

TOTALVIEW USERS GUIDE



APRIL 2003

VERSION 6.2

Copyright © 1998–1999, 2003 by Etnus Inc. All rights reserved.

Copyright © 1999–2003 by Etnus LLC. All rights reserved.

Copyright © 1996–1998 by Dolphin Interconnect Solutions, Inc.

Copyright © 1993–1996 by BBN Systems and Technologies, a division of BBN Corporation.

No part of this publication may be reproduced, stored in a retrieval system, or transmitted, in any form or by any means, electronic, mechanical, photocopying, recording, or otherwise without the prior written permission of Etnus LLC (Etnus).

Use, duplication, or disclosure by the Government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.227-7013.

Etnus has prepared this manual for the exclusive use of its customers, personnel, and licensees. The information in this manual is subject to change without notice, and should not be construed as a commitment by Etnus. Etnus assumes no responsibility for any errors that appear in this document.

TotalView and Etnus are registered trademarks of Etnus Inc.

All other brand names are the trademarks of their respective holders.



Book Overview

part I – Introduction

- 1 Discovering TotalView 3
- 2 Understanding Threads, Processes, and Groups..... 17

part II – Setting Up

- 3 Setting Up a Debugging Session..... 39
- 4 Setting Up Remote Debugging Sessions..... 73
- 5 Setting Up Parallel Debugging Sessions..... 91

part III – Using the GUI

- 6 Using TotalView’s Windows 145
- 7 Visualizing Programs and Data 159

part IV – Using the CLI

- 8 Seeing the CLI at Work..... 185
- 9 Using the CLI..... 193

part V – Debugging

- 10 Debugging Programs..... 211
- 11 Using Groups, Processes, and Threads..... 239
- 12 Examining and Changing Data..... 277
- 13 Examining Arrays..... 319
- 14 Setting Action Points 337

- Glossary** 385



Contents

About This Book

How to Use This Book	xvii
Using the CLI	xviii
Audience	xix
Conventions	xx
Note	xx
TotalView Documentation	xxi
Contacting Us	xxii

part I - Introduction

1 Discovering TotalView

First Steps	3
Starting TotalView	4
What About Print Statements?	5
Examining Data	6
Debugging Multiprocess and Multithreaded Programs	11
Supporting Multiprocess and Multithreaded Programs	12
Using Groups and Barriers	13
Introducing the CLI	14
What's Next	15

2 Understanding Threads, Processes, and Groups

A Couple of Processes	17
Threads	20
Complicated Programming Models	21
Kinds of Threads	23
Organizing Chaos	25

Creating Groups	29
Simplifying What You're Debugging	34

part II - Setting Up

3 Setting Up a Debugging Session

Compiling Programs	40
File Extensions	41
Starting TotalView	41
Initializing TotalView	43
Exiting from TotalView	46
Loading Executables	46
Loading Remote Executables	48
Attaching to Processes	49
Attaching Using the Unattached Page	50
Attaching Using File > New Program and dattach	51
Detaching from Processes	52
Examining Core Files	53
Viewing Process and Thread State	54
Attached Process States	55
Unattached Process States	56
Handling Signals	56
Setting Search Paths	59
Setting Command Arguments	61
Setting Input and Output Files	62
Setting Preferences	64
Setting Preferences, Options, and X Resources	69
Setting Environment Variables	70
Monitoring TotalView Sessions	71

4 Setting Up Remote Debugging Sessions

Starting the TotalView Debugger Server	73
Setting Single-Process Server Launch Options	74
Setting Bulk Launch Window Options	76
Starting the Debugger Server Manually	79
Using the Single-Process Server Launch Command	80
Bulk Server Launch on an SGI MIPS Machine	81
Bulk Server Launch on an IBM RS/6000 AIX Machine	83
Bulk Server Launch on an HP Alpha Machine	83

Disabling Autolaunch	84
Changing the Remote Shell Command	84
Changing the Arguments	85
Autolaunch Sequence	86
Debugging Over a Serial Line	86
Starting the TotalView Debugger Server	87
Starting TotalView on a Serial Line	88
Using the New Program Window	88

5 Setting Up Parallel Debugging Sessions

Debugging MPICH Applications	92
Starting TotalView on an MPICH Job	92
Attaching to an MPICH Job	94
MPICH P4 procgroup Files	96
Debugging HP Alpha MPI Applications	96
Starting TotalView on a HP Alpha MPI Job	96
Attaching to a HP Alpha MPI Job	97
Debugging HP MPI Applications	97
Starting TotalView on an HP MPI Job	97
Attaching to an HP MPI Job	98
Debugging IBM MPI (PE) Applications	99
Preparing to Debug a PE Application	99
Using Switch-Based Communication	99
Performing Remote Logins	100
Setting Timeouts	100
Starting TotalView on a PE Job	100
Setting Breakpoints	101
Starting Parallel Tasks	101
Attaching to a PE Job	102
Attaching from a Node Running poe	102
Attaching from a Node Not Running poe	102
Debugging QSW RMS Applications	103
Starting TotalView on an RMS Job	103
Attaching to an RMS Job	104
Debugging SGI MPI Applications	104
Starting TotalView on a SGI MPI Job	104
Attaching to an SGI MPI Job	105
Debugging Sun MPI Applications	105
Attaching to a Sun MPI Job	106
Displaying the Message Queue Graph	107

Displaying the Message Queue	109
Message Queue Display Overview	109
Using Message Operations	110
Diving on MPI Processes	111
Diving on MPI Buffers	111
Pending Receive Operations	111
Unexpected Messages	112
Pending Send Operations	112
MPI Debugging Troubleshooting	113
Debugging OpenMP Applications	113
Debugging OpenMP Programs	114
TotalView OpenMP Features	115
OpenMP Platform Differences	115
OpenMP Private and Shared Variables	117
OpenMP THREADPRIVATE Common Blocks	118
OpenMP Stack Parent Token Line	120
Debugging PVM and DPVM Applications	121
Supporting Multiple Sessions	121
Setting Up ORNL PVM Debugging	122
Starting an ORNL PVM Session	122
Starting a DPVM Session	123
Automatically Acquiring PVM/DPVM Processes	124
Attaching to PVM/DPVM Tasks	126
Reserved Message Tags	127
Cleanup of Processes	127
Debugging Shared Memory (SHMEM) Code	128
Debugging UPC Programs	129
Invoking TotalView	130
Viewing Shared Objects	130
Pointer to Shared	132
Parallel Debugging Tips	134
Attaching to Processes	134
General Parallel Debugging Tips	137
MPICH Debugging Tips	139
IBM PE Debugging Tips	140

part III - Using the GUI

6 Using TotalView's Windows

Using the Mouse Buttons	145
Using the Root Window	146
Using the Process Window	150
Diving into Objects	152
Resizing and Positioning Windows and Dialog Boxes	155
Editing Text	156
Saving the Contents of Windows	157

7 Visualizing Programs and Data

Displaying Your Program's Call Tree	159
Displaying Memory Statistics	161
Using the Visualizer to Display Array Data	163
How the Visualizer Works.....	164
Configuring TotalView to Launch the Visualizer	165
Visualizer Launch Command	166
Data Types That TotalView Can Visualize	167
Viewing Data	167
Visualizing Data Manually	168
Visualizing Data Programmatically	169
Using the Visualizer	170
Directory Window	171
Data Windows	172
Using the Graph Window	173
Displaying Graphs	174
Manipulating Graphs	176
Using the Surface Window	177
Displaying Surface Data	177
Manipulating Surface Data	179
Launching the Visualizer from the Command Line	180

part IV - Using the CLI

8 Seeing the CLI at Work

Setting the EXECUTABLE_PATH State Variable	185
Initializing an Array Slice	187
Printing an Array Slice	187

Writing an Array Variable to a File	189
Automatically Setting Breakpoints	189
9 Using the CLI	
Tcl and the CLI	193
The CLI and TotalView	194
The CLI Interface	195
Starting the CLI	196
Startup Example	197
Starting Your Program	198
CLI Output	200
"more" Processing	201
Command Arguments	201
Using Namespaces	202
Command and Prompt Formats	203
Built-In Aliases and Group Aliases	203
Effects of Parallelism on TotalView and CLI Behavior	204
Kinds of IDs	205
Controlling Program Execution	206
Advancing Program Execution	206
Action Points	207

part V - Debugging

10 Debugging Programs

Searching and Looking Up Program Elements	211
Searching for Text	212
Looking for Functions and Variables	212
Finding the Source Code for Functions	213
Resolving Ambiguous Names	214
Finding the Source Code for Files	215
Resetting the Stack Frame	216
Viewing the Assembler Version of Your Code	216
Editing Source Text	218
Manipulating Processes and Threads	219
Using the Toolbar to Select a Target	219
Stopping Processes and Threads	220
Updating Process Information	221
Holding and Releasing Processes and Threads	221

Examining Groups	223
Displaying Groups	224
Placing Processes into Groups	225
Starting Processes and Threads	225
Creating a Process Without Starting It	226
Creating a Process by Single-Stepping	227
Stepping and Setting Breakpoints	227
Using Stepping Commands	229
Stepping into Function Calls	230
Stepping Over Function Calls	230
Executing to a Selected Line	231
Executing to the Completion of a Function	232
Displaying Thread and Process Locations	232
Continuing with a Specific Signal	233
Deleting Programs	234
Restarting Programs	235
Checkpointing Programs and Processes	235
Setting the Program Counter	236
Interpreting Status and Control Registers	237

