane you want to close.

2. Choose Hide this source pane from the pop-up menu.

Using the Line-Number Region

The line-number region shows the line numbers associated with the source code displayed in the source window. FIGURE 2-4 shows a portion of a line-number region, with a breakpoint set.




FIGURE 2-4 Line-Number Region

You will see the following symbols in the line-number region:

- The > symbol in the line-number region in FIGURE 2-3 is the *execution pointer*. When a program is being executed, the execution pointer points to the next line to be executed for the most-active function call or to the call site for functions higher on the stack. If you move elsewhere in the source code, typing Control-X returns to the current execution point.
- A B displayed next to a line number indicates that all processes in the current pset have a breakpoint set at that line. A b is displayed next to a line number when some, but not all, processes in the current pset have a breakpoint set at that line. Methods for setting breakpoints are described in “Setting Breakpoints” on page 104.

Shift-click on the B or b in the line-number region to display the *event* associated with the breakpoint. See “Overview of Events” on page 93 for a discussion of events.

- A T is displayed next to a line number when all processes in the current pset have a tracepoint set at that line. A t is displayed next to a line number when some, but not all, processes in the current pset have a tracepoint set at that line. See “Tracing Program Execution” on page 109 for additional information.

Shift-click on \top or τ in the line-number region to display the *event* associated with the tracepoint. See “Overview of Events” on page 93 for a discussion of events.


- The `-` symbol is the *scope pointer*; it indicates the current source position (that is, the scope). The Prism environment uses the current source position to interpret names of variables. When you scroll through source code, the scope pointer moves to the middle line of the code that is displayed. Various Prism commands also change the position of the scope pointer.
- The `*` symbol is used when the current source position is the same as the current execution point; this happens whenever execution stops.

Note – If a breakpoint and tracepoint are both set at a given line, the breakpoint is indicated, but not the tracepoint.

Using the Command Window

The command window is the area at the bottom of the main window in which you type commands and receive Prism output.

The command window consists of two boxes: the command line at the bottom and the history region above it. FIGURE 2-5 shows a command window.



```
trace in procedure "main" at "julia.c":105
trace in procedure "main" at "julia.c":105
trace in procedure "main" at "julia.c":105
trace in procedure "main" at "julia.c":105
trace in procedure "main" at "julia.c":105
trace in procedure "main" at "julia.c":105
trace in procedure "main" at "julia.c":105
trace in procedure "main" at "julia.c":105
Program exited with 1
(prism) print ittr on ded
```

FIGURE 2-5 Command Window With History Region

You can resize the command window and scroll through it independently of the main window. If you do not intend to issue commands in the command window, you might want to make this window smaller so that you can display more code in the source window. If you use the command window frequently, you might want to make it bigger. If you change the size of the window, the new size is saved when you leave the Prism environment.

Using the Command Line

You can type in the command-line box whenever it is highlighted and an I-shaped cursor, called an *I-beam*, appears in it. See TABLE 2-3 for a list of keystrokes you can use in editing the command line. Press Return to issue the command. Type Control-C to interrupt execution of a command (or choose the Interrupt selection from the Execute menu).

You can issue multiple commands on the Prism command line, separating them with semicolons (;). There is one exception to this—you cannot follow a file name argument with a semicolon because the Prism environment parses it as part of the file name.

The Prism environment keeps the commands that you issue in a buffer. Type Control-P to display the previous command in this buffer. Type Control-N to display the next command in the buffer. You can then edit the command and issue it in the usual way.

During long-running commands (for example, when you have issued the `run` command to start a program executing), you may still be able to execute other commands. If you issue a command that requires that the current command complete execution, you receive a warning message and the Prism environment waits for the command to complete.

Using the History Region

Commands that you issue on the command line are echoed in the history region, above the command line. The Prism environment's response appears beneath the echoed command. The Prism environment also displays other messages in this area, as well as command output that you specify to go to the command window. Use the scroll bar at the right of this box to move through the display.

▼ To Specify the Maximum Number of Lines in the History Region

- **Type:**

```
(prism all) set $history = value
```

For example:

```
set $history = 2000
```

This limits the maximum number of lines in the history buffer to 2000. The default is 10,000.

The Prism environment uses memory in maintaining a large history region. A smaller history region, therefore, may improve performance and reduce the chances of the Prism environment running out of memory.

▼ To Select Text in the History Region

1. **Select text using one of these methods:**

- Double-click to select the word on which the mouse pointer is positioned.
- Triple-click to select the line on which the mouse pointer is located.
- Press the left mouse button and drag the mouse over the text to select it.

2. **Click the middle mouse button to paste the selected text into other text areas.**

▼ To Re-Execute a Command

1. **Triple-click on a line in the history region to select it.**

2. **Click the middle mouse button with the mouse pointer still in the history region.**

3. **Click the middle mouse button with the mouse pointer on the command line.**

The selected text appears on the command line but is not executed. This gives you an opportunity to edit the text before executing it.

Redirecting Output

▼ To Redirect Output to a File

- **Type:**

```
(prism all) where @ filename
```

For example,

```
(prism all) where @ where.output
```

This puts the output of the `where` command (a stack trace) into the file `where.output` in your current working directory within the Prism environment. This method works for most commands.

▼ To Redirect Output to a Window

- **Type:**

```
(prism all) where on window
```

This directs the output of the `where` command to the Prism environment window specified by the `window` argument. The argument `window` must be one of:

- `command` (abbreviated `com`) – This sends output to the command window; this is the default.
- `dedicated` (abbreviated `ded`) – This sends output to a window dedicated to output for this command. If you subsequently issue the same command and specify that output is to be sent to the dedicated window, this window will be updated. For example:

```
(prism all) list on ded
```

This displays the output of the `list` command in a dedicated window. Some commands that have equivalent menu selections display their output in the standard window for the menu selection.

- `snapshot` (abbreviated `sna`) – This creates a window that provides a snapshot of the output. If you subsequently issue the same command and specify that output is to be sent to the snapshot window, the Prism environment creates a separate window for the new output. The time each window was created is shown in its title. Snapshot windows let you save and compare outputs.
- `windowname` – This creates a window with a name you have created. *Windowname* appears in the title of the window. This is useful if you want a particular label for a window. For example, if you were doing a stack trace at line 22, you could issue this command:

```
(prism all) where on line 22
```

This labels the window with the location of the stack trace.

Note – The output of `edit`, `make`, and `sh` cannot be redirected. The output of `run` can, however, be redirected by `>` and other shell redirection conventions.

Logging Commands and Output

You can use the `log` command along with the `source` command to replay a session in the Prism environment. When replaying a Prism session, however, you must edit the log file to remove Prism output.

Use the log file for logging commands and output from within the Prism environment.

▼ To Specify the Name of a Log File.

- **Type:**

```
(prism all) log @ filename
```

The log file *filename* will be located in the current directory. This can be helpful in saving a record of a Prism session. For example,

```
(prism all) log @ prism.log
```

logs output to the file `prism.log`.

▼ To Append the Log to an Existing File

- **Type:**

```
(prism all) log @@ filename
```

▼ To Turn Off Logging

- **Type:**

```
(prism all) log off
```

Rerunning a Prism Session That Was Saved to a Log File

▼ To Rerun a Saved Prism Debug Session

- **Type:**

`(prism all) source filename`

where *filename* is the name of the log file.

Writing Expressions in the Prism Environment

While working in the Prism environment, there are circumstances in which you may want to write expressions that the Prism environment will evaluate. For example, you may want to print or display expressions or specify an expression as a condition under which an action is to take place. You can write these expressions in the language of the program you are working on. This section discusses additional aspects of writing expressions.

How the Prism Environment Chooses the Correct Variable or Procedure

Multiple variables and procedures can have the same name in a program. This can be a problem when you specify a variable or procedure in an expression. To determine which variable or procedure you mean, the Prism environment tries to resolve its name by using these rules:

1. Prism first tries to resolve the name using the scope of the current function. For example, if you use the name *x* and there is a variable named *x* in the current function or the current file, the Prism environment uses that *x*. The current function is ordinarily the function at the program's current stopping point, but you can change this, as described in "Choosing the Current File and Function" on page 88.
2. If this fails to resolve the name, the Prism environment goes up the call stack looking for the variable *x* in the caller of the current function, and then its caller, and so on, following the scoping and visibility rules of the current language.

3. If no match is found in any routine active on the stack, the Prism environment searches the static and global name spaces. If no match is found there, the Prism environment prints an error.
4. If the name is not found in the call stack, the Prism environment arbitrarily chooses one of the variables or procedures with the name in the source code. When the Prism environment prints out the information, it adds a message of the form “[using qualified name]”. Qualified names are discussed below.

Using Qualified Names

You can override the way that the Prism environment resolves names by *qualifying* the name.

A fully qualified name starts with a back-quotation mark, or back-quote, (```). The symbol farthest to the left in the name is the load object, followed optionally by the file, followed optionally by the procedure, followed by the variable name. Each element is preceded by a backquote (```). Examples of the Prism environment’s identifier syntax are shown in TABLE 2-5.

TABLE 2-5 Prism Identifier Syntax

Syntax	Description
<code>a</code>	Specifies the variable <code>a</code> in the current scope. An error is reported if no variable <code>a</code> exists in the current scope.
<code>`a</code>	Specifies the variable <code>a</code> in the global scope.
<code>``a</code>	Specifies the variable <code>a</code> in the global or file-static scope.
<code>`foo.c`a</code>	Specifies the variable <code>a</code> in file <code>foo.c</code> .
<code>`foo.c`foo`a</code>	Specifies the variable <code>a</code> in the procedure <code>foo</code> in the file <code>foo.c</code> .
<code>`foo`a</code>	Specifies the variable <code>a</code> in function <code>foo</code> (if <code>foo</code> is active).
<code>`a.out`foo.c`foo`a</code>	Specifies the variable <code>a</code> in function <code>foo</code> in file <code>foo.c</code> in load object <code>a.out</code> .

TABLE 2-5 Prism Identifier Syntax (Continued)

Syntax (Continued)	Description
<code>`a.out`foo.c`foo:line`a</code>	Specifies the variable <code>a</code> at line number <code>line</code> in function <code>foo</code> in file <code>foo.c</code> in load object <code>a.out</code> .
<code>`foo.x`foo.cc`Symbol::print:71`dummy</code>	Specifies the variable <code>dummy</code> at line number <code>71</code> in member function <code>print</code> of class <code>Symbol</code> in file <code>foo.cc</code> in load object <code>foo.x</code> .
<code>"foo.c":line</code>	Specifies the line number <code>line</code> in the file <code>foo.c</code> . Note the use of double quotes.

▼ To Display the Fully Qualified Name of a Variable

- **Type:**

`(prism all) which identifier`

This command displays the fully qualified name, as described below.

Partially qualified names do not begin with ```, but have a ``` in them. For example, `foo`a`

In this case, the Prism environment looks up the function name on the left first and picks the innermost symbol with that name that is visible from your current location. This is useful primarily in referring to variables in routines on the call stack.

Use the `whereis` command to display a list of all the fully qualified names that match the identifier you specify.

The Prism environment assigns its own names to variables in local blocks of C code. This approach disambiguates variable names, in case you reuse a variable name in more than one of these local blocks.

When debugging Fortran, the Prism environment attempts to be case-insensitive in interpreting names, but uses case to resolve ambiguities.

Using Fortran Intrinsic Functions in Expressions

The Prism environment supports the use of a subset of Fortran intrinsic functions in writing expressions; the intrinsics work for all languages that the Prism environment supports, except as noted below.

The intrinsics, along with the supported arguments, are:

- `ALL(logical array)` – Determines whether all elements are true in a logical array. This works for Fortran only.

- `ANY(logical array)` – Determines whether any elements are true in a logical array. This works for Fortran only.
- `CMPLX(numeric-arg, numeric-arg)` – Converts the arguments to a complex number. If the intrinsic is applied to Fortran variables, the second argument must not be of type `complex` or `double` (double-precision complex).
- `COUNT(logical array)` – Counts the number of true elements in a logical array. This works for Fortran only.
- `ILEN(I)` – Returns one less than the length, in bits, of the two’s-complement representation of an integer. If *I* is nonnegative, `ILEN(I)` has the value $\log_2(I + 1)$; if *I* is negative, `ILEN(I)` has the value $\log_2(-I)$.
- `IMAG(complex number)` – Returns the imaginary part of a complex number. This works for Fortran only.
- `MAXVAL(array)` – Computes the maximum value of all elements of a numeric array.
- `MINVAL(array)` – Computes the minimum value of all elements of a numeric array.
- `PRODUCT(array)` – Computes the product of all elements of a numeric array.
- `RANK(scalar or array)` – Returns the rank of the array or scalar.
- `REAL(numeric argument)` – Converts an argument to `real` type. This works for Fortran only.
- `SIZE(array)` – Counts the total number of elements in the array.
- `SUM(array)` – Computes the sum of all elements of a numeric array.

The intrinsics can be either uppercase or lowercase.

Using C Arrays in Expressions

The Prism environment handles arrays slightly differently from the way C handles them.

For example, in a C program, you might make the following declaration and use `a` in an expression:

```
int a[10];
```

The type of `a` converts from “array of `ints`” to “pointer to `int`”. Following the rules of C, therefore, a Prism command like this one should print a hexadecimal pointer value:

```
(prism all) print a + 2
```

Instead, it prints two more than each element of `a` (that is, `a[0] + 2`, `a[1] + 2`, etc.). This enables you to do array operations and use visualizers on C arrays in the Prism environment. The `print` command and visualizers are discussed in Chapter 5.

To get the C behavior, issue the command as follows:

```
(prism all) print &a + 2
```

Using Array-Section Syntax in C Arrays

You can use Fortran 90 array-section syntax when specifying C arrays. This syntax is useful, for example, if you want to print the values of only a subset of the elements of an array. The syntax is:

```
(lower-bound: upper-bound: stride)
```

where

- *lower-bound* – Specifies the lowest-numbered element you choose along a dimension; it defaults to 0.
- *upper-bound* – Specifies the highest-numbered element you choose along the dimension; it defaults to the highest-numbered element for the dimension.
- *stride* – Specifies the increment by which elements are chosen between the lower bound and upper bound; it defaults to 1.

You must enclose the values in parentheses (rather than brackets), as in Fortran. If your array is multidimensional, you must separate the dimension specifications with commas within the parentheses, again as in Fortran.

For example, if you have the following array:

```
int a[10][20];
```

you can issue the following command to print the values of elements 2–4 of the first dimension and 2–10 of the second dimension:

```
(prism all) print a(2:4,2:10)
```

Hints for Detecting NaNs and Infinities

The Prism environment provides expressions that you can use to detect NaNs (values that are “not a number”) and infinities in your data. These expressions derive from the way NaNs and infinities are defined in the IEEE standard for floating-point arithmetic.

▼ To Find Out if x Is a NaN

- **Use the expression:**

```
(x .ne. x)
```

For example, if x is an array, issue this command to print only the elements of x that are NaNs:

```
(prism all) where (x .ne. x) print x
```

The `print` command is discussed in Chapter 5.

Also, note that if there are NaNs in an array, the mean of the values in the array will be a NaN. The mean is available via the Statistics selection in the Options menu of a visualizer—see Chapter 5 for details.

▼ To Find Out if x Is an Infinity

- **Type:**

```
(prism all) (x * 0.0 .ne. 0.0)
```

Using Fortran 90 Generic Procedures

You can use Fortran 90 generic procedures in any Prism command or dialog box that asks for a procedure. If you do so, the Prism environment prompts you for the name(s) of the specific procedure(s) you want to use.

For example, you use the syntax `stop in procedure` to set a breakpoint in a procedure. If you use this syntax for a generic procedure in the Prism graphical interface, a dialog box like the one shown in FIGURE 2-6 is displayed.

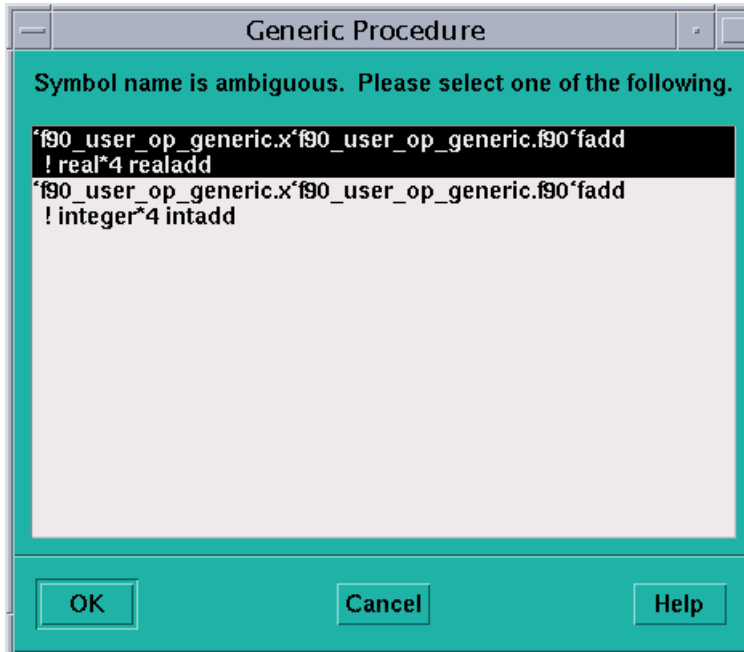


FIGURE 2-6 Generic Procedure Dialog Box

The commands-only interface of the Prism environment prompts you in a similar fashion, as shown below:

```
(prism all) stop in fadd
More than one identifier 'fadd'.
Select one of the following names:
0) Cancel
1) `f90_user_op_generic.x'f90_user_op_generic.f90 `fadd
! real*4 realadd
2) `f90_user_op_generic.x'f90_user_op_generic.f90 `fadd
! integer*4 intadd
>
```

If you choose 0 or press Return, the command is cancelled. If you choose another number, the Prism environment sets the breakpoint(s) in the specified procedure(s). For example, the following sets the breakpoint in the `fadd` procedure that uses `real*4 data`.

```
Select one of the following names:
0) Cancel
1) `f90_user_op_generic.x`f90_user_op_generic.f90`fadd
! real*4 realadd
2) `f90_user_op_generic.x`f90_user_op_generic.f90`fadd
! integer*4 intadd
> 1
(1) stop in fadd
(prism)
```

Issuing Solaris Commands

You can issue Solaris commands from within the Prism environment.

▼ To Issue Solaris Commands From Within the Prism Environment

- **Perform one of the following operations:**
 - From the menu bar – Choose the Shell selection from the Utilities menu. The Prism environment creates a Solaris shell that is independent of the Prism environment. You can issue Solaris commands from it just as you would from any Solaris shell. The type of shell that is created depends on the setting of your `SHELL` environment variable.
 - From the command window – Issue the `sh` command on the command line. With no arguments, it creates a Solaris shell. If you issue the `sh` command with arguments, the command is executed and the results are displayed in the history region. The `sh` command does not result in the creation of a new shell.

Note – For instructions on starting the default editor from the Prism command window, see “To Start the Default Editor on the Current Source File” on page 191.

Some Solaris commands have equivalents in the Prism environment, as described below.

Changing the Current Working Directory

Use the standard Solaris shell commands, such as `pwd`, `cd`, and `ls`, to manage your current working directory. By default, your current working directory within the Prism environment is the directory from which you started the Prism environment.

The Prism environment interprets all relative file names with respect to the current working directory. The Prism environment also uses the current working directory to determine which files to show in file-selection dialog boxes.

Setting and Displaying Environment Variables

You can set, unset, and display the settings of environment variables from within the Prism environment, just as you do in the Solaris environment.

Killing Processes Within the Prism Environment

▼ To Kill a Process or Job Running Within the Prism Environment

- **Type:**

```
(prism all) kill
```

▼ To Kill a Spawned Prism Session

- **Type:**

```
(prism all) kill
```

Issuing a `kill` command in the primary Prism session also kills all of the secondary Prism sessions.

Leaving the Prism Environment

▼ To Exit a Single-Job Prism Session

1. Perform one of the following:

- From the menu bar — Choose the Exit selection from the File menu. You will be asked if you are sure you want to exit. Click on OK if you're sure; otherwise, click on Cancel or press the Esc key to stay in the Prism environment.
- From the command window — Type the `quit` command on the command line. You will not be asked if you are sure you want to quit.

If you have created subprocesses, such as Solaris shells, while in the Prism environment, the Prism environment displays a message before exiting. For example, see FIGURE 2-7:

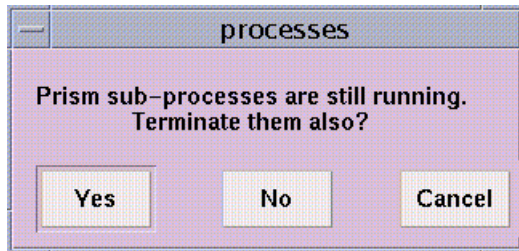


FIGURE 2-7 Subprocess Warning

2. Take one of the following actions:

- Choose Yes (the default) to leave the Prism environment and terminate the subprocesses.
- Choose No to leave the Prism environment without terminating the subprocesses.
- Choose Cancel to stay in the Prism environment.

▼ To Quit a Spawned Prism Session

- **Type**

```
(prism all) quit -all
```

Issuing the `quit` command without any options quits only the session in which it is invoked.

To quit all Prism sessions—in primary and secondary sessions—use `quit -all` in the primary Prism session.

Note – The `-all` option is valid only in the primary session.

If you brought the job into the Prism environment using the `attach` command, `quit -all` leaves the job's processes running and closes both the primary and secondary Prism sessions.

The Quit selection on the Prism File menu has the same effect as the `quit` command without the `-all` argument.

Loading and Executing a Program

This chapter describes how to load and run programs within the Prism environment. It covers the following topics:

- “Loading a Program” on page 49
- “Associating a Core File With a Loaded Program” on page 52
- “Attaching to a Running Message-Passing Process” on page 53
- “Detaching From a Running Process” on page 54
- “Executing a Program in the Prism Environment” on page 55
- “Using Psets in the Prism Environment” on page 61
- “Using Psets in Commands” on page 82
- “Using the Prism Environment With Sun MPI Client/Server Programs” on page 87
- “Choosing the Current File and Function” on page 88
- “Creating a Directory List for Source Files” on page 90

To use this chapter most effectively, you should already have an executable program that you want to run within the Prism environment. You can also develop a new program by calling up an editor within the Prism environment; see Chapter 6.

Loading a Program

Before you can execute or debug a program in the Prism environment, you must first load the program into the Prism environment. You can load only one program at a time.

As described in Chapter 2, you can load a program into the Prism environment by specifying its name as an argument to the `prism` command. If you don't use this method, you can load a program once you are in the Prism environment by using one of the methods discussed next.

▼ To Load a Program From the Menu Bar

1. Choose Open from the File menu or the tear-off region.

A dialog box like the one shown in FIGURE 3-1 appears:

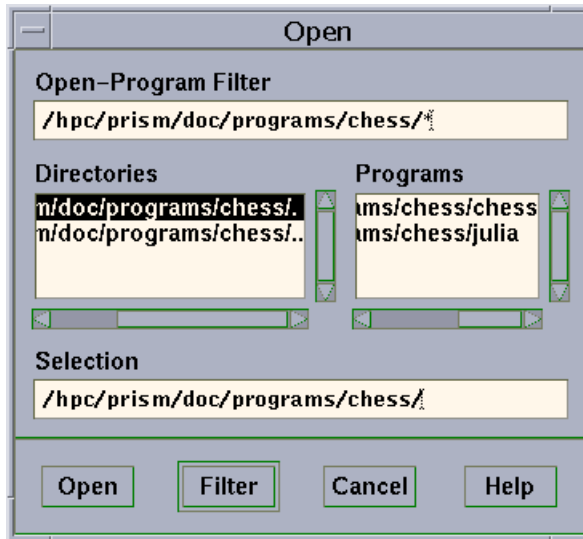


FIGURE 3-1 Open Program Filter

2. Double-click on the program name in the Programs scrollable list (if the name is listed).

Or, you can put its path name into the Selection box, then click on Open. To put the file's path name into the Selection box, you can either type it directly in the box or click on its name in the Programs list. The Programs list contains the executable programs in your current working directory.

Use the Open-Program Filter dialog box to control the display of file names in the Programs list; the box uses standard Solaris filters. For example, you can click on a directory in the Directories list if you want to change to that directory. But the Programs list does not update automatically to show the programs in the new directory. Instead, the filter changes to *directory-name**, indicating that all files in *directory-name* are to be displayed. Click on Filter to display the file names of the programs. Or simply double-click on the directory name in the Directories list to display the programs in the directory.

If you want to use a different filter, you can edit the Open-Program Filter box directly. For example, change it to *directory-name/prog** to display only programs beginning with *prog*.

3. Click on Cancel or press the Esc key if you decide not to load a program.

▼ To Load a Program From the Command Window

- **Type:**

`(prism all) load program`

Use the name of the executable program as its argument. For example, to load the program `myprogram`, enter

`(prism all) load myprogram`

What Happens When You Load a Program

Once a program has been successfully loaded, the following occurs:

- The program's name appears in the Program field in the main window.
- The source file containing the program's main function appears in the source window.
- The Load dialog box disappears (if it has been displayed).
- The status region displays the message `not started`.

You can now issue commands to execute and debug this program.

If the Prism environment cannot find the source file, it displays a warning message in the command window. Select Use from the File menu to specify other directories in which the Prism environment is to search; see “Creating a Directory List for Source Files” on page 90 for details.

▼ To Load Subsequent Programs

- **To load a new program, perform one of the following:**

- If you have a program already loaded and you want to switch to a new program, simply load the new program; the previously loaded program is automatically unloaded.
- If you want to start fresh with the current program, issue the `reload` command with no arguments. The currently loaded program is reloaded into the Prism environment.

Associating a Core File With a Loaded Program

You can have the Prism environment associate a core file with a program by specifying its file name after the name of the program on the `prism` command line.

▼ To Associate a Core File With a Loaded Program

- **Type:**

`(prism all) core corefile`

Where *corefile* is the name of the corresponding core file.

The Prism environment's `core` command is not available when using the Prism environment with message-passing programs. Instead, you must specify the name of the process core file from the Prism command line.

In either case, when a program failure results in a core dump, the Prism environment reports the error that caused the core dump and loads the program with a *stopped* status at the location where the error occurred. You can then work with the program within the Prism environment. You can, for example, examine the stack and print the values of variables. You cannot, however, continue execution from the current location.

▼ To Examine the Core File of a Local Process

You can examine a core file created by a message-passing program from within the Prism environment. However, because only one file will be examined, the Prism environment will be started in scalar mode. The procedure for examining a core file for a local process is as follows:

1. **Type:**

`% prism program corefile`

2. **Type:**

`(prism all) where`

This produces a stack trace.

3. Type:

(prism all) **print variable**

This lets you inspect the state of your process at the time the core dump was taken.

Note the following restrictions to this procedure:

- Because the Prism environment is started in scalar mode, you cannot use process sets (*psets*: pronounced “pea-sets”) or other features of MP Prism mode.
- You cannot issue any commands that require execution, such as `run`, `cont`, or `step`.
- You cannot change the values of variables with the `assign` command.
- You cannot use the `core` command to examine a core file once you have started the Prism environment in MP Prism mode. If multiple processes dumped core, the resulting core file may be overwritten and therefore invalid.

Attaching to a Running Message-Passing Process

You can attach the Prism environment to a MP process that has already started.

▼ To Attach to a Running Message-Passing Process

1. Obtain the job ID of the processes.

- If you are using the LSF environment, issue the `bjobs` command. You can also get the job ID from the `bsub` command when it starts the job.
- If you are using the CRE environment, issue the `mpps` command. You can also get the job ID from the `mprun` command when it starts the job.

The following is an example of `bjobs` output (LSF environment):

```
host4-0 54 =>bjobs
JOBID USER STAT QUEUE FROM_HOST EXEC_HOST JOB_NAME SUBMIT_TIME
15232 jay RUN hpc host4-0 host4-0 chess Sep 24 13:35
host4-1
```

2. Type:

% **prism -options [program | -] job_ID**

Note that *job_ID* is the ID of all the processes (not an individual process ID).

3. Use the `-n` (or `-np`, `-c`, `-p`) option when you request that the Prism environment attach to a job.

Without one of these options, the Prism environment assumes that the ID number is a process ID rather than a job ID.

For example, the following starts the Prism environment and attaches to the running processes in job 15232:

```
% prism -n 2 mpiprogram 15232
```

See the *LSF Batch User's Guide* for further information about `bjobs`. See the *Sun HPC ClusterTools User's Guide* for further information about `mpps`.

You attach to a single process of a message-passing program by specifying its process ID. If you do this, however, you won't be able to view or debug what is happening in the other processes.

▼ To Attach to Multiple Jobs When Starting Prism

When you launch the Prism environment, you can specify a list of jobs to attach. Note that if the list of job IDs belongs to the same executable, the Prism environment starts only one session.

- **Type:**

```
% prism -jid_list
```

Detaching From a Running Process

A job is automatically detached from the Prism environment if you quit it or run another program. You can detach from the job without leaving the Prism environment by issuing the `detach` command.

The Prism environment lets you detach only when all the processes in the job are stopped. The detach operation itself sets them all running again, outside control of the debugger.

Executing a Program in the Prism Environment

You start execution of a program in the Prism environment by issuing the `run` command or choosing the Run or Run (args) selection from the Execute menu. You can attach to a program that is already running by using the `attach` command. This is described in “Attaching to a Running Message-Passing Process” on page 53.

▼ To Run a Program

- **Perform one of the following:**

- From the menu bar – If you have no commandline arguments you want to specify, choose Run from the Execute menu. Execution starts immediately. By default, the Run selection is in the tear-off region.

If you have command-line arguments, choose Run (args) from the Execute menu. A dialog box is displayed in which you can specify any command-line arguments for the program. This dialog box is represented in FIGURE 3-1. If you have more arguments than will fit in the input box, your entries scroll to the left. Click on the Run button to start execution.




FIGURE 3-2 Run (args) Dialog Box

- From the command window – Type the `run` command, including any arguments to the program on the command line. You can abbreviate the command to `r`.

If you have already run the program, you can issue the `rerun` command to run it again. The same argument list you previously passed to the program will be used. In both cases, you can redirect input or output using the standard Solaris redirection operators (`<` `>`).

When the program starts executing, the status region displays the message `running`.

You can continue to interact with the Prism environment while a program is running, but many features will be unavailable. Unavailable selections are grayed out in menus. If you issue a command that cannot be executed while the program is running, it is queued until the program stops.

Program I/O

By default, the Prism environment creates a new window for a program's I/O. This window persists across multiple executions and program loads, giving you a complete history of your program's input and output. If you prefer, you can display I/O in the X terminal from which you invoked the Prism environment. See "Resource Descriptions" on page 206 for details.

Status Messages

The Prism environment displays status messages before, during, and after the execution of a program. FIGURE 3-1 lists the various status messages that might be displayed.

TABLE 3-1 Status Messages

Message	Meaning
<code>connected</code>	Prism has connected to other nodes to work on a message-passing program.
<code>connecting</code>	Prism is connecting to other nodes in order to work on a message-passing program.
<code>error</code>	Prism has encountered an internal error.
<code>initial</code>	Prism is starting up without a program loaded.
<code>interrupted</code>	The program has been interrupted.
<code>loading</code>	Prism is loading a program.
<code>not started</code>	The program is loaded but not yet started.

TABLE 3-1 Status Messages (Continued)

Message	Meaning
running	The program is running.
stopped	The program has stopped at a breakpoint or signal.
terminated	The program has run to completion and the process has gone away.

Stepping and Continuing Through a Program

When using the Prism environment to debug a multiprocess and/or multithreaded program, menu actions such as Step and Next apply to the processes (or threads) belonging to the current set of processes (or threads).

Waiting for and Interrupting Processes

When debugging multiprocess programs, it can be useful at times to be able to identify when a specific process or set of processes has stopped execution. It can also be useful to be able to interrupt execution of individual processes. The Prism environment meets these needs with the commands `wait` and `interrupt`.

▼ To Wait for a Specified Process or Set of Processes to Stop Execution

● Type:

`(prism) wait [argument]`

A process is considered to have stopped if it has entered the `done`, `break`, `interrupted`, or `error` state.

There are two versions of the `wait` command:

- Use the syntax `wait` or `wait every` to wait for every member of the specified pset to stop. If no pset is specified, the command applies to the current pset.

`(prism notx) wait every`

waits for every process in the pset `notx` to stop. The current process will be whatever it would normally be; see “The Current Pset” on page 75. This is the default behavior of the `wait` command.

- Use the syntax `wait any` to wait for any member of the specified pset to stop. If no pset is specified, the command applies to the current pset. When the first process stops, it becomes the current process of this pset.

`(prism all) wait any pset foo`

waits for the first process in pset foo to stop.

There are corresponding Wait Any and Wait Every selections in the Prism environment's Execute menu. They apply to the processes of the current set.

If you prefer to have the `step` and `next` commands wait for processes to finish executing before letting you issue other commands, you can issue them along with the `wait` command. For example,

```
(prism all) step; wait
```

executes the next line, then waits for all processes in the current pset to finish execution.

If you use this command sequence frequently, you can provide an alias for it through the `alias` command. The Prism environment provides the default alias `contw` for the following command sequence:

```
(prism all) cont; wait
```

▼ To End the Wait

- **Perform one of the following:**

- Type Control-C; this does not affect processes that are running.
- Choose Interrupt from the Execute menu; this stops processes that are running, in addition to ending the wait.

▼ To Interrupt the Execution of a Process or Set of Processes

- **Perform one of the following:**

- Type:

```
(prism all) interrupt
```

For example, the following interrupts execution of process 0:

```
(prism all) interrupt pset 0
```

If there is a predefined pset called `running`, type the following to interrupt all the processes in that pset:

```
(prism all) interrupt pset running
```

Using the `interrupt` command resets the predefined pset `interrupted` so that it includes the newly interrupted processes. Processes leave this pset when they continue execution.

- Select Interrupt from the Execute menu.

This interrupts all running processes that are in the current pset.

Execution Pointer

When using the Prism environment to debug a scalar program, the > symbol in the line-number region points to the next line to be executed. In a multiprocess or multithreaded program, there can be multiple execution points within the program. The Prism environment marks all the execution points for the processes in the current set with a > in the line-number region or a * character when the current source position is the same as the current execution point.

▼ To Display a Pop-Up Window Showing the Executing Process(es)

- **Shift-click on the execution pointer symbol.**

This shows the process(es) for which the symbol is the execution pointer.

▼ To Find Out Execution Status in the Current Process Set

- **Type:**

`(prism all) pstatus pset_qualifier`

This finds out the execution status of the pset specified by the *pset_qualifier* argument.

▼ To Find Out Execution Status in a Specified Process Set

- **Type:**

`(prism all) pstatus`

This finds out the execution status of processes in the current pset.

Rerunning Spawned Prism Sessions

▼ To Rerun a Multiple Job Session Within the Prism Environment

- **Type”**

`(prism all) rerun`

When issued in the primary Prism session, this command cleans up current secondary Prism sessions (shutting down both the secondary Prism sessions and the debuggee processes) before rerunning the currently loaded program.

The `rerun` command is not valid when issued in the secondary Prism sessions.

For more information about debugging multiple sessions within the Prism environment, see “Enabling Support for Spawned MPI Processes” on page 18.

Controlling Programs With the Commands-Only Interface

▼ Starting a Program

- **Type:**

`(prism all) run`

This starts a program using the commands-only interface. The program starts up in the background.

▼ Bringing a Program to the Foreground

- **Type:**

`(prism all) fg`

This brings the running program into the foreground. Note that you cannot execute Prism commands while the program is executing in the foreground.

▼ Sending a Program to the Background

- **Type:**

Ctrl-Z

This key sequence sends the running program to the background and regains the `(prism)` prompt.

▼ Quitting a Prism Debugging Session

- **Type:**

`(prism all) quit`

This command terminates the debugging session. Before quitting, the Prism environment kills the debugging process if it was started with `run`, or the Prism environment will detach from it if the program was previously attached.

Using Psets in the Prism Environment

The Prism environment enables you to view your program at the level of an individual process or individual thread.

Note – To view a program at the process level means to view the program at the level of the *main* thread.

You can use the Prism environment to view all of the processes (and threads), or subsets of the processes (and threads) that make up the program. For example, at times it may be useful to look at the status of process 0 (or thread 0.1). At other times, you might want to look at all processes or threads that have encountered an error, or at all processes or threads in which the value of a particular variable exceeds a certain number.

Such groups of processes or threads (psets) are often of interest because they have some useful characteristic in common. The Prism environment treats a pset as a unit: For example, you can use the name of a pset as a qualifier for many commands. The command is then executed for each process in the set.

You can view psets in the Psets window, as described in “Using the Psets Window” on page 62 and “Viewing Pset Contents From the Psets Window” on page 71.

Note – MPI rank numbers are used to identify processes in psets. In multithreaded programs, the Prism environment identifies threads numerically. For example, the first thread in process 0 is thread 0.1. Do not confuse these numbers with the Solaris process IDs (pids) assigned by the system to the message-passing processes.

The Prism environment provides predefined psets for certain groups of processes (and threads). For example, the set of all processes (and threads) in an error state is a predefined pset. You can also define your own psets, as described in “Defining

Psets” on page 66; for example, you can define a pset to be those processes (and threads) in which variable *x* is greater than 0. “To Delete Psets From the Psets Window” on page 75 describes how to delete psets.