11 Using Groups, Processes, and Threads

Defining the GOI, POI, and TOI	239
Setting a Breakpoint	241
Stepping (Part I)	241
Group Width	242
Process Width	243
Thread Width	243
Using "Run To" and duntil Commands	244
Using P/T Set Controls	245
Setting Process and Thread Focus	247
Process/Thread Sets	248
Arenas	249
Specifying Processes and Threads	250
The Thread of Interest (TOI)	250
Process and Thread Widths	251
Specifier Examples	253
Setting Group Focus	253
Specifying Groups in P/T Sets	255
Arena Specifier Combinations	257
'All' Does Not Always Mean <i>All</i>	259

Setting Groups	261
Using the 'g' Specifier: An Extended Example	263
Focus Merging	266
Incomplete Arena Specifiers	267
Lists with Inconsistent Widths	267
Stepping (Part II): Some Examples	268
Using P/T Set Operators	270
Using the P/T Set Browser	271
Using the Group Editor	275

12 Examining and Changing Data

Changing How Data Is Displayed	277
Displaying STL Variables	278
Changing Size and Precision	280
Displaying Variables	281
Displaying Program Variables	282
Displaying Variables in the Current Block	282
Browsing for Variables	283
Displaying Local Variables and Registers	284
Displaying Long Variable Names	286
Automatic Dereferencing	287
Displaying Areas of Memory	289
Displaying Machine Instructions	290
Closing Variable Windows	290
Diving in Variable Windows	291
Displaying Array of Structure Elements	293
Scoping and Symbol Names	294
Qualifying Symbol Names	295
Changing the Values of Variables	296
Changing the Data Type of Variables	297
Displaying C Data Types	298
Pointers to Arrays	299
Arrays	299
Typedefs	300
Structures	300
Unions	301
Built-In Types	302
Character Arrays (<string> Data Type)	304
Areas of Memory (<void> Data Type)	304
Instructions (<code> Data Type)	304

Type Casting Examples	304
Displaying the argv Array	305
Displaying Declared Arrays	306
Displaying Allocated Arrays	306
Working with Opaque Data	306
Changing the Address of Variables	306
Changing Types to Display Machine Instructions	307
Displaying C++ Types	307
Classes	307
Changing Class Types in C++	309
Displaying Fortran Types	310
Displaying Fortran Common Blocks	310
Displaying Fortran Module Data	310
Debugging Fortran 90 Modules	312
Fortran 90 User-Defined Types	314
Fortran 90 Deferred Shape Array Types	314
Fortran 90 Pointer Types	315
Displaying Fortran Parameters	315
Displaying Thread Objects	317

13 Examining Arrays

Examining and Analyzing Arrays	319
Displaying Array Slices	319
Using Slices and Strides	320
Using Slices in the Lookup Variable Command	322
Array Data Filtering	324
Filtering Array Data	324
Filtering by Comparison	325
Filtering for IEEE Values	326
Filtering By a Range of Values	327
Creating Array Filter Expressions	329
Using Filter Comparisons	329
Sorting Array Data	330
Obtaining Array Statistics	331
Displaying a Variable in All Processes or Threads	333
Diving in a Laminated Pane	334
Editing a Laminated Variable	335
Visualizing Array Data	335
Visualizing a Laminated Variable Window	336

14 Setting Action Points

Action Points Overview	337
Setting Breakpoints and Barriers	339
Setting Source-Level Breakpoints	340
Choosing Source Lines	340
Setting and Deleting Breakpoints at Locations	340
Displaying and Controlling Action Points	343
Disabling	343
Deleting	344
Enabling	344
Suppressing	344
Setting Machine-Level Breakpoints	345
Setting Breakpoints for Multiple Processes	346
Setting Breakpoints When Using fork()/execve()	347
Processes That Call fork()	348
Processes That Call execve()	348
Example: Multiprocess Breakpoint	349
Barrier Points	350
Barrier Breakpoint States	350
Setting a Barrier Breakpoint	351
Creating a Satisfaction Set	353
Hitting a Barrier Point	353
Releasing Processes from Barrier Points	353
Deleting a Barrier Point	353
Changes When Setting and Disabling a Barrier Point	354
Defining Evaluation Points and Conditional Breakpoints	354
Setting Evaluation Points	356
Creating Conditional Breakpoint Examples	356
Patching Programs	357
Conditionally Patching Out Code	357
Patching in a Function Call	358
Correcting Code	358
Interpreted vs. Compiled Expressions	358
Interpreted Expressions	359
Compiled Expressions	359
Allocating Patch Space for Compiled Expressions	360
Dynamic Patch Space Allocation	360
Static Patch Space Allocation	361
Using Watchpoints	363
Architectures	363

Creating Watchpoints	365
Displaying Watchpoints	366
Watching Memory	366
Triggering Watchpoints	367
Using Multiple Watchpoints	367
Data Copies	368
Using Conditional Watchpoints	368
Saving Action Points to a File	370
Evaluating Expressions	371
Writing Code Fragments	373
TotalView Variables	374
Built-In Statements	375
C Constructs Supported	377
Data Types and Declarations	377
Statements	378
Fortran Constructs Supported	378
Data Types and Declarations	379
Statements	379
Writing Assembler Code	380
Glossary	385
Index	401



About This Book

This book describes how to use TotalView®, a source- and machine-level debugger for multiprocess, multithreaded programs. The guide assumes that you are familiar with programming languages, the UNIX operating systems, the X Window System, and the processor architecture of the platform on which you are running TotalView.

You will be reading a user guide that combines information for two TotalView debuggers. One uses Motif to present windows and dialog boxes. The other runs in an xterm-like window and requires that you type commands. This book emphasizes the Motif interface, as it is easier to use. In addition, once you “see” what you can do, you will know what can be done using the command interface.

This book covers using TotalView on any platform.

How to Use This Book



The information in this book is presented in 5 parts.

■ I: Introduction

Here you’ll find an overview of some of TotalView’s features and an introduction to TotalView’s process/thread model. Please read this information. It’s easy reading and you’ll get a feel for what TotalView can do.

■ II: Setting Up

Most people don’t spend a lot of time in this section. Chapter 3 tells you what you need to know about configuring TotalView. Chapters 4 and 5 tell you how to get your programs running under TotalView’s control. Look at Chapter 4 if you’re having problems getting the TotalView Debugger Server (**tvdsvr**) running and if you’re reconfiguring how **tvdsvr** gets launched.

You will never need to read all of Chapter 5. Instead, go to the table of contents and find the section that has the information you need.

■ III: Using the GUI

The chapters in this section look at some of TotalView's windows and how you use them. You are also shown tools such as the **Visualizer** and the **Call Tree** that help you analyze what your program is doing.

■ IV: Using the CLI

The chapters in this section explain the conventions of using a command-line debugger and how to create Tcl macros.

■ V: Debugging

In many ways, most of what has preceded this part of the book is "introductory" material. Here is where you'll find out how to examine your program and its data. In this part, you'll find information on setting the action points that allow you to stop and monitor your program's execution.

Equally important, Chapter 11 is a detailed examination of TotalView's group, process, and thread model. The more you understand this model, the easier time you'll have debugging multiprocess and multithreaded programs.

Using the CLI

To use the CLI (Command Line Interface), you need to be familiar with and have experience debugging programs with the TotalView GUI. As CLI commands are embedded within a Tcl interpreter, you will get better results if you are familiar with Tcl. However, if you don't know Tcl, you will still be able to use the CLI, but you will lose the programmability features that Tcl gives. For example, CLI commands operate upon a set of processes and threads. You can save this set and apply it to commands based upon what you have saved.

You can obtain information on using Tcl at many bookstores, and you can also order these books from online bookstores. Two excellent books are

- Ousterhout, John K. *Tcl and the Tk Toolkit*. Reading, Mass.: Addison Wesley, 1997.
- Welch, Brent B. *Practical Programming in Tcl & Tk*. Upper Saddle River, N.J.: Prentice Hall PTR, 1997.

There is also a rich supply of resources available on the Web. The best starting point is www.tcltk.com.

The fastest way to gain an appreciation of the actions performed by CLI commands is to review Chapter 1 of the TOTALVIEW REFERENCE GUIDE, which contains an overview of CLI commands.

Audience

Many of you are very sophisticated programmers, having a tremendous knowledge of programming and its methodologies and almost all of you have used other debuggers and have developed your own techniques for debugging the programs that you write.

We know you are expert in your areas, whether it be threading, high-performance computing, client/server interactions, and the like. So, this book won't try to tell you about what you're doing. Instead, it tells you about TotalView.

As you will see, TotalView is a rather easy-to-use product. We can't tell you how to use TotalView to solve your problems because your programs are unique and complex, and we can't anticipate what you want to do. We also know you don't want to spend a lot of time reading about using TotalView. Consequently, you're not going to see a lot of quasi-procedural discussions that tell you what to put in dialog boxes. You already know what to do.

This book also doesn't spend a lot of time explaining what you do with a dialog box or the kinds of data you can type. If you want that information, you'll find it in the online Help. If you prefer, an HTML version of this information is available on our Web site. If you have purchased TotalView, you can also post this HTML documentation on your intranet.

Conventions

The following table describes the conventions used in this book:

TABLE I: Book Conventions

Convention	Meaning
[]	Brackets are used when describing parts of a command that are optional.
<i>arguments</i>	In a command description, text in italic represent information you type. Elsewhere, italic is used for emphasis. You won't have any problems distinguishing between the uses.
Dark text	In a command description, dark text represent keywords or options that you must type exactly as displayed. Elsewhere, it represents words that are used in a programmatic way rather than their normal way.
Example text	In program listings, this indicates that you are seeing a program or something you'd type in response to a shell or CLI prompt. If this text is in bold, it's indicating that what you're seeing is what you'll be typing. Bolding this kind of text is done only when it's important. You'll usually be able to differentiate what you type from what the system prints.
	This graphic symbol indicates that a feature is only available in the GUI. If you see it on the first line of a section, all the information in the section is just for GUI users. When it is next to a paragraph, it tells you that just the sentence or two being discussed applies to the GUI.
CLI EQUIVALENT:	The primary emphasis of this book is on the GUI. It shows the windows and dialog boxes that you use. This symbol tells you how to do the same thing using the CLI.

Note

This book discusses the TotalView GUI and the CLI simultaneously. You will see where something is done in the GUI and what the CLI equivalent is. But in many cases, we assume that you know which to use, and can fill in the details. For example, when a dialog box is discussed, this book does not mention that using the dia-

log box is something you do when using the GUI. In most cases, reading what the GUI does tells you what you need to know when using the CLI.

This book minimizes the amount of stuff you have to read since we assume you just want to use TotalView so that you can get the bugs out of your program.

TotalView Documentation

The following table describes other TotalView documentation:

TABLE II: TotalView Documentation

Title	Contents	Online			
		Help	HTML	PDF	Print
TotalView Reference Guide	Contains descriptions of CLI commands, how you run TotalView, and platform-specific information	✓	✓	✓	
TotalView QuickView	Presents what you need to know to get started using TotalView				✓
TotalView Commands	Defines all TotalView GUI commands	✓	✓	✓	
Creating Type Transformations	Tells how to create Tcl CLI macros that change the way structures and STL containers appear		✓	✓	
TotalView Installation Guide	Contains the procedures to install TotalView and the FLEX/mlicense manager	✓	✓	✓	
TotalView New Features	Tells you about new features added to TotalView	✓	✓	✓	
TotalView Release Notes	Lists known bugs and other information related to the current release	✓	✓	✓	
IBM Considerations	Briefly describes things you should know when run on IBM RS6000 machines	✓	✓	✓	
Linux Considerations	Briefly describes things you should know when running on Linux platforms	✓	✓	✓	

TABLE II: TotalView Documentation (cont.)

Title	Contents	Online Help	HTML	PDF	Print
Patching Platforms	Describes how to apply vendor supplied patches to operating systems and compilers	✓	✓	✓	
Platforms and System Requirements	Lists the platforms upon which TotalView runs and the compilers it supports	✓	✓	✓	

Contacting Us

Please contact us if you have problems installing TotalView, questions that are not answered in the product documentation or on our Web site, or suggestions for new features or improvements.

Our Internet E-Mail address for support issues is support@etnus.com

For documentation issues, the address is: documentation@etnus.com

Call: 1-800-856-3766 in the United States

(+1) 508-652-7700 worldwide

If you are reporting a problem, please include the following information:

- The *version* of TotalView and the *platform* on which you are running TotalView
- An *example* that illustrates the problem
- A *record* of the sequence of events that led to the problem

The TOTALVIEW RELEASE NOTES contains complete instructions on how to report problems.



Part I: Introduction

This part of the TOTALVIEW USERS GUIDE contains two chapters.

Chapter 1: Discovering TotalView

Presents an overview of what TotalView is and the ways in which it can help you debug programs. If you haven't used TotalView before, reading this chapter gives you a high-level understanding of what TotalView can do for you.

Chapter 2: Understanding Threads, Processes, and Groups

Defines TotalView's model for organizing processes and threads. While most programmers have an intuitive understanding of what their programs are doing, debugging multiprocess and multithreaded programs requires an exact knowledge of what's being done. This chapter begins a two-part look at TotalView's process/thread model. This chapter contains introductory information. Chapter 11: "Using Groups, Processes, and Threads" on page 239 contains information on using this model with TotalView commands.



Chapter 1

Discovering TotalView

The Etnus TotalView® debugger is a powerful, sophisticated, and programmable tool that allows you to debug, analyze, and tune the performance of complex serial, multi-processor, and multithreaded programs.

If you want to jump in and get started quickly, you should go to our Website at www.etnus.com and go to TotalView's "Getting Started" area.

Topics in this chapter are:

- *"First Steps"* on page 3
- *"Debugging Multiprocess and Multithreaded Programs"* on page 11
- *"Using Groups and Barriers"* on page 13
- *"Introducing the CLI"* on page 14
- *"What's Next"* on page 15

First Steps

Getting started with TotalView is similar to getting started using other debuggers:

- You use the `-g` option when compiling your program.
- You start your program under the debugger's control.
- You set breakpoints.
- You examine data.

And, the way you go about doing these things is just about the same. Where TotalView differs from what you're used to using is in its raw power, the breadth of commands available, and its native ability to handle multiprocess, multithreaded programs.

Starting TotalView

After execution begins—you'll probably have typed something like `totalview programname`—you'll see a five-pane window. (See Figure 1.)

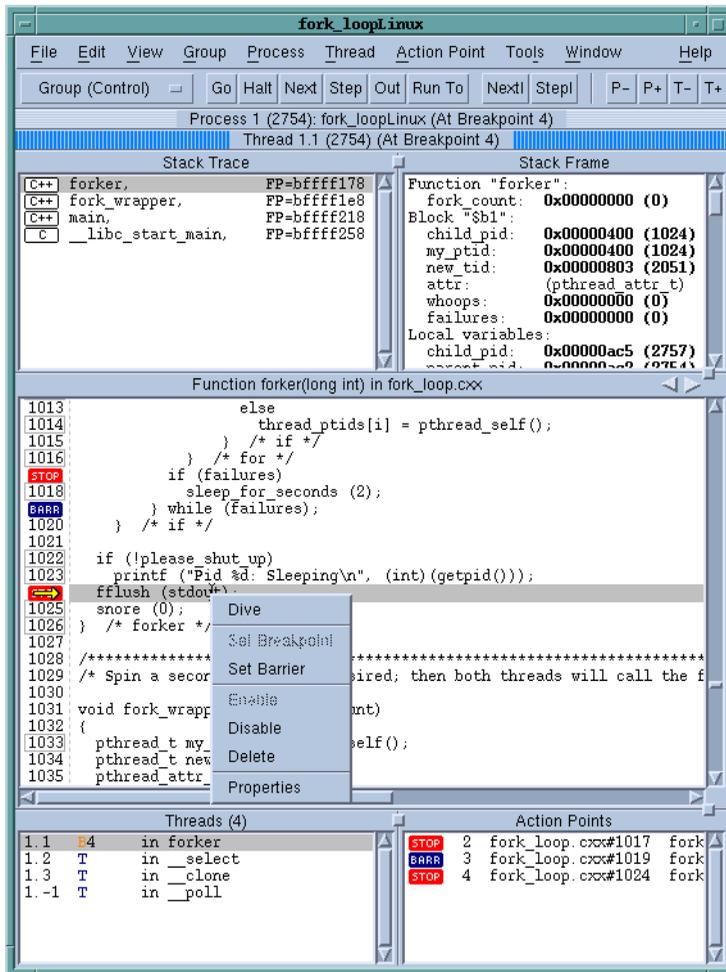


FIGURE 1: The Process Window

You can get going in several ways. Perhaps the easiest is to click on **Step** in the toolbar. This gets your program started, which means all the initialization stuff performed by the program gets done but no statements are executed. Alternatively, you could scroll your program to find where you want it to run to, select the line,

then click on **Run To** in the toolbar. Or you could click on the line number, which tells TotalView to create a breakpoint on that line, and then click **Go** in the toolbar.

If your program is large, and usually it will be, you can use the **Edit > Find** command to locate the line for you. Or, if you want to stop execution when your program reaches a subroutine, use the **Action Point > At Location** command to set a breakpoint before clicking on **Go**.

As you can see, you've got lots of choices. Unlike other debuggers, TotalView gives you choices that allow you to debug your program however you want to debug it.

What About Print Statements?

Most programmers learned to debug by using print statements. That is, you load your program with **printf()** or **PRINT** statements and then inspect what gets written. There's a problem with this. Every time you want to add a new statement, you've got to recompile your program. What's worse is that in a multiprocess, multi-threaded program, what you want printed may not print in the right order. While TotalView is much more sophisticated than this about showing your data (as you'll soon see), you can even use print statements simply.

In TotalView, breakpoints are called "action points". This is because they can be much more powerful than what you're used to. So, if you don't want to change the way you've been debugging, you can add your breakpoints easily. Figure 2 on page 6 shows the **Action Point Properties** Dialog Box. The easiest way to display this dialog box is to right-click on a line and then select **Properties** in the context menu. This menu is shown in Figure 1.