Some Prism commands can be directed to apply to specific psets by including pset qualifiers as arguments to those commands. If you don’t specify a pset as a qualifier to one of these commands, the command is executed on the *current* pset. The concept of the current pset is described in “The Current Pset” on page 75. “The Current Process” on page 78 describes the current process, which is a distinguished process (or thread) within a pset.

Note – In threaded programs, the Prism environment extends the notion of current process to refer to the current thread of a pset.

Using the Psets Navigator

You can navigate to any defined pset using the pull-down menu and arrow keys on the main MP Prism window. The pset navigator controls are shown at the bottom of FIGURE 3-1.




FIGURE 3-3 Pset Navigator Controls

Using the Psets Window

You can use the Psets window to view the current status of the processes in your program and to perform many of the actions associated with psets.

▼ To Display the Psets Window

- **Perform one of the following operations:**

- From the menu bar – Choose the Psets selection from the Debug menu.
- From the command window – Type

```
(prism all) show psets on dedicated
```


FIGURE 3-4 shows the Psets window for a nonthreaded 16-process message-passing program, including several user-defined psets.

FIGURE 3-5 shows the Psets window for a multithreaded program, including the predefined psets.

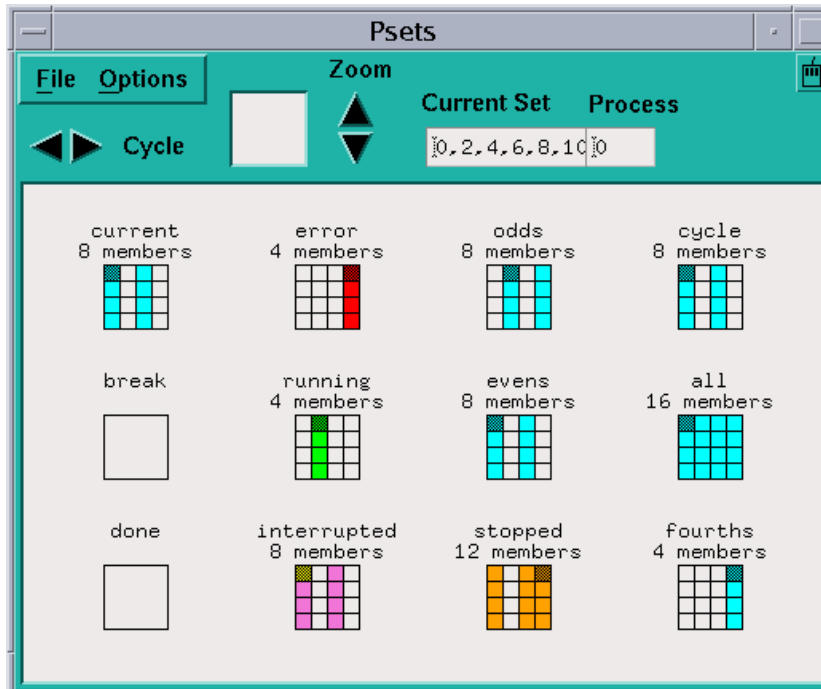


FIGURE 3-4 Psets Window (Nonthreaded)

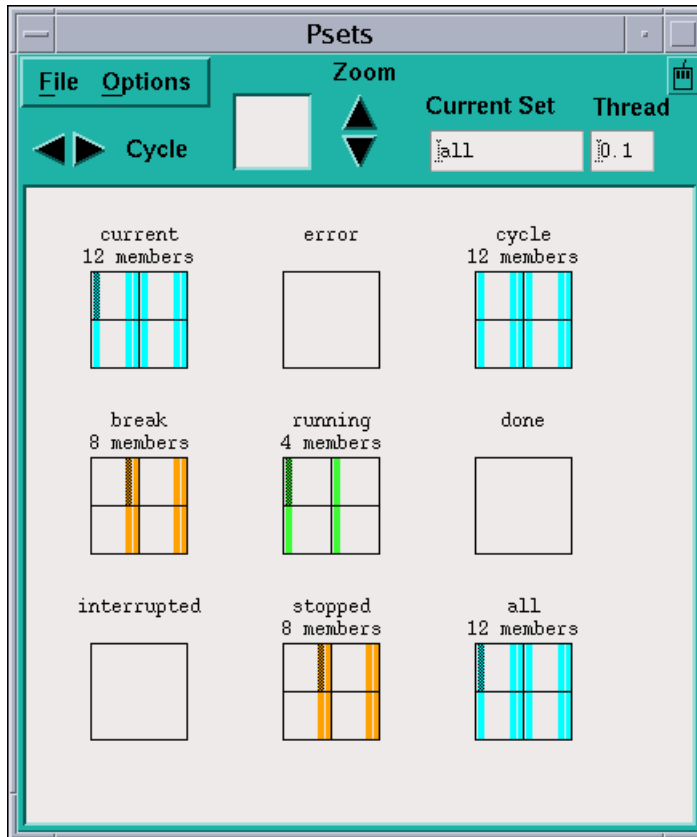


FIGURE 3-5 Psets Window (Threaded)

The various components of the window are described in detail in later sections. Here is a brief overview:

- The main area of the window shows psets and their members.
- In nonthreaded psets, processes that are members of a set are shown as colored (or black) cells within a rectangle that represents the entire set of processes that make up the message-passing program.
- Threads of a threaded pset are shown as colored stripes. By default, threads 2, 3, and 4 in all ranks belong to the *hide* set. These are auxiliary threads created by any program that is linked with `libthread.so` and are rarely interesting to a programmer. For further information about hiding threads, see “Hiding Threads From Psets” on page 81.
- The current process or thread for each pset is shown in a darker shade of the color used in the other squares—gray for noncolor monitors.

- The `current` pset is shown in the upper left corner of the window. The name of the `current` pset and the number of the current process are displayed in the small window on the upper right corner of the control panel.
- You can cycle through the `cycle` pset by clicking on the left and right arrows labeled `Cycle` at the top left of the control panel.
- If you have many psets and a large number of processes or threads, you can use the `Zoom` arrows to zoom in or out on particular psets. The box to the left of the arrows shows what part of the entire display you are seeing; you can drag the mouse through this box to pan through the display.
- You can view and change the `current` pset and current process or thread via the boxes at the top right of the window.
- The `Options` menu at the top left of the window lets you hide, display, create, and delete psets. See the discussions starting with “Defining Psets” on page 66 through “To Delete Psets From the Psets Window” on page 75.
- The `File` menu lets you close the `Psets` window.

Predefined Psets

The Prism environment provides the following predefined psets:

- `all` – Contains all the processes (and threads) in the program; it is the default pset at start-up. The `all` pset does not contain processes (or threads) that have terminated or were joined.
- `running` – Contains all processes (and threads) that are currently executing.
- `error` – Contains all processes (and threads) that have encountered an error.
- `interrupted` – Contains the processes (and threads) that were interrupted most recently by the user. See “Waiting for and Interrupting Processes” on page 57 for a discussion of the `interrupt` command and a further explanation of this pset.
- `break` – Contains the processes (and threads) that are currently stopped at breakpoints.
- `stopped` – Contains all processes (and threads) that are currently stopped. It is the union of the sets `error`, `interrupted`, and `break`.
- `done` – Contains all processes (and threads) that have terminated successfully. For *user* threads (not thread 1), the `done` set contains only *zombie* threads (threads that are unjoined). Once a thread is joined, it ceases to exist.

These sets are *dynamic*; that is, as a program executes, the Prism environment automatically adjusts the contents of each set to reflect the program’s current state.

In addition, there are two set names that have special meaning: `current` and `cycle`. These are discussed in “The Current Pset” on page 75 and “The `cycle` Pset” on page 80, respectively.

Defining Psets

You can create psets in the Prism environment. This section describes the syntax of pset creation. This syntax provides a convenient shorthand method for entering complicated pset specifications.

Psets can be constructed using one or more of the following syntactical building blocks:

- *An individual process (or thread) number.*
- *The name of an existing pset.* The new pset will have the same definition as the existing set.
- *A list of process (or thread) numbers.* Separate the numbers with commas. Use a colon between two process (or thread) numbers to indicate a range. Use a second colon to indicate the stride to be used within this range.
- *A union, difference, or intersection of psets.* To specify the union, use the symbol `+`, `|`, or `||`. To specify the difference, use the minus sign (`-`). To specify the intersection, use the symbol `&`, `&&`, or `*`.
- *A snapshot of a pset expression.* In a multithreaded program, use the `snapshot (pset_expression)` argument to define a pset with a constant value that could otherwise change during program execution. For more information about the snapshot intrinsic, see “Using Snapshots of Unbounded Psets in Commands” on page 85.

Note that you can use predefined psets to define new psets. Except for pset `all`, when you use a predefined pset to define a new pset, the Prism environment uses the instantaneous value of the predefined pset. Thus, even if the predefined pset changes, the user-defined pset remains unchanged until the user forces reevaluation with a Prism command, such as `eval pset`.

▼ To Specify a Pset as an Argument to a Command

- **Type:**

```
(prism all) command pset pset_specifier
```

Put the *pset_specifier* clause at the end of the command line (but before an *on window* clause, if any).

```
(prism all) print x pset error
```

prints the values of the variable `x` in the predefined pset `error`. (See “Visualizing Multiple Processes” on page 168 for a discussion of printing variables in the Prism environment.)

▼ To Specify a Pset as a Subset of a Pset Clause

- **Perform one of the following:**

- Specify an individual process number. An individual process can constitute a pset.


```
(prism all) print x pset 0
```

 prints the value of `x` in process 0. If the program is multithreaded, it prints the value of `x` in all threads in process 0.
- Specify an individual thread number. An individual thread can constitute a pset.


```
(prism all) print x pset 0.1
```

 prints the value of `x` in process 0, thread 1.
- Specify the name of a pset. Name a pset using the `define pset` command, as described in “Naming Psets” on page 68.


```
(prism all) print x pset foo
```

 prints `x` in the processes you have defined to be members of pset `foo`.
- Specify a list of process numbers. Separate the numbers with commas.


```
(prism all) print x pset 0, 4, 7
```

 prints `x` in processes 0, 4, and 7.

Ranges and strides are allowed. Use a colon between two process numbers to indicate a range. Use a second colon to indicate the stride to be used within this range.

```
(prism all) print x pset 0:10
```

 prints `x` in processes 0 through 10. And


```
(prism all) print x pset 0:10:2
```

 prints `x` in processes 0, 2, 4, 6, 8, and 10.

You can also combine comma-separated process numbers and range specifications.

```
(prism all) print x pset 0, 1, 3:5, 8
```

 prints `x` in processes 0, 1, 3, 4, 5, and 8.
- Specify a union, difference, or intersection of psets. To specify the union of two psets, use the symbol `+`, `|`, or `||`.


```
(prism all) print x pset 0:2 + 8:10
```

 prints `x` in processes 0, 1, 2, 8, 9, and 10.


```
(prism all) print x pset foo | bar
```

 prints `x` in processes that are members of either pset `foo` or pset `bar`.

The Prism environment evaluates the pset expression from left to right. If a process is a member of the first part of the expression, it is not evaluated in the rest of the expression. In the example above, if a process is a member of `foo`, its value of `x` is printed; the Prism environment does not test for its membership in `bar`.

- Specify the difference of two psets by using a minus sign.

```
(prism all) print x pset stopped - foo
```

prints `x` in all processes that are stopped except those belonging to the pset `foo`.

- Specify the intersection of two psets, using the `&`, `&&`, or `*` symbol.

```
(prism all) print x pset foo & bar
```

prints `x` in processes that are members of both pset `foo` and pset `bar`. If a process returns `false` for the first part of the expression, it is not evaluated further. In the example above, if a process is not a member of `foo`, the Prism environment does not test for its membership in `bar`; it won't be printed in any case.

- Specify a condition to be met. Put braces around an expression that evaluates to true or false in each process. Processes in which the expression is true are part of the set.

```
(prism all) print x pset {y > 1}
```

prints `x` in processes where `y` is greater than 1. Likewise,

```
(prism all) print x pset all - {y == 1}
```

prints `x` in all processes except those in which `y` is equal to 1.

- Membership in some psets can change based on the current state of your program; such a pset is referred to as *variable*. See “To Evaluate Variable Psets” on page 70 to learn how to update the membership of a variable pset.

For this syntax to work, the variable must be active in all processes in which the expression is evaluated. If the variable is not active in a process, you get an error message and the command is not executed. To ensure that the command is executed, use the intrinsic `isactive` in the pset definition. The expression `isactive(variable)` returns `true` if `variable` is on the stack for a process or is a global.

Thus, you could use this syntax to ensure that `x` is printed:

```
(prism all) print x pset stopped && {isactive(x)}
```

Naming Psets

You can assign a name to a pset. This is convenient if you plan to use the set frequently in your Prism session.

Use the syntax described above in “Defining Psets” on page 66 to specify the pset. You can use any name except the names that the Prism environment predefines; see “Predefined Psets” on page 65. The name must begin with a letter; it can contain any alphanumeric character, plus the dollar sign (\$) and underscore (_).

▼ To Name a Pset

- **Do one of the following:**

- From the Psets window – Choose Define Set from the Options menu. A dialog box is displayed that prompts for the name and definition of the pset. Click on Create to create the pset.
- From the command line – Issue the `define pset` command.

To create a pset called `odd`, which contains the odd-numbered processes between 1 and 31, type:

```
(prism all) define pset odd 1:31:2
```

To create a pset from the first thread in process one, type:

```
(prism all) define pset gui_thread 1.1
```

To create a pset from the second thread in process one, type:

```
(prism all) define pset io_thread 1.2
```

To create a pset from an expression that takes the intersection of all processes and threads and subtracts the two psets defined in the two previous examples, type:

```
(prism all) define pset workers (all.all - gui_thread -  
io_thread)
```

To create a pset that consists of all processes in which `x` is not equal to 0, enter

```
(prism all) define pset xon {x .NE. 0}
```

Note that `x` must be active in all processes for this syntax to work. As described above, you can use the intrinsic `isactive` to ensure that `x` is active in the processes that are evaluated. The following example creates a variable pset whose contents will change based on the value of `x`:

```
(prism all) define pset xon {isactive(x) && (x .NE. 0)}
```

Variable psets are discussed in a later section.

Finally, note that all processes must be stopped for this syntax to work. To ensure that the definition applies only to stopped processes, use the following syntax:

```
(prism all) define pset xon stopped && {isactive(x) && (x .NE.  
0)}
```

Dynamic, user-defined psets are deleted when you reload a program. To get a list of these psets before reloading, issue the command `show psets`. You can then use this list to help reissue the `define pset` commands. See “Viewing Pset Contents From the Psets Window” on page 71 for more information about `show psets`.

The Prism environment evaluates the membership of a variable pset when it is defined. If no processes meet the condition, the Prism environment prints appropriate error messages, but the set is defined.

▼ To Evaluate Variable Psets

- **Type:**

```
(prism all) eval pset psetname
```

The following example evaluates the membership of the pset `xon`:

```
(prism all) eval pset xon
```

This causes the display for the pset to be updated in the Psets window.

Note that this evaluation fails if one of the following is true:

- Processes are running that need to be polled in evaluating the pset.
- The pset's definition contains a variable that is not active in any of the processes being polled.

For example, type the following command:

```
(prism all) define pset foo {x > 0}
```

You must be certain that all processes are stopped and that `x` is active on all processes when you type the command:

```
(prism all) eval pset foo
```

For greater assurance that the evaluation will succeed, use the following syntax:

```
(prism all) define pset foo stopped && {isactive(x) && (x > 0)}
```

The extra information provided ensures that the evaluation takes place only in processes that are stopped and in which `x` is active.

If an evaluation fails, the membership of the pset remains what it was before you issued the `eval pset` command.

You can use the `eval pset` command in event actions; see “Events Taking Pset Qualifiers” on page 101.

Combining Named Psets and Pset Expressions

You can use combinations of named psets and pset expressions to isolate processes and threads of interest, as shown in TABLE 3-2:

TABLE 3-2 Examples of Pset Composition

pset	Contains
<code>pset 1.3</code>	Thread 3 in process 1
<code>pset 1:10.3</code>	Thread 3 in processes 1 to 10
<code>pset 1.1, 2.2:5</code>	Process 1, thread 1 and process 2, threads 2, 3, 4, and 5
<code>pset 1.all</code>	All threads in process 1
<code>pset 1</code>	All threads in process 1
<code>pset .4</code>	Thread 4 in all processes. Same as all.4
<code>pset 1,2.(3,4)</code>	All threads in process 1, threads 3 and 4 in process 2
<code>pset 1,2.3,4</code>	All threads in processes 1 and 4, thread 3 in process 2
<code>pset {isactive(var) && var == 1}</code>	All threads in which the variable <i>var</i> is on the stack for a process (or is a global) and has value 1

Each of the following lines specify the same pset:

```
pset {var_i == 3} . { var_j == 4}
pset {var_i == 3} & { var_j == 4}
pset {var_i == 3 && var_j == 4}
```

Viewing Pset Contents From the Psets Window

The easiest way to view the contents of psets is to use the Psets window.

By default, the window displays the current pset (which starts out being the predefined pset `all`) and the psets `break`, `running`, and `error`. When you create a new pset via the `define pset` command, that set is also displayed automatically.

The processes within a pset are numbered starting at the upper left, increasing from left to right and then jumping to the next row. You can display information about them as follows:

- Shift-click on a cell to view the Prism ID number of the process it represents.
- Shift-click elsewhere in the pset rectangle (for example, on a border) to display all the ID numbers of the processes in the pset.
- Shift-middle-click on a cell to view the process's Solaris pid and the hostname of the node on which it is running.
- Shift-middle-click elsewhere in the rectangle to display the entire list of pids and hostnames for the processes in the pset.

▼ To Display a Pset

1. **Choose the Show selection from the Options menu in the Psets window.**

This displays a list of psets; the predefined psets are at the top, followed by any user-defined set names.

2. **Click on a set name, and that set is displayed in the window.**

▼ To Hide a Pset

1. **Choose the Hide selection from the Options menu.**

This displays the list of predefined and user-defined psets.

2. **Click on a set name to remove that set from the display.**


Note that hiding a pset doesn't otherwise affect its status; it still exists and can be used in commands.

Note also that the Show and Hide submenus include the choices All Sets and all. The All Sets choice refers to all psets; the all choice refers to the predefined pset `all`.

▼ To View Psets Not Shown in the Display Window

1. **Use the navigator rectangle (between the Cycle and the Zoom arrows) to pan through the psets.**

The gray box in the rectangle shows the position of the display area relative to all the psets that are to be displayed:



2. **Either drag the box or click at a spot in the rectangle.**

The box moves to that spot, and the display window shows the psets in this area of the total display.

To display more psets at the same time, click on the Zoom up arrow to the right of the navigator rectangle. This raises the zoom factor, increasing the size of the boxes that represent the psets. Clicking on the Zoom down arrow decreases the size of these boxes.

Viewing Pset Contents From the Command Line

▼ To Print the Contents of the Specified Pset

- **Type:**

```
(prism all) show pset [psetname]
```

For example, this command might produce this subsequent response:

```
(prism all) show pset stopped  
The set contains the following processes: 0:3.
```

The `show pset` command is discussed further in “To Find Out the Current Pset” on page 76.

The `show psets` command displays the contents and status of all psets, as shown by the following sample output:

```
(prism all)show psets
foo:
  definition = 0:31:2
  members = 0,2,4,6,8,10,12,14,16,18,20,22,24,26,28,30
  current process = 0
break:
  definition = break
  members = nil
  current process = (none)
done:
  definition = done
  members = 0:31
  current process = 0
interrupted:
  definition = interrupted
  members = nil
  current process = (none)
error:
  definition = error
  members = nil
  current process = (none)
running:
  definition = running
  members = nil
  current process = (none)
stopped:
  definition = stopped
  members = nil
  current process = (none)
current:
  definition = 6, 9, 12
  members = 6,9,12
  current process = 6
cycle:
  definition = 6, 9, 12
  members = 6,9,12
  current process = 6
all:
  definition = all
  members = 0:31
  current process = 12
```

Deleting Psets

You can delete named psets that you have defined. You cannot delete any predefined pset except `cycle`; see “The cycle Pset” on page 80.

▼ To Delete Psets From the Psets Window

● Perform one of the following:

- From the Psets window – Select Delete from the Options menu. This displays a list of psets that you can delete. Click on the name of the pset you want to delete. If it is currently displayed in the Psets window, it disappears.
- From the command line – Issue the `delete pset` command, using a pset qualifier to specify the name of a user-defined pset. For example, this deletes the pset named `xon`:

```
(prism all) delete pset xon
```

See “Events Taking Pset Qualifiers” on page 101 for a discussion of the effects of deleting a pset on events that have been defined to affect the members of that set.

The Current Pset

The command syntax described in “Defining Psets” on page 66 lets you apply a command to a specific pset. If you don’t use this syntax, the command is applied to the *current* pset; *current* is a predefined pset name in the Prism environment. In addition, many graphical actions in the Prism environment apply only to the members of the current set.

Note that you cannot change the current pset to one that has no members. If you try to do so, nothing will happen in the Psets window and you will get a message like the following in the history region of the command window:

```
Cannot set current pset to running -- it is empty.
```

When a program is first loaded, the current pset is the default pset, `all`.

▼ To Change the Current Pset

● Perform one of the following:

- From the Psets window – There are several ways of changing the current pset via the Psets window:
 - If the set is displayed in the Psets window, simply double-click anywhere in its display (for example, on its name or in the box beneath its name).

- Select Set Pset from the Options menu. This displays a list of psets. Click on the name of the set you want to be current.
- Edit the name of the pset in the box below Current Set, then press Return.

When you change the current set, the new name appears in the Current Set box in the Psets window, and the current set shown at the top left of the psets area changes to reflect the contents of the new set.

- From the command line, type:

```
(prism all) pset pset_specifier
```

For example, the following changes the current pset to foo:

```
(prism all) pset foo
```

You can also use the `pset` command with the pset-specification syntax described in “Defining Psets” on page 66.

▼ To Find Out the Current Pset

- **Perform one of the following:**

- Look for the name in the Current Set box in the Psets window.
- Look in the status region in the Prism environment’s main window.
- Type

```
(prism all) pset
```

This displays the current set.

▼ To List the Processes in the Current Pset

- **Type:**

```
(prism all) show pset
```

The Psets window also displays the processes in the current pset. In this example, issuing the `show pset` command at the Prism commandline causes the Prism Environment to list the processes in the current pset:

```
(prism foo) show pset
pset 'current' is defined as 'foo'.
The set contains the following processes: 1,2.
```

Current Pset and Predefined Psets

“Predefined Psets” on page 65 describes predefined psets—sets like `running`, `stopped`, and `interrupted` whose contents the Prism environment automatically updates during execution of the program.

If you specify a predefined pset as the current pset, it becomes a *static* pset that consists of the processes that are members of the predefined set at the time it is made the current set. To make this clear, the (prism) prompt changes to list the processes that are members of this static set. In this example, if processes 0, 1, and 13 are the only processes that are stopped, the pset command changes the Prism commandline prompt to list the processes 0, 1, and 13:

```
(prism all)pset stopped
(prism 0:1, 13)show pset
The current set was created by evaluating the pset
'stopped' once at the time when it became the current set.
The set contains the following processes: 0:1, 13.
```

Output of the show pset command is explicit under these circumstances.

Issuing the pset command with no arguments displays the same information.

Note that the (prism) prompt can become quite long if there are many processes in a current pset derived from a dynamic pset. By default, the prompt length is limited to 25 characters. You can change this default by issuing the set command with the \$prompt_length variable, specifying the maximum number of characters to appear in the pset part of the prompt. In this example, the command shortens the prompt long_pset_name to long_pset:

```
(prism long_pset_name)set $prompt_length=9
(prism long_pset)
```

Current Pset and Variable Psets

“Defining Psets” on page 66 describes how to create variable psets—user-defined psets whose membership can change in the course of program execution.

▼ To Update the Membership of a Variable Pset

- **Type:**

```
(prism all) eval pset
```

If you make a variable pset your current set, its membership is determined by the most recent eval pset command you have executed for the set.

The Current Process

Each pset has a current process, which serves as the scoping point for Prism commands. By default, the process with the lowest rank is the current process. If that process is threaded, its lowest numbered thread is the current process.

The Prism environment uses the current process and current thread in several ways:

- The Prism environment's source window displays the source that is executing in the current process or current thread.
- The Prism environment centers the Where graph around the call stack of the current process or current thread.
- The Prism environment uses the current process or current thread to resolve variable names.

▼ To Change the Current Process

● Perform one of the following:

- From the Psets window – Do either of the following:
 - Click on the cell representing the process in the displayed pset. The cell turns a darker shade of the color used for the other processes.
 - To change the current process in the current pset, you can also edit the number in the box under Process (or Thread if the loaded program is a threaded program) at the top right of the window, then press Return.
- From the command line – Issue the `process` command. This is illustrated in the following example:

```
(prism all) process 2  
The current process is now 2.
```

The syntax of the `process` command includes both process number and thread ID, as shown below:

```
(prism all) process process_number.thread_ID
```

process_number is the rank of the process that you want to make the current process. *thread_ID* is the identifying number of the thread in that process that you want to be the current thread.

If you do not specify the *thread_ID* value, it defaults to the lowest numbered thread on the specified process. This is illustrated in the following example:

```
(prism all) pset 1:4:2.2:3  
(prism 1:4:2.2:3)process 3  
(prism 1:4:2.2:3)
```


This example replaces the `pset all` with a `pset` consisting of the processes and threads: 1.2, 1.3, 3.2, and 3.3. In other words, the `pset` command specifies that the new `pset` will consist of a subset of processes in the range of 1 through 4. Within this range, the subset is restricted to threads 2 and 3 on those processes that are selected by a process stride of 2.

When the `pset` command completes execution, the default current process is 1.2. This is the process with the lowest rank and the lowest numbered thread on that process.

The `process` command on the second line changes the current process to 3.2.

▼ To Print the Current Process of the Current Pset

- **Type:**

(prism all) `process`

Scope in the Prism Environment

When using the Prism environment to debug a message-passing program, the scope of the current process determines the scope for resolving the names of variables.

If a command applies to a `pset` other than the current set, the Prism environment uses the scope of that set's current process.

It is possible that other members of the `pset` will have different scopes from that of the current process, or that its scope level will not even exist in these processes. In these cases, you receive an error message when you try to issue a command (for example, `print` or `display`) that requires a consistent scope. To solve the problem, you can do one of the following:

- Restrict your `pset` so that it contains only members with the same scope.
- If the current process's scope level does not exist in other processes in the set, you can use the `up` command to move up its call stack to a point where it has a scope level that does exist in the other processes.
- If different processes in the set have different scopes, you can issue the `up` and `down` commands as needed to ensure that they all have the same scope.

Commands such as `pset` and `process` that affect scope print the current function when you issue them.

The `cycle` Pset

In debugging a message-passing program, you may often want to look, in turn, at each process within a pset. The `cycle` pset provides you with a convenient way of doing this.

▼ To Create a `cycle` Pset out of an Existing Pset

- **Type:**

```
(prism all)define pset cycle psetname
```

If *psetname* is dynamic, the `cycle` pset is statically fixed when you create it. You can then cycle through each process in this pset to examine it in turn.

By default, the `cycle` pset is equivalent to the `current` pset. For more information about the `define pset` command, see “Defining Psets” on page 66.

For example, the following copies `foo` into the `cycle` pset:

```
(prism all)define pset cycle foo
```

Note that changing the `cycle` pset erases any previous cycling information. For example, if you do the following:

1. Make `foo` the current set and cycle partway through it.
2. Make `bar` the current set.
3. Once again make `foo` the current set.

Then you start at the beginning again when you cycle through the members of `foo`.

▼ To Cycle Through the Processes in a `cycle` Pset

- **Perform one of the following:**

- From the Psets window:

Use the Cycle arrows at the top left of the window to cycle through the members of the `cycle` set.

Click on the right arrow to cycle up through the members of the set; click on the left arrow to cycle down through the members.

- From the command line:

Type:

```
(prism all) cycle
```

This has the same effect as clicking on the right cycle arrow in the Psets window.

In a nonthreaded program, the `cycle` command sets the current process to the next one in the current pset. In a threaded program, it sets the current process to be the next valid thread on the current process. If the `cycle` command is entered when the current thread is the last thread in the process, it steps to the next process in the pset.

In the following example, this Prism session defines a pset, makes it the current set, and then cycles through its members:

```
(prism all) define pset foo 0:3  
(prism all) pset foo  
(prism foo) cycle  
(prism 1) cycle  
(prism 2) cycle  
(prism 3) cycle  
(prism 0)
```

- From the Source-Window pop-up menu:

Choose Cycle from the Source-Window pop-up menu. This advances to the next member of the `cycle` pset.

Cycle Visualizer Window

The Prism environment includes a Cycle window type for visualizing data. When you print a variable's value to the Cycle window, the value changes to that of the variable in the new process whenever you cycle through the members of the `cycle` pset. For more information, see "Visualizing Multiple Processes" on page 168.

Hiding Threads From Psets

The `pset` command takes two thread-specific options, `-hide` and `-unhide`. These options control membership in a set of *hidden* threads.

Hidden threads never appear in any pset, and debugging commands are never sent to them regardless of the definition of the current set. Once hidden, those threads are represented by empty stripes in the Psets window and Where graph. By default, the

set of hidden threads consists of threads 2, 3, and 4 in all ranks. These are auxiliary threads created by any program that is linked with `libthread.so` and are rarely interesting to a programmer.

The following procedures are valid only when debugging a multithreaded program.

▼ To Hide Threads From Psets

- **Type:**

```
(prism all)pset -hide pset_expression
```

The Prism environment evaluates *pset_expression* and adds the result to the set of hidden threads.

▼ To Make Hidden Threads Available to Psets Again

- **Type:**

```
(prism all)pset -unhide pset_expression
```

The Prism environment evaluates *pset_expression* and subtracts the result from the set of hidden threads.

▼ To Show Currently Hidden Threads

- **Type:**

```
(prism all)pset -hide
```

Using Psets in Commands

As mentioned at the beginning of this chapter, you can specify pset qualifiers with several Prism commands. The following commands can take a pset as a qualifier:

- *address/*
- *assign*
- *call*
- *catch*
- *cont, contw*
- *display*
- *ignore*
- *interrupt*
- *lwps*
- *next, nexti*
- *print*
- *pstatus*

- return, stepout
- step, stepi
- stop, stopi
- sync, syncs
- thread, threads
- trace, tracei
- wait
- whatis
- where

▼ To Use a Pset Qualifier

- **Type:**

`(prism) command options pset_qualifier [on window]`

A command with a pset qualifier applies only to the processes in the set. If you omit the qualifier, the command applies to the processes in the current set.

For example, the following sets a breakpoint at line 12 for the processes in pset error:

```
(prism all) stop at 12 pset error
```

The following displays the Where graph for processes 0 through 10:

```
(prism all) where pset 0:10 on dedicated
```

See “Displaying the Where Graph” on page 113 for a description of the Where graph.

The following creates a trace event for the members of the current pset:

```
(prism all) trace at 12 if x > 10
```

Note that this last command applies only to the members of the current pset. To apply it to all processes, use the syntax:

```
(prism all) trace at 12 if x > 10 pset all
```

Many commands cannot logically take a pset qualifier. You get an error message if you try to issue one of these commands with a pset qualifier.

Using Unbounded Psets in Commands

When running threaded programs in the Prism environment, you can encounter *unbounded* psets. An unbounded pset is one that contains the value of `all` in the thread-part of a pset specifier. The membership of such psets varies unpredictably. The term *unbounded* distinguishes such psets from those whose membership varies deterministically, which are referred to as *variable* psets (see “Naming Psets” on page 68).

For example:

```
pset 3.all
```

The size of such an unbounded pset is not constant, since it contains all threads created during the life of the program. The size of this set will change as threads are created and destroyed.

Pset expressions that omit specifying the thread-part imply all threads, so that pset 2 means pset 2.all, and pset all means pset all.all, both of which are unbounded sets.

If a pset expression includes one or more unbounded psets, it is also unbounded.

Note – The use of all in only the process-part of a pset specifier does not create an unbounded set. The Prism environment creates a constant number of processes at startup, taken from the number of processes you specify when you start the Prism environment with a `-n` (or `-np`) argument. For example, pset all.1 is a bounded set.

The Prism environment places several restrictions on the use of unbounded psets. You cannot use an unbounded pset as the context for an event specification or a `wait every` command. For an overview of information about event specifications, see “Overview of Events” on page 93.

In the following example, both of these uses of the `wait every` command are illegal:

```
(prism all) wait every pset all
...
(prism all) pset all
(prism all) wait every
```

Similarly, you may not use unbounded sets as the context for the `stop` or `trace` commands when these commands contain actions. Examples:

```
(prism all) stop in foo {print x} pset all ; illegal
(prism all) stop in foo pset all ; legal, does not contain an action
```

The Prism environment handles the psets that apply to the `wait every`, `stop`, and `trace` commands in a similar manner. When using a constant (bounded) pset, the Prism environment records the membership of the pset when the command is issued. When using an unbounded pset, the Prism environment reevaluates the pset each time the command executes. In the following example, `foo` is an unbounded pset.

```
(prism all) stop at 10 pset foo
```

Each time a thread executes line 10, the Prism environment reevaluates pset `foo`, and stops the thread if it is a member of `foo`.

Using Snapshots of Unbounded Psets in Commands

The Prism environment enables you to control the contents of psets derived from unbounded sets of threads. You can specify a constant membership of such a pset by capturing snapshots of the unbounded sets. You capture snapshots of unbounded sets by including the `snapshot` (*pset_expression*) argument with any command that takes a pset qualifier.

Here is an example of how the contents of unbounded psets can vary:

```
(prism all)pset
The current set was created by evaluating the Pset
'all' once at the time when it became the current set.
The current set is defined as 'all'.
The set contains threads 0:3.1
(prism all)define pset all1 all - 1.1
(prism all)show pset all1
Pset 'all1' is defined as 'all - 1.1'.
The set contains the following threads: (0,2,3).1.
```

After running the program for a while, the membership of `all` and `all1` both change, as shown in the following example:

```
(prism all)show pset all
The set contains the following threads: 0:2.(1,5,6).
(prism all)show pset all1
Pset 'all1' is defined as 'all - 1.1'.
The set contains the following threads: (0,2).(1,5,6), 1.(5,6).
```

▼ To Create a Bounded Pset From an Unbounded Pset

- **Type:**

```
(prism all) command (pset_name) pset snapshot (expression)
```

In the following example, a snapshot pset, called `snap1`, is created. `snap1` contains a subset of the threads in `pset all`, as defined by the snapshot expression.

```
(prism all)pset
The current set was created by evaluating the Pset
'all' once at the time when it became the current set.
The set contains threads: 0:2.1.
(prism all)define pset snap1 snapshot ( all - 1.1 )
(prism all)show pset snap1
Pset 'snap1' is defined as 'snapshot ( all - 1.1 )'.
The set contains the following threads: (0,2).1.
```

Then, after running the program for a while, the membership of `(all - 1.1)` and `snap1` differ, as shown in this example:

```
(prism all)show pset all
The set contains the following threads: 0:2.(1,5,6).
(prism all)show pset snap1
Pset 'snap1' is defined as 'snapshot ( all - 1.1 )'.
The set contains the following threads: (0,2).1.
```

However, you can force the update of the membership of `pset snap1` by issuing the `eval pset` command. For example,

```
(prism all)eval pset snap1
(prism all)show pset snap1
Pset 'snap1' is defined as 'snapshot ( all - 1.1 )'.
The set contains the following threads: (0,2).(1,5,6), 1.(5,6)
```

The following example shows a situation in which using an unbounded pset, `all`, generates an error. Note that, in a threaded program, `all` is equivalent to the unbounded set of `all.all`, which is the union of all processes and all threads. The use of the snapshot argument, however, avoids that error.

```
(prism all - 1.1)stop in func {print 1 } pset all
Currently, dynamic psets are not allowed in events.
Action is dropped from event 3 because of dynamic pset all
(3) stop in func pset all
(prism all - 1.1)stop in func {print 2 } pset snapshot(all)
(4) stop in func { print 2 } pset snapshot(all)
```


Referring to Nonexistent Thread Identifiers

You cannot establish a current pset that contains non-existent threads. The `pset` command in the following example is wrong because it specifies a thread (5) that has not been created yet.

```
(prism all)show pset all
The set contains the following threads: (0:3).1
(prism all)pset all.5
```

However, in certain contexts, such as setting breakpoints, you may use a pset qualifier containing nonexistent threads, as shown in the following example:

```
(prism all)stop in foo pset all.5
(prism all)
```

Using the Prism Environment With Sun MPI Client/Server Programs

You can use a Prism session to debug more than one Sun MPI job at a time. To debug a child or client program, it is necessary to launch an additional Prism session. If the child program is spawned using calls to `MPI_Comm_spawn()` or `MPI_Comm_spawn_multiple()`, Prism can (if enabled) debug the child program as well.

However, if an MPI job connects to another job, the current Prism session has control only of the parent or server MPI job. It cannot debug the children or clients of that job. This might occur, for example, when an MPI job sets up a client/server connection to another MPI job with `MPI_Comm_accept` or `MPI_Comm_connect`.

With the exception of programs using calls to `MPI_Comm_spawn()` or `MPI_Comm_spawn_multiple()`, to use the Prism environment to debug a Sun MPI program, the program must be written in the SPMD (single program, multiple data) style—that is, all processes that make up a Sun MPI program must be running the same executable.

Note – `MPI_Comm_spawn_multiple` can create multiple executables with only one job id; therefore, you can use the Prism environment to debug jobs with different executables that have been spawned with this command.

Choosing the Current File and Function

The Prism environment uses the concepts of *current file* and *current function*.

The current file is the source file currently displayed in the source window. The current function is the function or procedure displayed in the source window. You might change the current file or function if, for example, you want to set a breakpoint in a file that is not currently displayed in the source window, and you don't know the line number at which to set the breakpoint.

In addition, changing the current file and current function changes the scope used by the Prism environment for commands that refer to line numbers without specifying a file, as well as the scope used by the Prism environment in identifying variables; see “How the Prism Environment Chooses the Correct Variable or Procedure” on page 37 for a discussion of how the Prism environment identifies variables. The scope pointer (-) in the line-number region moves to the current file or current function to indicate the beginning of the new scope.

▼ To Change the Current File

- **Perform one of the following:**

- From the menu bar – Choose the File selection from the File menu. A window is displayed (shown in FIGURE 3-6), listing in alphabetical order the source files that make up the loaded program. Click on one, and it appears in the Selection box. Then click on OK and the source window updates to display the file. Or simply double-click, rapidly, on the source file. You can also edit the file name in the Selection box.

Note – The File window displays only files compiled with the `-g` switch.

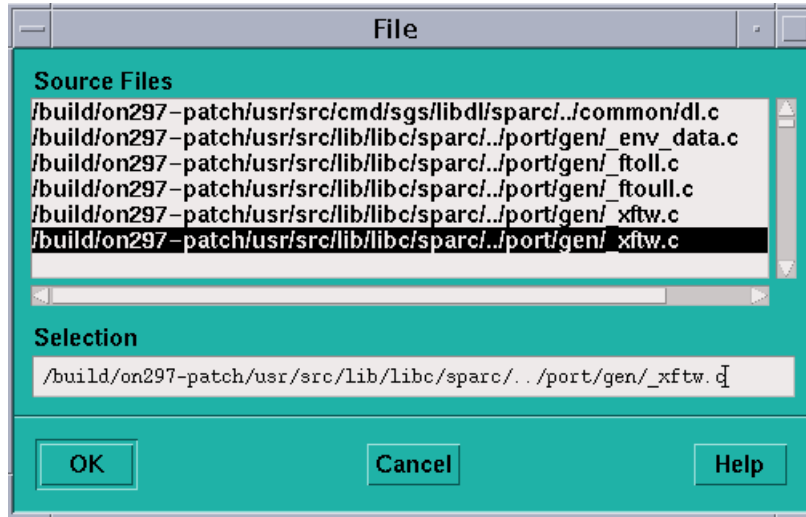


FIGURE 3-6 File Window

- From the command window – Issue the `file filename` command. The source window updates to display the file.

▼ To Change the Current Function or Procedure

- **Perform one of the following:**

- From the menu bar – Choose the Func selection from the File menu. A window is displayed, listing the functions in the program in alphabetical order. (Fortran procedure names are converted to all lowercase.) Click on one, and it appears in the Selection box. Then click on OK, and the source window updates to display the function. Or simply double-click on the function name in the list. You can also edit the function name in the Selection box.

By default, the Func window displays only functions in files compiled with the `-g` switch. To display all functions in the program, click on the Select All Functions button. The button then changes to Show `-g` Functions; click on it to return to displaying only the `-g` functions.

- From the command window – Issue the `func` command with the name of a function or subroutine as its argument. The source window updates to display the function.
- From the source window – Select the name of the function in the source window by dragging the mouse over it while pressing the Shift key. When you let go of the mouse button, the source window is updated to display the definition of this function.

Note – Include only the function name, not its arguments.

Note that if the function you choose is in a different source file from the current file, changing to this function also has the effect of changing the current file.

Creating a Directory List for Source Files

If the Prism environment cannot find a file—because you moved it or for some other reason—you can explicitly add its directory to the Prism environment’s search path.

▼ To Add a Directory to the Search Path

- **Perform one of the following:**
 - From the menu bar – Select Use from the File menu. This displays a dialog box, as shown in FIGURE 3-7. To add a directory, type its path name in the Directory box, then click on Add. To remove a directory, click on it in the directory list, then click on Remove.

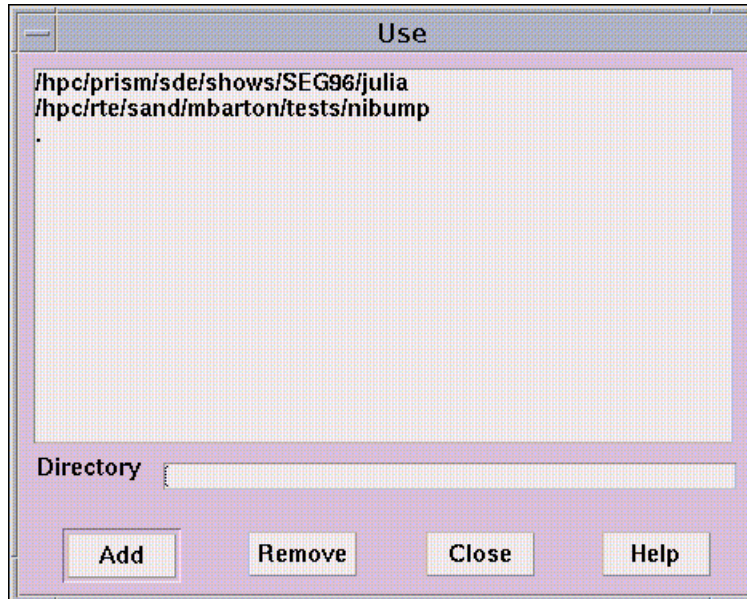


FIGURE 3-7 Use Dialog Box

- From the command window – Issue the `use` command on the command line. Specify a directory as an argument; the directory is added to the front of the search path. Issue `use` with no arguments to display the list of directories to be searched.

Note – No matter what the contents of your directory list are, the Prism environment searches for the source file first in the directory in which the program was compiled.

Debugging a Program

This chapter discusses how to debug message-passing programs in the Prism environment. It also describes how to use *events* to control the execution of a program.

Note that many principles that apply to debugging serial programs also apply to debugging message-passing programs. However, debugging a message-passing program can be considerably more complex than debugging a serial program, since you are in effect debugging multiple individual programs concurrently. The Prism environment's concept of psets lets you focus your debugging efforts on the processes that are of particular interest.

For information about debugging serial programs, see Appendix C.

This chapter is organized into the following sections:

- “Overview of Events” on page 93
- “Using the Event Table” on page 95
- “Setting Breakpoints” on page 104
- “Tracing Program Execution” on page 109
- “Displaying and Moving Through the Call Stack” on page 112
- “Combining Debug and Optimization Options” on page 120
- “Debugging Spawned Sun MPI Processes” on page 122
- “Examining the Contents of Memory and Registers” on page 125

Overview of Events

A typical approach to debugging is to stop the execution of a program at different points so that you can perform various actions, such as checking the values of variables. You stop execution by setting a *breakpoint*. If you perform a *trace*, execution stops, then automatically continues.

In the Prism environment, breakpoints and traces are referred to as *events*. Before the execution of a program begins, you can specify what events are to take place during execution. When an event occurs:

1. The execution pointer moves to the current execution point.
2. A message is printed in the command window.
3. If you specified that an action was to accompany the event, it is performed. An example of this might be to print a variable's value.
4. If the event is a trace, execution then continues. If it is a breakpoint, execution does not resume until you explicitly order it to.

The Prism environment provides various ways of creating these events—for example, by issuing commands or by using the mouse in the source window. “Setting Breakpoints” on page 104 describes how to create breakpoint events; “Tracing Program Execution” on page 109 describes how to create trace events. “Using the Event Table” on page 95 describes the *Event Table*, which provides a unified method for listing, creating, editing, and deleting events.

See “Events Taking Pset Qualifiers” on page 101 for a discussion of events in the Prism environment.

You can define events so that they occur:

- When the program reaches a certain point in its execution, such as at a specified line or function.
- When the value of a variable changes – For example, you can define an event that tells the Prism environment to stop the program when *x* changes value. This kind of event is sometimes referred to as a *watchpoint*. It slows execution considerably, since the Prism environment has to check the value of the variable after each statement is executed.
- At every line or assembly-language instruction.
- Whenever a program is stopped – An example of this would be to define an event that tells the Prism environment to print the value of *x* whenever the program stops.
- So that it occurs *only* if a specified condition is met – For instance, you can tell the Prism environment to stop at line 25 if *x* is not equal to 1. Like watchpoints, this kind of event slows execution.
- So that it occurs *only* after a condition has been met a specified number of times.

You can include one or more Prism commands as actions that are to take place as part of the event. One example of this would be to define an event that tells the Prism environment to stop at line 25, print the value of *x*, and do a stack trace.

Using the Event Table

The Event Table provides a unified method for controlling the execution of a program. Creating an event in any of the ways discussed later in this chapter adds an event to the list in this table. You can also display the Event Table and modify its contents directly by:

- Adding new events
- Deleting existing events
- Editing existing events

To display the Event Table, select Event Table from the Events menu.

This section describes the general process of using the Event Table.

Description of the Event Table

FIGURE 4-1 shows the Event Table.

The top area of the Event Table is the *event list*—a scrollable region in which events are listed. When you execute the program, the Prism environment uses the events in this list to control execution. Each event is listed in the format you would use to enter it as a command in the command window. It is prefaced by an ID number assigned by the Prism environment, which is 1 in the FIGURE 4-1 example.

The middle area of the Event Table is a series of fields that you fill in when editing or adding an event; not all of the fields are relevant to every event. The fields are:

- ID – This is an identification number associated with the event. You cannot edit this field.

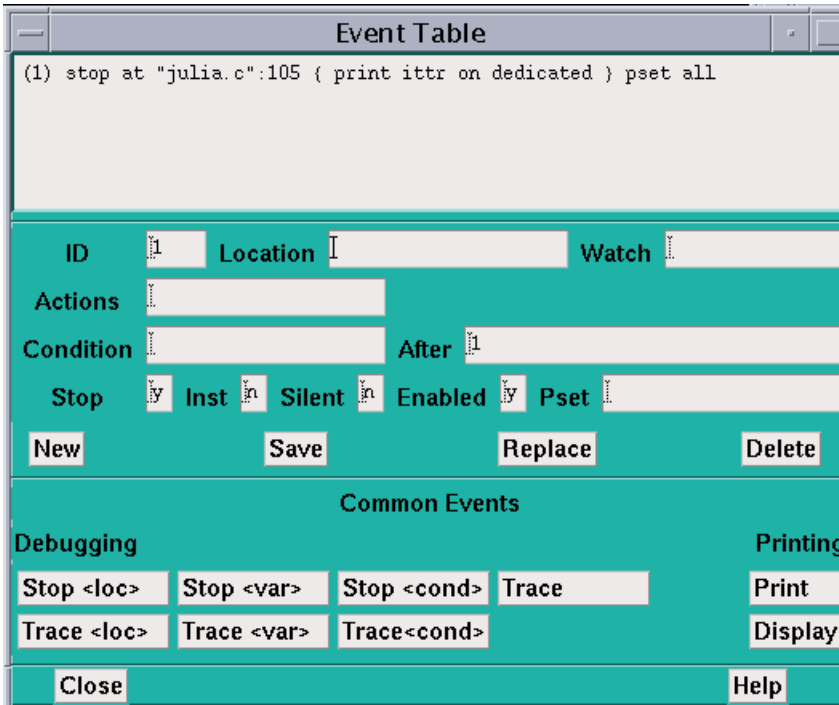


FIGURE 4-1 Event Table

- Location – Use this field to specify the location in the program at which the event is to take place. Use the syntax *"filename":line-number* to identify the source file and the line within this file. If you specify only the line number, the Prism environment uses the current file. There are also three keywords you can use in this field:
 - Use *eachline* to specify that the event is to take place at each line of the program; this is the default.
 - Use *eachinst* to specify that the event is to take place at each assembly-language instruction.
 - Use *stopped* to specify that the event is to take place whenever the program stops execution.
- Watch – Use this field to specify a variable or expression whose value is to be watched; the event takes place if the value of the variable or expression changes. (If the variable is an array or a parallel variable, the event takes place if the value of any element changes.) This slows execution considerably.

- **Actions** – Use this field to specify the action(s) associated with the event. Most Prism commands can be used in this field. Separate multiple commands with semicolons. (The commands that you cannot include in the Actions field are `attach`, `core`, `detach`, `load`, `return`, `run`, and `step`.)
- **Condition** – Use this field to specify a logical condition that must be met if the event is to take place. The logical condition can be any language expression that evaluates to true or false. See “Writing Expressions in the Prism Environment” on page 37 for more information. Specifying a condition slows execution considerably, unless you also specify a location at which the condition is to be checked.
- **After** – Use this field to specify how many times a triggering condition is to be met before the event is to take place. The Event Table updates during execution to show the current count (that is, how many times are left for the triggering condition to be met before the event is to take place). Once the event takes place, the count is reset to the original value. The default setting is 1, and the event takes place each time the condition is met. See “Overview of Events” on page 93 for a discussion of triggering conditions.
- **Stop** – Use this field to specify whether or not the event is to halt execution of the program. Putting a `y` in this field creates a breakpoint event; putting an `n` in this field creates a trace event.
- **Inst** – Set this field to `y` to cause a disassembled assembly-language instruction to be displayed when the event occurs. Specify `n` to prevent such displays.
- **Silent** – Use this field to specify whether or not the event is to cause a message to appear in the command window when it occurs.
- **Enabled** – Use this field to specify whether the event is enabled. Putting an `n` in this field disables the event; it still exists, but it does not affect program execution.
- **Pset** – Use this field to specify the intended pset (for events that take pset qualifiers).

The buttons beneath these fields are for use in creating and deleting events; they are described below.

The area headed **Common Events** contains buttons that provide shortcuts for creating certain standard events.

Click on **Close** or press the **Esc** key to cancel the Event Table window.

Adding an Event

You can either add an event explicitly, editing field by field, or you can use the **Common Events** buttons to automatically fill in some of the fields for you. You can add an event from the beginning if it is not similar to any of the categories covered by the **Common Events** buttons.

▼ To Add an Event, Editing Field by Field

- 1. Click on the New button.**

All values currently in the fields are cleared.

- 2. Fill in the relevant fields to create the event.**

- 3. Click on the Save button to save the new event.**

It appears in the event list.

▼ To Add an Event, Using Common Events Buttons

- 1. Click on the button for the event you want to add—for example, Print.**

This fills in certain fields and highlights all fields that you need to fill in.

- 2. Fill in the highlighted field(s).**

You can also edit other fields if you like.

- 3. Click on Save to add the event to the event list.**

Most of these Common Events buttons are also available as separate selections in the Events menu. This lets you add one of these events without having to display the entire Event Table. The menu selections, however, prompt you only for the field(s) you must fill in. You cannot edit other fields.

Individual Common Events buttons are discussed throughout the remainder of this guide.

You can also create a new event by editing an existing event; see “Editing an Existing Event” on page 99.

Deleting an Existing Event

You can delete events using the Event Table or the Delete selection from the Events menu.

▼ To Delete an Existing Event

- 1. Click on the line representing the event in the Event Table or move to it with the up and down arrow keys.**

This causes the components of the event to be displayed in the appropriate fields beneath the list.

2. Click on the Delete button.

You can also select Delete from the Events menu to display the Event Table. You can then follow the procedure described above.

Deleting a breakpoint at a program location also deletes the B in the line-number region at that location.

Editing an Existing Event

You can edit an existing event to change it or to create a new event similar to it.

▼ To Edit an Existing Event

1. Click on the line representing the event in the event list or move to it with the up or down arrow keys.

This causes the components of the event to be displayed in the appropriate fields beneath the list.

2. Edit these fields.

You can, for example, change the Location field to specify a different location in the program.

3. Save the newly edited event.

Click on the Save button to save the new event *in addition to* the original version of the event; it is given a new ID and is added to the end of the event list. Clicking on Save is a quick way of creating a new event similar to an event you have already created.

Disabling and Enabling Events

You can disable and enable events. When you disable an event, the Prism environment keeps it in the event list, but it no longer affects execution. You can subsequently enable it when you once again want it to affect execution. This can be more convenient than deleting events and then redefining them.

▼ To Disable an Event

● Perform one of the following:

- From the Event Table – The Event Table has an Enabled field. By default, there is a y in this field, meaning that the event is enabled. Click on the field and change the y to an n to disable the event. The event remains in the event list, but is

labeled (disabled). You can then edit the event as described in “Editing an Existing Event” on page 99 and change the field back to a `y` to enable the event once again.

- From the command line – Issue the `disable` command to disable an event. Use the event’s ID as the argument. You can obtain this ID from the event list in the Event Table or by issuing the `show events` command.

In the following example, the sequence of commands displays the event list, then disables an event, and then redisplay the event list:

```
(prism all) show events
(1) trace
(2) when stopped { print board }
(prism all) disable 1
event 1 disabled
(prism all) show events
(1) trace (disabled)
(2) when stopped { print board }
```

▼ To Enable an Event

- **Type:**

```
(prism all) enable event_ID
```

This re-enables *event_ID*.

Saving Events

Events that you create for a program are automatically maintained when you reload the same program during a Prism session. This saves you the effort of redefining these events each time you reload a program.

Note these points:

- The Prism environment prints a warning message if it can’t maintain an event. This would happen, for example, if an event is supposed to occur at a source line that no longer exists.
- Changing a program can also change the meaning of events. A breakpoint set at line 32, for example, may still be a valid event, but it may not be the event you want if you have deleted lines earlier in the program.
- Disabled events become enabled when a program is reloaded.
- Events are deleted when you leave the Prism environment.

▼ To Save Events to a File

You can use Prism commands to save your events to a file and then execute them from the file rather than interactively.

1. Redirect the output to a file.

For example, the following redirects the list of events to the file `primes.events`:

```
(prism all) show events @ primes.events
```

2. Edit `primes.events` to remove the ID number at the beginning of each event.

This leaves you with a list of Prism commands.

3. Type:

```
(prism all) source primes.events
```


This reads in and executes the commands from `primes.events`.

Events Taking Pset Qualifiers

Events in the Prism environment can take a pset qualifier.

▼ To Specify a Pset Qualifier

- Type the pset name in the Pset field in the Event Table, as shown in FIGURE 4-2.



The image shows a screenshot of the Prism Environment's Event Table. The table has several rows and columns. The first row contains 'ID', 'Location', and 'Watch'. The second row is 'Actions'. The third row is 'Condition', with an 'After' field to its right. The fourth row is 'Stop', with checkboxes for 'Inst', 'Silent', and 'Enabled', and a 'Pset' field containing the text 'cycle'. At the bottom of the table are four buttons: 'New', 'Save', 'Replace', and 'Delete'.

FIGURE 4-2 Pset Field in the Prism Environment's Event Table

If you do not supply a pset qualifier, the event applies to the current pset.

In the following example, the current pset is `all`.

```
(prism all) stop in receive pset notx
```

Because the pset `notx` is specified, this command sets a breakpoint in the `receive` routine for the processes in the set `notx`. Each process in `pset notx` stops when it reaches this routine. It is possible, of course, that some processes may never reach this routine. This might become an issue when you include actions in an event.

The following command stops execution for any process in the current pset if the process's value for the variable `x` is greater than 10.

```
(prism all) stop if x > 10
```

Because no other pset was specified in this example, this event applies to the current pset, which is `all`. The Prism environment evaluates the expression in the condition locally—that is, separately for each process. Similarly, if `a` and `b` are arrays, the following command stops execution for a process in the current set if the sum of the values of `a` in that process is greater than the sum of the values of `b`:

```
(prism all) stop if sum(a) > sum(b)
```

All processes that are stopped at breakpoints are members of the predefined pset `break`.

▼ To Continue All the Processes in a Pset

- **Type:**

```
(prism all) cont
```

The following command causes the processes in pset `notx` to continue running:

```
(prism all) cont pset notx
```

Events and Dynamic Psets

If you use a dynamic pset as a qualifier for an event, its membership is evaluated when you issue the command defining the event. Thus, the following command creates a breakpoint only in the processes that are interrupted at the time the command is issued:

```
(prism all) stop at 10 pset interrupted
```

If no processes are currently interrupted, you will receive an error message.

One result of this is that you cannot define events that involve dynamic psets before the program starts execution.

Events and Variable Psets

If you specify a user-defined variable pset as a qualifier, its membership is determined by the most recent `eval pset` command issued for that pset.

As is the case with dynamic psets, you cannot define events that involve variable psets before the program starts execution.

Actions in Events

Events in the Prism environment can take action clauses. For example, the following action clause prints `x` for the pset `foo` when the members of `foo` are stopped at line 10:

```
(prism all) stop at 10 {print x} pset foo
```

Note – Associating an action with an event forces a global synchronization at the breakpoint or tracepoint. In the example above, every process in pset `foo` must stop at line 10 before `x` can be printed. If a member does not stop at line 10, the action never takes place. In a trace event, all processes in the pset must stop at the specified place and synchronize; the action then takes place, and the processes automatically continue execution.

You can include an `eval pset` command as an event action. For example, this evaluates the pset `sending` when all the members of the current pset are stopped in `send`:

```
(prism all) stop in send {eval pset sending}
```

You receive error messages if it is impossible to evaluate membership in a pset. This would happen, for example, if a variable in the set definition is not active.

Note these limitations in using event actions:

- You cannot include the following commands that manipulate psets:
 - `define pset`
 - `delete pset`
 - `process`
 - `pset`
- You cannot include a pset qualifier in the action. The command in the action clause takes its pset from the pset of the event.
- You cannot include commands that affect program execution; these are
 - `cont` and `contw`
 - `run`
 - `step` and `stepi`
 - `next` and `nexti`
 - `wait`
- You cannot include the `load`, `reload`, `return`, and `core` commands.
- You cannot use an unbounded pset as the context for an event specification. For information about unbounded psets, see “Using Unbounded Psets in Commands” on page 83.

▼ To Display Events by Process

- **Type:**

```
(prism all) show events (processnumber)
```

This displays all events associated with the specified process.

Issuing `show events` with no arguments has its standard behavior. That is, it prints out all events, as shown in the following example:

```
(prism all) show events  
(1) trace  
(2) when stopped { print board }  
(prism all) disable 1  
event 1 disabled  
(prism all) show events  
(1) trace (disabled)  
(2) when stopped { print board }
```

Events and Deleted Psets

If you create an event that applies to a particular pset and subsequently delete the pset, the event continues to exist. Its printed representation, however, is changed so that it shows the processes that were members of the pset at the time you deleted the set.

Setting Breakpoints

A *breakpoint* stops execution of a program when a specific location is reached, if a variable or expression changes its value, or if a certain condition is met. This section describes the methods available in the Prism environment for setting a breakpoint.

You can set a breakpoint in the following ways:

- By using the line-number region
- By using the Event Table and the Events menu
- By issuing the command `stop` or `when` in the command window

The line-number region is easiest for setting simple breakpoints. However, the other two methods give you greater flexibility, such as in setting up a condition under which the breakpoint is to take place.

In all cases, an event is added to the list in the Event Table. If you delete the breakpoint using any of the methods described in this section, the corresponding event is deleted from the event list. If you set a breakpoint at a program location, a B appears next to the line number in the line-number region.

Note – Secondary (spawned) Prism sessions do not inherit breakpoints set within primary Prism sessions.

Using the Line-Number Region

To use the line-number region to set a breakpoint, the line at which you want to stop execution must appear in the source window. If it does not, you can scroll through the source window (if the line is in the current file) or use the File or Func selection from the File menu to display the source file you are interested in.

▼ To Set a Breakpoint in the Line-Number Region

- 1. Position the mouse pointer to the right of the line numbers.**

The pointer turns into a B.

- 2. Move the pointer next to the line at which you want to stop execution.**

- 3. Left-click the mouse.**

A B is displayed, indicating that a breakpoint has been set for that line.

A message appears in the command window confirming the breakpoint, and an event is added to the event list.

The source line you choose must contain executable code; if it does not, you receive a warning in the command window, and no B appears where you clicked.

- 4. Shift-click on the letter in the line-number region to display the complete event (or events) associated with it.**

See “Using the Line-Number Region” on page 31 for more information on the line-number region.

▼ To Delete Breakpoints Using the Line-Number Region

- Left-click on the B that represents the breakpoint you want to delete.**

The B disappears; a message appears in the command window, confirming the deletion.

What Happens in a Split Source Window

As described in “Moving Through the Source Code” on page 28, you can split the source window to display source code and the corresponding assembly code.

You can set a breakpoint in either pane. The **B** appears in the line-number region of both panes, unless you set the breakpoint at an assembly code line for which there is no corresponding source line.

Deleting a breakpoint from one pane of the split source window deletes it from the other pane as well.

Using the Event Table and the Events Menu

▼ To Set a Breakpoint Using the Event Table

1. Select **Stop <loc>** or **Stop <var>** from the Events menu.

These choices are also available as Common Events buttons within the Event Table itself; see “Adding an Event” on page 97.

2. Perform one of the following:

- **Stop <loc>** prompts for a location at which to stop the program. If you want to set the breakpoint at a particular line in the program, enter the line number. Or, to set the breakpoint at the first line of a function or procedure, enter the function or procedure name instead.

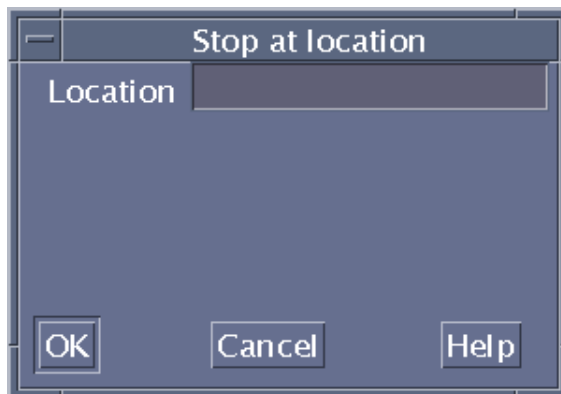


FIGURE 4-3 Stop <loc> Dialog Box

- Stop <var> prompts for a variable name. The program stops when the variable's value changes. The variable can be an array, in which case execution stops whenever any element of the array changes. This slows execution considerably.
- Stop <cond> (available as a Common Events button on the Event Table) prompts for a condition, which can be any expression that evaluates to true or false. The program stops when the condition is met. This slows execution considerably.

See “Writing Expressions in the Prism Environment” on page 37 for more information on expressions.

You can also use the Event Table to create combinations of these breakpoints; for example, you can create a breakpoint that stops at a location if a condition is met. In addition, you can use the Actions field of the Event Table to specify the Prism commands that are to be executed when execution stops.

▼ To Delete Breakpoints Using the Event Table

- **Perform one of the following:**

- From the Events menu, select Delete.
- From the Event Table, click on the Delete button.

For more information about deleting events, see “Deleting an Existing Event” on page 98.

Setting a Breakpoint Using Commands

▼ To Set a Breakpoint Using Commands

- **Type:**

`(prism all) stop`

or

`(prism all) when`

The `when` command is an alias for the `stop` command.

The syntax of the `stop` command is also used by the `stopi`, `when`, `trace`, and `tracei` commands. The general syntax for these commands is:

`command [variable | at line | in func] [if expr] [{cmd[: cmd...]}] [after n]`

where

- *command* – Is the name of a command, which can be `stop`, `stopi`, `when`, `trace`, or `tracei`.

- *variable* – Is the name of a variable. The command is executed if the value of the variable changes. If the variable is an array, an array section, or a parallel variable, the command is executed if the value of any element changes. This form of the command slows execution considerably. You cannot specify both a variable and a program location.
- *line* – Specifies the line number where the stop or trace is to be executed. If the line is not in the current file, use the format:
at *filename:line-number*
- *func* – Is the name of the function or procedure in which the stop or trace is to be executed.
- *expr* – Is any language expression that evaluates to true or false. This argument specifies the logical condition, if any, under which the stop or trace is to be executed. For example, the following expression evaluates to true whenever variable *a* is greater than 1.

```
if a .GT. 1
```

This form of the command slows execution considerably, unless you combine it with the *at line* syntax. See “Writing Expressions in the Prism Environment” on page 37 for more information on writing expressions in the Prism environment.

- *cmd* – Is any Prism command (except *attach*, *core*, *detach*, *load*, *return*, *run*, or *step*). This argument specifies the actions, if any, that are to accompany the execution of the *stop* or *trace*. For example, `{print a}` prints the value of *a*. If you include multiple commands, separate them with semicolons.
- *n* – Is an integer that specifies how many times a triggering condition is to be reached before the *stop* or *trace* is executed; see “Overview of Events” on page 93 for a discussion of triggering conditions. This is referred to as an *after count*. The default is 1. Once the *stop* or *trace* is executed, the count is reset to its original value. Note that if there is both a condition and an after count, the condition is checked first.

The first option listed (specifying the location or the name of the variable) must come first on the command line. The other options, if you include them, can be in any order.

For the *when* command, you can use the keyword *stopped* to specify that the actions are to occur whenever the program stops execution.

When you issue the command, an event is added to the event list. If the command sets a breakpoint at a program location, a *B* appears in the line-number region next to the location.

Examples of the `stop` Command

To stop execution the tenth time in function `foo` and print `a`, type:

```
(prism all) stop in foo {print a} after 10
```

To stop at line 17 of file `bar` if `a` is equal to 0, type:

```
(prism all) stop at "bar":17 if a == 0
```

To stop whenever `a` changes, type:

```
(prism all) stop a
```

To stop the third time `a` equals 5, type:

```
(prism all) stop if a .eq. 5 after 3
```

To print `a` and do a stack trace every time the program stops execution, type:

```
(prism all) when stopped {print a; where}
```

▼ To Set a Breakpoint Using Machine Instructions

● Type:

```
(prism all) stopi at machine-address
```

For example, the following command stops execution at address 1000 (hex):

```
(prism all) stopi at 0x1000
```

The history region displays the address and the machine instruction. The source pointer moves to the source line being executed.

▼ To Delete Breakpoints Using the Command Window

1. Type:

```
(prism all) show events
```

This prints out the event list. Each event has an ID number associated with it.

2. Type:

```
(prism all) delete ID [ID ...]
```

List the ID numbers of the events you want to delete; separate multiple IDs with one or more blank spaces. For example, this deletes the events with IDs 1 and 3. Use the argument `all` to delete all existing events:

```
delete 1 3
```

Tracing Program Execution

You can trace program execution by using the Event Table or Events menu or by issuing commands. All methods add an event to the Event Table. If you trace a source line, the Prism environment displays a `T` next to the line in the line-number region.

Tracing is essentially the same as setting a breakpoint, except that execution continues automatically after the breakpoint is reached. When tracing source lines, the Prism environment steps into procedures if they were compiled with the `-g` option; otherwise it steps over them as if it had issued a `next` command.

Using the Event Table and Events Menu

▼ To Trace Program Execution Using the Event Table and the Events Menu

- **Select Trace, Trace <loc>, or Trace <var> from the Events menu.**

These choices are also available as Common Events buttons within the Event Table itself.

- `Trace` displays source lines in the command window before they are executed.
- `Trace <loc>` prompts for a source line. The Prism environment displays a message immediately prior to the execution of this source line.
- `Trace <var>` prompts for a variable name. A message is printed when the variable's value changes. The variable can be an array, an array section, or a parallel variable, in which case a message is printed any time any element changes. This slows execution considerably.
- `Trace <cond>` (available as a Common Events button) prompts for a condition, which can be any expression that evaluates to true or false; see “Writing Expressions in the Prism Environment” on page 37 for more information on writing expressions. The program displays a message when the condition is met. This also slows execution considerably.

For variations of these traces, you can create your own event in the Event Table. You can also use the Actions field to specify Prism commands that are to be executed along with the trace.

▼ To Delete Traces

- **Choose the Delete selection from the Events menu, or use the Delete button in the Event Table.**

For more information about deleting existing events, see “Deleting an Existing Event” on page 98.

Using the Command Window

▼ To Trace Program Execution Using Commands

- **Type:**

```
(prism all) trace
```

Issuing `trace` with no arguments causes each source line in the program to be displayed in the command window before it is executed.

The `trace` command uses the same syntax as the `stop` command; see “Setting a Breakpoint Using Commands” on page 107. For example:

To trace and print `a` on every source line, type:

```
(prism all) trace {print a}
```

To trace line 17 if `a` is greater than 10, type:

```
(prism all) trace at 17 if a .GT. 10
```

In addition, the Prism environment interprets these two commands as being the same:

```
(prism all) trace line-number
```

```
(prism all) trace at line-number
```

▼ To Trace Machine Instructions

- **Type:**

```
(prism all) tracei address
```

When tracing machine instructions, the Prism environment follows all procedure calls down. The `tracei` command has the same syntax as the `stop` command; see “Setting a Breakpoint Using Commands” on page 107.

The history region displays the address and the machine instruction. The execution pointer moves to the next source line to be executed.

▼ To Delete Traces

1. **Type:**

```
(prism all) show events
```

This obtains the ID associated with the trace.

2. **Type:**

```
(prism all) delete ID
```

For further information, see “Setting a Breakpoint Using Commands” on page 107.

Displaying and Moving Through the Call Stack

The *call stack* is the list of procedures and functions currently active in a program. The Prism environment provides you with methods for examining the contents of the call stack.

See “Displaying the Where Graph” on page 113 for a discussion of displaying the call stack graphically in the Prism environment.

▼ To Display the Call Stack

- **Perform one of the following:**
 - From the menu bar – Select Where from the Debug menu. The Where window is displayed; see FIGURE 4-4. The window contains the call stack. It is updated automatically when execution stops or when you issue commands that change the stack.

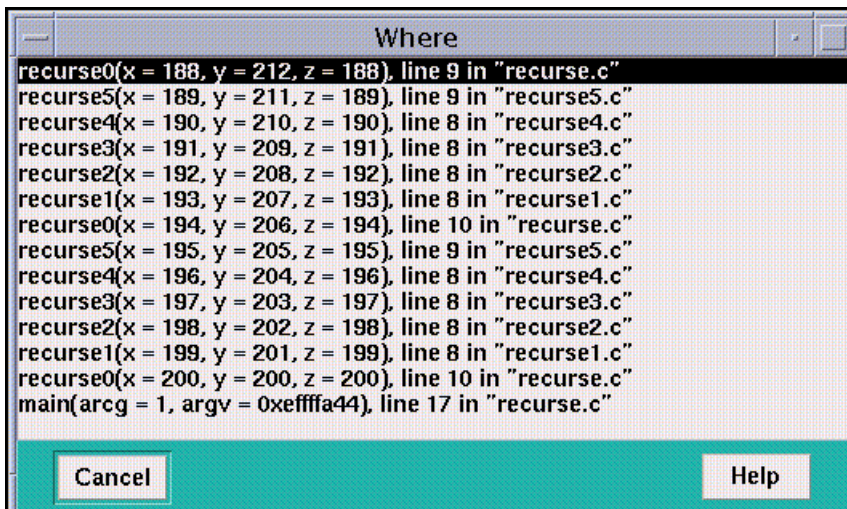


FIGURE 4-4 Where Window

- From the command window – Type `where` on the Prism command line. If you include a number, it specifies how many active procedures are to be displayed; otherwise, all active procedures are displayed in the history region.
- From the command window – Type `where on snapshot` on the Prism command line to put the history-region output into a window.

Values of arguments in displayed procedures are shown in the default radix, which is decimal unless you change it via the `set $radix` command; see “To Change the Default Radix” on page 131.

Moving Through the Call Stack

Moving *up* through the call stack means heading toward the main procedure. Moving *down* through the call stack means heading toward the current stopping point in the program.

Moving through the call stack changes the current function and repositions the source window at this function. It also affects the scope that the Prism environment uses for interpreting the names of variables in expressions and commands. See “Scope in the Prism Environment” on page 79 for more information.

▼ To Move Through the Call Stack

- **Perform one of the following:**
 - From the menu bar – Choose Up or Down from the Debug menu. Up moves up one level in the call stack; Down moves down one level. These selections are available by default in the tear-off region.
 - From the command window – Enter `up` or `down` on the command line to move up or down one level. To move more than one level, specify an integer argument.
 - From the Where window – If the Where window is displayed, clicking on a function in it changes the stack level to make that function current.

Displaying the Where Graph

Selecting Where from the Debug menu displays the call stacks for the program being debugged. A multiprocess program can have multiple call stacks, one for each process. A threaded program can have a separate stack for each thread in each process.

To show the relationships among these call stacks, the Prism environment provides a *Where graph*; this window displays a snapshot of the dynamic call graph of the program. The Where graph displays information about all processes that are not running.

▼ To Display the Where Graph

- **Perform one of the following:**

- From the menu bar – Choose Where from the Debug menu.
- From the command line – Type `where` on dedicated.

A window like the one shown in FIGURE 4-5 is displayed.