After clicking on the **Evaluate** radio button, you can add any code you want to a breakpoint. Because there's code associated with this breakpoint, it is now called an "eval point." Here's where TotalView does things a little differently. When the eval point is reached, TotalView executes the code, which prints the value of *i*. Eval points do exactly what you tell them to do. In this case, because you didn't tell TotalView to stop executing, it keeps on going. In other words, you don't have to stop program execution just to see data. You can, of course, tell TotalView to stop. Figure 3 on page 7 shows two evaluation points that stop execution. (One of them does something else as well.)

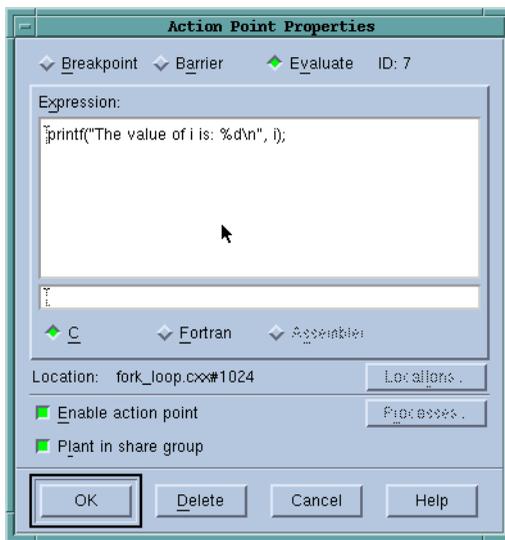


FIGURE 2: **Action Point Properties Dialog Box**

The one in the foreground uses a programming language statements and a built-in TotalView function to stop a loop every 100 iterations. It also prints what the value of *i* is. The one in the background just stops the program every 100 times a statement gets executed.

Evaluation points even allow you to patch your programs and route around code that you want replaced. The evaluation point shown in Figure 4 on page 8 tells TotalView to execute three statements and then skip to line 658.

Examining Data

Programmers use print statements as an easy way to examine data. They usually do this because the debugger doesn't have sophisticated ways of showing your data. In contrast, Chapter 12, "Examining and Changing Data," on page 277 and Chapter 13, "Examining Arrays," on page 319 explain how you can display data values with TotalView. In addition, Chapter 7, "Visualizing Programs and Data," on page 159 describes how to visualize your data in a graphical way.

Because data is difficult to see, the Stack Frame Pane (the pane in the upper right corner of the Process Window, which is show in Figure 1 on page 4) has a list of all

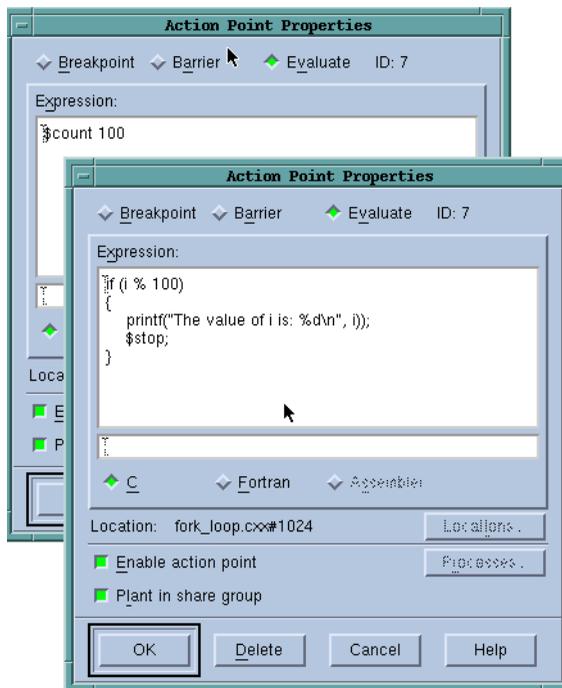


FIGURE 3: **More Conditions**

variables that exist in your current routine. If the value is simple, you can see it in this pane.

If it isn't, just dive on the variable to get more information.

NOTE “Diving” is something you can do almost everywhere in TotalView. What happens depends on where you are. To dive on something, position the cursor on the item and click your middle mouse button. If you have a two-button mouse, you can double-click your left mouse button.

Diving on a variable tells TotalView to display a window containing information about the variable. (As you read this manual, you'll come across many other kinds of diving.)

Notice that some of the values in the Stack Frame Pane are in bold type. This lets you know that you can click on the value and then edit it.

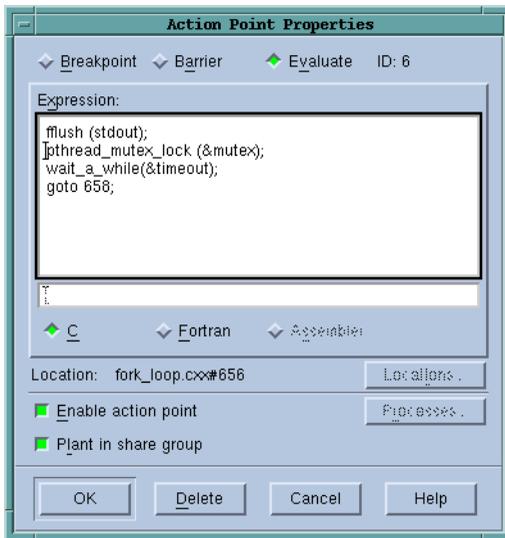


FIGURE 4: Patching Using an Evaluation Point

Figure 5 on page 9 shows two Variable Windows. One was created by diving on a structure and the other by diving on an array. Because the data displayed in a Variable Window may not be simple, you can also dive on data in the Variable Window. Normally, this information is displayed in the same window. However, you can use the **View > Dive Anew** command to display this information in a separate window.

If the data being displayed is a pointer, diving on the variable chases the pointer and then displays the data that is being pointed to. In this way, you can follow linked lists. Notice the forward- and backward-facing arrows in the upper right corner of the Variable Windows. Selecting them lets you “undive” and “redive.” For example, if you’re following a pointer chain, clicking the left-pointing arrow takes you back to where you just were. Clicking the right-pointing arrow takes you “forward” to the place you previously dove on.

Many arrays have copious amounts of data. Consequently, TotalView has a variety of ways to simplify how it should display this data.

The window in the background of Figure 6 on page 9 shows a basic *slice* operation. This operation tells TotalView that it should only display array elements whose positions are named by the slice. The foreground window combines a *filter* with a slice.

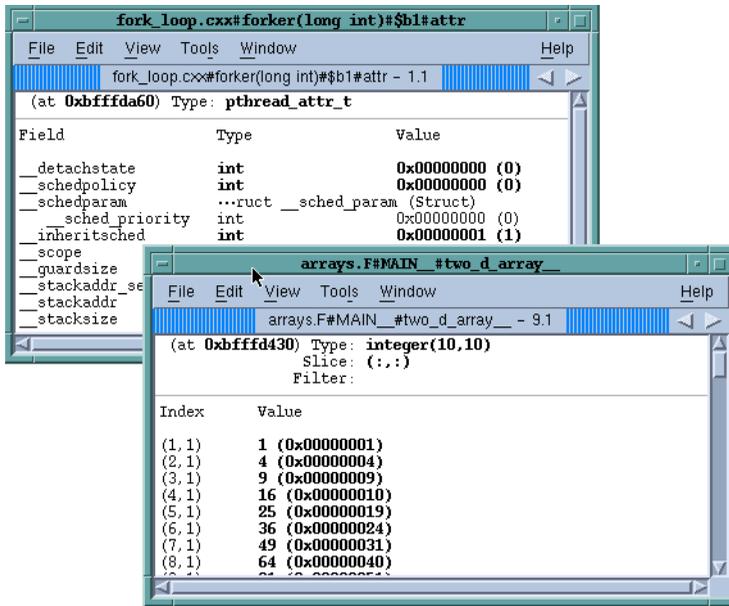


FIGURE 5: Two Variable Windows

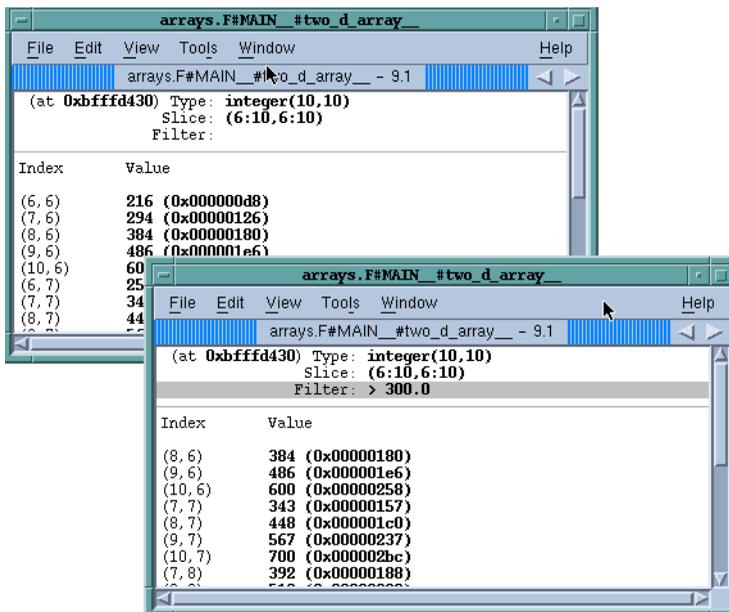


FIGURE 6: Two More Variable Windows

In this case, the filter says “of the array elements that could be displayed, only display the elements whose value is greater than 300.”