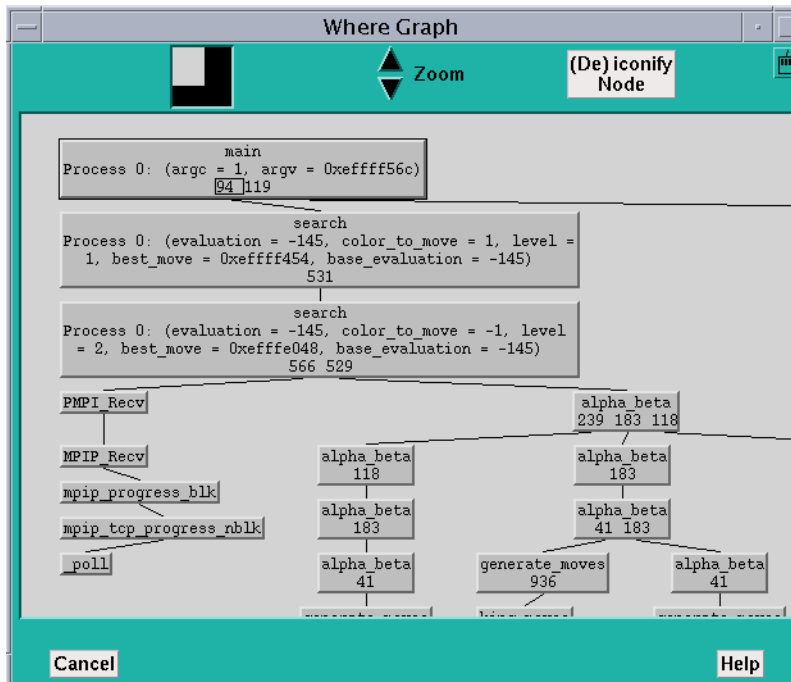


FIGURE 4-5 Where Graph

The Where graph centers on the current process of the current pset. That is, the processes related to it are lined up in a single column. In FIGURE 4-5, process 0 is the current process. If you change the current process, the Where graph rearranges itself. The default zoom level of the Where graph shows the arguments for the current process.

The line numbers at the bottom of each box indicate where processes branch.

▼ To Display Processes Containing a Specific Function in Their Call Stacks

- **Shift-click in each function's box.**

This displays a pop-up window showing the numbers of the processes with this function in their call stack, along with their arguments.

Panning and Zooming in the Where Graph

As FIGURE 4-6 shows, the Where graph can get quite large, so the Prism environment provides methods for panning through it and zooming in and out.

The white box in the navigator rectangle at the top of the window shows the position of the display area relative to the entire Where graph.

▼ To Move the Position Displayed in the Where Graph

- **Perform one of the following:**
 - Drag the box.
 - Click at a spot in the navigator.

The box moves to that spot, and the window shows the Where graph in this area of the total display.

▼ To Display More of the Where Graph

- **Click on the Zoom down arrow to the right of the navigator.**

This reduces the size of the boxes representing the functions and removes information. FIGURE 4-6 shows the Where graph of FIGURE 4-5, zoomed out one level. Note that the information about the current process's arguments is gone.

As you zoom further out, the Where graph removes the line numbers, and one more level after that removes the function names, leaving only boxes connected by lines.

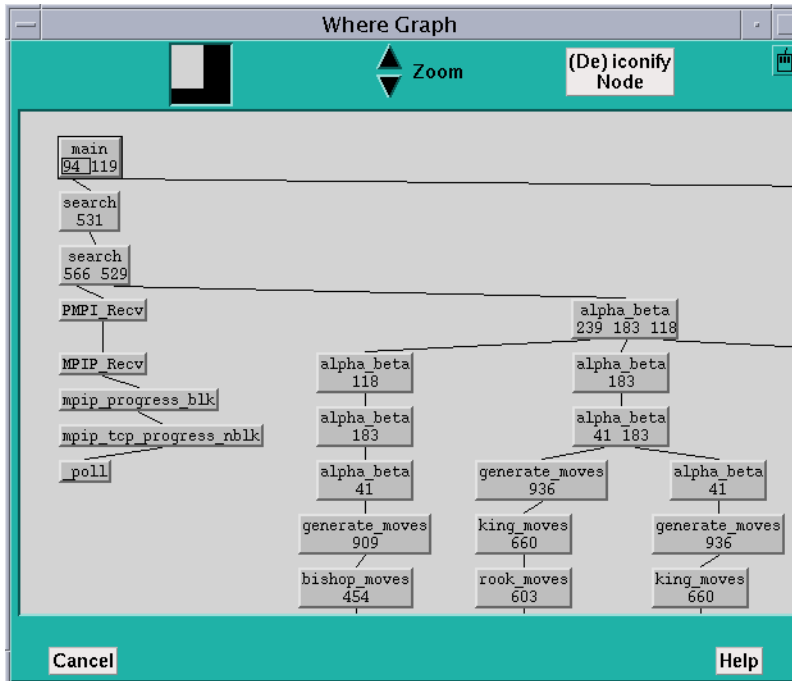


FIGURE 4-6 Where Graph, Zoomed Out One Level

▼ To Display Additional Information About a Box in the Where Graph

- Shift-click on a box to display information about it.

If your program is multithreaded, its call stacks are not rooted at `main`. Thus, at maximum zoom, the Where graph displays the call stacks as multiple trees, as shown in FIGURE 4-7.

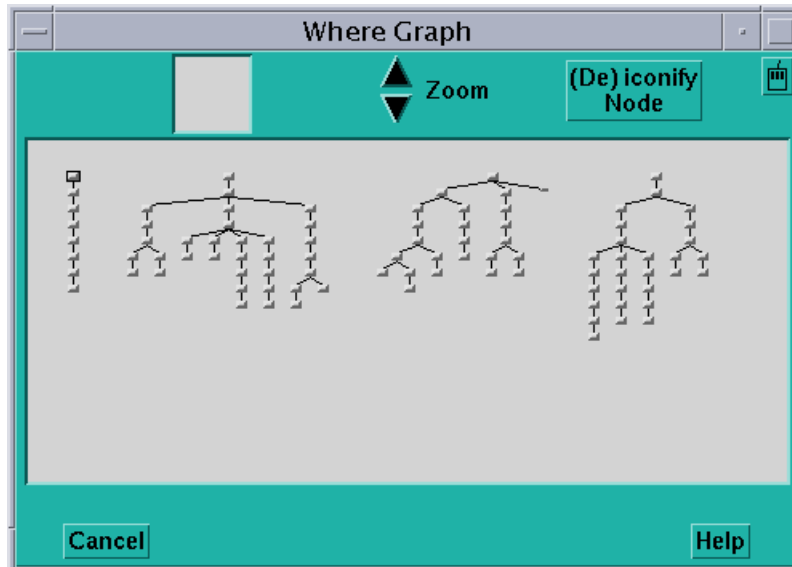


FIGURE 4-7 Where Graph, Zoomed Out to the Maximum

▼ To Increase the Size of the Where Graph's Function Boxes

- Click on the Zoom up arrow.

This increases the size of the function boxes and includes more information in them. FIGURE 4-8 shows the Where graph of FIGURE 4-5, zoomed in. In this case, the Where graph shows, for each function, the processes that have that function in their call stack. As in the Psets window, the processes are represented as bitmaps of cells, with numbering starting at the upper left, increasing from left to right and then jumping to the next row.

If your Where graph displays a threaded program, you can zoom in to the level shown in FIGURE 4-9.

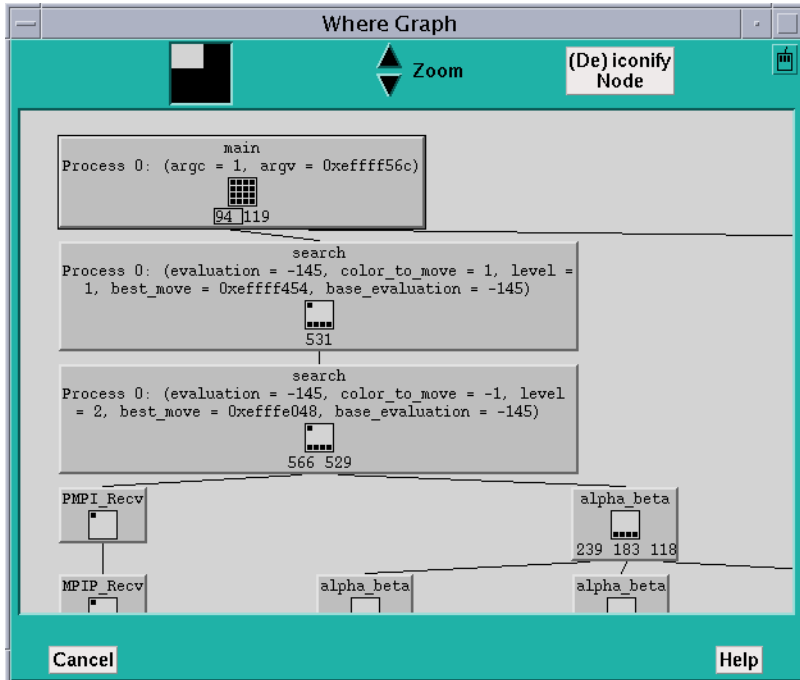


FIGURE 4-8 Where Graph, Zoomed In

Zooming in another level shows all arguments for all processes.

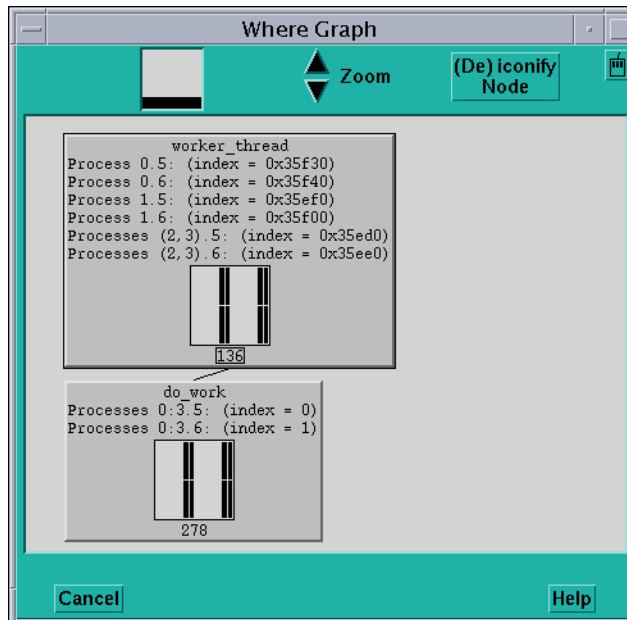


FIGURE 4-9 Where Graph of a Threaded Program, Zoomed in to Show Thread Stripes

▼ To View Information About Individual Threads

- **Shift-click on the individual stripes.**

This displays information about the corresponding threads.

▼ To Shrink Selected Portions of the Where Graph

You can shrink selected portions of the Where graph. This is useful if you want to see the overall structure of the graph but also want to focus on certain functions.

- **Perform one of the following:**
 - Middle-click on a function to iconify it and all of its children. Middle-click on an iconified function to re-expand it and its children to the current zoom level.
 - Alternatively, you can click on the (De) iconify Node button next to the Zoom arrows at the top of the Where graph. This changes the mouse pointer to a target. You can then left-click on a function to iconify it and its children. If it is already iconified, left-clicking on it will re-expand it and its children. To cancel the operation, left-click anywhere outside of the boxes surrounding the functions.

▼ To Move Through the Where Graph

When you first display the Where graph, the `main` function is highlighted.

- **Left-click on a function to highlight it. Or, move through the Where graph using the keyboard:**
 - Use the up arrow key to move to the parent of the highlighted function.
 - If line numbers are visible in the highlighted function, by default the leftmost number is selected by having a box drawn around it. Use the left and right arrows to select other line numbers in the function. You can then use the down arrow key to highlight the function called at the selected line.

▼ To Make a Function the Current Pset

- **Press the spacebar while in the Where graph.**

The following actions occur:

- The current function changes to the function that is highlighted in the Where graph.
- The highlighted function in the source window is displayed.
- A new current pset is created, with the same name as the function and containing the processes with this function in their call stack. The current process of this current set is the lowest-numbered process in the set.

Combining Debug and Optimization Options

When you use the Prism environment on programs that have been compiled with optimization options, Prism commands behave differently and the visibility of variables in the optimized programs changes.

Interpreting Interaction Between an Optimized Program and the Prism Environment

When the control flow is inside a routine that has been compiled with both `-g` and an optimization option (a debuggable optimized routine), the `next` and `step` commands change their behavior:

- next steps out of the current routine and stops in the next debuggable routine that differs from the original routine.
- step stops in the next debuggable routine (including recursive calls of the original routine).

You can set breakpoints using the `stop at` command inside debuggable optimized routines only at the first line of such a routine. If the routine name is `foo` and the first instruction in `foo` is `ADDR_INSTR`, then the breakpoint is set as if you had used `stop in foo` or `stopi at ADDR_INSTR`.

Note that the following commands are unaffected:

- `nexti`
- `stepi`
- `stopi`

When either `return` or `stepout` is used to return control flow to a debuggable optimized routine, the Prism environment assumes that the current position is at the first line of the current routine. The Prism environment makes the same assumption when the source file position is updated as a result of `up` or `down` commands that result in a debuggable optimized routine.

Accessing Variables in Optimized Routines

Due to the effects of optimization on variable locations in executable programs that have been compiled with optimization, the Prism environment cannot access all variables at all times.

The accessibility of variables can be defined by whether the variables can be used in expressions that require the right value of the variable (such as `print X` or `call foo(X)`) or the left value of the variable (such as `assign X=1`).

The limits of accessibility can be described by the flow of control in an optimized program. When the flow of control is in a routine compiled with both `-g` and an optimization flag, the following conditions apply:

- If the control flow is at the first machine instruction of the routine (which has not yet been executed), then all global variables and the routine's arguments are accessible. No other local variable is accessible.
- If the first machine instruction of the current routine has already been executed, then only the global variables are accessible. No local variable is accessible.

The following commands can use only accessible variables:

- `assign`
- `call`
- `display`
- `dump`

- print
- trace
- tracei
- varsave
- when
- where

The `where` command reports all active stack frames that have a stack pointer. The `where` command does not report routines that have no frame pointer and routines that have been inlined.

Note – The `where` stack displays values only for accessible arguments and ‘???’ for all others.

Debugging Spawned Sun MPI Processes

When debugging Sun MPI jobs that spawn other Sun MPI jobs, you should be especially careful to ensure that Sun MPI or Prism processes do not exit while other processes depend on communicating with them.

For example, suppose MPI job `foo` spawns MPI job `bar`, job `foo` uses `MPI_Send` to communicate with a process in job `bar`, and job `bar` uses `MPI_Recv` to handle a message from job `foo`.

If you are debugging both jobs in the Prism environment and you issue the Prism `quit` command in the primary Prism session (`foo`) before the process in `foo` calls the `MPI_Send` function, then job `foo` will exit. However, `bar` (which you are still debugging in a secondary Prism session) cannot continue past the `MPI_Recv` call, because `foo` has already exited.

If you issue a `quit -all` command in the primary Prism session while debugging a job that has many deeply nested `MPI_Comm_spawn` calls, it may not terminate all spawned secondary Prism sessions. To terminate a secondary debug session, you must manually issue the `quit` command in the secondary Prism session(s).

Debugging Spawned Sessions Using the Commands-Only Interface

When the Prism environment is started with the `-CX` option, it opens new X terminal windows in response to the spawning of new processes. It labels a new window with the title `about:jid`, where `jid` is the job ID of the spawned process.

You must set the `DISPLAY` variable if you debug programs with calls to `MPI_Comm_spawn()` or `MPI_Comm_spawn_multiple()`, even when launching the Prism environment with the commands-only interface. For more information about the commands-only interface, see Appendix A.

Prism Commands With Special Functions in Spawned Sessions

Several Prism commands perform special functions in spawned Prism sessions.

- `attach` - The Prism `attach` command enables you to attach to an executable without issuing a prior `load` command. You can simply attach to the process ID or job ID, as follows:

```
(prism all) attach jid
```

The `attach` command will clean up the current session before attaching to the `jid` (job ID) specified in the command.

The `attach` command does not accept multiple job IDs.

However, if the specified job ID is a result of an `MPI_Comm_spawn_multiple()`, multiple Prism sessions get created.

- `detach` - The `detach` command only applies to the Prism session where it is invoked. If you issue the `detach` command in a primary session, it is not propagated down to secondary sessions.
- `run` and `rerun` - When you issue the `run` or `rerun` commands in the primary Prism session, the Prism environment will clean up all the secondary Prism sessions. That is, the Prism environment will shut down the secondary Prism sessions and the debuggees.

The `run` and `rerun` commands are not valid in the secondary Prism sessions.

- `kill` - If you issue a `kill` command in a primary Prism session, the command propagates to the secondary Prism sessions. That is, the Prism environment shuts down the secondary Prism sessions and the debuggees.
- `quit` - The `quit` command does not propagate down to the secondary sessions unless you issue the command with the `-all` option.

To quit all Prism sessions, you must type:

```
(prism all) quit -all
```

If the job was run by the primary Prism session, the command `quit -all` kills the debuggees in the primary as well as the secondary Prism sessions and close all the Prism sessions.

If you attached to the job in the primary Prism session, `quit -all` leaves the debuggees running and close all the Prism sessions.

The `-all` option is valid only in the primary Prism session.

For convenience, you can add the `quit` command with the `-all` option to the tear-off region of the Prism graphical interface. For example,

```
(prism all) pushbutton quitall "quit -all"
```

This creates a button labeled `quitall` in the tear-off region.

Error Conditions Arising From Spawned Sessions

TABLE 4-1 lists and explains error messages that may be displayed when error conditions are encountered in debugging spawned processes.

TABLE 4-1 Error Messages Related to Debugging of Spawned Processes

Error Message	Description
Command not allowed in spawned prisms	The Prism environment displays this error message when a user attempts to issue a <code>run</code> , <code>rerun</code> , or <code>quit -all</code> command in a secondary Prism session.
Timed out waiting for spawned prism	The Prism environment displays this message when system conditions prevent a secondary Prism session from starting and it successfully communicates its status to the primary Prism session. In such a situation, quit the primary Prism session and all secondary sessions by issuing a <code>quit -all</code> command in the primary Prism session. Then repeat the debugging session with <code>follow_spawn</code> disabled.
Nodal startup failed in spawned prism	The Prism environment displays this error message when the Prism executable on a specific node fails to start due to a system error on that node. Quit and repeat the debugging session with <code>follow_spawn</code> disabled.

TABLE 4-1 Error Messages Related to Debugging of Spawned Processes (*Continued*)

Error Message	Description
Timed out connecting to parent prism	The Prism environment displays this error message in the secondary Prism session if it fails to connect to the primary Prism session and exchange status information with it.
Failed to spawn new prism	The Prism environment displays this error message in the primary Prism session if it fails to spawn either a new Prism session to debug a newly spawned Sun MPI debuggee or if additional Sun MPI jobs have been specified in the prism command line.
Could not continue stopped processes	The Prism environment displays this error message when it fails to change debuggee process status from the stopped state to the running state. In such cases, partial debugging of Sun MPI jobs with different executables (debugging only some of the processes) cannot usually be continued because the debugged processes cannot communicate with the non-debugged ones.
<code>follow_spawn</code> requires the MPI library to be linked in	The Prism environment displays this error message when you attempt to use spawn-related Prism commands, such as <pre>set follow_spawn = on set debug_spawn_aout = list</pre> while debugging a non-MIMD executable—that is, an executable that does not have the Sun MPI library linked in.

For more information about using the Prism environment with Sun MPI programs that issue calls to `MPI_Comm_spawn()` or `MPI_Comm_spawn_multiple()`, see “Enabling Support for Spawned MPI Processes” on page 18.

Examining the Contents of Memory and Registers

You can issue commands in the command window to display the contents of memory addresses and registers.

▼ To Display Memory

- **Specify the address on the command line, followed by a slash (/).**

The following displays the memory contents at address 10000 (hex).

```
(prism all) 0x10000/
```

If you specify the address as a period, the Prism environment displays the contents of the memory address immediately following the one printed most recently.

Specify a symbolic address by preceding the name with an &. For example, this prints the contents of memory for variable *x*:

```
(prism all) &x/
```

The address you specify can be an expression made up of other addresses and the operators +, -, and indirection (unary *). For example, this prints the contents of the location 100 addresses above address 0x1000:

```
(prism all) 0x1000+100/
```

After the slash you can specify how memory is to be displayed. TABLE 4-2 lists the supported memory address formats.

TABLE 4-2 Memory Address Formats

Format	Description
d	Print a short word in decimal
D	Print a long word in decimal
o	Print a short word in octal
O	Print a long word in octal
x	Print a short word in hexadecimal
X	Print a long word in hexadecimal
b	Print a byte in octal
c	Print a byte as a character
s	Print a string of characters terminated by a null byte
f	Print a single-precision real number
F	Print a double-precision real number
i	Print the machine instruction

The initial format is *x*. If you omit the format in your command, you get either *x* (if you haven't previously specified a format) or the format you specified previously.

You can print the contents of multiple addresses by specifying a number after the slash (and before the format). For example, this displays the contents of eight

memory locations starting at address 0x1000:

```
(prism all) 0x1000/8x
```

These contents are displayed as hexadecimal long words.

▼ To Display the Contents of Registers

You can examine the contents of registers in the same way that you examine the contents of memory.

- **Specify a register by preceding its name with a dollar sign.**

For example, this prints the contents of the `f0` register:

```
(prism all) $f0/
```

Specify a number after the slash to print the contents of multiple registers. For example, this prints the contents of registers `f0`, `f1`, and `f2`:

```
(prism all) $f0/3
```

The order in which the registers are displayed is that shown in TABLE 4-3.

You can also specify a format, as described above. The format specifier controls the display of the output; it doesn't affect how much of the register contents is displayed. Thus, this displays three registers:

```
$f0/3X
```

The output is displayed as hexadecimal longwords. The following table shows the names and descriptions of UltraSPARC™ registers

TABLE 4-3 UltraSPARC Registers

Name	Register
<code>\$g0-\$g7</code>	Global registers (64 bits)
<code>\$o0-\$o7</code>	Output registers (64 bits)
<code>\$l0-\$l7</code>	Local registers
<code>\$i0-\$i7</code>	Input registers
<code>\$psr</code>	Processor state register
<code>\$pc</code>	Program counter
<code>\$npc</code>	Next program counter
<code>\$y</code>	Y register
<code>\$wim</code>	Window invalid mask
<code>\$tbr</code>	Trap base register

TABLE 4-3 UltraSPARC Registers (*Continued*)

Name	Register
\$f0-\$f31	Floating-point registers
\$fsr	Floating status register (64 bits)
\$f0f1-\$f62f63	Floating-point registers
\$xg0-\$xg7	Upper 32 bits of \$g0-\$g7 (SPARC V8 plus only, or higher)
\$xo0-\$xo7	Upper 32 bits of \$o0-\$o7 (SPARC V8 plus only, or higher)
\$xfsr	Upper 32 bits of \$fsr (SPARC V8 plus only, or higher)
\$fprs	Floating-point registers state (SPARC V8 plus only, or higher)
\$tstate	Trap state register (SPARC V8 plus only, or higher)
\$fp	Frame pointer (synonym for \$i6)
\$sp	Stack pointer (synonym for \$o6)

Visualizing Data

This chapter describes how to examine the values of variables and expressions in your program. In addition, it describes how to find out the type of a variable and change its values.

This chapter consists of the following sections:

- “Overview of Data Visualization” on page 129
- “Choosing the Data to Visualize” on page 131
- “Working With Visualizers” on page 137
- “Saving, Restoring, and Comparing Visualizers” on page 153
- “Visualizing Structures” on page 157
- “Printing the Type of a Variable” on page 163
- “Printing Pointers as Array Sections” on page 166
- “Visualizing Multiple Processes” on page 168
- “Visualizing MPI Message Queues” on page 171
- “Displaying MPI Message Queue Text” on page 182
- “Displaying and Visualizing Sun S3L Arrays” on page 183

Overview of Data Visualization

You can visualize either variables (including arrays, structures, pointers, etc.) or expressions. In addition, you can provide a *context*, so that the Prism environment handles the values of data elements differently, depending on whether they meet the conditions you specify.

Printing and Displaying

The Prism environment provides two general methods for visualizing data:

- Printing data shows the value(s) of the data at a specified point during program execution.
- Displaying data causes its value(s) to be updated every time the program stops execution.

Printing or displaying to the history region of the Command window of the Prism graphical user interface shows the numeric or character values of the data in standard fashion.

Printing or displaying to a graphical window creates a *visualizer* window, which provides you with various options for representing the data.

Visualization Methods

The Prism environment provides the following methods for choosing what to print or display:

- Choosing the Print or Display selection from the Debug menu in the menu bar (see “To Print or Display a Variable or Expression at the Current Program Location” on page 131)
- Selecting text within the source window (see “To Print or Display From the Source Window” on page 132)
- Adding events to the Event Table (see “To Print or Display From the Event Table” on page 134)
- Issuing commands from the Command window (see “To Print or Display From the Command Window” on page 134)

In all cases, choosing Display adds an event to the event list, since displaying data requires an action to update the values each time the program is stopped. Note that, since Display updates automatically, the only way to keep an unwanted display window from reappearing is to delete the corresponding display event.

You create print events through the Event Table and the Events menu.

▼ To Change the Default Radix

- **Type:**

```
(prism all) set $radix = number
```

where *number* can be 2 (binary), 8 (octal), or 16 (hexadecimal). For example, this changes the default representation to hexadecimal:

```
(prism all) set $radix = 16
```

By default, the Prism environment prints and displays values as decimal numbers.

You can override the default for an individual print or display operation. See “To Print or Display From the Command Window” on page 134 and “Using the Options Menu” on page 140.

The default setting also affects the display of argument values in procedures in the call stack; see “To Display the Call Stack” on page 112.

Data Visualization Limits

Note these points in visualizing data:

- You cannot print or display any variables after a program finishes execution.
- Visualizers do not deal correctly with Fortran adjustable arrays. The size is determined when you create a visualizer for such an array. Subsequent updates to the visualizer will continue to use the original size, even though the size of the array may have changed since the last update. This will result in incorrect values in the visualizer. Printing or displaying values of an adjustable array in the Command window or to a new window will work, however.

Choosing the Data to Visualize

This section describes the methods for printing and displaying data.

▼ To Print or Display a Variable or Expression at the Current Program Location

1. **Perform one of the following:**

- To print a variable or expression at the current program location, choose Print from the Debug menu. It is, by default, also in the tear-off region.

- To display a variable or expression every time execution stops, choose Display from the Debug menu. The display begins at the current location in the program.
- When you choose Print or Display, a dialog box appears; FIGURE 5-1 shows an example of the Print dialog box.




FIGURE 5-1 Print Dialog Box

2. In the Expression box, enter the variable or expression whose value(s) you want to print.

Text selected in the source window appears as the default; you can edit this text.

The dialog boxes also allow selection of the window in which the values are to appear:

- You can specify that the values are to be printed or displayed in a standard window dedicated to the specified expression. The first time you print or display the data, the Prism environment creates this window. If you print data, and subsequently print it again, this standard window is updated. This is the default choice for both Print and Display.
- You can create a separate *snapshot* window for printing or displaying values. This is useful if you want to compare values between windows.
- You can print out the values in the Command window.
 - Click Print or Display to print the values of the specified expression at the current program location.
 - Click Cancel or press the Esc key to close the window without printing or displaying.

▼ To Print or Display From the Source Window

1. Select the variable or expression by dragging over it with the mouse or double-clicking on it.

To print without displaying the menu, press the Shift key while selecting the variable or expression.

2. Right-click the mouse to display a pop-up menu.

3. Click Print in this menu.

This displays a snapshot visualizer containing the value(s) of the selected variable or expression at that point in the program's execution.

4. Click Display.

This displays a visualizer that is automatically updated whenever execution stops.

Note – The Prism environment prints the correct variable when you choose it in this way, even if the scope pointer sets a scope that contains another variable of the same name.

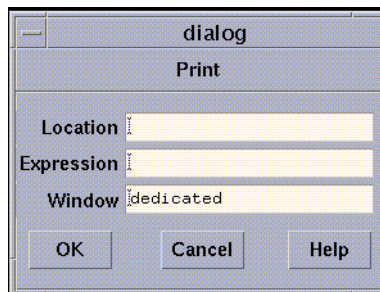
▼ To Print or Display From the Events Menu

1. Select Print on the Events menu.

You can use the Events menu to define a print or display event that is to take place at a specified location in the program.

2. Fill out the fields in the Print dialog box.

The Print dialog box prompts for the variable or expression whose value(s) are to be printed, the program location at which the printing is to take place, and the name of the window in which the value(s) are to be displayed.



Screenshot of the Events menu Print dialog box. Buttons are OK, Cancel, and Help.

FIGURE 5-2 Print Dialog Box

Predefined window names are *dedicated*, *snapshot*, and *command*. You can also specify custom names. The default window is *dedicated*. See “To Redirect Output to a File” on page 35 for a discussion of these names.

3. Click OK.

The event is added to the Event Table. When the location is reached in the program, the value(s) of the expression or variable is printed.

The Display dialog box is similar, but it does not prompt for a location; the display visualizer will update every time the program stops execution.

▼ To Print or Display From the Event Table

You can use the Event Table to define a print or display event that is to take place at a specified location in the program.

- **Click Print or Display in the Common Events buttons to create an event that will print or display data.**
 - If you click Print, the Location and Actions fields are highlighted. Put a program location in the Location field. Complete the print event in the Actions field, specifying the variable or expression, and the window in which it is to be printed. For example,

```
(prism all) print d2 on dedicated
```
 - If you click Display, the Location field displays `stopped` and the Actions field displays `print on dedicated`. Complete the description of the print event, as described above. The variable or expression you specify is then displayed whenever the program stops execution.

▼ To Print or Display From the Command Window

- **Perform one of the following:**
 - Type:

```
(prism all) print
```

This prints the value(s) of a variable or expression from the Command window.
 - Type:

```
(prism all) display
```


This displays the value(s).

The `display` command prints the value(s) of the variable or expression immediately and creates a display event so that the values are updated automatically whenever the program stops.

The commands have the following syntax:

```
[where (expression)] command variable[, variable ...]
```

The optional `where (expression)` sets the context for printing the variable or expression.

`command` is either `print` or `display`, and `variable` is the variable or expression to be displayed or printed.

Redirection of output to a window via the `on window` syntax works slightly differently for `display` and `print` from the way it works for other commands; see “To Redirect Output to a File” on page 35 for a discussion of redirection. Separate windows are created for each variable or expression that you print or display. Thus, the commands displayed in the following example create three windows, each of which is updated separately:

```
(prism all) display x on dedicated as colormap  
(prism all) display y/4 on dedicated as histogram  
(prism all) display [0:128:2]z on dedicated as text
```

▼ To Print or Display the Contents of a Register

- **Type:**

```
(prism all) print $name
```

or

```
(prism all) display $name
```

where `name` is the *name* of the register of interest. For example, this prints the contents of the program counter register:

```
(prism all) print $pc
```

See “To Display the Contents of Registers” on page 127 for a list of register names supported by the Prism environment.

▼ To Set the Context

- **Type:**

```
(prism all) where (expression) print variable
```

or

```
(prism all) where (expression) display variable
```

You can precede the `print` or `display` command with a `where` statement that can make elements of a variable or array *inactive*. Inactive elements are not printed in the Command window; “Overview of Data Visualization” on page 129 describes how they are treated in visualizers. Making elements of a variable or array inactive is referred to as *setting the context*.

The expression must evaluate to true or false for every element of the variable or array being printed.

For example, this prints only the values of *i* that are greater than 0:

```
(prism all) where (i .gt. 0) print i
```

You can use certain Fortran intrinsics in the `where` statement. For example, this prints the element of *a* that has the largest value:

```
(prism all) where (a .eq. maxval(a)) print a
```

This is equivalent to the `MAXLOC` intrinsic function. See “Writing Expressions in the Prism Environment” on page 37 for more information on writing expressions in the Prism environment.

Note that setting the context affects only the printing or displaying of the variable. It does not affect the actual context of the program as it executes.

▼ To Specify the Radix

- **Type:**

(prism all) **print**/*radix variable*

or

(prism all) **display**/*radix variable*

radix can be b (binary), d (decimal), x (hexadecimal), or o (octal).

For example, this prints the binary representation of `pvar1` in the Command window:

```
(prism all) print/b pvar1
```

The following example displays the hexadecimal values of `pvar2` in a dedicated window:

```
(prism all) display/x pvar2 on dedicated
```

The default radix is decimal, unless you have used the `set $radix` command to change it; see “To Change the Default Radix” on page 131.

Working With Visualizers

The window that contains the data being printed or displayed is called a *visualizer*. FIGURE 5-3 shows a visualizer for a three-dimensional array.

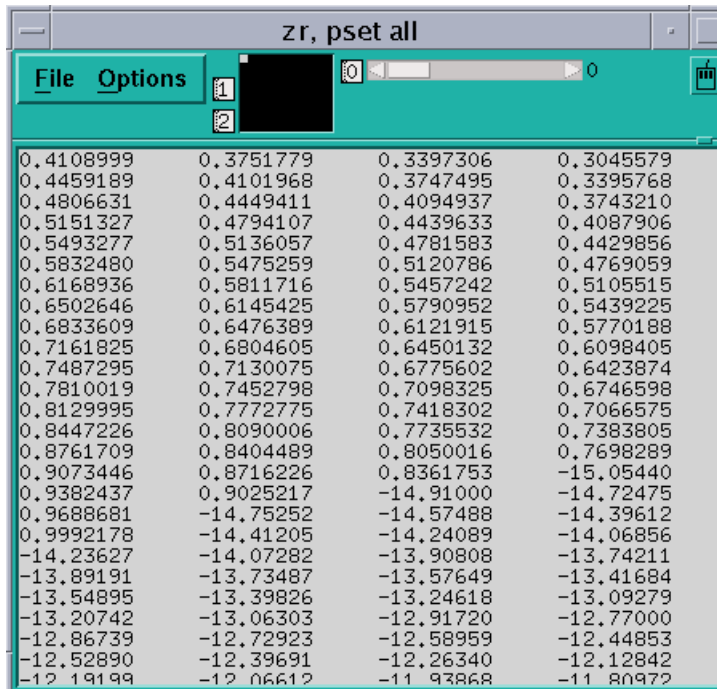


FIGURE 5-3 Visualizer for a Three-Dimensional Array

The visualizer consists of two parts: the *data navigator* and the *display window*. There are also File and Options pull-down menus.

The *data navigator* shows which portion of the data is being displayed and provides a quick method for moving through the data. The appearance of the data navigator depends on the number of dimensions in the data. It is described in more detail in “Using the Display Window in a Visualizer” on page 139.

The *display window* is the main part of the visualizer. It shows the data, using a representation that you can choose from the Options menu. The default is `text`; that is, the data are displayed as numbers or characters. FIGURE 5-3 is a Text visualizer. The display window is described in more detail in “Using the Options Menu” on page 140.

The File menu lets you save, update, or cancel the visualizer. The Options menu, among other things, lets you change the way values are represented. Both menus are described more fully later in this section.

Using the Data Navigator in a Visualizer

The data navigator helps you move through the data being visualized. It has different appearances, depending on the number of dimensions in your data.

Note – If your data is a single scalar value, there is no data navigator.

For one-dimensional arrays and parallel variables, the data navigator is the scroll bar to the right of the data. The numbers to the right of the buttons for the File and Options menus indicate the coordinates of the first element that is displayed. The elevator in the scroll bar indicates the position of the displayed data relative to the entire data set.

For two-dimensional data, the data navigator is a rectangle in the shape of the data, with the axes numbered. The white box inside the rectangle indicates the position of the displayed data relative to the entire data set. You can either drag the box or click at a spot in the rectangle. The box moves to that spot, and the data displayed in the display window changes.

For three-dimensional data, the data navigator consists of a rectangle and a slider, each of which you can operate independently. The value to the right of the slider indicates the coordinate of the third dimension. Changing the position of the bar along the slider changes which two-dimensional plane is displayed out of the three-dimensional data.

For data with more than three dimensions, the data navigator adds a slider for each additional dimension.

▼ To Change the Axes

You can change the way the visualizer lays out your data by changing the numbers that label the axes.

- 1. Click in the box surrounding the number.**

It is highlighted, and an I-beam appears.

- 2. Type in the new number of the axis; you don't have to delete the old number.**

The other axis number automatically changes. For example, if you change axis 1 to 2, axis 2 automatically changes to become axis 1.

Using the Display Window in a Visualizer

The display window shows the data being visualized.

In addition to using the data navigator to move through the data, you can drag the data itself relative to the display window by holding down the left mouse button; this provides finer control over the display of the data.

To find out the coordinates and value of a specific data element, click it while pressing the Shift key. Its coordinates are displayed in parentheses, and its value is displayed beneath them. If you have set a context for the visualizer, you also see whether the element is active or inactive. Drag the mouse with the Shift key pressed, and you see the coordinates, value, and context of each data element over which the mouse pointer passes.

You can resize the visualizer to display more or less data either horizontally or vertically.

▼ To Use the File Menu

1. Click on File to pull down the File menu.

2. Perform one of the following:

- Choose Update from this menu to update the display window for this variable, using the value(s) at the current program location. See “Updating and Closing the Visualizer” on page 152 for more information on updating a visualizer.
- Choose Save or Save As to save the visualizer’s values to a file. See “To Save the Values of a Variable” on page 153 for more information.
- Choose Diff or Diff With to compare the visualizer’s values with values stored in a file. See “To Compare the Data” on page 155 for more information.
- Choose Snapshot to create a copy of the visualizer, which you can use to compare with later updates.
- Choose Close to cancel the visualizer.