While slicing and filtering let you reduce the amount of data that TotalView will display, there are many times when you want to see the shape of the data. If you select the **Tools > Visualize** command, TotalView shows a graphic representation of the information in the Variable Window. Here’s an example:

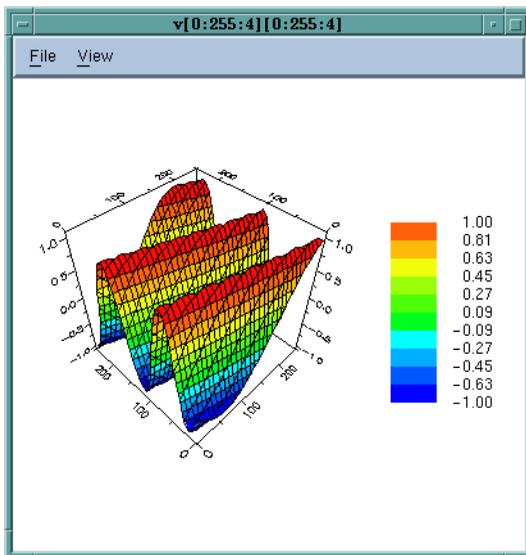


FIGURE 7: **Array Visualization**

There’s yet another way to look at data. TotalView’s watchpoints let you see when a variable’s data changes. This works in a different way than other action points. A watchpoint stops execution whenever a data value changes no matter what instruction changed the data. That is, if you change data from 30 different statements, the watchpoint stops execution whenever any of the 30 make a change. In contrast, other action points do something right before a line in your source executes. To create a watchpoint, select the **Tools > Watchpoint** command from any Variable Window.

Debugging Multiprocess and Multithreaded Programs

When your program creates processes and threads, TotalView can automatically bring them under its control. If the processes are already running, they too can be acquired. You don't need to have multiple debuggers running. TotalView is enough.

The processes that your program creates can be local or remote. Both are presented to you in the same way. The only difference between debugging a single-process program and a multiprocess, multithreaded program is that you gain the ability to display these additional threads and processes in Process Windows. You can display them in the current Process Window or display them in another window. As always, there are several ways to do it.

TotalView's Root Window (see Figure 8), which is automatically displayed after you start TotalView, contains an overview of everything being debugged, so diving on a process or a thread listed in the Root Window takes you quickly to the information you want to see. If you need to debug processes that are already running, the **Unattached** Page lets you dive on other processes you own. After diving on them, they can also be debugged.

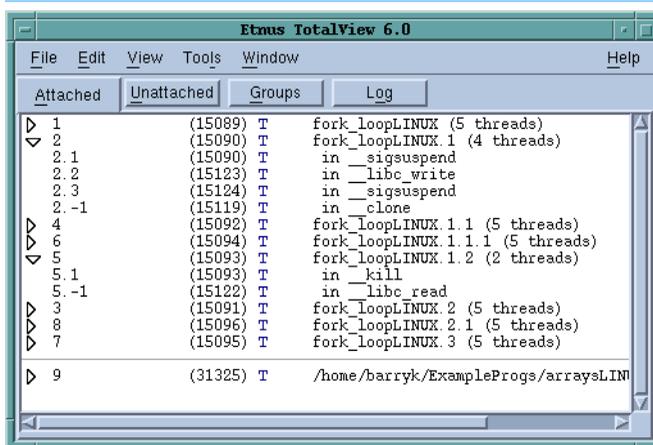


FIGURE 8: The Root Window

In the Process Window, you can switch between processes and threads by clicking the process and thread switching buttons in the toolbar. These are the buttons on the right side of the toolbar in Figure 9.

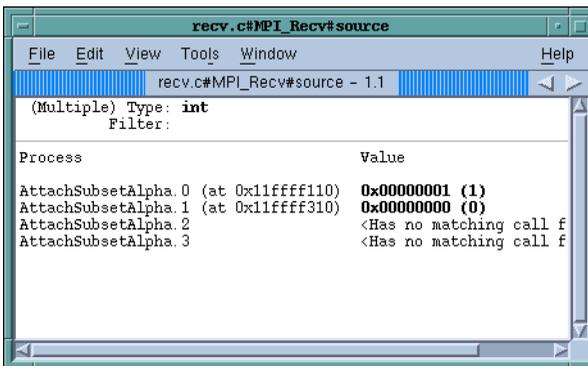
FIGURE 9: **Process and Thread Switching Icons**

Every time you click on one of these buttons, TotalView switches. The switching order is the order in which you see things in the Root Window.

In many cases, you'll be using one of the popular parallel execution models. TotalView supports MPI and MPICH, OpenMP, ORNL PVM (and HP Alpha DPVM), SGI shared memory (shmem) and UPC. You could be using threading in your programs. Or your programs can be compiled using products provided by your hardware vendor or third-party programs such as those from KAI and the Free Software Foundation (the GNU compilers).

Supporting Multiprocess and Multithreaded Programs

TotalView's laminated data view lets you see data values across threads or across processes. Here's an example of what you'll see after you use the Variable Window's **View > Laminate** command:

FIGURE 10: **A Laminated Variable Window**

When debugging MPI programs, using the **Tools > Message Queue Graph** to display the message queues graphically is the quickest way to see what is going on. (See Figure 11 on page 13.)

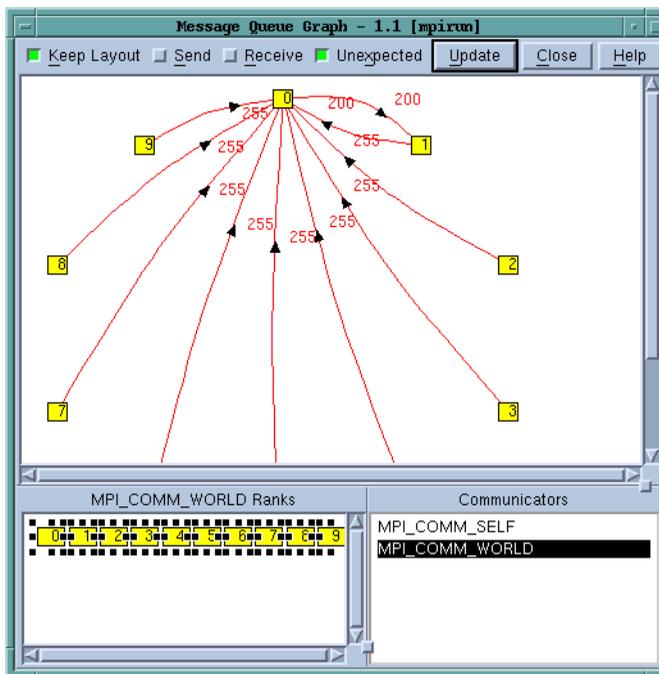


FIGURE 11: **A Message Queue Graph**

Clicking on the boxed numbers tells TotalView to place the associated process into a Process Window. Clicking on a number next to the arrow tells TotalView to display more information about that message queue.

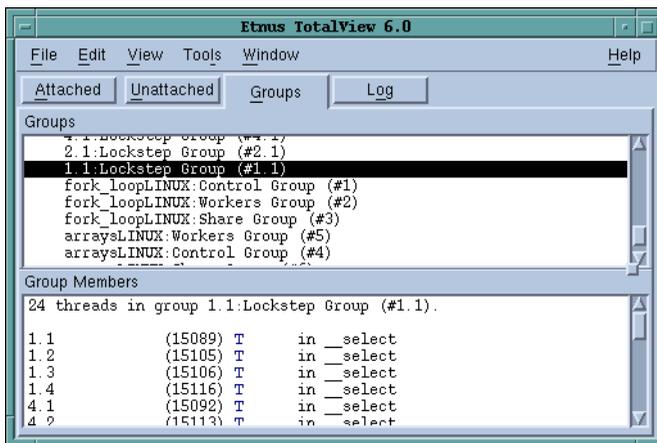
As you go through this book, you'll find many more examples.

Using Groups and Barriers

When running a multiprocess and multithreaded program, TotalView tries to automatically place your executing processes into different groups. While you can always individually stop, start, step, and examine any thread or process, TotalView lets you perform these actions on groups of threads and processes. In most cases, you'll be doing the same kinds of operations on the same kinds of things. The two pulldown menus on the toolbar let you indicate what the target of your action will be. See Figure 12 on page 14.

FIGURE 12: **Toolbar with Pulldown**

For example, if you are debugging an MPI program, you'd probably set the pull-down to **Process Workers**. The reasons for setting them like this are in Chapter 11. The definition of what these groups actually are is in Chapter 2. The **Groups** Page of the Root Window (Figure 13) shows you the processes and threads that are in a group.

FIGURE 13: **The Root Window's Group Page**

Introducing the CLI

TotalView is programmable. The CLI, the TotalView Command Line Interface, contains an extensive set of commands that you can type into a command window. These commands are embedded in a version of the Tcl command interpreter. When you open a CLI window, you can enter any Tcl statements that you could enter in any version of Tcl. You can also enter commands that have been added that allow you to debug your program. Because these debugging commands are native to

TotalView's Tcl, you can also use Tcl to manipulate the program being debugged. This means that you can use the CLI to create your own commands or perform any kind of repetitive operation. For example, the **dbreak 1038** command sets a breakpoint at line 1038. Here's an example that uses a debugging command and a Tcl command to set breakpoints at three lines:

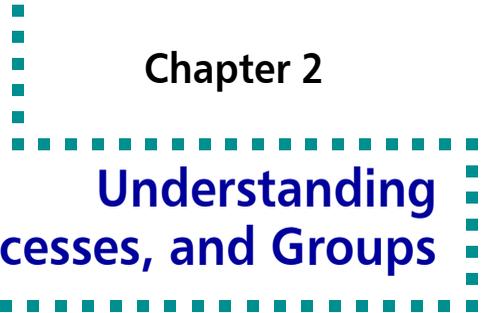
```
foreach i {1038 1043 1045} {  
    dbreak $i  
}
```

You'll find information about the CLI scattered throughout this book. CLI Commands are described in Chapter 2 of the TOTALVIEW REFERENCE GUIDE.

What's Next

This chapter has presented just a few of TotalView's highlights. The rest of this book tells you more about all of TotalView features, both the ones mentioned here and those not yet discussed.

All TotalView documentation is available on our Web site at <http://www.etnus.com/Support/docs> in PDF and HTML formats. In addition, this information is also contained within TotalView's online Help.



Chapter 2

Understanding Threads, Processes, and Groups

While the specifics of how multiprocess, multithreaded programs execute differ greatly from one hardware platform to another, from one operating system to another, and from one compiler to another, all share some general characteristics. This chapter defines how TotalView looks at processes and threads.

This chapter presents the concepts of thread, process, and group. Chapter 11, *"Using Groups, Processes, and Threads"* on page 239 is a more exacting and comprehensive look at these topics.

Topics in this chapter are:

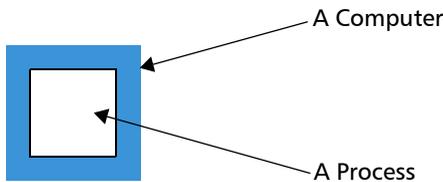
- *"A Couple of Processes"* on page 17
- *"Threads"* on page 20
- *"Complicated Programming Models"* on page 21
- *"Kinds of Threads"* on page 23
- *"Organizing Chaos"* on page 25
- *"Creating Groups"* on page 29
- *"Simplifying What You're Debugging"* on page 34

A Couple of Processes

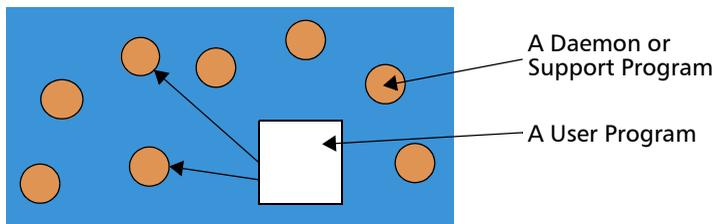


When programmers write single-threaded, single-process programs, they can almost always answer the question "Do you know where your program is?" These kind of programs are rather simple, looking something like what's shown in Figure 14 on page 18.

If you use any debugger on these kinds of programs, you can almost always figure out what's going on. Before the program begins executing, you set a breakpoint, let the program run until it hits the breakpoint, and then inspect variables to see their values. If you suspect there's a logic problem, you can step the program through its statements, seeing what happens and where things are going wrong.

FIGURE 14: **A Uniprocessor**

What is actually occurring, however, is a lot more complicated since a number of programs are always executing on your computer. For example, your computing environment could have daemons and other support programs executing, and your program can interact with them. (See Figure 15.)

FIGURE 15: **A Program and Daemons**

These additional processes can simplify your life because your program no longer has to do everything itself. It can hand off some tasks and not have to focus on how the work will get done.

Figure 15 assumes that the application program just sends requests to a daemon. This architecture is very simple. More typical is the kind of architecture shown in Figure 16 on page 19. Here, an E-mail program is communicating with a daemon on one computer. After receiving a request, this daemon sends data to an E-mail daemon on another computer, which then delivers the data to another mail program.

This architecture assumes that the jobs are disconnected and that they do not need to cooperate. This model has one program handing off work to another. After the handoff, the programs do not interact. While this is a useful model for many kinds of computation, a more general model allows a program to divide its work into smaller jobs, and parcel them out to other computers. This model relies on programs on other machines to do some of the first program's work. To gain any advantage, however, the work a program parcels out

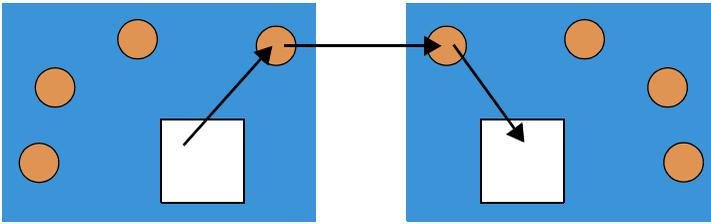


FIGURE 16: Mail Using Daemons

must be work that it doesn't need right away. In this model, the two computers act more or less independently. And, because the first computer doesn't have to do all the work, the program can complete its work faster. (See Figure 17.)

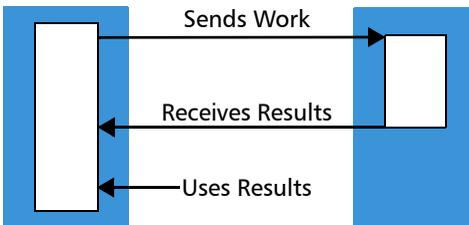


FIGURE 17: Two Computers Working on One Problem

Using more than one computer doesn't mean that less computer time is being used. Overhead due to sending data across the network and overhead for coordinating multiprocessing always means more work is being done. It does mean, however, that your program finishes sooner than if only one computer were working on the problem.

Here is one of the problems with this model: how does a programmer debug what's happening on the second computer? One solution is to have a debugger running on each computer. The TotalView solution to this debugging problem is better. It places a server on all remote processor as they are launched. These servers then communicate with the "main" TotalView. This debugging architecture gives you one central location from which you can manage and examine all aspects of your program.

NOTE You can also have TotalView attach to programs that are already running on other computers.

In all cases, it is far easier to write your program so that it only uses one computer at first. After you've got it working, you can split it up so it uses other computers. It is likely that any

problems you find will occur in the code that splits up the program or in the way the programs manipulate shared data, or in some other area related to the use of more than one thread or process. This assumes, of course, that it is practical to write your program as a single-process program. For some algorithms, executing a program on one machine means that it will take weeks to execute.

Threads

The daemon programs discussed in the previous section are owned by the operating system. They perform a variety of activities from managing computer resources to providing standard services such as printing.

If operating systems can have many independently executing components, why can't a program? Obviously, it can and there are various ways to do this. One programming model splits the work off into somewhat independent tasks within the same process. This is the *threads* model. (See Figure 18.) This figure also shows, for the last time, the daemon processes that are executing. From now on, just assume that they are there.

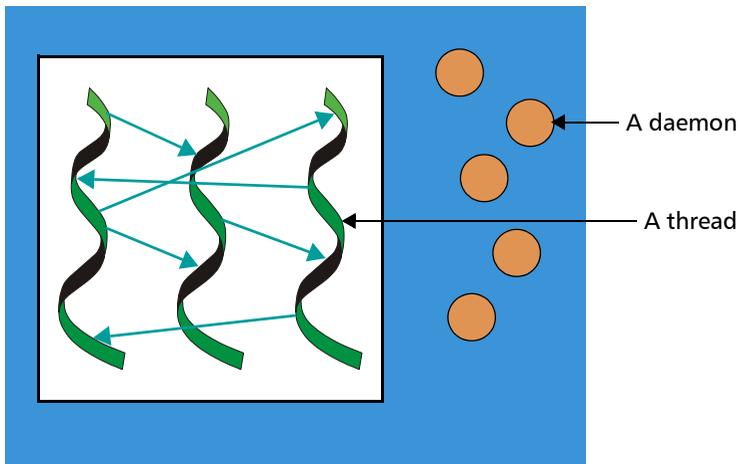


FIGURE 18: **Threads**

In this computing model, a program (the main thread) creates threads. If they need to, these newly created threads can also create threads. Each thread executes relatively independently from other threads. You can, of course, program them to share data and to synchronize how they execute.

The debugging problem here is similar to the problem of processes running on different machines. In both, a debugger must intervene with more than one executing entity.

NOTE There's not a lot of difference between a multithreaded or a multiprocess programs when you are using TotalView. Except for operating system support, the way in which TotalView displays process information is very similar to how it displays thread information.

Complicated Programming Models

While most computers being sold today have one processor, high-performance computing uses computers that have more than one processor. And as hardware prices decrease, this model is starting to become more widespread. Having more than one processor means that the threads model shown in Figure 18 changes to look something like what's shown in Figure 19.