Using the Options Menu

Click on Options to pull down the Options menu. See FIGURE 5-4.

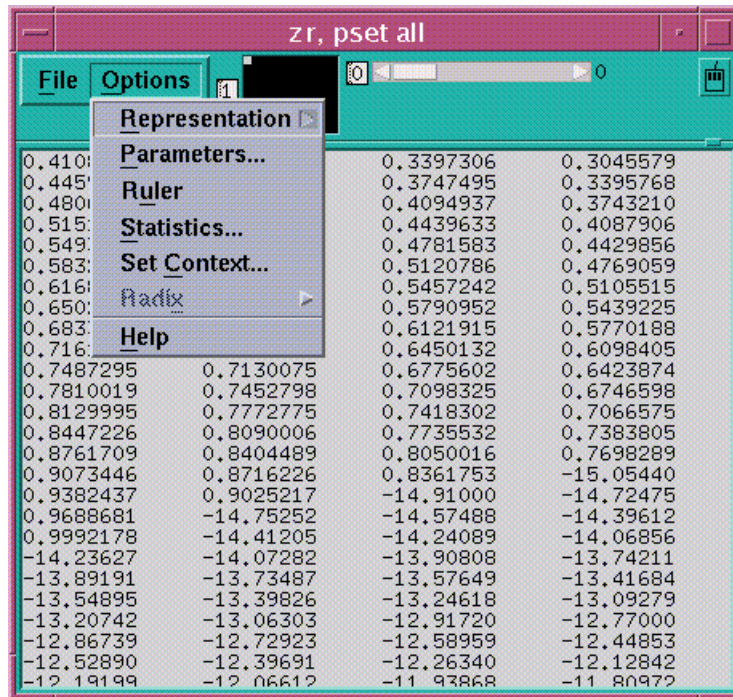


FIGURE 5-4 Options Menu in a Visualizer

▼ To Choose the Representation

- **Choose Representation from the Options menu.**

This displays another menu that gives the choices for how the values are represented in the display window. The choices are described below.

- Choose Text to display the values as numbers or letters. This is the default. For information about changing the default, see “To Change the Default Representation” on page 147.
- Choose Histogram to display the values of an array or parallel variable in a histogram. See FIGURE 5-5 for an example.

The vertical axis displays the number of data points; the horizontal axis displays the range of values. The Prism environment divides up this range evenly in creating the histogram bars. It prints summary data above the histogram.

Shift-click on a histogram bar to display the range and number of data points it represents.

Note that the histogram represents all the values of the variable, not just those shown in the two-dimensional slice of data that happens to be displayed in other representations.

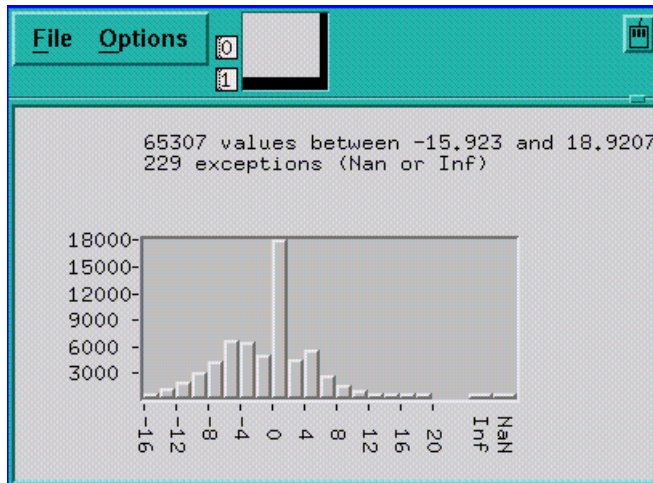


FIGURE 5-5 Histogram Visualizer

- Choose Dither to display the values as a shading from black to white. Groups of values in a low range are assigned more black pixels; groups of values in a high range are assigned more white pixels. This has the effect of displaying the data in various shades of gray. FIGURE 5-6 shows a two-dimensional dither visualizer. The lighter area indicates values that are higher than values in the surrounding areas; the darker area indicates values that are lower than surrounding values.
- You can left-click on a histogram visualizer bar to get a pop-up window, showing its contents.

For complex numbers, the Prism environment uses the modulus.




FIGURE 5-6 Dither Visualizer

- Choose Threshold to display the values as black or white. By default, the Prism environment uses the mean of the values as the threshold; values less than or equal to the mean are black, and values greater than the mean are white. FIGURE 5-7 shows a threshold representation of a three-dimensional array. For complex numbers, the Prism environment uses the modulus.

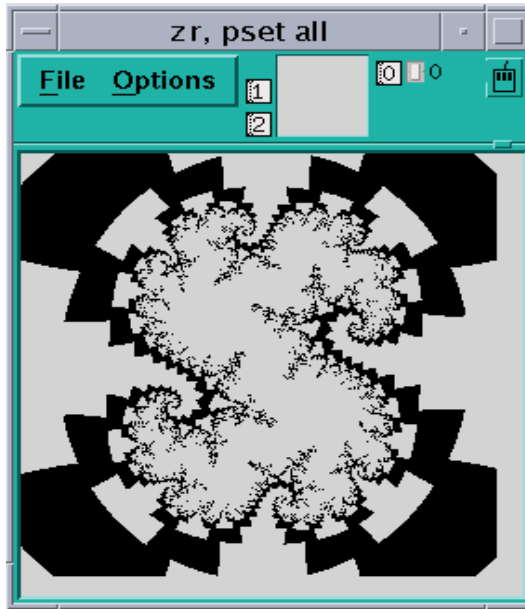


FIGURE 5-7 Threshold Visualizer

- Choose Colormap (if you are using a color workstation) to display the values as a range of colors. By default, the Prism environment displays the values as a continuous spectrum from blue (for the minimum value) to red (for the maximum value). You can change the colors that the Prism environment uses; see “Changing Colors” on page 213.

For complex numbers, the Prism environment uses the modulus.

- Choose Graph to display values as a graph, with the index of each array element plotted on the horizontal axis and its value on the vertical axis. A line connects the points plotted on the graph. This representation is particularly useful for one-dimensional data, but can be used for higher-dimensional data as well; for example, in a two-dimensional array, graphs are shown for each separate one-dimensional slice of the two-dimensional plane.

FIGURE 5-8 shows a graph visualizer for a one-dimensional slice of an array.




FIGURE 5-8 One-Dimensional Graph Visualizer

- Choose Surface (if your data has more than one dimension) to render the three-dimensional contours of a two-dimensional slice of data. In the representation, the two-dimensional slice of data is tilted 45 degrees away from the viewer, with the top edge farther from the viewer than the bottom edge. The data values rise out of this slice. FIGURE 5-9 is an example.

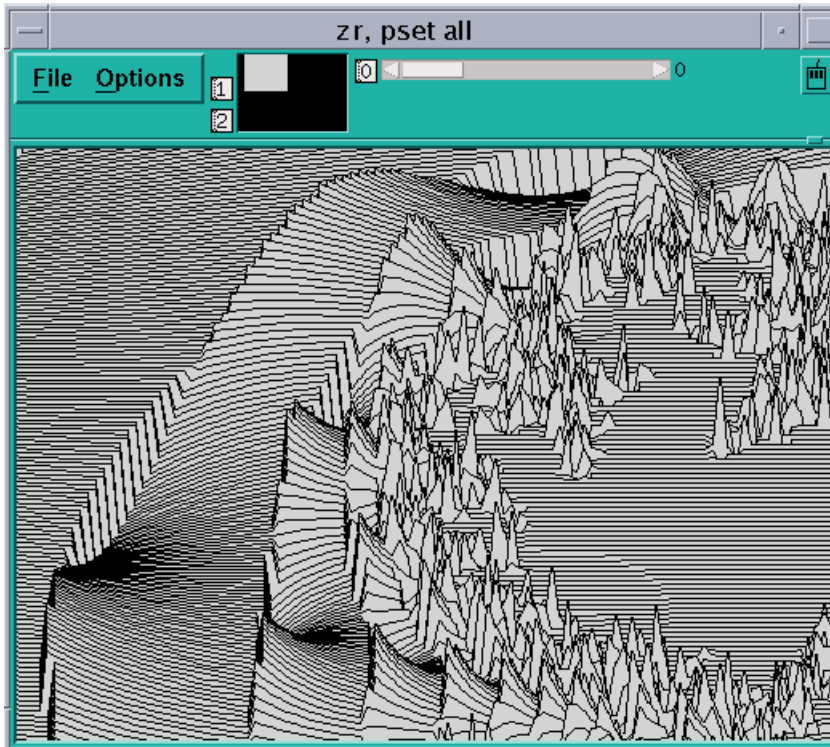


FIGURE 5-9 Surface Visualizer

Note – If there are large values in the top rows of the data, they may be drawn off the top of the screen. To see these values, flip the axes as described earlier in this section, so that the top row appears in the left column.

- Choose Vector to display data as vectors. The data must be a Fortran complex or double-complex number, or a pair of variables to which the `CMPLEX` intrinsic function has been applied (see “Using Fortran Intrinsic Functions in Expressions” on page 39). The complex number is drawn showing both magnitude and direction. The length of the vector increases with magnitude. Because direction is difficult to see for smaller vectors, the minimum vector length is five pixels. By default, the lengths of all vectors scale linearly with magnitude, varying between the minimum and maximum vector lengths. FIGURE 5-10 shows a vector visualizer.




FIGURE 5-10 Vector Visualizer

▼ To Change the Default Representation

- Type:

```
(prism all) set $viz = "representation"
```

By default, the Prism environment creates new visualizers with a Text representation. Set the Prism variable `viz` to control the representation that is applied when a new visualizer is created. Possible values are:

- Text
- Histogram
- Dither
- Threshold
- Colormap
- Graph
- Surface
- Vector

▼ To Set Parameters

- **Choose Parameters from the Options menu.**

This displays a dialog box in which you can change various defaults that the Prism environment uses in setting up the display window; see FIGURE 5-11. If a parameter is grayed out or missing, it does not apply to the current representation.

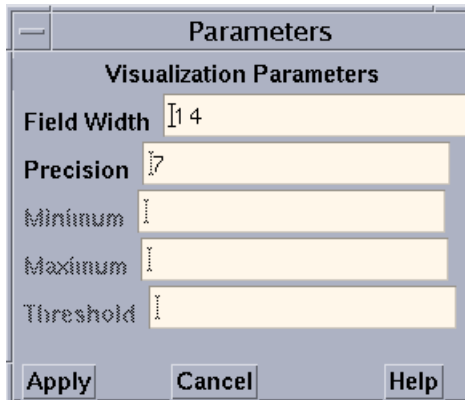


FIGURE 5-11 Visualization Parameters Dialog Box

The parameters for all representations except the histogram representation are:

- **Field Width** – Type a value in this box to change the width of the field that the Prism environment allocates to every data element.

For the text representation, the field width specifies the number of characters in each column. If a number is too large for the field width you specify, dots are printed instead of the number.

For dither, threshold, colormap, and vector representations, the field width specifies how wide (in pixels) the representation of each data element is to be. By default, dither, threshold, and colormap visualizers are scaled to fit the display window. Note, however, that for dither visualizers, the gray shading may be more noticeable with a smaller field width.

For the graph representation, the field width specifies the horizontal spacing between elements.

For the surface representation, it specifies the spacing of elements along both directions of the plane.

- **Precision** – Type a value in this box to change the precision with which the Prism environment displays real numbers in a text visualizer. The precision must be less than the field width. By default, the Prism environment prints the values of doubles with 16 significant digits, and floating-point values with seven significant

digits. You can change this default by issuing the `set` command with the `$d_precision` variable (for doubles) or the `$f_precision` variable (for floating-point values). For example,

```
(prism all) set $d_precision = 11
```

sets the default precision for doubles to 11 significant digits.

- **Minimum and Maximum** – For colormap representations, use these variables to specify the minimum and maximum values that the Prism environment is to use in assigning color values to the data elements. Data elements that have values below the minimum and above the maximum are assigned default colors.

For graph, surface, and vector representations, these parameters represent the bottom and top of the range that is to be represented. Values below the minimum are shown as the minimum; values above the maximum are shown as the maximum.

By default, the Prism environment uses the entire range of values for all these representations.

- **Threshold** – For threshold representations, use this variable to specify the value at which the Prism environment is to change the display from black to white. Data elements whose values are at or below the threshold are displayed as black; data elements whose values are above the threshold are displayed as white. By default, the Prism environment uses the mean of the data as the threshold.

The parameters for the histogram representation are:

- **Bar Width** – Specifies the width in pixels of each histogram bar (except for the bars representing infinities and NaNs, which must be wide enough to fit the `Inf` or `NaN` label underneath). The default is 10 pixels.
- **Bar Height** – Specifies the height in pixels of the largest histogram bar. The default is 100 pixels.
- **Minimum** – Specifies the minimum value to be included in the histogram. By default the actual minimum value is used.
- **Maximum** – Specifies the maximum value to be included in the histogram. By default the actual maximum value is used.

If you specify a different minimum or maximum, values below the minimum or above the maximum are not displayed in the histogram but are counted as outliers instead; the number of outliers is displayed above the histogram.

- **Max Buckets** – Specifies the number of “buckets” into which values are to be poured—in other words, the number of histogram bars to be used. The default is 30. (The Prism environment might use fewer to make the horizontal labels come out evenly.)

▼ To Display a Ruler

- **Choose Ruler from the Options menu.**

This toggles the display of a ruler around the display window.

The ruler is helpful in showing which elements are being displayed. FIGURE 5-12 shows a three-dimensional threshold visualizer with the ruler displayed.

In the surface representation, the ruler cannot indicate the coordinates of elements in the vertical axis, since they change depending on the height of each element. However, you can press the Shift key and left-click to display the coordinates and value of an element.




FIGURE 5-12 Threshold Visualizer With a Ruler

▼ To Display Statistics

- **Choose Statistics from the Options menu.**

This displays a window containing statistics and other information about the variable being visualized.

The window contains:

- The name of the variable
- Its type and number of dimensions

- The total number of elements the variable contains and the total number of active elements, based on the context you set within the Prism environment (see the next section for a discussion of setting the context)
- The variable's minimum, maximum, and mean; these statistics reflect the context you set for the visualizer

FIGURE 5-13 gives an example of the Statistics window.




FIGURE 5-13 Statistics for a Visualizer

For complex numbers, the Prism environment uses the modulus.

▼ To Use the Set Context Dialog Box

● Choose Set Context from the Options menu.

In this dialog box, you can specify which elements of the variable are to be considered active and which are to be considered inactive. Active and inactive elements are treated differently in visualizers:

- In text, graph, surface, and vector visualizers, inactive elements are grayed out.
- In colormap visualizers, inactive elements by default are displayed as gray. You can change this default; see “Changing Colors” on page 213.
- Context has no effect on dither and threshold visualizers.

FIGURE 5-14 shows the Set Context dialog box.



FIGURE 5-14 Set Context Dialog Box

By default, all elements of the variable are active; this is the meaning of the keyword “everywhere” in the text-entry box. To change this default, you can either edit the text in the text-entry box directly or click on the Where button to display a menu. The choices in the menu are everywhere and other:

- Choose `everywhere` to make all elements active.
- Choose `other` to erase the current contents of the text-entry box. You can then enter an expression into the text-entry box.

In the text-entry box, you can enter any valid expression that will evaluate to true or false for each element of the variable.

The context you specify for printing does not affect the program’s context; it just affects the way the elements of the variable are displayed in the visualizer.

Click on Apply to set the context you specified. Click on Cancel or press the Esc key to close the dialog box without setting the context.

▼ To Change the Radix

1. Choose Radix from the Options menu.
2. Choose one of the items from the submenu: Decimal, Hex, Octal, and Binary.

Updating and Closing the Visualizer

If you created a visualizer by issuing a `display` command, it automatically updates every time the program stops execution.

If you created the visualizer by issuing a `print` command, its display window is grayed out when the program resumes execution and the values in the window are outdated.

▼ To Update Values

- Choose Update from the visualizer’s File menu.

▼ To Close the Visualizer

- Choose Close from the File menu, or press the Esc key.

Saving, Restoring, and Comparing Visualizers

You can save the values of a variable or expression to a file. You can subsequently visualize these values and compare them with the values in another visualizer, such as the same variable later in the run or the same variable during a separate execution of the program. This provides a convenient way of spotting changes in the values of a variable.

▼ To Save the Values of a Variable

You can save the values of a variable or expression to a file for later use.

- **Perform one of the following:**
 - From the command line – Use the command `varsave` to save the values of a variable or expression to a file.

Its syntax is:

```
varsave "filename" expression
```

where *filename* is the name of the file to which the data is to be saved, and *expression* is the variable or expression whose values are to be saved.

For example, this saves the values of the variable `alpha` in the file `alpha.data` in your current working directory within the Prism environment:

```
(prism all) varsave "alpha.data" alpha
```

The following saves the results of the expression `alpha*2` in the file with the path name `/u/kathy/alpha2.data`:

```
(prism all) varsave "/u/kathy/alpha2.data" alpha*2
```

- From a visualizer – Use the Save or Save As selection from a visualizer's File menu to save the visualizer's values to a file.

If you choose Save As, a dialog box appears in which you can specify the name of the file to which the values are to be saved; see FIGURE 5-15.




FIGURE 5-15 Saving a Visualizer's Data to a File

The highlighted directory is the current working directory. If you want to put the file there, simply type its name in the Save As box and click on OK.

If you want to put the file in another directory, click on the directory. (The parent directories of the current working directory are shown above it in the Directories list; its subdirectories are listed beneath it.) This displays the subdirectories of the directory you clicked on. You can traverse the directory structure in this manner until you find the directory in which you want to put the file, or you can simply type the entire path name in the Save As box.

Choose the Save selection to save the values in the file you most recently specified. If you have not specified a file, the values are saved in a file called `noname.var` in your current working directory.

▼ To Restore the Data

This intrinsic brings values you have saved to a file back into the Prism environment.

- **Type:**

`(prism all) command varfile("filename")`

where *filename* is the name of the file that contains the values you want to restore.

Note – The `varfile` intrinsic is not available for use with message-passing programs.

You can use the `varfile` intrinsic anywhere you could have used the original variable or expression that you saved to a file. For example, if you saved `x`:

```
(prism all) varsave "x.var" x
```

then the command

```
(prism all) print varfile("x.var")
```

is equivalent to

```
(prism all) print x
```

Note that this enables you to save a variable's values, then print them during a later Prism session without having a program loaded or running.

▼ To Compare the Data

You can compare a variable or expression whose values have been saved in a file with another version of the variable or expression. This comparison could take place later in the same run of the program, during a subsequent run, or even during a second, simultaneous Prism session.

You can also compare the values with those of another variable, as long as both variables have the same base type (that is, you can't compare integers with floating-point numbers).

● Perform one of the following:

- From the command line – Type:

```
(prism all) print - varfile ("filename")
```

or

```
(prism all) display - varfile ("filename")
```

This performs a comparison between two versions of a variable or expression.

For example, if you saved `x` in the file `x.var`:

```
(prism all) varsave "x.var" x
```

then this command prints the difference between the current and saved values of `x`:

```
(prism all) print x - varfile("x.var")
```

If an element is printed as 0, it is the same in both versions. If it is nonzero, its value is different in the two versions.

- From a visualizer – Select Diff or Diff With from a visualizer's File menu.

This performs a comparison between the visualizer's values and the values stored in a file.

Select Diff With to choose the file containing the values. It displays a dialog box like the one shown in FIGURE 5-16.

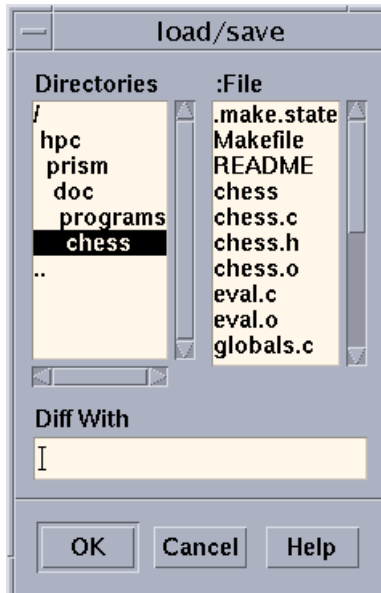


FIGURE 5-16 Diff With Dialog Box

The dialog box has the same format as the Save As dialog box described in “To Save the Values of a Variable” on page 153. It lists the files found in your current working directory in the Prism environment. Click on a file name, then click OK to choose the file. Or type a file name in the Diff With text-entry box and click OK.

Choose Diff to compare the visualizer's values to those in the most recently specified file. If no file has been specified, values are compared to those in the file `noname.var` in your current working directory in the Prism environment.

Once you have specified a file via Diff or Diff With, the Prism environment creates a new visualizer that displays the difference in values between the visualizer and the file. If an element's value in the new visualizer is 0, the value is the same in both versions. If it is nonzero, it is different in the two versions.

You can work with this visualizer as you would any visualizer. For example, you can change the representation and display summary statistics.

Visualizing Structures

If you print a pointer or a structure (or a structure-valued expression) in a window, a *structure visualizer* appears.

FIGURE 5-17 shows an example of a structure visualizer.

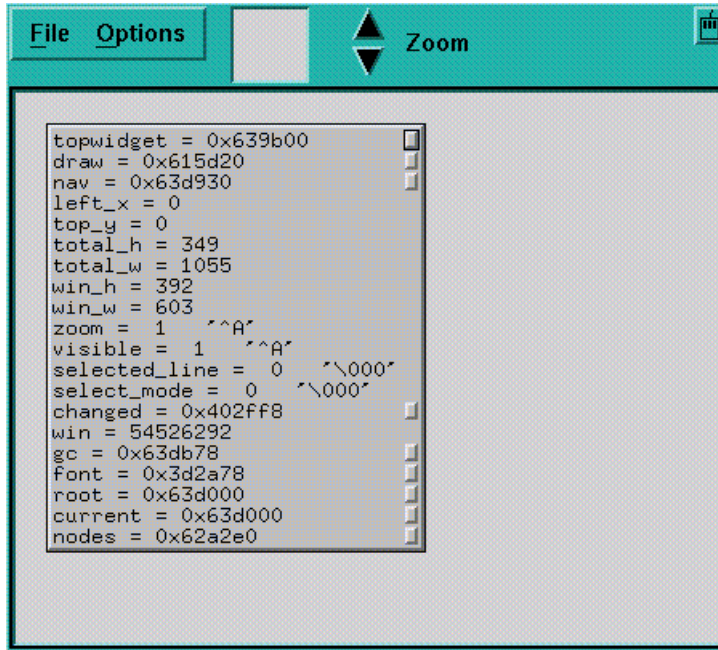


FIGURE 5-17 Structure Visualizer

The structure you specified appears inside a box; this is referred to as a *node*. The node shows the fields in the structure and their values. If the structure contains pointers, small boxes appear next to them; they are referred to as *buttons*. Left-click on a node to select it. Use the up and down arrow keys to move between buttons of a selected node.

You can perform various actions within a structure visualizer, as described below.

Expanding Pointers

You can expand scalar pointers in a structure to generate new nodes. (You cannot expand a pointer to a parallel variable.)

▼ To Expand a Single Pointer

- **Perform one of the following:**

- With the mouse – Left-click a button to expand the pointer. For example, clicking the button next to the `nav` field in FIGURE 5-17 changes the visualizer as shown in FIGURE 5-18.
- From the keyboard – Use the right arrow key to expand and visit the node pointed to by the current button. If the node is already expanded, pressing the right arrow key simply visits the node. Use the left arrow key to visit the parent of a selected node.

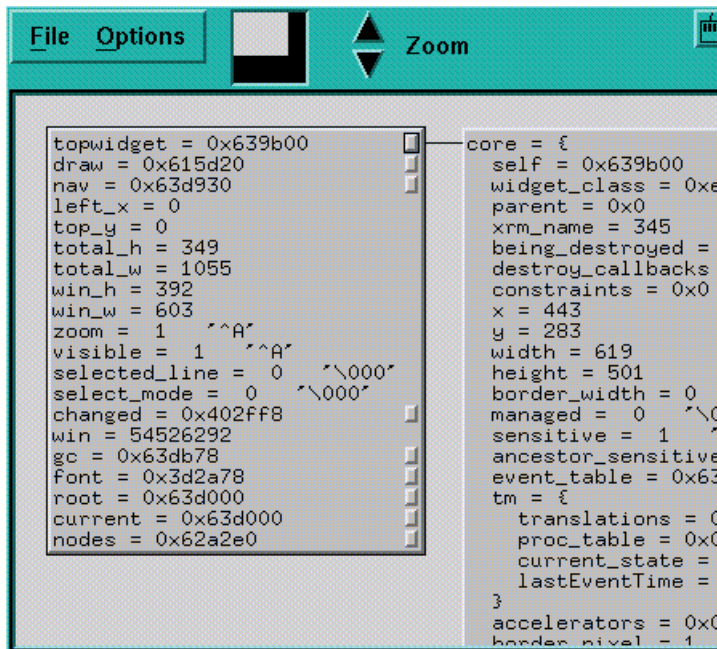


FIGURE 5-18 Structure Visualizer, With One Pointer Expanded

▼ To Expand All Pointers in a Node

- **Perform one of the following:**

- With the mouse – Double-click or Shift-left-click on the node.

- From the keyboard – Press the Shift key along with the right arrow key.
- From the Options menu – Click Expand. The cursor turns into a target; move the cursor to the node you are interested in and left-click.

▼ To Expand All Pointers Recursively From the Selected Node on Down

- **Perform one of the following:**
 - With the mouse – Triple-click or Control-left-click on the node.
 - From the keyboard – Press the Control key and the right arrow key.
 - From the Options menu – Click Expand All. The cursor turns into a target; move the cursor to the node you are interested in and left-click.

▼ To Pan and Zoom

- **Perform one of the following:**
 - Left-click and drag through the data navigator or the display window to pan through the data.
 - Left-click on the Zoom arrows to “zoom” in and out on the data.
 - Click on the down arrow to zoom out and see a bird’s-eye view of the structure. Click on the up arrow to get a closeup.
 - Left-click on a node in a zoomed-out structure visualizer to pop up a window showing the full contents of the node.

For information about navigating through visualizers, see “Using the Data Navigator in a Visualizer” on page 139 and “Using the Display Window in a Visualizer” on page 139.

FIGURE 5-19 shows part of a complicated structure visualizer after zooming out.

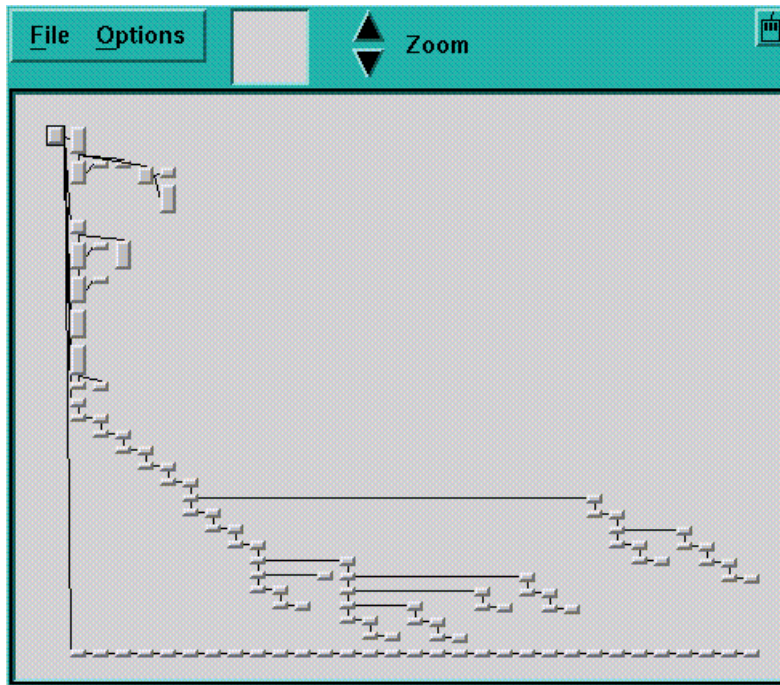


FIGURE 5-19 Zooming Out in a Structure Visualizer

The selected node is centered in the display window whenever you zoom in or out.

▼ To Delete Nodes

- **To delete a node (except the root node):**
 - With the mouse – Middle-click on a node (except the root node).
 - From the Options menu – Click Delete. The cursor turns into a target; move the cursor to the node you want to delete and left-click.

Deleting a node also deletes its children (if any).

More About Pointers in Structures

Note the following about pointers in structure visualizers:

- Null pointers have “ground” symbols next to them.
- If you have previously expanded a pointer, it has an arrow next to its button; you cannot expand the pointer again. This prevents infinite loops on circular data structures.

- A pointer containing a bad address has an X drawn over its button.

Augmenting the Information Available for Display

You can provide a special function for each of your data types that makes additional information available to the Prism environment. This enables the Prism environment to more accurately display the contents of structures with that data type.

For C or C++ union types, you can identify which member of the union is valid. For a pointer within a structure, you can specify that the pointer's target is an array of elements, rather than a single element, and you can further specify the length of the array.

You must embed these specifications within a special function that is compiled and linked with your program being debugged. The function has the following form:

```
void prism_define_typename (typename *ptr);
```

where *typename* is the tag name of one of your structure data types. Thus, you can define one such function for each of your data types. When the Prism environment displays a variable of this type, it checks whether an augmentation function is defined in the program. If so, the Prism environment calls the function, passing a pointer to the instance of the structure being displayed. Your function can then look at the instance to choose valid union members and to size dynamic arrays.

You communicate this information back to the Prism environment with the following call, which is defined in `/opt/SUNWhpc/include/prism.h`:

```
void prism_add_array(char *member_name, int len);
```

This call specifies that the pointer named `member_name` points to an array of length `len`. The pointer's name, `member_name`, is the name of one of the members of the structure, as found in the structure's C or C++ declaration. The results are undefined if `member_name` is not a pointer.

The following call specifies that the member named `name` is of type union, and of all the members of this union, only `valid_member` is to be displayed.

```
void prism_add_union(char *name, char *valid_member);
```

Both `name` and `valid_member` are names as found in the C or C++ declarations of structs or unions.

▼ To Augment the Information That the Structure Visualizer Displays

- **Link your program with the library `libprism.so` in `/opt/SUNWhpc/lib`.**

Sun HPC ClusterTools software currently supplies both static (`libprism.a`) and shared versions (`libprism.so`) of this library. Use the shared version.

Note – To prevent the occurrence of unresolved references in applications shipped to your customers, ensure that you remove all references to `libprism.so` in your production code.

Assume that data in the declaration below is a dynamic array:

```
struct Vector {
    int len;
    int *data;
};
```

The function you write looks like this:

```
#include "prism.h"
void prism_define_Vector(struct Vector *v)
{
    prism_add_array("data", v->len);
}
```

Assume that the member `type` discriminates the union value in this example:

```
enum Type {INT, DOUBLE};
struct Value {
    enum Type type;
    union {
        int i;
        double d;
    } value;
};
```

The function you write would look like this:

```
#include "prism.h"
void prism_define_Value(struct Value *val)
{
    if (val->type == INT)
        prism_add_union("value", "i");
    else
        prism_add_union("value", "d");
}
```

There are no restrictions on the number or order of calls to `prism_add_union` and `prism_add_array`.

▼ To Update and Close a Structure Visualizer

1. Update the structure visualizer by left-clicking Update in the File menu.

This updates a structure visualizer. When you do this, the root node is reread; the Prism environment attempts to expand the same nodes that are currently expanded. (The same thing happens if you reprint an existing structure visualizer.)

2. Close the structure visualizer by left-clicking Close in the File menu.

Printing the Type of a Variable

The Prism environment provides several methods for finding out the type of a variable.

▼ To Print the Type of a Variable From the Menu Bar

1. Select Whatis from the Debug menu.

The Whatis dialog box appears; it prompts for the name of a variable.

2. Click on Whatis.

This displays the information about the variable in the Command window.

3. Click on Type.

The Prism environment treats *name* as a type name.

▼ To Print the Type of a Variable From the Source Window

1. Select a variable by double-clicking it or by dragging over it while pressing the left mouse button.

2. Hold down the right mouse button.

A pop-up menu appears.

3. Choose **WhatIs** from this menu.

Information about the variable appears in the Command window.

▼ To Print the Type of a Variable From the Command Window

● Type:

(prism all) **whatIs** [*type*] *variable*

If you specify a type (`struct`, `class`, `enum`, or `union`) before the name of the variable, the Prism environment treats *variable* as a type name. The type keywords resolve ambiguities where there are types and variables with the same name.

What Is Displayed

The Prism environment displays information about the variable in the Command window. In this example, the `whatIs` command is used with variable `primes`:

```
whatIs primes
logical primes(1:999)
```

▼ To Modify Visualizer Data

- **Type:**

```
(prism all) assign variable = value
```

This assigns new values to a variable or an array.

For example,

```
(prism all) assign x = 0
```

assigns the value 0 to the variable `x`. You can put anything on the left-hand side of the statement that can go on the left-hand side in the language you are using—for example, a variable or a Fortran array section.

If the right-hand side does not have the same type as the left-hand side, the Prism environment performs the proper type coercion.

Changing the Radix of Data

▼ To Change the Radix of a Value

- **Type:**

```
(prism all) value = base
```

This changes the radix of a value in the Prism environment. The value can be a decimal, hexadecimal, or octal number. Precede hexadecimal numbers with `0x`; precede octal numbers with `0` (zero). The base can be `D` (decimal), `X` (hexadecimal), or `O` (octal). The Prism environment prints the converted value in the Command window.

For example, to convert 100 (hex) to decimal, issue this command:

```
(prism all) 0x100=D
```

The Prism environment responds:

```
256
```

▼ To Print the Names and Values of Local Variables

- **Type:**

```
(prism all) dump routine
```

To print the names and values of all local variables in a function or procedure, specify its name. If you omit the function name, `dump` uses the current function. If you specify a period, `dump` prints the names and values of all local variables in the functions in the stack.

Printing Pointers as Array Sections

The Prism environment enables you to print simple arrays by section. The following examples assume the following declarations and code:

```
double da[]={0.1,1.1,2.1,3.1,4.1,5.1,6.1,7.1,8.1,9.1,10.1};
double *pd=da;
int a[]={0,1,2,3,4,5,6,7,8,9,10};
int *pa=a;

int *par[10];
int **ppi=par;
void *ptr=(void*)da;
...

for(i=0;i<10;i++)
    par[i]=&a[9-i];

<----- assume that the program is stopped here -----
...

```

▼ To Print an Array by Section

- **Type**

(prism all) **print** *array-name*[*section-specifier*]

For example,

```
(prism) print a[1:5:2]
a[1:5:2] =
(1:3) 1 3 5

```


▼ To View a Pointer as a One-Dimensional Array

- **Type**

```
(prism all) print pointer[section_specifier]
```

Specify a section when printing the pointer.

For example:

```
(prism all) print pa[1:5:2]  
pa[1:5:2] =  
(1:3) 1 3 5
```

▼ To Dereference an Array of Pointers

- **Type**

```
(prism all) *pointer[section_specifier]
```

If the array element is a pointer, then the Prism environment enables you to dereference the section.

For example,

```
(prism all) *par[1:5:2] =  
(1:3) 8 6 4
```

▼ To Cast Pointers

- **Type**

```
(prism all) print ((type*) pointer) [section_specifier]
```

For example,

```
(prism all) print ((double*)ptr)[1:4:2]
((double*)ptr)[1:4:2] =
(1:2) 1.1000000000000000      3.1000000000000000
```

Currently, the Prism environment supports only one level of dereferencing.

Assuming this declaration:

```
int **appi[2];
```

The Prism environment does *not* support:

```
(prism all) print **(appi[0:1])
```

Although the Prism environment allows one level of dereferencing for sections, the Prism environment does not support indexing. Thus, the Prism environment allows:

```
(prism all) print *par[1:5:2]
```

but the Prism environment does *not* allow:

```
(prism all) print par[1:5:2][0]
```

Visualizing Multiple Processes

When you print or display an object in the Prism environment, the data are shown for all processes in the pset you specify. If you do not include a pset qualifier, the current pset will be the source. Choosing the Print or Display selection from the Debug menu prints or displays data for processes in the `current` pset.

If there is only one process in the pset, the visualizer that is displayed is no different from the visualizer you would see in the scalar mode of the Prism environment.

If the pset contains more than one process, the Prism environment adds a dimension to the visualizer. The extra dimension represents the processes in the set. For example, if the variable is scalar, the Prism environment displays a one-dimensional array that represents the value of the variable in each process. If you are printing a one-dimensional array, the Prism environment uses a two-dimensional visualizer.

For C programs, axis 0 represents the processes. For Fortran 77 programs, the highest-numbered axis represents the processes.

The Prism environment can aggregate data from multiple processes only if the expression has the same size and number of dimensions in each process; if it doesn't, the Prism environment prints an error message.