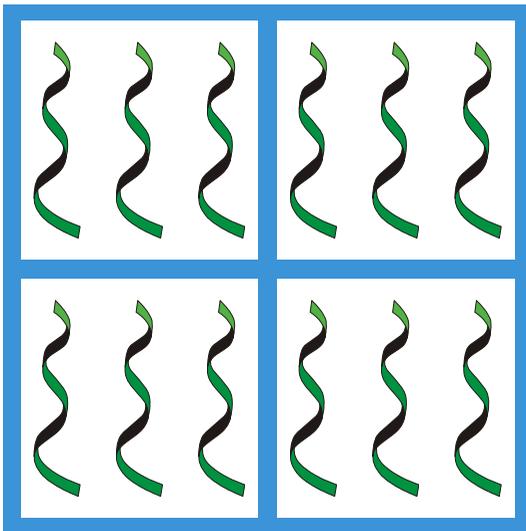


FIGURE 19: **Four Processor Computer**

This figure shows four linked processors in one computer, each of which has three threads. This architecture is an extension to the model that links more than one computer together. Its advantage is that the processor doesn't need to communicate with other processors over a network as it is completely self-contained.

The next step, of course, is to join many multiprocessor computers together. Figure 20 shows five computers, each having four processors with each processor running three threads. If this figure is showing the execution of one program, then the program is using 60 threads.

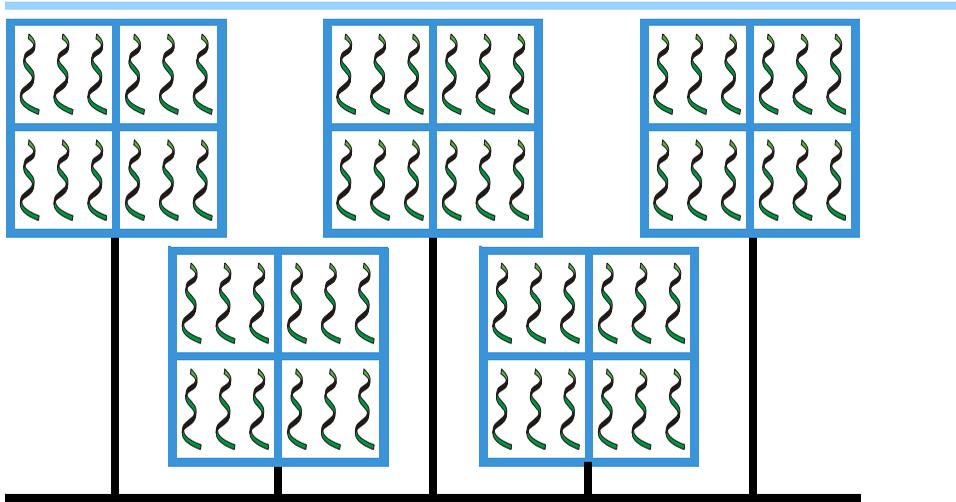


FIGURE 20: **Four-Processor Computer Networks**

This figure depicts only processors and threads. It doesn't have any information about the nature of the programs and threads or even if the programs are copies of one another or represent different executables.

At any time, it is next to impossible to guess which threads are executing and what a thread is actually doing. To make matters worse, many multiprocessor programs begin by invoking a process such as **mpirun** or IBM's **poe** whose function is to distribute and control the work being performed. In this kind of environment, a program (or the program in a library) is using another program to control the workflow across processors.

When there are problems in this scenario—and there are always problems—traditional debuggers and solutions are helpless. As you will see, TotalView, on the other hand, organizes this mass of executing procedures for you and lets you distinguish between threads and processes that the operating system uses from those that your program uses.

Kinds of Threads

All threads aren't the same. Figure 21 shows a program with three threads.

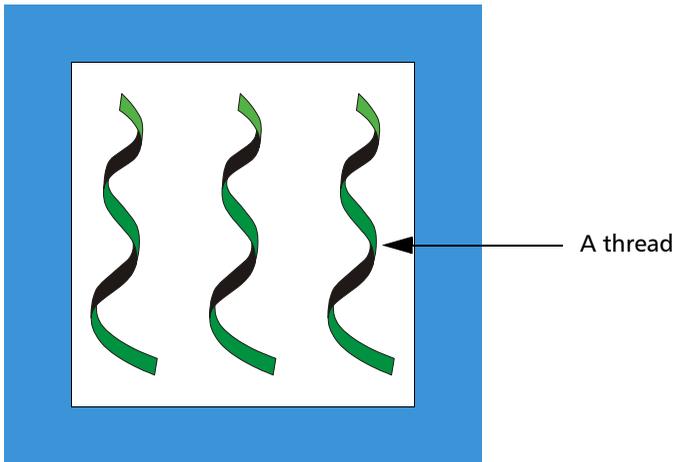


FIGURE 21: **Threads**

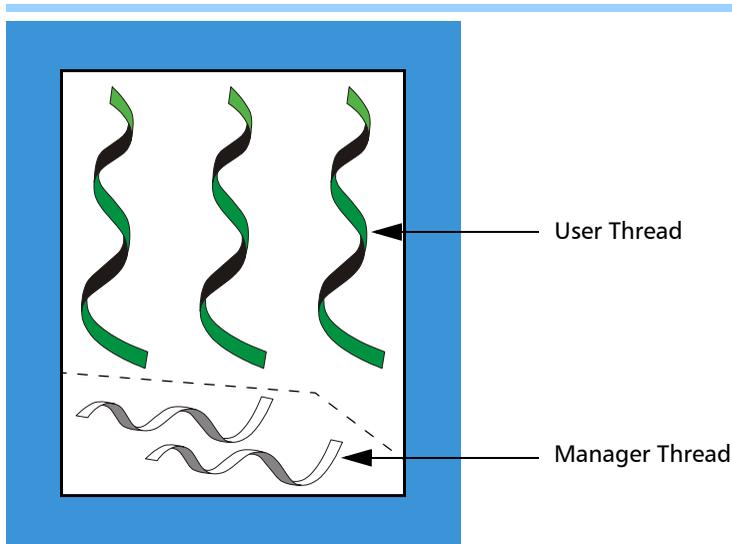
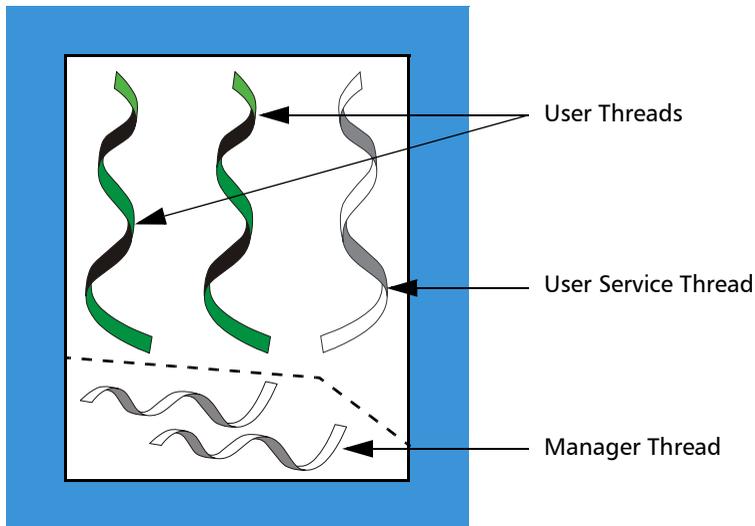
For the moment, assume that all of these threads are *user threads*; that is, they are threads that perform some activity that you've programmed.

NOTE Many computer architectures have something called "user mode", "user space," or something similar. "User threads" means something else. Without trying to be rigorous, the TotalView definition of a "user thread" is simply a unit of execution created by a program.

Because they are created by your program to do the work of your program, they are called *worker threads*.

Other threads can also be executing within the process. For example, the threads that are part of the operating environment are manager threads. A *manager thread* is a thread that your environment or operating system adds to your program to help it get work done. In Figure 22 on page 24, the horizontal threads at the bottom are user-created manager threads.

Things would be nice and easy if this was all there was to it. Unfortunately, all threads are not created equal and all threads do not execute equally. In most cases, a program also creates manager-like threads. As these user-created manager threads are designed to perform services for other threads, they can also be called *service threads*. (See Figure 23 on page 24.)

FIGURE 22: **User Threads and Service Threads**FIGURE 23: **User, Service, and Manager Threads**

These service threads are, of course, also worker threads. They are called different things just to keep the different kinds of things that they do separate. As an example, this could be

a thread whose sole function is to send data to a printer in response to a request from the other two threads.

One reason you need to know which of your threads are service threads is that a service thread performs different kinds of activities from your other threads. Because their activities are different, they are usually developed separately and, in many cases, are not involved with the fundamental problems being solved by the program. The code that sends messages between processes is far different than the code that performs fast Fourier transforms. For example, a service thread that queues and dispatches messages sent from other threads may have bugs, but the bugs are different than the rest of your code and you can deal with them separately from the bugs that occur in non-service user threads.

In contrast, your user threads are the agents performing the program's work, and their interactions are where the action is. Being able to distinguish between the two kinds of threads means that you can focus on the threads and processes that are actively participating in an activity, rather than on threads in the background performing subordinate activities.

So, while this figure shows five threads, most of your debugging effort will focus on just two threads.

Organizing Chaos

While it is possible to debug programs that are running thousands of processes and threads across hundreds of computers by individually looking at each, this is clearly impractical. The only workable approach is to organize your processes and threads into groups and then debug your program by using these groups. In other words, in a multiprocess, multithreaded program, you are most often not programming each process or thread individually. Instead, most high-performance computing programs perform the same or similar activities on different sets of data.