In the example shown in FIGURE 5-20, the variable `board` is an 8x8 array (representing a chess board); the current pset contains four processes. Therefore, the Prism environment displays a three-dimensional visualizer. Axis 0 represents the processes. The figure shows the values of `board` in the first process in the set. You drag the white bar in the slider portion of the data navigator to display the values in the other processes in the set. (Note that, for a two-dimensional Fortran array, where axis 3 would represent the processes, you might want to rearrange the display axes so that axis 3 is on the slider. You can do this by clicking in the box to the left of the slider and changing the number to a 3.)

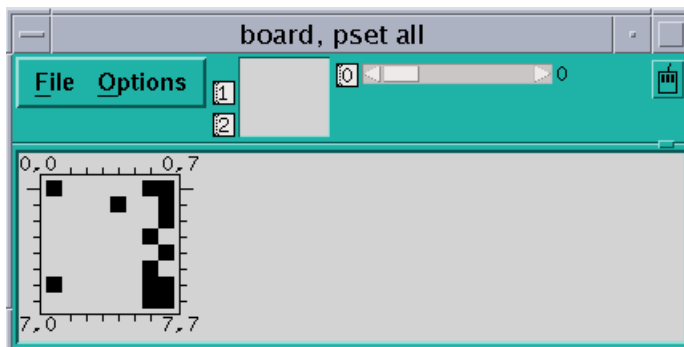


FIGURE 5-20 Visualizer in the Prism Environment (Threshold Representation)

▼ To Find Out the Value and Process Number for an Element

- **Shift-click on the element.**

Printing to the history region or in the commands-only mode of the Prism environment works the same way. Axis 0 represents the processes. Here is a portion of the history-region output for the data shown below:

```
(prism all) print board
board =
process 0
(0,0,0:4)  4  1  0  3  0
(0,0,5:7) -1  0 -4
(0,1,0:4)  2  1  0  0  0
(0,1,5:7)  0 -1  0
(0,2,0:4)  3  1  0  0  0
(0,2,5:7)  2 -1 -3
(0,3,0:4)  5  0  0  0 -1
(0,3,5:7)  0  0 -5
(0,4,0:4)  4  0  0 -2  0
(0,4,5:7)  0  0 -6
(0,5,0:4)  0  1  0  0  0
(0,5,5:7)  0 -1  0
(0,6,0:4)  0  1  0  0  0
(0,6,5:7)  0 -1  0
(0,7,0:4)  6 -1  0  0  0
(0,7,5:7)  0 -1 -4
process 1
(1,0,0:4)  4  1  0  3  0
(1,0,5:7) -1  0 -4
(1,1,0:4)  2  1  0  1  0
...
```

The elements of axis 0 do not necessarily correspond to the numbers of the processes they represent. For example, if you were visualizing a variable in pset (1, 3, 5, 7), element 0 of axis 0 would represent process 1, element 1 would represent process 3, and so forth.

The Prism environment provides a `Cycle` visualizer window you can use to display the values of a variable in the `cycle` pset. See “The cycle Pset” on page 80 for more information.

▼ To Open a Cycle Visualizer Window

- **Type**

(prism all) `print variable on cycle`

The Prism environment displays a window containing the value of *variable* in the current process of the `current` pset. If you then issue the `cycle` command or otherwise cycle through the members of the `cycle` pset, this window automatically updates to display the value of *x* in the next member of the set. This provides a convenient way of examining a particular variable in a series of processes.

Visualizing MPI Message Queues

The Prism MPI queue visualizer allows you to examine message queues created by your Sun MPI program. The visualizer shows you the status of messages generated by nonblocking send and receive routines that have not been reaped by a call to `MPI_Test` or `MPI_Wait`.

By showing you the state of the queue, detailing the messages that have not completed, the Prism environment gives you clues regarding where your program's logic can be tuned.

The Prism queue visualizer also shows you unexpected receive routines, indicating performance or correctness problems:

- **Performance** – An unexpected receive indicates the receipt of a message before a posted matching receive; you might receive an extra copy of the message.
- **Correctness** – An unexpected receive can arise due to an intended receive's not having been posted or having been posted incorrectly, such as with the wrong tag. The program could deadlock due to errors.

In addition to viewing the status of messages, you can also view the contents of the messages themselves to ensure that the correct data was transmitted.

Note – The Prism environment does not display blocking sends and receives on message queues. If a blocking routine such as an `MPI_Send` hangs your program, you can use the Prism environment to display a stack backtrace to find the problem, showing the `MPI_Send` present on the stack. The Prism environment also does not display MPI generalized requests.

▼ To Launch the MPI Queue Visualizer

- **Choose the MPI Msgs selection under the Prism Debug menu.**

This selection is available only when a program linked to the Sun MPI library has been loaded into the Prism environment. See FIGURE 5-21 for an example.

Each row of messages displayed in the message queue window corresponds to a process rank, numbered from zero. The following sections describe how each part of the MPI queue visualizer window affects the display of messages.

▼ To Select the Queue to Visualize

- **Choose an item from the View menu.**

This selects the queues to visualize. You can view three classes of MPI queues for each rank:

- Posted Sends
- Posted Receives
- Unexpected Receives

You can view queues only when a rank has stopped. Otherwise, the visualizer displays the label `running` for that rank. Prism reevaluates the queue every time the rank stops.

▼ To Zoom Through Levels of Message Detail

- **Click the Zoom buttons to navigate through four levels of message detail.**

By default, the MPI queue visualizer opens at zoom level three. Examples of the zoom levels are illustrated in FIGURE 5-21 through FIGURE 5-24.

FIGURE 5-21 shows a single pixel per message. This zoom level is useful when examining very large MPI jobs.

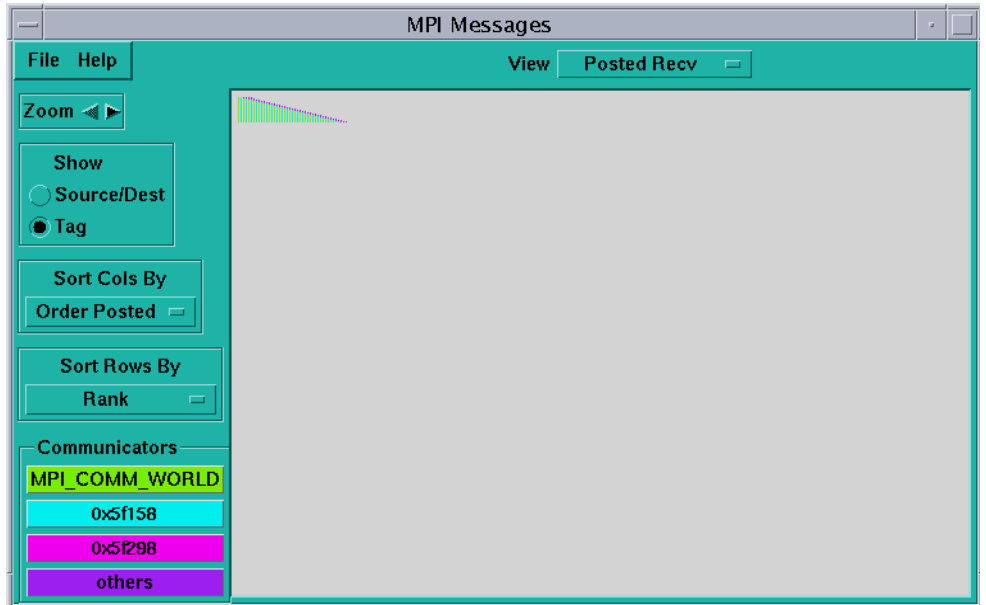


FIGURE 5-21 Queue Visualizer at Zoom Level One

FIGURE 5-22 shows a simple box per message. The size of the box increases with the size of the message.

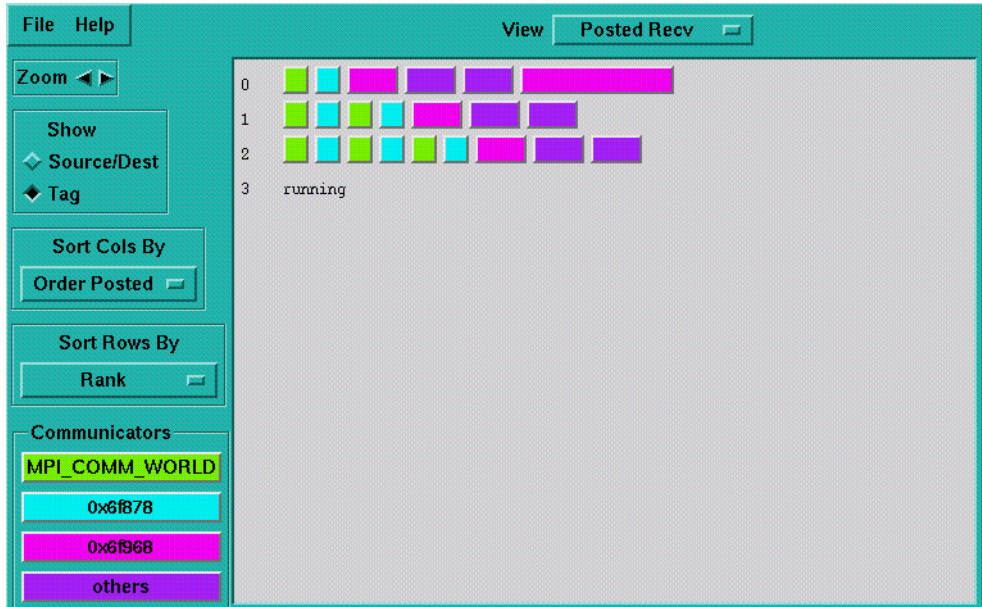


FIGURE 5-22 Queue Visualizer at Zoom Level Two

FIGURE 5-23 shows a single label on a message. Clicking the buttons on the Show menu toggles the labels. Label choices are Source/Destination and Tag.

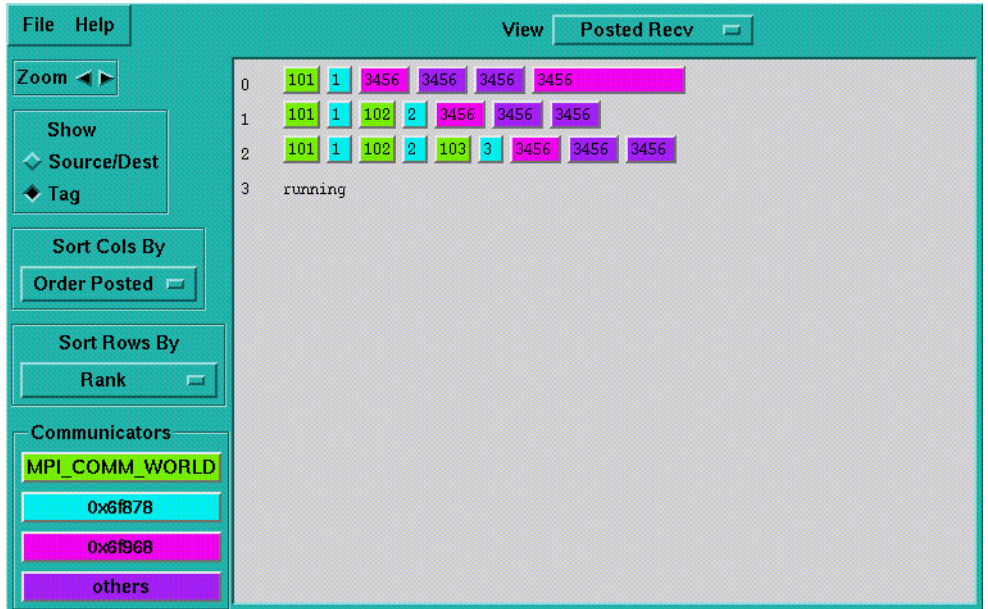


FIGURE 5-23 Queue Visualizer at Zoom Level Three

FIGURE 5-24 shows the entire message.

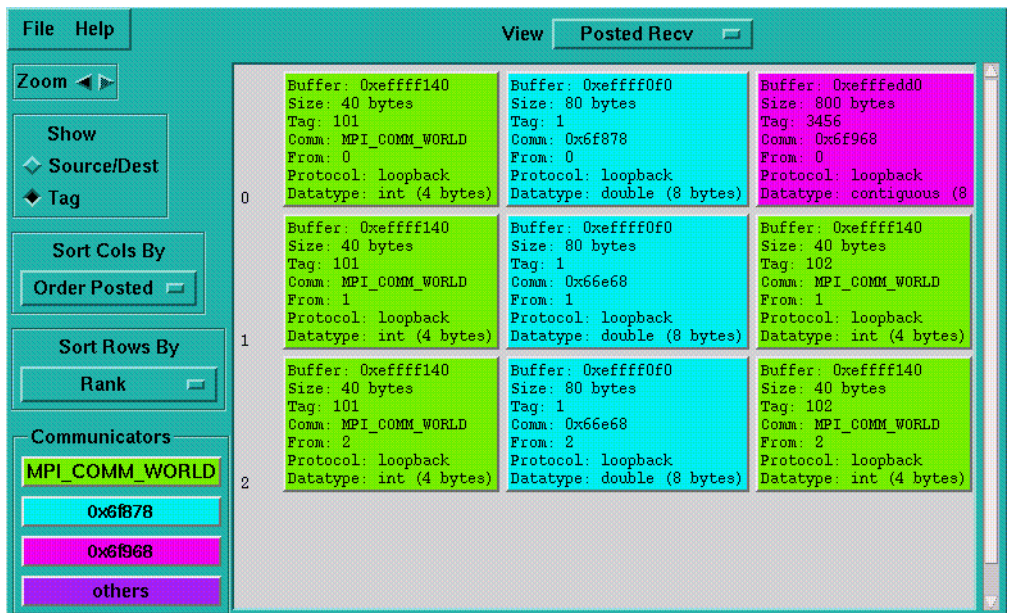


FIGURE 5-24 Queue Visualizer at Zoom Level Four

▼ To Control the Values of Message Labels

- **Click the toggle buttons under Show on the MPI queue visualizer.**

This controls the value of the message labels:

- Select Source/Dest to show the source or destination rank for the message.
- Select Tag to show the MPI tag of the message.

Clicking the Show toggle buttons affects the display of messages at zoom level three only.

▼ To Sort Messages

- **Choose selections from the Sort Rows By and Sort Cols By option menus.**

These sort messages by row or by column using the criteria listed in TABLE 5-2 and TABLE 5-1.

TABLE 5-1 Row Sort Criteria

Sort Criteria	Description
Rank	Sort rows from the smallest to the largest process rank (the default).
Message Count	Sort by the number of messages posted.
Message Volume	Sort by the sum of the sizes, in bytes, of all messages for each rank.

TABLE 5-2 Column Sort Criteria

Sort Criteria	Description
Order posted	Sort messages by the order in which messages are posted by the MPI program, with the earliest posted on the left. This is the default. The use of MPI send operations for transmitting large messages might queue and dequeue the message several times once the rendezvous begins. In such cases, the posted order seen in the visualizer will, at some point, no longer match the programmatic order. At present there is no way to distinguish such messages.
Source/Destination	Sort by the source rank for receives and the destination rank for sends.
Tag	Sort by the messages' tag values.
Size	Sort by size in bytes, from small to large.
Communicator	Sort by communicator address.
Protocol	Group messages sent with the same transport protocol. Protocols are loopback, shared memory, RSM, and TCP.

The MPI queue visualizer does not scale message labels to exactly correspond to the size of the messages they represent. Collections of labels of small messages can appear disproportionately large when compared to the label of a single, very large message, even when the number of bytes in the single large message exceeds the total size of the collection of smaller messages.

▼ To Display Message Fields

- **Click an individual message.**

This opens the Message dialog box, shown in FIGURE 5-25.

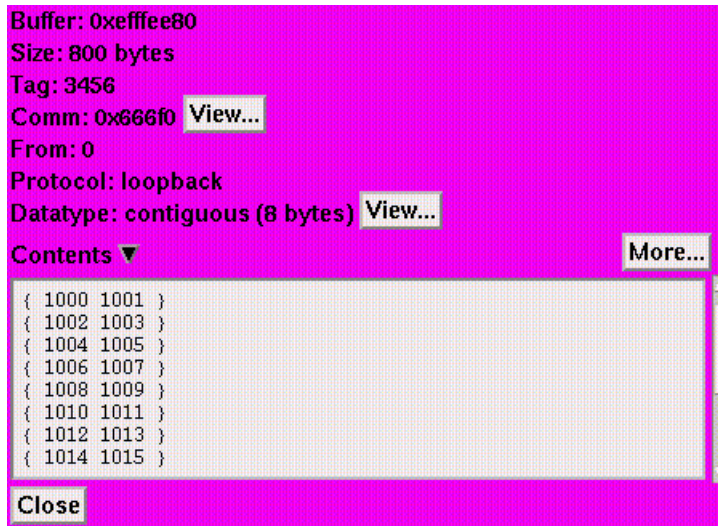


FIGURE 5-25 Message Dialog Box

Interpreting Message Dialog Fields

The fields in the Message dialog box are described in TABLE 5-3.

TABLE 5-3 Message Dialog Box Fields

Field Name	Description
Buffer	The address of the message.
Size	The length (in bytes) of the message.
Tag	The MPI tag argument passed in the call to post the message.
Comm	The name of the MPI communicator in which the message belongs, or the communicator's address if it is unnamed. Click on the Communicator View button to display the Communicator dialog box.
To	The rank of the destination of the message. Prism displays this field only for posted sends.
From	The rank of the sender of the message. Prism displays this field only for posted receives or unexpected receives.

TABLE 5-3 Message Dialog Box Fields (*Continued*)

Field Name	Description
Protocol	The implementation method by which the message has been sent. Possible values are: loopback, shared memory, RSM, and TCP.
Datatype	The MPI data type of the message, with the size of a single data type element in bytes. If this is a user-defined data type, click on the Data Type View button to display the Data Type dialog box. See “To Display Data Types” on page 181 for more information about the Data Type dialog box. The View button is available only for user-defined data types.
Contents	The contents of the message. Click on the triangular button to open or close the contents area. Click on More repeatedly to scroll through the message.

When the Message dialog box displays a posted receive, it displays the value of the buffer address as null (indicating that no buffer has been allocated) and disables the Contents button.

When the Message dialog box displays an unexpected receive, it shows the delivered message with no data type. This characteristic is due to MPI design, since a posted receive declares the data type. Here too, the Contents button is disabled, and the visualizer displays the value of the buffer address as null.

Displaying Communicator Data

The Prism environment displays MPI Communicators in the Communicators region of the MPI queue visualizer window. The visualizer does not display all the communicators that have been created in an MPI program; rather, it displays only communicators referenced by currently posted messages. Thus, if no messages are visible, the visualizer displays no communicators.

The Prism environment displays as many as three distinct communicators. Each communicator is color coded, and messages are drawn using the color of their communicator. If more than three communicators are present, then the excess are grouped together under a single color labeled Others.

▼ To Change Communicator Colors

- **Set the following X resources in the Prism application defaults file:**
 - `Prism.commlColor`
 - `Prism.comm2Color`
 - `Prism.comm3Color`

- Prism.commOtherColor

For information about modifying values in the Prism applications defaults file, see “Changing Prism Environment Defaults” on page 209.

▼ To Display Communicator Data

- **Press any of the Communicator buttons.**

This reveals the Communicator dialog box.

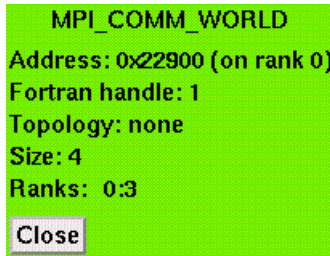


FIGURE 5-26 Communicator Dialog Box

FIGURE 5-26 shows the Communicator dialog box, which includes the information described in TABLE 5-4.

TABLE 5-4 Communicator Dialog Box Fields

Field Name	Description
Name	Name of the communicator.
Address	The address of the communicator
Fortran handle	The Fortran identifier for the communicator, if defined. Built-in communicators such as MPI_COMM_WORLD have predefined Fortran handles. Other communicators are assigned a Fortran handle only if they are used in a Fortran subroutine.
Topology	The options are: Cartesian – Communicators created using MPI_Cart_create Graph – Communicators created using MPI_Graph_create None – All others

TABLE 5-4 Communicator Dialog Box Fields (Continued)

Field Name	Description
Size	The number of ranks
Remote size	Shown only for intercommunicators; the size of the remote group (the number of ranks). For information about intercommunicators, see the <code>MPI_Intercomm_create</code> man page.
Ranks	The list of ranks. This may be annotated with job identifiers if the communicator was created via an MPI client/server rendezvous or an <code>MPI_Spawn</code> . The ranks displayed for a communicator are relative to <code>MPI_COMM_WORLD</code> , rather than relative to the communicator's parent.

▼ To Display Data Types

- Click on the **Datatype View** button in a Message dialog box.

FIGURE 5-27 shows the Data Type dialog box.

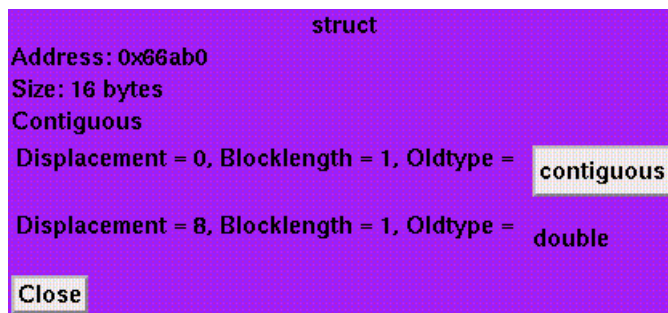


FIGURE 5-27 Data Type Dialog Box

The fields in the Data Type dialog box are described in TABLE 5-5.

TABLE 5-5 Data Type Dialog Box Fields

Field name	Description
Type	The description of the data type, such as <code>struct</code> or <code>contiguous</code>
Address	The address of the corresponding <code>MPI_Datatype</code> object in the MPI program

TABLE 5-5 Data Type Dialog Box Fields (Continued)

Field name	Description
Size	The size in bytes of a single element of this data type
Contiguous	An indication that the bytes of this data type are contiguous and may be sent or received without any intermediate packing or unpacking. If the data type is not contiguous, the label changes to Non-contiguous
Additional information	Additional information that is specific to the data type, representing arguments that were passed to MPI to create the data type. This can include offsets, block sizes, pointers to other data types, and so forth. In this example, Displacement, Blocklength, and Oldtype refer to arguments the programmer used when creating the MPI <code>struct</code> data type. Click on buttons that name other data types to display the Data Type dialog box for that other type

Displaying MPI Message Queue Text

You can use the `mpimsgs` command to display (in text format) message queues created by a Sun MPI program. The Prism environment displays the messages in the output window, sorting the messages by rank. The fields of each message are displayed, including message size, tag, to (or from), comm (communicator), protocol, and data type.

Specify `send` to view Posted Sends, `recv` to view Posted Receives, or `urecv` to view Unexpected Receives. Specify the `verbose` option to display more details about the communicator and to display the contents of the message. For example,

```
(prism all) mpimsgs send verbose pset foo
```


The following example shows a typical message with the verbose option enabled,

```
Queues for Rank 0:

** 6 Posted Sends:
  # 0: size = 40
      tag = 101
      to = 0
      comm:
          name = MPI_COMM_WORLD
          fortran handle = 1
          topology = none
          size = 4
          ranks = 0:3
          protocol = loopback
          data type = int
          contents:

0 1 2 3 4 5 6 7 8 9
```

Displaying and Visualizing Sun S3L Arrays

In a multiprocess Sun MPI program, a parallel array is an array whose elements may be distributed among the processes of the program. The Prism environment can extract the global dimensionality and distribution information from these arrays and manipulate them as single entities. For the purpose of this discussion, arrays that are not distributed (arrays that belong in their entirety to a single process) are referred to as *regular* arrays.

Sun S3L's parallel array syntax is based on *array handles*, which define the properties of the parallel array.

By default, the Prism environment recognizes an array handle as a simple variable. In Fortran 77 and Fortran 90, the array handle is a variable of type `integer*8`. In C, the array handle is type `S3L_array_t`.

The following examples assume the code in CODE EXAMPLE 5-1:

CODE EXAMPLE 5-1 S3L Array Demonstration Program

```
c
c Copyright (c) 1998, by Sun Microsystems, Inc.
c All rights reserved
c
c   program test_prism_s3l
c   include 's3l/s3l-f.h'
c
c In f77 programs, s3l arrays are integer*8
c
c   integer*8 a
c
c   integer*4 ext(2),local(2),ier
c
c Initialize the S3L library and the prism/s3l interface.
c
c   call s3l_init(ier)
c
c
c Declare a parallel S3L array of size 2 x 3,
c with the second dimension distributed.
c
c   ext(1) = 2
c   ext(2) = 3
c
c   local(1) = 1
c   local(2) = 0
c
c   call s3l_declare(a,2,ext,S3L_float,local,S3L_USE_MALLOC,ier)
c
c
c Initialize the array randomly by using S3L_rand_lcg
c
c
c   call s3l_rand_lcg(a,123456,ier)
c   w = 1.0
c
c free the resources associated with the parallel S3L array
c
c   call s3l_free(a,ier)
c
c finalize the S3L library.
```

CODE EXAMPLE 5-1 S3L Array Demonstration Program (Continued)

```
call s3l_exit(ier)
c
stop
end
```

Note that, before using the `type` command, the `whatis` command reports that the Sun S3L array handle, `a`, has been declared an `integer*8`.

▼ To Display the Data Type of an Array Handle

- **Type**

```
(prism all) whatis array-handle
```

This shows that the array handle, `a`, is a variable of type `integer*8`:

```
(prism all) whatis a
integer*8 a
```

▼ To Create an S3L Parallel Array

- **Type**

```
(prism all) type data-type array-handle
```

This identifies `array-handle` as a Sun S3L parallel array and specifies its basic data type. Basic data types are `int`, `float`, `double`, `complex8`, and `complex16`.

The example below executes the `type` command, associating the Sun S3L handle, `a`, with the data type, `float`, the same type used to declare the element type of the Sun S3L array in your program:

```
(prism all) type float a
"a" defined as "float a"
```

▼ To Display and Visualize Sun S3L Parallel Arrays

- **Type**

```
(prism all) print array_handle
```

or

```
(prism all) display array_handle
```

At this point, the Prism environment recognizes `a` as a Sun S3L array. You could now use the Prism environment's `print` command to display the values of `a`, as illustrated in the following example:

```
(prism all) print a
a =
(0:1,0) 0.000000      1.000000
(0:1,1) 0.100000      1.100000
(0:1,2) 0.200000      1.200000
```

In all respects, you could use `a` as you would use any array in the Prism environment. In the following example, you use `a` as an array variable:

```
(prism all) assign a=9
(prism all) print a
a =
(0:1,0) 9.000000      9.000000
(0:1,1) 9.000000      9.000000
(0:1,2) 9.000000      9.000000
```

Sun S3L arrays are distributed across multiple processes, with all processes having the same global view of the array. Since `a` has the same content for all processes, the Prism environment prints the values of the array only once.

However, when the Prism environment prints a regular array, it prints the values of the array separately for each process or pset. For regular arrays, the values of the array can differ in each process, since every process has its own copy.

In the following example, `larr` is a regular array with each process having its own copy of the array:

```
(prism all) print larr
larr =
Pset 0
(1:2,1,1) 0.000000      1.000000
(1:2,2,1) 0.1000000    1.100000
(1:2,3,1) 0.2000000    1.200000
Pset 1
(1:2,1,2) 0.000000      1.000000
(1:2,2,2) 0.1000000    1.100000
(1:2,3,2) 0.2000000    1.200000
(prism all) assign larr=larr*5 pset 0
(prism all) print larr
larr =
Pset 0
(1:2,1,1) 0.000000      5.000000
(1:2,2,1) 0.5000000    5.500000
(1:2,3,1) 1.000000      6.000000
Pset 1
(1:2,1,2) 0.000000      1.000000
(1:2,2,2) 0.1000000    1.100000
(1:2,3,2) 0.2000000    1.200000
```

The Prism environment also prints expressions involving Sun S3L parallel arrays only once, unless they include a variable in the user program. The following example shows the S3L array, `a`, in an expression that does not include any other variables. Since the values of `a` are the same for all processes, its contents are printed only once.

```
(prism all) print 10* a - 1
10* a - 1 =
(0:1,0) -1.0000000000000000      9.0000000000000000
(0:1,1) 0.0000000000000000      10.00000023841858
(0:1,2) 1.000000029802322      11.00000047683716
(prism all) print a
a =
(0:1,0) 0.000000      1.000000
(0:1,1) 0.1000000    1.100000
(0:1,2) 0.2000000    1.200000
```

However, if you use a Sun S3L array in an expression that includes a variable, the Prism environment replicates the array on each process and then evaluates the array separately on each process. The following example adds the variable *w* to *a*. The Prism environment prints the results for both processes.

```
(prism all) print w
w =
(1:2) 0          0
(prism all) print a+w
a+w =
Pset 0
(0:1,0,1) 0.0000000000000000    1.0000000000000000
(0:1,1,1) 0.1000000014901161    1.100000023841858
(0:1,2,1) 0.2000000029802322    1.200000047683716
Pset 1
(0:1,0,2) 0.0000000000000000    1.0000000000000000
(0:1,1,2) 0.1000000014901161    1.100000023841858
(0:1,2,2) 0.2000000029802322    1.200000047683716
```

▼ To Visualize the Layouts of S3L Parallel Arrays

- **Type**

```
(prism all) print layout arrayname
```

This returns the numbers of the nodes on which the data elements of an S3L array are located. In the following example, the `print layout` command is used with array *a*:

```
(prism all) print layout a
layout (a) =
a =
(0:1,0) 0 0
(0:1,1) 0 0
(0:1,2) 1 1
```

where *a* is an S3L array. You can use the Fortran 90 array-section syntax described in “Using Array-Section Syntax in C Arrays” on page 41 to specify a range of elements within an S3L array.

▼ To Print or Display an S3L Array Using the `layout Intrinsic`

- **Type**

`(prism all) print layout (arrayname) on window as representation`

This creates a visualizer that is the same size and shape as S3L array *arrayname*. The visualizer displays the rank of the process that is holding each value.

Note that you can specify any visualizer representation—for example, text, dither, or colormap—to display the layout graphically.

Editing and Compiling Programs

This chapter discusses editing and compiling source code while in the Prism environment. It is organized into the following sections:

- “Editing Source Code” on page 191
- “Using the make Utility” on page 192

Editing Source Code

The Prism environment provides an interface to the editor of your choice so that you can edit source code without having to leave the environment in which you compile and debug the code.

▼ To Start the Default Editor on the Current Source File

- **Perform one of the following:**
 - From the menu bar – Select Edit from the Utilities menu.
 - From the Prism command window – Type

```
(prism all) edit [file-name | function-name]
```

You can specify which editor the Prism environment is to call by using the Customize window to set a Prism resource; see “Changing Prism Resource Defaults” on page 204. If this resource has no setting, Prism uses the setting of your `EDITOR` environment variable. Otherwise, the Prism environment uses the default editor listed in the Customize window.

The editor is invoked on the current file, as displayed in the source window. Some editors start at the current execution point. Other editors start at the beginning of the file.

After the editor has been created, it runs independently. This means that changes you make in the current file are not reflected in the source window. To update the source window, you must recompile and reload the program. You can do this using the Make selection from the Utilities menu, as described below.

Using the make Utility

The Prism environment provides an interface to the standard Solaris tool `make`. The `make` utility lets you automatically recompile and relink a program that is broken up into different source files. See your Solaris documentation for an explanation of `make` and makefiles.

Creating the Makefile

Create the makefile as you normally would. Within the Prism environment, you can choose the Edit selection from the Utilities menu to bring up a text editor in which you can create the file; see “Editing Source Code” on page 191.

Using the Makefile

After you have made changes in your program, you can run `make` to update the program.

The Prism environment uses the standard Solaris `make` utility, `/usr/ccs/bin/make`, unless you specify otherwise. You do this by using the Customize utility to change the setting of a Prism resource; see “Changing Prism Resource Defaults” on page 204.

▼ To Run `make` From the Menu Bar

- **Perform the following:**

- Choose Make from the Utilities menu.

A window like the one shown in FIGURE 6-1 appears.

- Edit the fields in the `make` window, if necessary. The window prompts for the names of the makefile, the target file(s), the directory in which the makefile is located, and other arguments to `make`. If a file is loaded, its name is in the Target box, and the directory in which it is located is in the Directory box. You can change these if you like.

If you leave the Makefile or the Target box empty, `make` uses a default. See your Solaris documentation for a discussion of these defaults. If you leave the Directory box empty, `make` looks for the makefile in the directory from which you started the Prism environment.

You can specify any standard `make` arguments in the Other Args box.

The dialog box also asks if you want to reload after the make. Answering Yes automatically reloads the newly compiled program into the Prism environment if the make is successful. If you answer No, the program is not reloaded. Yes is the default.

- To cancel the make while it is in progress, click the Cancel button. If a make is not in progress, clicking Cancel closes the window.
- View the output from `make` in the box at the bottom of the Make window. Subsequent makes use the same window, unless you start a new make while a previous make is still in progress.



FIGURE 6-1 The `make` Window

▼ To Run `make` From the Command Window

- **Type**

`(prism all) make arguments`

You can specify any arguments that are valid in the Solaris version of `make`.

Getting Help

This chapter describes how to obtain information about the Prism environment and other Sun products available at your site. It is organized into the following sections:

- “Prism Online Help Systems” on page 195
 - “Obtaining Online Documentation” on page 197
-

Prism Online Help Systems

▼ To Get Help in the Prism Environment

- **Perform one of the following:**
 - Select an entry from the Help menu in the menu bar. The Help menu provides help on several major topics. See “Choosing Selections From the Help Menu” on page 196.
 - Select an entry from the Help menus and Help button in windows and dialog boxes. These Help menus and the Help button provide instructions for using these screen areas. Pressing the F1 key in a window or dialog box also displays a help screen.
 - Use the command-line help. The syntax of command-line help is
`(prism all) help commandname`
Command-line help provides information about commands you can issue from the command window.

Using the Browser-Based Help System

The Prism environment displays its help files using your World Wide Web browser. The default browser is Netscape™, although your system administrator can change this.

To specify the HTML browser you want to use for the graphical mode of the Prism environment, set the Prism environment resource `Prism.helpBrowser` to the executable name of the browser. For detailed information about customizing this feature of the Prism environment, see “Specifying a Different Browser for Displaying Help” on page 219.

If you don't have a browser running, the Prism environment starts one. If you have a browser currently running as you use the Prism environment, by default the Prism environment displays the help information in that browser. You can change this behavior using the `Prism.helpUseExisting` resource. For detailed information about customizing this feature of the Prism environment, see “Specifying a Different Browser for Displaying Help” on page 219.

Note – See “Setting Up Your Working Environment” on page 10 for important information about setting up your environment for the Prism environment's use of your default browser to display the Prism environment's online help files.

Choosing Selections From the Help Menu

The Help menu provides information in a variety of ways. You can choose:

- Using Help to display an overview of the Help system.
- Overview to display an overview of the features of the Prism environment.
- Glossary to display a list of terms used in the Prism environment. You can click on a term to find out more about it.
- Commands Reference to display a list of Prism commands. You can click on a command's link marker to obtain its reference description.
- Tutorial to display a tutorial that will teach you the basics of the Prism environment.

Getting Help on Using the Mouse

Some Prism windows include an icon of a mouse:



Click this icon to display information about using the mouse in the window.

Obtaining Help From the Command Window

▼ To Obtain Help From the Command Window

- **Type**

`(prism all) help commands`

This displays a list of Prism commands and editing key combinations.

- **Type**

`(prism all) help command-name`

This displays help on that command.

- **Type**

`(prism all) help`

This displays a brief message about how to use command-line help.

Obtaining Online Documentation

The Prism environment's documentation is available both in print and in the PDF and PostScript™ formats. Prism also comes with a Solaris-style manual page.

Viewing Manual Pages

▼ To Obtain a Manual Page

- **Choose the Man Pages selection from the Doc menu.**

This brings up `xman`, a standard X program for viewing manual pages; `xman` operates independently of the Prism environment.

Help for `xman` appears in the `xman` window, as shown in FIGURE 7-1. You can use `xman` to view any Solaris manual pages available on your Sun system.

Note – If `xman` is not available on your system, you will not be able to use this feature.

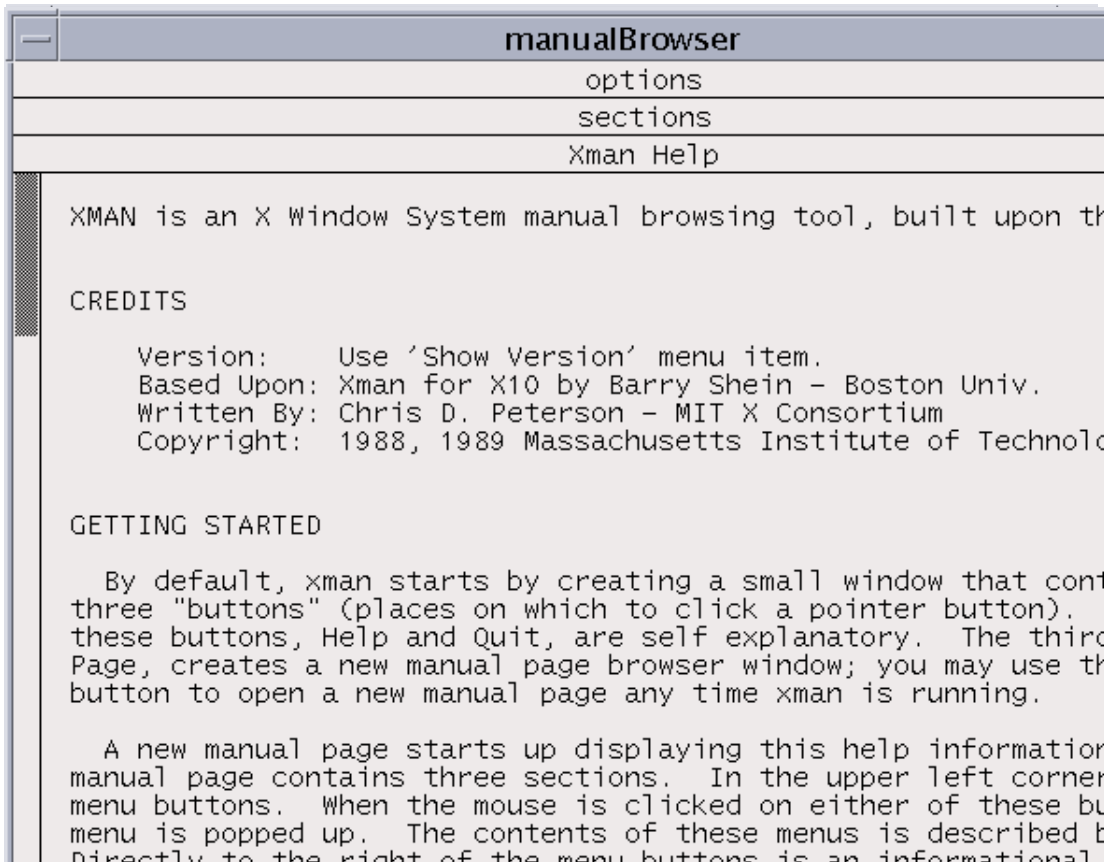


FIGURE 7-1 xman Window

Customizing the Prism Programming Environment

This chapter discusses ways in which you can change various aspects of the Prism environment's appearance and the way the Prism environment operates. This discussion is organized into the following sections:

- “Initializing the Prism Environment” on page 199
- “Using the Tear-Off Region” on page 201
- “Creating Aliases for Commands and Variables” on page 203
- “Changing Prism Resource Defaults” on page 204
- “Changing Prism Environment Defaults” on page 209

Initializing the Prism Environment

Use the `.prisminit` file to initialize the Prism environment when you start it up. You can put any Prism commands into this file. When the Prism environment starts, it executes these commands, echoing them in the history region of the command window.

When starting up, the Prism environment first looks in the current directory for a file called `.prisminit`. If the file is there, the Prism environment uses it. If the file isn't there, the Prism environment looks for it in your home directory. If the file isn't in either place, the Prism environment starts up without executing a `.prisminit` file.

The `.prisminit` file is useful if there are commands that you always want to execute when starting the Prism environment. For example,

- If you always want to log command output, put a `log` command in the file; see “Logging Commands and Output” on page 36.

- If you want to use your own aliases for Prism commands, put the appropriate `alias` commands in the file; see “Creating Aliases for Commands and Variables” on page 203.

Note that you don’t need to put `pushbutton` or `tearoff` commands into the `.prisminit` file, because changes you make to the tear-off region are automatically saved when you leave the Prism environment; see “Customizing MP Prism Mode” on page 200.

In the `.prisminit` file, the Prism environment interprets lines enclosed between C-style comment characters, `/*` and `*/`, as comments. If `\` is the final character on a line, the Prism environment interprets it as a continuation character.

Customizing MP Prism Mode

Using the `.prisminit` file, you can reserve commands in your `.prisminit` file exclusively for debugging multiprocess programs by bracketing the commands with `#ifdef MP` and `#endif`. For example, the following command sequence defines `c` to aliases in the scalar and MP modes of the Prism environment and sets the initial `pset` to 0 (zero) in the MP Prism mode.

```
alias c cont
#ifdef MP
pset 0
alias c "cont; wait every"
#endif
```

To provide this feature, the Prism environment must preprocess the `.prisminit` file; by default it does not do this.

▼ To Force the Prism Environment to Preprocess the `.prisminit` File

- **Change the setting of the Prism resource `Prism.cppPath`, and specify the path to your C preprocessor as its setting.**

Typically, this setting is `/lib`. Thus, you would set the resource as follows:

```
Prism.cppPath: /lib
```

See “Changing Prism Environment Defaults” on page 209 for information on setting the Prism environment’s resources.

Note, however, that the commands-only mode of the Prism environment is not aware of the settings of Prism resources such as `Prism.cppPath`, unless the settings are contained in the system-wide Prism `app-defaults` file.

Using the Tear-Off Region

You can place frequently used menu selections and commands in the tear-off region below the menu bar. They become buttons that you can click on to execute functions. FIGURE 8-1 shows the buttons that are in this region by default.



FIGURE 8-1 Tear-Off Region

Putting menu selections and commands in the tear-off region lets you access them without having to pull down a menu or issue a command from the command line.

Changes you make to the tear-off region are saved when you leave the Prism environment; see “Where the Prism Environment Stores Your Changes” on page 208 for more information.

Adding Menu Selections to the Tear-Off Region

You can add menu selections to the tear-off region from either the menu bar or the command line.

▼ To Add a Menu Selection to the Tear-Off Region

- **Perform one of the following:**
 - From the menu bar – Enter tear-off mode by choosing Tear-off from the Utilities menu. A dialog box appears that describes tear-off mode; see FIGURE 8-2.

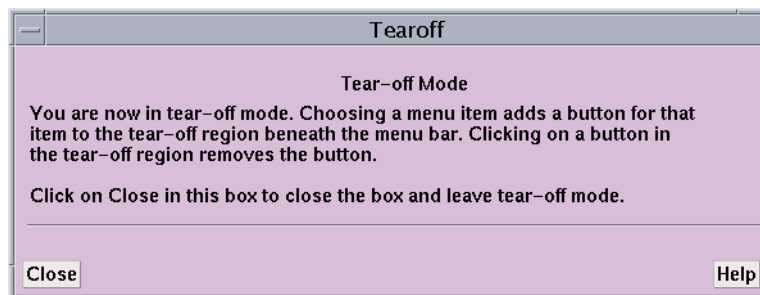


FIGURE 8-2 Tear-Off Region Dialog Box

While the dialog box is on the screen, choosing any selection from a menu adds a button for this selection to the tear-off region. Clicking a button in the tear-off mode removes that button.

If you fill up the region, you can resize it to accommodate more buttons. To resize the region, drag the small resize box at the bottom right of the region.

Click Close or press the Esc key while the mouse pointer is in the dialog box to close the box and leave tear-off mode.

When you are not in tear-off mode, clicking on a button in the tear-off region has the same effect as choosing the equivalent selection from a menu.

- From the command window - Use the `tearoff` and `untearoff` commands from the command window to add menu selections to and remove them from the tear-off region, respectively. Put the selection name in quotation marks. Case does not matter, and you can omit spaces and the ellipsis (...) that indicates the selection displays a window or dialog box. If the selection name is ambiguous, put the menu name in parentheses after the selection name. For example,

```
(prism all) tearoff "print (events)"
```

adds a button for the Print selection from the Events menu to the tear-off region.

Adding Prism Commands to the Tear-Off Region

▼ To Add a Command to the Tear-Off Region

- **Type**

```
(prism all) pushbutton label command
```

The *label* must be a single word. The *command* can be any valid Prism command, along with its arguments.

For example,

```
(prism all) pushbutton printa print a on dedicated
```

adds a button labeled `printa` to the tear-off region. Clicking on it executes the command `print a on dedicated`.

Creating Aliases for Commands and Variables

The Prism environment provides commands that let you create alternative names for commands, variables, and expressions.

▼ To Create an Alias for a Prism Command

- **Type**

```
(prism all) alias new-name command
```

For example,

```
(prism all) alias ni nexti
```

makes `ni` an alias for the `nexti` command. The Prism environment provides some default aliases for common commands. Issue `alias` with no arguments to display a list of the current aliases.

▼ To Remove an Alias

- **Type**

```
(prism all) unalias new-name
```

For example,

```
(prism all) unalias ni
```

removes the alias created above.

▼ To Set Up an Alternative Name for a Variable or Expression

- **Type**

```
(prism all) set variable = expression
```

For example,

```
(prism all) set alan = annoyingly_long_array_name
```

abbreviates the annoyingly long array name to `alan`. You can use this abbreviation subsequently in your program to refer to this variable. Use the `unset` command to remove a setting. For example,

```
(prism all) unset alan
```

removes the setting created above.

Changes you make via `alias` and `set` last for your current Prism session. To make them permanent, you can add the appropriate commands to your `.prisminit` file, which is described in “Initializing the Prism Environment” on page 199.

Changing Prism Resource Defaults

Many aspects of the Prism environment’s behavior and appearance—for example, the colors it displays on color workstations and the fonts it uses for text—are controlled by the settings of *Prism resources*. The default settings for many of these resources appear in the file `Prism` in the X11 `app-defaults` directory for your system. Your system administrator can change these system-wide defaults. You can override these defaults in two ways:

- For many defaults, you can use the *Customize* selection from the *Utilities* menu to display a window in which you can change the settings. This section describes this method.
- A more general method is to add an entry for a resource to your X resource database, as described in the next section. Using the *Customize* utility is much more convenient, however.

▼ To Launch the Prism Customize Utility

- **Choose Customize from the Utilities menu.**

This displays the window shown in FIGURE 8-3.

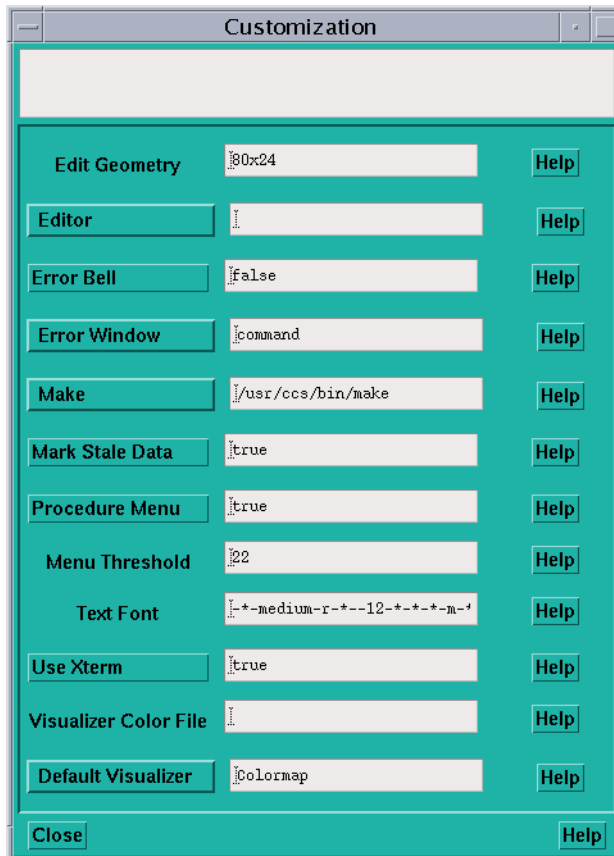


FIGURE 8-3 Customize Window

Changing a Resource Setting

On the left of the Customize window are the names of the resources. Next to each resource is a text-entry box that contains the resource's setting (if any). To the right of the fields are Help buttons. Clicking on a Help button or anywhere in the text-entry field displays help about the associated resource in the box at the top of the window.

▼ To Set a Value for a Prism Resource

- **Perform one of the following:**
 - For Edit Geometry, Menu Threshold, Text Font, and Visualizer Color File, enter the setting in the resource's text-entry box.

- For Editor, Error Window, and Make, left-click on the button labeled with the resource's name. This displays a menu of choices for the resource. Clicking on one of these choices displays it in the resource's text-entry box. For Editor and Make, you can also enter the setting directly in the text-entry box.
- For Error Bell, Procedure Menu, Mark Stale Data, and Use Xterm, there are only two possible settings, true and false; clicking on the button labeled with the resource's name toggles the current setting.

Whenever you make a change in a text-entry box, Apply and Cancel buttons appear to the right of it. Click on Apply to save the new setting; it takes effect immediately. Click on Cancel to cancel it; the setting changes back to its previous value.

▼ To Close the Customize Window

- Click on Close or press the Esc key.

Resource Descriptions

The following list summarizes the X Windows system resources that can be modified through the Customize window.

- `Edit Geometry` - Use this resource to specify the X geometry string for the editor created by the Edit and Email selections from the Utilities menu. The geometry string specifies the number of columns and rows and, optionally, the left and right offsets from the corner of the screen. The Prism environment's default is 80x24 (that is, 80 rows and 24 columns). See your X documentation for more information on X geometries.
- `Editor` - Use this resource to specify the editor that the Prism environment is to invoke when you choose the Edit selection from the Utilities menu. Click on the Editor box to display a menu of possible choices. If you leave this field blank, the Prism environment uses the setting of your `EDITOR` environment variable to determine which editor to use.
- `Error Bell` - Use this resource to specify how the Prism environment is to signal errors. Choosing `true` tells the Prism environment to ring the bell of your workstation. Choose `false` to have the Prism environment flash the screen instead; this is the Prism environment's default.
- `Error Window` - Use this resource to tell the Prism environment where to display the Prism environment's error messages. Choose `command` to display them in the command window; this is the Prism environment's default. Choose `dedicated` to send the messages to a dedicated window; the window is updated each time a new message is received. Choose `snapshot` to send each message to a separate window.

- **Make** – Use this resource to tell the Prism environment which `make` utility to use when you choose the Make selection from the Utilities menu. The Prism environment’s default is the standard Solaris `make` utility, `/usr/ccs/bin/make`. Click on the Make box to display a menu of possible choices.
- **Mark Stale Data** – Use this resource to tell the Prism environment whether to mark stale data with distinctive stripes. Choose `true` (the default) to have the Prism environment draw diagonal lines over the data; choose `false` to leave the visualizer’s appearance unchanged.
- **Procedure Menu** – Use this resource to specify whether a menu is to be displayed when you set a breakpoint in a Fortran 90 generic procedure. If you choose `true` (the default), a menu of possible procedures is displayed. You can then choose the procedure(s) in which the breakpoint is to be set. Choose `false` if you want breakpoints to be set automatically in all the generic procedures.
- **Menu Threshold** – Use this resource to specify the maximum number of procedures that are to be displayed in a menu when you perform an action on a Fortran 90 generic procedure, such as setting a breakpoint. The default is 22. Enter 0 to indicate that there should be no maximum. If the number of procedures exceeds the specified threshold, you are prompted to either enter the procedure name or display the menu.
- **Text Font** – Use this resource to specify the name of the X font that the Prism environment is to use for histogram labels and text in visualizers. The default, `8x13`, is a 12-point fixed-width font. To list the fonts available on your system, issue the Solaris command `xlsfonts`. Specifying a font much larger than the default can cause display problems, because the Prism environment doesn’t resize windows and buttons to accommodate the larger font.
- **Use Xterm** – Use this resource to tell the Prism environment what to do with the program’s I/O. Specify `true` (the Prism environment default) to tell the Prism environment to create an Xterm in which to display the I/O. Specify `false` to send the I/O to the Xterm from which you started the Prism environment.
- **Visualizer Color File** – Use this resource to tell the Prism environment the name of a file that specifies the colors to be used in `colormap` visualizers. If you leave this field blank, the Prism environment uses gray for elements whose values are not in the context you specify. For elements whose values are in the context, it uses black for values below the minimum, white for values above the maximum, and a smooth spectral map from blue to red for all other values.
- **Default Visualizer** – Use this resource to tell the Prism environment which representation you want to use as your initial representation when you display data in a visualizer. If you leave this field blank, the Prism environment uses `Text` for the initial representation.

If you specify a visualizer color file, that file must be in ASCII format. Each line of the file must contain three integers between 0 and 255 that specify the red, green, and blue components of a color.

The first line of the visualizer color file must contain the color that is to be displayed for values that fall below the minimum you specify in creating the visualizer. The next-to-last line must contain the color for values that exceed the maximum. The last line must contain the color used to display the values of elements that are not in the context specified by the user in a `where` statement. The Prism environment uses the colors in between to display the values falling between the minimum and the maximum. See TABLE 8-1 for an example.

TABLE 8-1 Sample Visualizer Colors

Red	Green	Blue
0	0	0
255	0	0
255	255	0
0	255	0
0	255	255
0	0	255
255	0	255
255	255	255
100	100	100

Like the default settings, this file specifies black for values below the minimum, white for values above the maximum, and gray for values outside the context. But the file reverses the default spectral map for other values: from lowest to highest, values are mapped red-yellow-green-cyan-blue-magenta.

Where the Prism Environment Stores Your Changes

The Prism environment maintains a file called `.prism_defaults` in your home directory. In it, the Prism environment keeps:

- Changes you make to the Prism environment via the `Customize` utility
- Changes you make to the tear-off region
- Changes you make to the size of the panes within the main Prism window

Do not attempt to edit this file; make all changes to it through the Prism environment itself. If you remove this file, you get the default configuration the next time you start the Prism environment.

Changing Prism Environment Defaults

As mentioned in the previous section, you can change the settings of many Prism resources either by using the `Customize` utility or by adding them to your X resource database. This section describes how to add a Prism resource to your X resource database.

An entry in the X resource database has the form:

resource-name: *value*

where *resource-name* is the name of the Prism resource, and *value* is the value to which it is set. TABLE 8-2 lists the Prism resources.

TABLE 8-2 Prism Resources

Resource	Use
<code>Prism.comm1Color</code>	Specifies the color of the first communicator displayed in the MPI queue visualizer.
<code>Prism.comm2Color</code>	Specifies the color of the second communicator displayed in the MPI queue visualizer.
<code>Prism.comm3Color</code>	Specifies the color of the third communicator displayed in the MPI queue visualizer.
<code>Prism.commOtherColor</code>	Specifies the color of the fourth communicator displayed in the MPI queue visualizer.
<code>Prism.cppPath</code>	Specifies the path to your C preprocessor.
<code>Prism.dialogColor</code>	Specifies the color for dialog boxes.
<code>Prism.editGeometry</code>	Specifies the size and placement of the editor window.
<code>Prism.editor</code>	Specifies the editor to use.
<code>Prism.errorBell</code>	Specifies whether the error bell is to ring.
<code>Prism.errorWin</code>	Specifies the window to use for error messages.
<code>Prism*fontList</code>	Specifies the font for labels, menu selections, and so on.

TABLE 8-2 Prism Resources (*Continued*)

Resource	Use
<code>Prism.graphBgColor</code>	Specifies the background color of all graphics windows, such as the structure browser, Where graph, and visualizer.
<code>Prism.graphFillColor</code>	Specifies the interior fill color for objects in graphics windows that have 3-D shadow borders.
<code>Prism.helpBrowser</code>	Specifies the browser to use for displaying help.
<code>Prism.helpUseExisting</code>	Specifies whether to use a currently running browser for displaying help.
<code>Prism.mainColor</code>	Specifies the main background color for Prism.
<code>Prism.make</code>	Specifies the make utility to use.
<code>Prism.markStaleData</code>	Specifies how Prism is to mark stale data in visualizers.
<code>Prism.procMenu</code>	Specifies whether a menu is displayed when setting a breakpoint in a Fortran 90 generic procedure.
<code>Prism.procThresh</code>	Changes the maximum number of specific procedures automatically shown when performing an action on a Fortran 90 generic procedure.
<code>Prism.spectralMapSize</code>	Specifies the size of the default spectral color map for color visualizers.
<code>Prism.textBgColor</code>	Specifies the background color for widgets containing text.
<code>Prism.textFont</code>	Specifies the text font to use for certain labels.
<code>Prism.textManyFieldTranslations</code>	Specifies the keyboard translations for dialog boxes that contain several text fields.
<code>Prism.textMasterColor</code>	Specifies the color used to highlight the master pane in a split source window.
<code>Prism.textOneFieldTranslations</code>	Specifies the keyboard translations for dialog boxes that contain one text field.
<code>Prism.useXterm</code>	Specifies whether to use a new Xterm for I/O.

TABLE 8-2 Prism Resources (*Continued*)

Resource	Use
<code>Prism.vizColormap</code>	Specifies the colors to be used in colormap visualizers.
<code>Prism.vizRepresentation</code>	Specifies the initial representation to be used when displaying data in visualizers.
<code>Prism*XmText.fontList</code>	Specifies the text font to use for most running text.

Note that the defaults mentioned in the following sections are the defaults for the Prism environment as it was shipped; your system administrator may have changed these in the Prism environment's file in your system's `app-defaults` directory.

Note also that the commands-only mode of the Prism environment is not aware of the settings of any Prism resources, unless they are contained in the Prism environment's `app-defaults` file. This matters only for the resource `Prism.cppPath`.

Adding Prism Resources to the X Resource Database

The X resource database keeps track of default settings for programs running under X. Use the `xrdb` program to add a Prism resource to this database.

▼ To Add Resource Settings to the X Resource Database

- **Perform one of the following:**

- Use the `-merge` option to specify the resource and its setting from the standard input. For example, type the following command to specify a default editor:

```
% xrdb -merge
Prism.editor: emacs
```

- Put resource settings in a file, then merge the file into the database. For example, if your changes are in `prism.defs`, you could issue this command:

```
% xrdb -merge prism.defs
```

Note – You must include the `-merge` option. Otherwise, what you type replaces the contents of your database. The new settings take effect the next time you start the Prism environment.

▼ To Signal the End of Input

- **Press the Ctrl-D keys.**

Consult your X documentation for more information about `xrdb`.

Specifying the Editor and Its Placement

▼ To Specify an Editor and Its Placement

- **Change the following:**

- The setting of the `Prism.editor` resource

This resource specifies which editor is to be invoked when you choose Edit from the Utilities menu or issue the corresponding command.

- The setting of the `Prism.editGeometry` resource

This resource specifies the X geometry string to the selected editor. The geometry string specifies the number of columns and rows and the left and right offsets from the corner of the screen.

You can also change the settings of these resources using the `Customize` window; see “Changing Prism Resource Defaults” on page 204 for more information.

Specifying the Window for Error Messages

▼ To Specify the Window for Error Messages

- **Change the setting of the `Prism.errorwin` resource.**

This resource specifies the window the Prism environment is to use for error messages. The predefined values are `command`, `dedicated`, and `snapshot`. You can also specify your own name for the window.

You can also change the setting of this resource via the `Customize` utility; see “Changing Prism Resource Defaults” on page 204 for more information.

Changing the Text Fonts

You may need to change the fonts that the Prism environment uses. This will be necessary, for example, if the default fonts used by the Prism environment are not available on your system. Use the resources described below to make this change.

▼ To List the Names of the Fonts Available on Your System

- **Type**

- % `xlsfonts`

▼ To Specify a Different Prism Font

You should try to substitute a font that is about the same size as the default value of the Prism environment. Substituting a font that is much larger can cause display problems, since the Prism environment does not resize windows and buttons to accommodate larger fonts.

- **Edit the `Prism.textFont` resource.**

This setting specifies the resource that the Prism environment is to use in displaying the labels of histograms and text in visualizers. By default, the Prism environment uses a 12-point, fixed-width font for this text.

You can also change the setting of this resource via the `Customize` utility; see “Changing Prism Resource Defaults” on page 204 for more information.

- **Change the setting of the `Prism*XmText.fontList` resource.**

This setting specifies the font used for most of the running text in the Prism environment, such as code in the source window. By default, the Prism environment uses a 12-point, fixed-width font for this text.

- **Change the setting of the `Prism*fontList` resource.**

This setting specifies the font used for everything else, such as menu selections, pushbuttons, and list items. By default, the Prism environment uses a 14-point Helvetica font for this text.

Changing Colors

The Prism environment provides several resources for changing the default colors it uses when running on a color workstation.

▼ To Change the Colors Used for Colormap Visualizers

- **Change the setting of the `Prism.vizColormap` resource.**

This setting specifies a file containing the colors to be used. You can also change the setting of this resource via the `Customize` utility; see “Changing Prism Resource Defaults” on page 204. See “Resource Descriptions” on page 206 for a discussion of how to create a visualizer color file.

- **Change the setting of the `Prism.spectralMapSize` resource.**

This setting specifies how large the default spectral color map is to be for colormap visualizers. The default is 100 entries. If this many entries causes problems on your workstation, use this resource to specify fewer entries. To set the default to 50, for example, set the resource in your X resource database as follows:

```
Prism.spectralMapSize: 50
```

▼ To Change the Prism Environment’s Standard Colors

- **Change the setting of the `Prism.dialogColor` resource.**

This setting specifies the background color of dialog boxes.

- **Change the setting of the `Prism.textBgColor` resource.**

This setting specifies the background color for text in buttons, dialog boxes, and so on. Note that this setting overrides the setting of the X toolkit `-bg` option.

- **Change the setting of the `Prism.textMasterColor` resource.**

This setting specifies the color used to highlight the master pane when the source window is split.

- **Change the setting of the `Prism.graphFillColor` resource.**

This setting specifies the interior fill color for objects in graphics windows that have 3-D shadow borders.

- **Change the setting of the `Prism.graphBGColor` resource.**

This setting specifies the background color of all graphics windows, such as the structure browser, Where graph, and visualizer.

- **Change the setting of the `Prism.mainColor` resource.**

This setting specifies the color used for just about everything else.

The defaults are:

```
Prism.dialogColor: Thistle
Prism.textBgColor: snow2
Prism.textMasterColor: black
Prism.graphFillColor: grey
Prism.graphBGColor: light grey
Prism.mainColor: light sea green
```

▼ Changing the Colors of MPI Communicators in the MPI Queue Visualizer

- **Change the setting of the `Prism.comm1Color` resource.**

This setting specifies the color of the first communicator displayed in the MPI queue visualizer.

- **Change the setting of the `Prism.comm2Color` resource.**

This setting specifies the color of the second communicator displayed in the MPI queue visualizer.

- **Change the setting of the `Prism.comm3Color` resource.**

This setting specifies the color of the third communicator displayed in the MPI queue visualizer.

- **Change the setting of the `Prism.commOtherColor` resource.**

This setting specifies the color of the fourth communicator displayed in the MPI queue visualizer.

The defaults are:

```
Prism.comm1Color: chartreuse2
Prism.comm2Color: cyan2
Prism.comm3Color: magenta2
Prism.commOtherColor: purple
```

Changing Keyboard Translations

You can change the keys and key combinations that the Prism environment translates into various actions. In general, doing this requires an understanding of X and Motif programming. You may be able to make some changes, however, by reading this section and studying the defaults in the Prism environment's file in your system's `app-defaults` directory.

Changing Keyboard Translations in Text Widgets

▼ To Change Keyboard Translations for Dialog Boxes With a Single Text Field

- **Change the settings of the `Prism.textOneFieldTranslations` resource.**

This controls default keyboard translations for dialog boxes that contain only one text field. Its default definition is:

```
Prism.textOneFieldTranslations:
<Key>osfDelete: delete-previous-character()
  <Key>osfBackSpace: delete-previous-character()
    Ctrl<Key>u: erase_to_beginning()
    Ctrl<Key>k: erase_to_end()
    Ctrl<Key>d: delete_char_at_cursor_position()
    ctrl<Key>f: move_cursor_to_next_char()
    Ctrl<Key>h: move_cursor_to_prev_char()
    Ctrl<Key>b: move_cursor_to_prev_char()
    Ctrl<Key>a: move_cursor_to_beginning_of_text()
    Ctrl<Key>e: move_cursor_to_end_of_text()
```

The definitions with `osf` in them are special Motif keyboard symbols.

▼ To Change Keyboard Translations for Dialog Boxes With Several Text Fields

- **Change the settings in the `Prism.textManyFieldTranslations` resource.**

Its default definition is:

```
Prism.textManyFieldTranslations:
  <Key>osfDelete: delete-previous-character()
  <Key>osfBackSpace: delete-previous-character()
  <Key>Return: next-tab-group()
  <Key>KP_Enter: next-tab-group()
  Ctrl<Key>u: erase_to_beginning()
  Ctrl<Key>k: erase_to_end()
  Ctrl<Key>d: delete_char_at_cursor_position()
  Ctrl<Key>f: move_cursor_to_next_char()
  Ctrl<Key>h: move_cursor_to_prev_char()
  Ctrl<Key>b: move_cursor_to_prev_char()
  Ctrl<Key>a: move_cursor_to_beginning_of_text()
  Ctrl<Key>e: move_cursor_to_end_of_text()
```

If you make a change to any field in one of these resources, you must copy all the definitions.

Changing General Motif Keyboard Translations

The Prism environment uses the standard Motif translations that define the general mappings of functions to keys. They are shown below.

```
*defaultVirtualBindings:
  osfActivate : <Key>Return
  osfAddMode : Shift <Key>F8
  osfBackSpace : <Key>BackSpace
  osfBeginLine : <Key>Home
  osfClear : <Key>Clear
  osfDelete : <Key>Delete
  osfDown : <Key>Down
  osfEndLine : <Key>End
  osfCancel : <Key>Escape
  osfHelp : <Key>F1
  osfInsert : <Key>Insert
  osfLeft : <Key>Left
  osfMenu : <Key>F4
  osfMenuBar : <Key>F10
  osfPageDown : <Key>Next
  osfPageUp : <Key>Prior
  osfRight : <Key>Right
  osfSelect : <Key>Select
  osfUndo: <Key>Undo
  osfUp : <Key>Up
```

▼ To Change a General Motif Keyboard Translation

- **Change its entry in the `*defaultVirtualBindings` resource.**

For example, if your keyboard does not have an F10 key, you could edit the `osfMenuBar` line and substitute another function key.

Note the following points in changing this resource:

- All entries in the resource must be included in your resource database if you want to change any of them. Otherwise, the omitted entries are undefined.
- The entries in this resource apply to all Motif-based applications. If you want your changes to apply only to the Prism environment, change the first line of the resource to `Prism*defaultVirtualBindings`.

Changing Xterm Use With I/O

By default, the Prism environment creates a new Xterm for input to and output from a program.

▼ To Prevent the Prism Environment From Creating New I/O Windows

- **Set the `Prism.useXterm` resource to `false`.**

This setting causes I/O to go to the Xterm from which you invoked the Prism environment. You can also change the setting of this resource via the `Customize` utility; see “Changing Prism Resource Defaults” on page 204.

Changing the Way the Prism Environment Signals an Error

By default, the Prism environment flashes the command window when there is an error. You can, instead, use the bell for signaling errors.

▼ To Force the Prism Environment to Ring the Bell on Errors

- **Perform one of the following:**
 - Set the resource `Prism.errorBell` to `true`.
 - Change the setting of the `Prism.errorBell` resource using the `Customize` utility; see “Changing Prism Resource Defaults” on page 204.

Changing the make Utility to Use

By default, the Prism environment uses the standard Solaris `/usr/ccs/bin/make` utility. You can use a different `make` utility by changing the setting of the `Prism.make` resource.

▼ To Specify an Alternative Make Utility

- **Perform one of the following:**

- Change the setting of the resource `Prism.make`.

This resource specifies the path name of another version of `make` to use.

- Change the setting of the `Prism.make` resource using the `Customize` utility; see “Changing Prism Resource Defaults” on page 204.

Changing How the Prism Environment Treats Stale Data in Visualizers

By default, the Prism environment prints diagonal lines over *stale* data in visualizers. Data are considered stale when the program has continued execution beyond the spot where the data were collected.

▼ To Prevent the Prism Environment From Depicting Stale Data With Diagonal Lines

- **Perform one of the following:**

- Change the setting of the resource `Prism.markStaleData` to `false`.
- Change the setting of the `Prism.markStaleData` resource using the `Customize` utility; see “Changing Prism Resource Defaults” on page 204.

Specifying a Different Browser for Displaying Help

There are several resources you can use to affect the way help is displayed.

By default, the graphical mode of the Prism environment uses the Netscape browser to display help information; see “Using the Browser-Based Help System” on page 196.

▼ To Specify an Alternative HTML Browser for Displaying Online Help

- **Set the `Prism.helpBrowser` resource to the executable name of the other browser.**

The name of the browser must be on your path. The graphical mode of the Prism environment supports HTML browsers, such as Netscape. You can include in the setting any browser-specific options that you want passed to the browser when the Prism environment starts it up.

These options do not take effect if the Prism environment uses an existing browser. If you already have a browser running when you request help from the Prism environment, the Prism environment displays the help information in this browser.

▼ To Force the Prism Environment to Start a New Help Browser

- **Set the resource `Prism.helpUseExisting` to `false`.**

This forces the Prism environment to start a new browser.

To restore the default behavior, set `Prism.helpUseExisting` to `true`.

Changing the Way the Prism Environment Handles Fortran 90 Generic Procedures

There are two resources you can use to change the way the Prism environment handles Fortran 90 generic procedures.

By default, the Prism environment displays a menu (in the commands-only mode of the Prism environment) or a dialog box when you attempt to set a breakpoint in a Fortran 90 generic procedure.

▼ To Suppress the Display of Menus or Dialog Boxes When Setting Breakpoints in Fortran 90 Generic Procedures

- **Perform one of the following:**

- Change the setting of the Prism resource `Prism.procMenu` to `false`.

This setting specifies that the Prism environment is to set the breakpoint in every one of these procedures, without displaying a menu or dialog box.

- Change the setting of the resource `Prism.procMenu` using the `Customize` utility; see “Changing Prism Resource Defaults” on page 204.

By default, the commands-only interface of the Prism environment displays a maximum of 22 procedures on a menu when you attempt to perform an action (like setting a breakpoint) on a Fortran 90 generic procedure. If there are more than this number of specific procedures, the Prism environment asks you whether you want to specify the name of a specific procedure or to view a menu.

▼ To Display a Different Maximum Number of Fortran 90 Generic Procedures

- **Change the setting of the `Prism.procThresh` resource.**

This specifies a different maximum number of procedures. Set the resource to 0 (zero) to specify that there is to be no maximum.

Troubleshooting

This chapter discusses ways in which you can recognize and avoid potential difficulties when using the Prism environment. It is organized into the following sections:

- “Avoid Using the `-xs` Compiler Option” on page 223
- “Keep `.o` Files after Compilation” on page 223
- “Expect a Pause After Issuing the First run Command” on page 224
- “Monitor Your Use of Color Resources” on page 224
- “Expect Only Stopped Processes to Be Displayed in the Where Graph” on page 224
- “Use Only the MP Mode of the Prism Environment to Load MPI Programs” on page 225
- “Verify That `/opt/SUNWlsf/bin` Is in Your PATH” on page 225

Avoid Using the `-xs` Compiler Option

Loading code compiled with the `-xs` option can require long load times. The Prism environment does not require that you compile code with the `-xs` option.

Keep `.o` Files after Compilation

If you have not used `-xs` during compiling, do not move or delete the `.o` files of the program that you want to load into the Prism environment. If you move or delete `.o` files, the Prism environment can find no debugging information for the functions in those files, even though the final executable was compiled with the `-g` option.

Expect a Pause After Issuing the First `run` Command

The multiprocess mode of the Prism environment (*MP Prism*) may pause for an unexpectedly long time after you issue the `run` command. During this pause, the user interface is unresponsive. This pause is unavoidable and is due to the delay caused while loading either the LSF or CRE environment. The `run` command goes to completion.

Monitor Your Use of Color Resources

The Prism environment may issue messages indicating that it needs additional color resources. For example,

```
Can't allocate color for snow2
```

When that happens, shut down any unnecessary color applications and try again.

To reduce the likelihood of exhausting color resources, you can launch the Prism environment with the `-install` argument. This creates a private colormap for the Prism environment at startup.

Expect Only Stopped Processes to Be Displayed in the Where Graph

The Prism environment does not show all processes in the Where graph. The Where graph shows only the stacks of stopped processes.

Use Only the MP Mode of the Prism Environment to Load MPI Programs

Attempting to use the scalar mode of the Prism environment to run an MPI program can cause the Prism environment to abort the process and issue messages such as the following:

```
[unknown MPI_COMM_WORLD unknown] ERROR in MPI_Init:  
unclassified error: RTE_Init_lib:  
Job must be submitted to CRE: No such job  
Aborting.
```

To run an MPI program, you must launch the MP mode of the Prism environment. You launch it by specifying a number of processes to run. Use the `-n` option to specify the number of processes. For example,

```
% prism -n 4 a.out
```

launches the MP mode of the Prism environment and loads `a.out`.

Verify That `/opt/SUNWlsf/bin` Is in Your PATH

If LSF is your default runtime environment and if the directory containing LSF executables is not set in your `PATH` variable, attempting to launch the MP mode of the Prism environment will fail. For example,

```
hpc-450-3 44 =>prism -n 0 &  
[1] 26614  
hpc-450-3 45 =>/opt/SUNWhpc/bin/prism: bsub: not found  
[1] Exit 1 prism -n 0
```

MP Prism Environment Freezes on Startup

If you attempt to run the message-passing Prism environment in the background without specifying the `-n` flag to `mprun`, the Prism process hangs.

Once hung, you can either kill the Prism process manually from another terminal window using the `mpkill` command or put the Prism process in the foreground in your shell with the `fg` command, thereby releasing it from the freeze.

To run an MP Prism process in the background, add the `-n` flag to the `mprun` command line.

For example,

```
% mprun -n -np 4 prism &
```

Commands-Only Mode of the Prism Environment

You can run the Prism environment in a commands-only mode, without the graphical interface. This is useful if you don't have access to a terminal or workstation running X. All of the functionality of the Prism environment is available in commands-only mode except features that require graphics (for example, visualizers). See “Specifying the Commands-Only Option” on page 228.

If you are using an X-terminal, you can also run a commands-only mode of the Prism environment that lets you redirect the output of certain commands to X windows. This may be preferable to users who are used to a command-line interface for debugging but want to take advantage of some of the Prism environment's graphical features. See “Running the Commands-Only Mode of the Prism Environment From an X-terminal: the `-CX` Option” on page 230.

For further information on individual commands, read the sections of the main body of this guide that deal with the commands, and read the reference descriptions in the *Prism Reference Manual*.

This appendix discusses the following topics:

- “Specifying the Commands-Only Option” on page 228
- “Issuing Commands” on page 228
- “Useful Commands” on page 229
- “Leaving the Commands-Only Mode of the Prism Environment” on page 230
- “Running the Commands-Only Mode of the Prism Environment From an X-terminal: the `-CX` Option” on page 230

Specifying the Commands-Only Option

To enter commands-only mode, specify the `-C` option on the `prism` command line. You can also include other arguments on the command line; for example, you can specify the name of a program, so that the Prism environment comes up with that program loaded. X toolkit options are, of course, meaningless. See “Starting the Prism Environment” on page 11 for more information on command-line options.

When you have issued the command

```
% prism -C a.out
```

you receive this prompt:

```
(prism all)
```

You can issue most Prism commands at this prompt, except for commands that apply specifically to the graphical interface; these include `pushbutton`, `tearoff`, and `untearoff`.

Issuing Commands

You operate in the commands-only mode of the Prism environment just as you do when issuing commands on the command line in the graphical mode of the Prism environment; output appears below the command you type, instead of in the history region above the command line. You cannot redirect output using the `on window` syntax. You can, however, redirect output to a file using the `@ filename` syntax.

The commands-only mode of the Prism environment supports the editing key combinations supported by the graphical mode of the Prism environment, plus some additional combinations. Here is the entire list:

- Ctrl-A – Moves to the beginning of the line.
- Ctrl-B (or Ctrl-H) – Moves back one character.
- Ctrl-C – Interrupts execution.
- Ctrl-D – Deletes the character under the cursor.
- Ctrl-E – Moves to the end of the line.
- Ctrl-F – Moves forward one character.
- Ctrl-J (or Ctrl-M) – Signals done with input (equivalent to pressing the Return key).

- Ctrl-K – Deletes to the end of the line.
- Ctrl-L – Refreshes the screen.
- Ctrl-N – Displays the next command in the commands buffer.
- Ctrl-P – Displays the previous command in the commands buffer.
- Ctrl-U – Deletes to the beginning of the line.

When printing large amounts of output, the commands-only mode of the Prism environment displays a `more?` prompt after every screen of text. Answer `y` or simply press the Return key to display another screen; answer `n` or `q`, followed by another Return, to stop the display and return to the `(prism)` prompt.

You can adjust the number of lines the Prism environment displays before issuing the `more?` prompt by issuing the `set` command with the `$page_size` variable that specifies the number of lines you want displayed. For example, issue this command to display 10 lines at a time:

```
(prism all) set $page_size = 10
```

Set the `$page_size` to 0 to turn the feature off; the Prism environment will not display a `more?` prompt.

Useful Commands

This section describes some commands that are especially useful in the commands-only mode of the Prism environment.

Use the `list` command to list source lines from the current file. For example,

```
(prism all) list 10, 20
```

prints lines 10 through 20 of the current file.

Use the `show events` command to print the events list. Use the `delete` command to delete events from this list.

Use the `set` command with the `$print_width` variable to specify the number of items to be printed on a line. The default is 1.

Leaving the Commands-Only Mode of the Prism Environment

Issue the `quit` command to leave the commands-only mode of the Prism environment and return to your Solaris prompt.

Running the Commands-Only Mode of the Prism Environment From an X-terminal: the `-CX` Option

Issue the `prism` command with the `-CX` option from an Xterm to start up an instance of the commands-only mode of the Prism environment that lets you redirect the output of certain commands to X windows. The information presented earlier in this chapter about the commands-only mode of the Prism environment also applies to this version, except that this version lets you redirect output using the `on window` syntax.

You can redirect the following output to X windows:

- Visualizers (including structure visualizers) - `print` or `display` command
- Where graph (MP Prism environment only) - `where` command
- Psets window (MP Prism environment only) - `show psets` command

To redirect the output, issue the appropriate command with the `on dedicated` or `on snapshot` syntax, just as you would in the graphical mode of the Prism environment. For example, this command displays a visualizer for `x` in a dedicated window:

```
(prism all) print x on dedicated
```

You can specify the type of the visualizer as well, by adding `as type` after the `on window` argument. For example:

```
(prism all) print x on dedicated as colormap
```

In addition, you can display help windows from within windows that you pop up in this way.

C++ and Fortran 90 Support

This appendix identifies the particular features of C++ and F90 programs that the Prism environment is able to debug and those it does not support. This discussion is organized into the following sections:

- “C++ Support in the Prism Environment” on page 231
 - “Fortran 90 Support in the Prism Environment” on page 234
-

C++ Support in the Prism Environment

The Prism environment provides limited support for debugging C++ programs.

- “Fully Supported C++ Features” on page 231
- “Partially Supported C++ Features” on page 233
- “Unsupported C++ Features” on page 234

Fully Supported C++ Features

This section describes the C++ program features that, with few exceptions, are supported by the Prism environment.

Data Members in Methods

You can simply type `print member` to print a data member when in a class method.

C++ Linkage Names

You can set breakpoints using the `stop in` command with functions having either C or C++ linkage (mangled) names.

Methods of a Class

You can use the Prism environment `stop in`, `func`, and `list` commands with methods of a class.

```
(prism all) stop in class_name::method_name
(prism all) func class_name::method_name
(prism all) list class_name::method_name
```

Class Member Variables

The Prism environment supports assignment to class member variables.

Variables of Class Type and Template Classes

You can use the `whatis` and `print` commands with variables of class type and template classes.

`this` Identifier

The Prism environment recognizes the `this` identifier in C++ methods. Its value also appears in stack backtraces.

Overloaded Method Names

The Prism environment allows you to set breakpoints in overloaded method names. A list pops up from which you can select the correct method.

Template Functions

The Prism environment allows you to set breakpoints in template functions. A list pops up from which you can select the correct function.

Scope Operator in the Prism Environment's Identifier Syntax

The Prism environment's identifier syntax recognizes the C++ scope operator, `::`. For example:

```
(prism all) whereis dummy  
variable: `symbol.x`symbol.cc`Symbol::print:71`dummy
```

Partially Supported C++ Features

This section describes the C++ program features that, with significant limitations, are supported by the Prism environment.

Casts

The Prism environment recognizes casting a class pointer to the class of a base type only for single-inheritance relationships. For example, the Prism environment recognizes the following cast syntax when printing variable P:

```
(prism all) print (struct class_name *) P  
(prism all) print (class class_name *) P  
(prism all) print (class_name *) P
```

Static Class Members

You can print static class members when the current scope is a class method. You cannot print static class members when not in class scope. For example, the following command will fail if you issue it outside of the scope of *class_name*:

```
(prism all) print class_name::var_name
```

Breakpoints in Methods

You cannot use a method name that has some forms of non-C identifier syntax to set a breakpoint. For example, this fails with a syntax error:

```
(prism all) stop in class_name::operator+
```

You must instead use `stop at line` syntax. These method names are correctly identified in a stack trace, however.

Unsupported C++ Features

This section identifies the C++ programming features that are not supported by the Prism environment.

Inlined Methods Used in Multiple Source Files

Using the Prism environment, you cannot set a breakpoint in an inlined method that is used in multiple source files. Only one of the several debuggable copies of the inlined function gets the breakpoint.

Calling C++ Methods

The Prism environment does not support calling C++ methods, using any syntax.

Variables of Type Reference

The Prism environment does not support printing variables of type reference, such as `int &xref`. Also, variables of type reference appear as `(unknown type)` in stack traces.

Fortran 90 Support in the Prism Environment

The Prism environment provides limited support for debugging F90 programs. This support is described in three sections, as follows:

- “Fully Supported Fortran 90 Features” on page 234
- “Partially Supported Fortran 90 Features” on page 239
- “Unsupported Fortran 90 Features” on page 240

Fully Supported Fortran 90 Features

This section describes the F90 program features that, with few exceptions, are supported by the Prism environment.

Derived Types

With the exception of constructors, the Prism environment supports derived types in Fortran 90. For example, given these declarations:

```
type point3
  integer x,y,z;
end type point3
type(point3) :: var,var2;
```

you can use Prism commands with these Fortran 90 variables:

```
(prism all) print var
(prism all) whatis var
(prism all) whatis point3
(prism all) assign var=var2
(prism all) print var%x
(prism all) assign var%x = 70
```

Generic Functions

The Prism environment fully supports generic functions in Fortran 90. For example, given the generic function `fadd`, declared as follows:

```
interface fadd
  integer function intadd(i, j)
    integer*4, intent(in) :: i, j
  end function intadd
  real function realadd(x, y)
    real, intent(in) :: x, y
  end function realadd
end interface
```

you can use Prism commands with these Fortran 90 generic functions:

```
(prism all) p fadd(1,2)
(prism all) whatis fadd
(prism all) stop in fadd
```

In each case, the Prism environment asks you which instance of `fadd` your command refers to. For example:

```
(prism all) whatis fadd
More than one identifier 'fadd'.
Select one of the following names:
0) Cancel
1) `f90_user_op_generic.exe`f90_user_op_generic.f90`fadd
! real*4 realadd
2) `f90_user_op_generic.exe`f90_user_op_generic.f90`fadd
! integer*4 intadd
> 1
real*4 function fadd (x, y)
(dummy argument) real*4 x
(dummy argument) real*4 y
```

Simple Pointers

In addition to the standard assignment operator (`=`), the Prism environment supports the new Fortran 90 pointer assignment operator `=>`. For example:

```
program pnode
type node
integer x,y
type(node), pointer :: next
end type node
type(node), target :: n1,n2,n3
type(node), pointer :: pn1, pn2
...
pn1 => n1
pn2 => n2
i = 0
end
```

The following examples assume that a breakpoint has been set at the last statement, `i = 0`, and show how the Prism environment supports Fortran 90 pointers:

- `print pn1` - Prints the value pointed to by `pn1`, in this case `n1`.
- `print pn1%x` - Prints the value of the member `x` in the object pointed to by `pn1` (in this case `n1%x`).
- `assign pn1%x = 3` - Assigns `n1%x = 3`.
- `assign pn1=n3` - Assigns `n3` to the value pointed to by `pn1` (this has the same effect as `assign n1=n3`).

- `assign pn1=>n3` - Makes `pn1` point to `n3`.
- `assign pn1=>pn2` - Makes `pn1` point to the same object as `pn2`.

Interactive Examples of Support for Fortran 90 Pointers

If `pn1` does not point to any value, an attempt to access it will result in an error message:

```
(prism all) p pn1
Fortran variable is not allocated/associated.
```

You can find the state of a pointer using the `what is` command. Assume `pn1` has not been associated:

```
(prism all) what is pn1
node pn1 ! unallocated f90 pointer
```

Assume `pn1` has been associated with a value:

```
(prism all) what is pn1
node pn1 ! f90 pointer
```

Pointers to Arrays

The Prism environment supports pointers to arrays in the same way that it supports simple pointers. The Fortran 90 language constraints apply. For example, Fortran 90 allows pointer assignment between pointers to arrays. Assignment to arrays having different ranks is not allowed.

For example, given these declarations:

```
real, dimension(10), target :: r_arr1
real, dimension(20), target :: r_arr2
real, dimension(:), pointer :: p_arr1,p_arr2
```

you can use Prism commands with these Fortran 90 pointers to arrays:

```
(prism all) print p_arr1
(prism all) what is p_arr2
(prism all) assign p_arr1 => r_arr1
(prism all) assign p_arr1(1:2) = 7
```

Pointers to Sections of an Array in Fortran 90

The Prism environment does not handle Fortran 90 pointers to array sections correctly. For example,

```
array_ptr => some_array(1:10:3)
```

The Prism environment will print some elements of the array, although it will not print the correct elements or the correct number of elements.

Allocatable Arrays

The Prism environment supports allocatable arrays in the same way that it supports pointers to arrays. Fortran 90 support includes the Prism commands `print` and `what is`. The Prism environment also supports slicing and striding Fortran 90 allocatable arrays. For example, to print a section of allocatable array `alloc_array`, type:

```
(prism all) print alloc_array(1:30:2)
```

Fortran 90 language constraints apply. For example, Fortran 90 allows allocating or deallocating memory for an allocatable array but does not allow an allocatable array to point to another object. Therefore, the Prism environment does not recognize pointer assignment, `=>`, to allocatable arrays.

Array Sections and Operations on Arrays

The Prism environment supports Fortran 90 operations on arrays or array sections, and assignment to continuous sections of arrays.

```
(prism all) assign a=b+c  
(prism all) assign a(3:7)=b(2:10:2)+c(8:8)
```

Masked Array Operations

The Prism environment supports Fortran 90 masked `print` statements:

```
(prism all) where (arr>0) print arr
```


Variable Attributes

The Prism `what is` command shows variable attributes. These attributes include allocated and associated attributes for pointers, or the (function variable) attribute displayed for a `RESULT` variable in Fortran 90.

Consider, for example, this declaration:

```
function inc(i) result(j)
  integer i;
  integer k;
  integer j;
  k = i+1 j = k
end function inc
```

In this situation, the `what is` command displays the function variable attribute of `j`:

```
(prism all) what is j
(function variable) integer*4 j
```

Partially Supported Fortran 90 Features

This section describes the Fortran 90 program features that, with significant limitations, are supported by the Prism environment.

User-Defined Operators

The Prism environment views user-defined operators as functions. If a new operator `.my_op.` appears in a Fortran 90 program, then the Prism environment cannot deal with the operator `.my_op.` as an operator, but it can deal with the function `my_op`,

viewed as a generic function. You cannot use operators named * (or +, or any other keyword operator), but you can stop in functions that are used to define such operators. For example:

```
interface operator(.add_op.)
  integer function int_add(i, j)
    integer*4, intent(in) :: i, j
  end function int_add
  real function real_add(x, y)
    real, intent(in) :: x, y
  end function real_add
end interface
```

In this example, the Prism environment does not support debugging the user-defined function `.add_op`:

```
(prism all) print 1 .add_op. 2
```

However, the Prism environment supports the function `add_op`:

```
(prism all) print add_op(1,2)
```

A list pops up, allowing you to choose which `add_op` to apply.

Internal Procedures

The following commands can take internal procedure names as arguments:

- `stop in`
- `whatis`

If there are several procedures with the same name, a list pops up from which to select the desired procedure.

Supported Ininsics

The Prism environment supports the same intrinsics in Fortran 90 that it supports in Fortran 77. See “Using Fortran Intrinsic Functions in Expressions” on page 39.

Unsupported Fortran 90 Features

This section identifies the F90 programming features that are not supported by the Prism environment.

Derived Type Constructors

The Prism environment does not support constructors for derived types.

```
type point3
  integer x,y,z;end
type point3
type(point3) :: var,var2;
```

The Prism environment does support assignment to derived types, however. For example:

```
(prism all) assign var = var2
```

Although Fortran 90 allows the use of constructors, the Prism environment does not support them. The following example is not supported:

```
(prism all) assign var = point3(1,2,3)
```

Generic Functions

Assume the generic function is defined in the current module as in this example:

```
interface fadd
  integer function intadd(i, j)
    integer*4, intent(in) :: i, j
  end function intadd
  real function realadd(x, y)
    real, intent(in) :: x, y
  end function realadd
end interface
```

In this situation, only references to the `fadd` are supported, but references to specific functions that define `fadd` are not. For example:

```
(prism all) whatis intadd
prism: "intadd" is not defined in the scope
`f90_user_op_generic.exe`f90_user_op_generic.f90`main`
```

Pointer Assignment Error Checking

The error checking involved by the semantics of the `=>` operator is not fully supported. If your program causes an illegal pointer assignment, the Prism environment might not issue any error, and the behavior of the program will be undefined.

Printing Array-Valued Functions

The Prism environment does not print the result of an array-valued function.

The Scalar Mode of the Prism Environment

When viewing serial programs, the Prism environment behaves differently than it does when viewing multiprocess programs. In this situation, the Prism environment operates in *scalar mode*.

The scalar mode of the Prism environment does not support psets, since pset-related features require multiple processes or threads. This appendix provides descriptions of other differences between the MP mode and the scalar mode of the Prism environment. This discussion is organized into the following sections:

- “Starting the Prism Environment” on page 243
- “Stepping and Continuing Through a Serial Program” on page 244
- “Viewing the Call Stack” on page 245

Starting the Prism Environment

▼ To Launch the Prism Environment in Scalar Mode

- **Type:**

`% prism program`

This starts the Prism environment for a nonthreaded single-process program, using the scalar mode of the Prism environment.

Do not launch the Prism environment as an argument to the `bsub` command (LSF) or the `mprun` command (CRE). It creates redundant instances of the Prism environment. For information on `bsub`, see the *LSF Batch User's Guide*. For information about `mprun`, see the *Sun MPI Programming and Reference Guide*.

You can specify other options on the `prism` command line. For example, you can specify the `-C` option to bring up the Prism environment with the commands-only interface, or the `-CX` option (from an Xterminal) to bring it up with the commands-only interface, but be able to send the output of certain commands to X windows.

Stepping and Continuing Through a Serial Program

When operating on a serial program, the scalar mode of the Prism environment (like most other debuggers) waits for a `step`, `next`, or `cont` command to finish executing before letting you issue most other commands.

Execution Pointer

In the scalar mode of the Prism environment, the `>` symbol in the line-number region points to the next line to be executed; see “Using the Line-Number Region” on page 31. In a message-passing program, there can be multiple execution points within the program. The MP mode of the Prism environment marks all the execution points for the processes in the current set by a `>` in the line-number region (or a `*` if the current source position is the same as the current execution point). Shift-click on this symbol to display a pop-up window that shows the process(es) for which the symbol is the execution pointer.

Attaching to a Running Serial Process

As described in “Attaching to a Running Serial Program” on page 16, you can load a running process into the Prism environment by specifying the name of the executable program and the process ID of the corresponding running process on the command line of the Prism environment.

You can also attach to a running process from within the Prism environment.

Note – To attach to the running process of a serial program, the process must be running on the same node as the Prism environment.

▼ To Attach To a Running Process From Within the Prism Environment

1. Find out the process's ID by issuing the Solaris command `ps`.
2. Load the executable program for the process into the Prism environment.
3. Issue the `attach` command on the command line of the Prism environment, using the process's ID as the argument.

With either method of attaching to the process, the process is interrupted; a message is displayed in the command window giving its current location, and its status is stopped. You can then work with the program in the Prism environment as you normally would. The only difference in behavior is that it does not display its I/O in a special X-terminal window; see “Program I/O” on page 56.

To detach from a running process, issue the command `detach` from the command line of the Prism environment. The process continues to run in the background from the point at which it was stopped in the Prism environment; it is no longer under the control of the Prism environment. Note that you can detach any process in the Prism environment via the `detach` command, not just processes that you have explicitly attached.

Note – Use the `kill` command to terminate the process or job (rather than releasing it to run in the background) currently running within the Prism environment.

Viewing the Call Stack

In the scalar mode of the Prism environment, choosing Where from the Debug menu displays the call stack for the program; see “To Display the Call Stack” on page 112. Note that a multiprocess or multithreaded program can have multiple call stacks, one for each process or thread. To show the relationships among these call stacks, the MP mode of the Prism environment provides a Where graph. For information about the Where graph in the MP mode of the Prism environment, see “Displaying the Where Graph” on page 113.

Index

SYMBOLS

' , 200
* , 59, 244
.prism_defaults, 208
.prisminit, 23, 199, 200
/ command, 28
/**/, 200
/bin/make, 192, 207
>, 59, 244
? command, 28

A

accessibility of variables, 121
 commands, 121
adjustable arrays, printing, 131
alias command, 200, 203
aliases, creating, 203
ALL intrinsic function, 39
all pset, 72
ANY intrinsic function, 40
app-defaults file, 204, 216
arrow keys, 25
 using to scroll through source window, 28
assembly code, displaying in split source
 window, 30
assign command, 121, 165
 not available when examining node core files, 53
attach command, 22, 53, 54, 245
 cannot be used in actions field, 97

augmenting data type information, 161

B

base
 changing for a specific value, 165
 changing the default, 131
 changing via the Options menu, 152
 specifying in print or display command, 137
bjobs command, 53
break pset, 65, 102
breakpoints
 deleting, 105, 107, 109
 setting, 104
 using commands to set, 107
 using the event table and Events menu to
 set, 106
browser, default for displaying help, 219

C

C++ support, 231
 calling C++ methods, 234
 cast syntax, 233
 class member variables, 232
 class methods, 231
 class scope, 233
 inlined methods, 234
 linkage (mangled) names, 232
 method names, 233
 methods of a class, 232

- overloaded method names, 232
- template classes, 232
- template functions, 232
- variables of class type, 232
- variables of type reference, 234
- call command, 121
- call stack
 - displaying, 112
 - moving through, 113
- CDE, 10
- changes, where Prism stores, 208
- Client/Server programs
 - debugging, 87
- CMPLX intrinsic function, 40, 146
- Colormap visualizers, 144
 - minimum and maximum values of, 149
- colormap visualizers, 6
- colormap visualizers, changing the size of the
 - default spectral color map, 214
- colors, changing Prism's standard, 214
- command line, 32
 - using, 33
- command window, 5
 - using, 32
- commands
 - adding to the tear-off region, 202
 - issuing, 26
 - issuing multiple, 33
 - logging, 36
 - setting up alternative names for, 203
- Commands Reference selection, 196
- commands-only mode, 227
- Common Events buttons, 97, 134
- compiler options, combining, 10
- compilers, supported, 10
- compiling and linking, 10
 - from within Prism, 192
- complex numbers, 142, 151
- cont command
 - in MP Prism, 244
- context
 - setting via print or display command, 136
- contw command, 58
 - cannot be used in event actions, 103
- core command, 52
 - cannot be used in actions field, 97

- not available in MP Prism, 53
- core files
 - associating with loaded programs, 52
 - working with, 22
- COUNT intrinsic function, 40
- Ctrl-A, 26, 228
- Ctrl-B, 26, 228
- Ctrl-C, 26, 33, 228
 - ending a wait in MP Prism, 58
- Ctrl-D, 26, 228
- Ctrl-E, 26, 228
- Ctrl-F, 26, 228
- Ctrl-H, 228
- Ctrl-J, 228
- Ctrl-K, 26, 229
- Ctrl-L, 229
- Ctrl-M, 228
- Ctrl-N, 33, 229
- Ctrl-P, 33, 229
- Ctrl-U, 26, 229
- Ctrl-X, 28
- Ctrl-Z, 60
- current execution point, returning to, 28
- current file, 88
- current function, 88
 - changing, 89
 - changing via the Where graph, 120
- current process, 78
- current pset, 75
 - and dynamic psets, 76
 - and variable psets, 77
 - changing via the Where graph, 120
 - setting, 75
- current working directory, changing and
 - printing, 45
- Customize selection, 204
- Customize utility, using, 204
- cycle command, 81, 171
- cycle pset, 80, 171
- Cycle window, 81, 170

D

- data navigator, 6

- data navigator, using, 139
- data type information, augmenting, 161
- dbx, 27
- dedicated window, 35, 133
- define pset command, 69
 - cannot be used in event actions, 103
- delete command, 111, 229
- delete pset command, 75
 - cannot be used in event actions, 103
- Delete selection, 99, 110
- detach command, 54, 245
 - cannot be used in actions field, 97
- disable command, 100
- display command, 121, 134
 - redirecting output to X window, 230
 - specifying the radix in, 137
 - with varfile intrinsic, 155
- Display dialog box, 134
- DISPLAY environment variable, 11
- Display selection (Debug menu), 132
 - in MP Prism, 168
- display window, using, 139
- displaying
 - difference from printing, 129
 - from the command window, 134
 - from the event table, 134
- Dither visualizers, 142
- done pset, 65
- Down selection, 113
- dump command, 121, 165, 166

E

- eachinst keyword, 96
- eachline keyword, 96
- edit geometry, 206
- Edit selection, 206, 212
- editing source code, 191
- EDITOR environment variable, 191, 206
- editor, specifying default, 212
- effects of optimization, 121
- enable command, 100
- environment variables, setting and displaying, 45
- error bell, 206

- error messages, specifying window for, 212
- error pset, 65
- error window, 206
- errors, Prism's behavior after, 218
- eval pset command, 70, 77, 102
- event list, 95, 109
- Event Table
 - description of, 95
 - using, 95
- events
 - adding, 97
 - and deleted psets, 104
 - deleting, 98
 - disabling, 99
 - editing, 99
 - enabling, 99
 - maintaining across reloads, 100
 - saving, 100
- Events menu, 98
- execution pointer, 31
 - in MP Prism, 59, 244
- expressions, writing in Prism, 37

F

- F1 key, 25, 195
- file command, 89
- File menu in visualizers
 - Diff and Diff With selections, 155
 - Save and Save as selections, 153
- File menu in visualizers, using, 140
- File selection, 88, 89, 105
- focus, 25
- fonts, changing the default, 213
- Fortran 90 generic procedures
 - changing the way Prism handles, 220
 - using, 42
- Fortran 90 support
 - allocatable arrays, 238
 - print command, 238
 - whatis command, 238
 - array sections, 238
 - array valued functions, 242
 - derived types, 235, 241
 - Fortran 77 intrinsics, 240

- generic functions, 235, 241
- internal procedures, 240
- masked array operations, 238
- pointer assignment, 236
 - allocatable arrays, 238
- pointer assignment error checking, 242
- pointers to arrays, 237
- slicing and striding arrays, 238
- user defined operators, 239
- variable attributes, 239
- `what is` command, 237

Fortran intrinsic functions, 39

`func` command, 89

Func selection, 29, 89, 105

function definition, displaying in the source window, 29

functions, choosing the correct, 37

G

- `-g` compiler option, 10
- Glossary selection, 196
- Graph visualizers, 144
 - minimum and maximum of, 149

H

- help system
 - overview of, 6
 - using, 196
- Histogram visualizers, 141
 - parameters for, 149
- history region, 32
 - changing the default length of, 34
 - using, 33

I

- I/O, 56
 - specifying the Xterm for, 207, 218
- `ILEN` intrinsic function, 40
- `IMAG` intrinsic function, 40
- infinities, detecting, 41
- initialization file, 23

- `interrupt` command, 57
- Interrupt selection, 33
 - ending a wait in MP Prism, 58
 - in MP Prism, 58
- interrupted pset, 58, 65
- inactive intrinsic, 68, 69

K

- keyboard accelerators, 27
- keyboard alternatives to the mouse, 25
- `kill` command, 45, 245

L

- languages supported in Prism, 10
- `layout` intrinsic, 188, 189
- layouts, visualizing, 188
- leaving Prism, 46
- line-number region, 5, 31
- `list` command, 229
- `load` command, 51
 - cannot be used in actions field, 97
- Load selection, 50
- loading a program, 49
- local variables, printing names and values of, 165
- location cursor, 25
- `log` command, 199
- logging commands and output, 36

M

- `make` command, 193
- `make` utility, 192, 207
- Man Pages selection, 197
- manual pages, viewing, 197
- Mark Stale Data, 207
- `MAXLOC` intrinsic function, 136
- `MAXVAL` intrinsic function, 40
- memory
 - examining the contents of, 125
- menu bar, 4
 - using, 27

- menu threshold
 - for Fortran 90 generic procedures, 207
- message queues, text display, 182
- message queues, visualizing, 171 to 182
 - communicator colors, 179
 - communicator data, 179
 - communicator dialog box, 180
 - Data Type dialog box, 181
 - label values, 176
 - Message dialog box, 177
 - nonblocking sends and receives, 171
 - sort criteria, 176
 - stopped ranks, 172
 - unexpected receives
 - correctness problems, 171
 - performance problems, 171
 - zoom levels, 172
- Meta key, 26
- MINVAL intrinsic function, 40
- Motif keyboard translations, changing, 217
- mouse
 - getting help on using, 196
 - using, 24
- MP Prism
 - attaching in, 53, 54, 244
 - commands-only version, 60
 - customizing, 200
 - executing a program in, 55
 - scope in, 79
 - shortening prompt, 77
 - visualizing data in, 168
- MPI queues. See message queues
- MPI SPMD style requirement, 87
- MPI_Comm_spawn, 18
- MPI_Comm_spawn_multiple, 18, 87
- mpimsgs command, 182

N

- names
 - resolving, 37
- NaNs, detecting, 41
- Netscape, 219
- next command
 - in MP Prism, 244

O

- online documentation, 197
 - obtaining in commands-only Prism, 229
- optimization, effects of, 121
- Options menu in visualizers, using, 140
- output
 - logging, 36
 - redirecting
 - in -CX version of Prism, 230
- Overview selection, 196

P

- parallel array, 183
- print command, 122, 134, 231, 232, 238
 - redirecting output to X window, 230
 - specifying the radix in, 137
 - with varfile intrinsic, 155
- Print dialog box, 132
- Print selection (Debug menu), 131
 - in MP Prism, 168
- Print selection (Events menu), 133
- printing
 - changing the default precision for, 149
 - difference from displaying, 129
 - from the command window, 134
 - from the event table, 134
 - from the source window, 29, 132
 - specifying the number of items to be printed on a line, 229
- Prism
 - commands-only, 228 to 230
 - entering, 11
 - initializing, 199
 - languages supported in, 10
 - leaving, 46
 - look and feel of, 3
 - overview of, 1
- prism command
 - C option, 228
 - CX option, 230
- Prism defaults
 - changing, 209
- Prism resources
 - table of, 209

- Prism*defaultVirtualBindings resource, 218
- Prism*fontList resource, 209
- Prism*XmText.fontList resource, 211
- Prism.comm1Color resource, 209
- Prism.comm2Color resource, 209
- Prism.comm3Color resource, 209
- Prism.commOtherColor resource, 209
- Prism.cppPath, 200
- Prism.cppPath resource, 209
- Prism.dialogColor resource, 209
- Prism.editGeometry resource, 209, 212
- Prism.editor resource, 209, 212
- Prism.errorBell resource, 209, 218
- Prism.errorwin resource, 209, 212
- Prism.graphBGColor resource, 210
- Prism.graphFillColor resource, 210
- Prism.helpBrowser resource, 210, 220
- Prism.helpUseExisting resource, 210
- Prism.mainColor resource, 210
- Prism.make resource, 210
- Prism.markStaleData resource, 210, 219
- Prism.procMenu resource, 210, 220
- Prism.procThresh resource, 210, 221
- Prism.spectralMapSize resource, 210
- Prism.textBgColor resource, 210
- Prism.textFont resource, 210, 213
- Prism.textManyFieldTranslations resource, 210, 216
- Prism.textMasterColor resource, 210
- Prism.textOneFieldTranslations resource, 210, 216
- Prism.useXterm resource, 210, 218
- Prism.vizColormap resource, 211
- prism_add_array function, 161
- prism_define_typename function, 161
- procedure menu
 - for Fortran 90 generic procedures, 207
- process command, 78, 79
 - cannot be used in event actions, 103
- process, attaching to running, 22
- process, loading a running, 22
- processes
 - interrupting, 57

- spawned, 18
- waiting for, 57
- PRODUCT intrinsic function, 40
- programs
 - loading into Prism, 49
 - reloading into Prism, 51
 - rerunning, 56
- pset command, 76, 77, 79
 - cannot be used in event actions, 103
 - hide option, 81
 - unhide option, 81
- pset keyword, 66
- pset qualifier, 83
 - cannot be used in event actions, 103
- psets
 - cycling through the members of, 80
 - defining, 66
 - deleting, 75
 - dynamic, 65
 - and events, 102
 - and the current pset, 76
 - naming, 68
 - predefined, 65
 - snapshots of unbounded, 85
 - syntax for defining, 66
 - threads, 61
 - unbounded, 83
 - using, 61
 - using in commands, 82
 - variable, 68, 83
 - and events, 102
 - and the current pset, 77
 - evaluating membership in, 70
 - viewing the contents of, 71
- Psets selection, 62
- Psets window, 71
 - changing the current pset via, 75
 - using, 62
 - zooming in, 73
- pstatus command, 59
- pushbutton command, 200, 202, 228

Q

- qualified names, 38
- quit command, 46, 230

Quit selection, 46

R

radix

- changing for a specific value, 165
- changing the default, 131
- changing via the Options menu, 152
- specifying in `print` or `display` command, 137

RANK intrinsic function, 40

REAL intrinsic function, 40

registers

- examining the contents of, 125, 135

reload command, 51

rerun command, 56

resolving names, 37

return command

- cannot be used in actions field, 97

Run (args) selection, 55

Run button, 55

run command, 55

- cannot be used in actions field, 97

Run selection, 55

running pset, 65

S

S3L parallel arrays, 183

array handle, 183

data types, 183

visualizing layouts of, 188

scope

in MP Prism, 79

scope pointer, 32

set command, 204

`$d_precision` and `$f_precision`
variables, 149

`$history` variable, 34

`$page_size` variable, 229

`$print_width` variable, 229

`$prompt_length` variable, 77

`$radix` variable, 113

sh command, 44

Shell selection, 44

show events command, 100, 104, 107, 109, 111,
229

show pset command, 73, 75, 76, 77

show psets command, 62, 69, 74
redirecting output to X window, 230

SIZE intrinsic function, 40

snapshot window, 35, 133

snapshots, unbounded psets, 85

source code

- editing, 191
- moving through, 28

source command, 101

source files

- creating a directory list for, 90

source window, 5

scrolling, 28

splitting, 29

using, 28

Spawned MPI processes

debugging, 18

selecting, 19

special function, `prism_define_typename`, 161

status messages, 56

status region, 5

step command

- cannot be used in actions field, 97
- in MP Prism, 57, 244

Stop (loc), 106

Stop (var), 107

stop command, 232

stopi command, 107, 109

stopped keyword, 96

stopped pset, 65

structures

- visualizing, 157
- augmenting data type information, 161
- in commands-only Prism, 230

SUM intrinsic function, 40

Sun MPI Client/Server programs

using MP Prism with, 87

Surface visualizers, 145

minimum and maximum of, 149

T

- Tab, 25
- task ID, 53
- tearoff command, 200, 202, 228
- Tear-off dialog box, 202
- tear-off region, 5, 201, 202
- Tear-off selection, 201
- text font, 207
- Text visualizers, 141
 - precision of, 148
- text visualizers, 6
- text widgets, changing keyboard translations
 - in, 216
- text, selecting in source window, 29
- this identifier, 232
- threads, 61
 - hidden, 81
 - libmpi_mt library, 13, 14
 - libthread library, 13, 14
 - unbounded psets, 83
- Threshold visualizers, 143
 - threshold of, 149
- threshold visualizers, 6
- Trace (cond), 110
- Trace (loc), 110
- Trace (var), 110
- trace command, 107, 111, 122
- Trace selection, 110
- tracei command, 107, 122
- traces
 - deleting, 111
 - in MP Prism
 - requirement that processes synchronize, 103
- tracing program execution, 109
- Tutorial selection, 196

U

- unalias command, 203
- UNIX commands, issuing, 44
- unset command, 204
- untearoff command, 202, 228
- Up selection, 113
- use command, 91

- Use selection, 18, 51, 90
- Using Help selection, 196

V

- variables
 - choosing the correct, 37
 - comparing values of, 155
 - printing the type of, 163
 - restoring the values of from a file, 154
 - saving the values of to a file, 153
 - setting up alternative names for, 203
- variables, accessibility of, 121
- varsave command, 122, 153
- Vector visualizers, 146
 - minimum and maximum of, 149
- visualization parameters, 148
- visualizer color file, 207
 - creating, 207
- visualizers, 6, 137
 - closing, 152
 - comparing values in, 155
 - displaying a ruler for, 150
 - displaying from the source window, 29
 - field width of, 148
 - in MP Prism, 168
 - saving, restoring, and comparing, 153
 - setting the context for, 151
 - statistics for, 150
 - structure, 157
 - treatment of stale data in, 207
 - types of, 141
 - updating, 152
 - working with, 137
- visualizing layouts, 188

W

- Wait Any selection, 58
- wait command, 57
 - any argument, 57
 - every argument, 57
- Wait Every selection, 58
- watchpoint, 94
- what is command, 164, 237, 238, 239

- Whatis selection, 163
- when command, 122
- where command, 113, 122
 - in MP Prism
 - redirecting output to X window, 230
- Where graph, 113, 245
 - moving through, 120
 - panning and zooming in, 115
 - shrinking portions of, 119
 - view information about threads, 119
 - visualizing in commands-only Prism, 230
- Where selection, 112
 - in MP Prism, 113, 245
- Where window, 112, 113
- whereis command, 39
- which command, 39
- windowing environments, supported
 - Common Desktop Environment (CDE), 10
 - OpenWindows, 10

X

- X resource database, adding Prism resources to, 211
- X toolkit command-line options, 22
- X Window System, 1
- xman, 197
- xrdb, 211
- Xterm
 - specifying for I/O, 218