While TotalView cannot know your program's architecture, it can make some intelligent guesses based on what your program is executing and where the program counter is. Using this information, TotalView automatically organizes your processes and threads into the following predefined "groups":

- **Control Group:** All the processes that a program creates. These processes can be local or remote. If your program uses processes that it did not create, TotalView places them in separate control groups. For example, a client/server program has two distinct execut-

bles that run independently of one another. Each would be in a separate control group. In contrast, processes created by `fork()` are in the same control group.

- **Share Group:** All the processes within a control group that share the same code. In most cases, your program will have more than one share group. Share groups, like control groups, can be local or remote.
- **Workers Group:** All the worker threads within a control group. These threads can reside in more than one share group.
- **Lockstep Group:** All threads that are at the same PC (program counter). This group is a subset of a workers group. A lockstep group only exists for stopped threads. By definition, all members of a lockstep group are within the same workers group. That is, a lockstep group cannot have members in more than one workers group or more than one control group.

The first two groups in the above list only contain processes, and the last two only contain threads. Notice that "same code" means that the processes have the same executable file name and path.

TotalView lets you manipulate processes and threads individually and by groups. In addition, you can create your own groups and manipulate a group's contents (to some extent).

NOTE Not all operating systems let you individually manipulate threads.

Figure 24 on page 27 shows a processor running five processes (ignoring daemons and other programs not related to your program) and the threads within the processes. The figure shows a control group and two share groups within this control group.

The elements in this figure are as follows:

CPU	The one outer square. All elements in the drawing operate within one CPU.
Processes	The five white inner squares represent processes being executed by the CPU.
Control Group	The large rounded rectangle that surrounds the five processes. This drawing shows one control group. This diagram doesn't indicate which process is the main procedure.
Share Groups	The two smaller rounded rectangles having white dashed lines surround processes in a share group. This drawing shows two share groups within one control group. The three processes in the first share group have the same executable. The two processes in the second share group share a second executable.

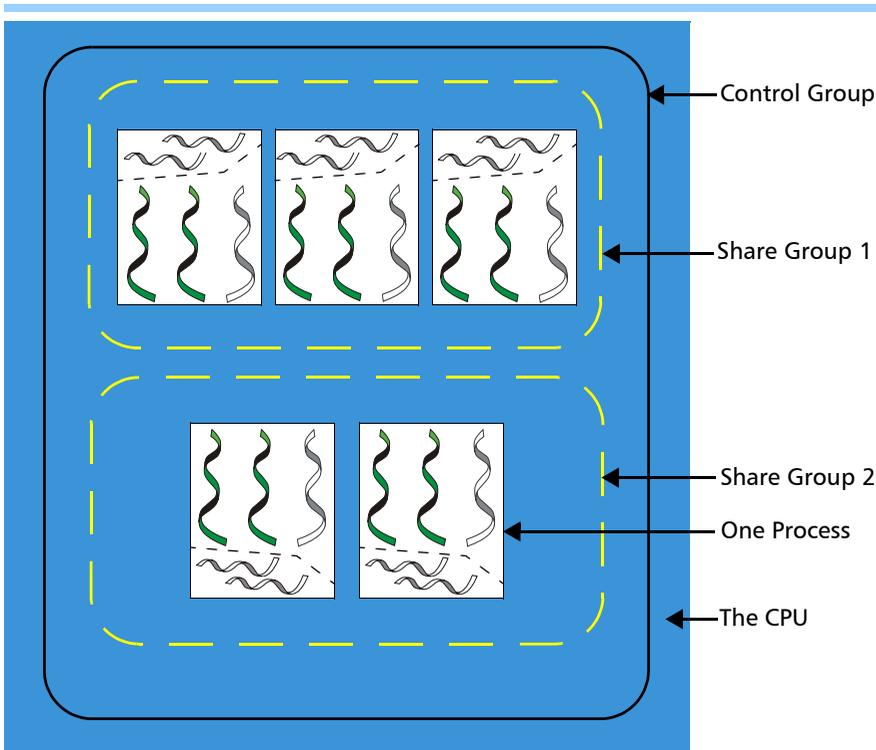


FIGURE 24: **Five Processors and Processor Groups (Part 1)**

The control group and the share group only contain processes. In contrast, the workers group and the lockstep group only contain threads. Figure 25 on page 28 show how TotalView organizes the threads in Figure 24. As you can see, this figure adds the workers group and two lockstep groups.

NOTE The control group is not shown as it encompasses everything in Figure 25.

The elements in this figure are as follows:

Workers Group All nonmanager threads within the control group make up the workers group. Notice that this group includes service threads.

Lockstep Group Each share group has its own lockstep groups. Figure 25 shows two lockstep groups, one in each share group.

If other threads are stopped, this picture indicates that they are not participating in either of these two lockstep groups. Recall that a

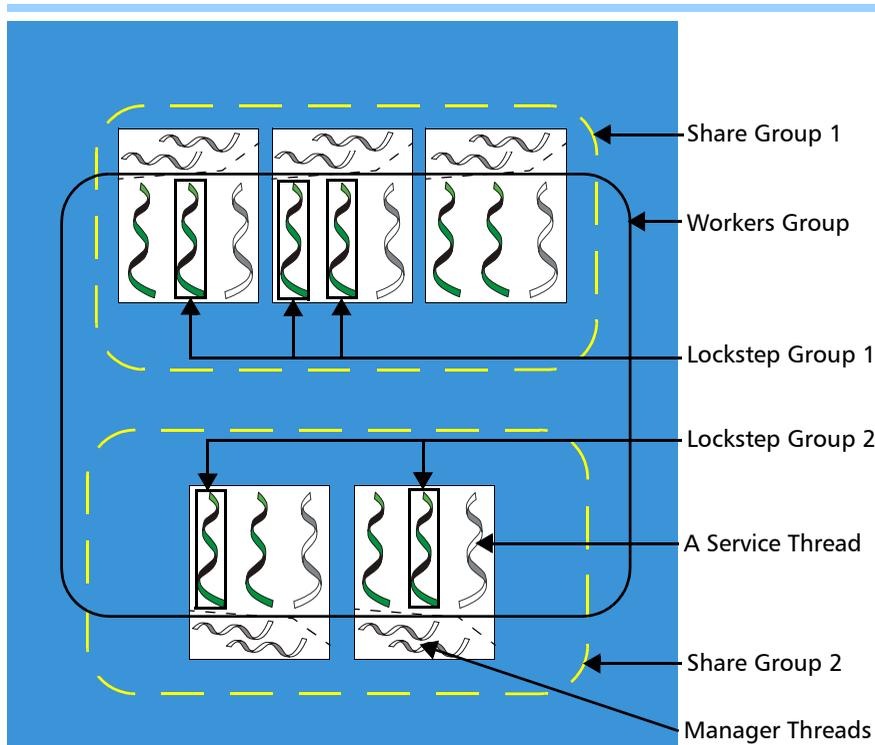


FIGURE 25: **Five Processors and Processor Groups (Part 2)**

stopped thread is always in a lockstep group. (It's OK if a lockstep group has only one member.)

Service Threads Each process has one service thread. While a process can have any number of service threads, this figure only shows one.

Manager Threads The only threads that are not participating in the workers group are the ten manager threads.

Figure 26 extends Figure 25 to show the same kinds of information executing on two processors.

Figure 26 differs from Figure 25 in that it has ten processes executing within two processors rather than five processes within one processor. Although the number of processors has changed, the number of control and share groups is unchanged. This is not to say that the number of groups could not be different. It's just that they are not in this example.

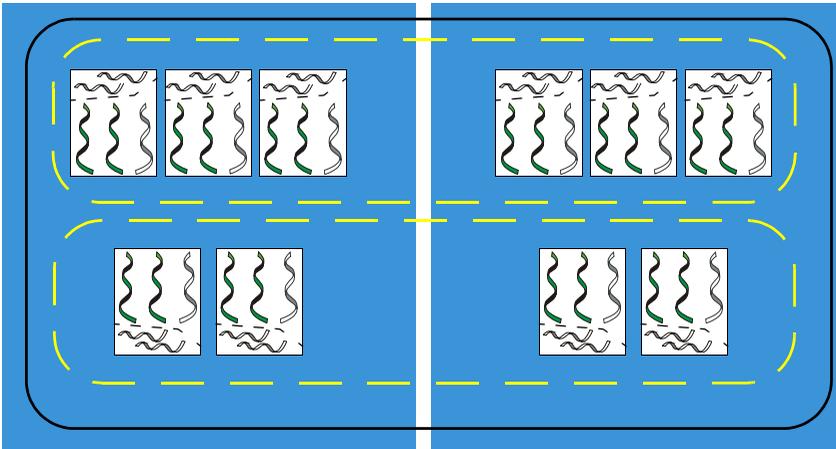


FIGURE 26: **Five Processes and Their Groups on Two Computers**

Creating Groups

TotalView places processes and threads in groups as your program creates them. The exception is the lockstep groups that are created or changed whenever a process or thread hits an action point or is stopped for any reason. While there are many ways in which this kind of organization can be built, the following steps indicate the beginning of how this might occur:

- 1 TotalView and your program are launched and your program begins executing. (See Figure 27.)

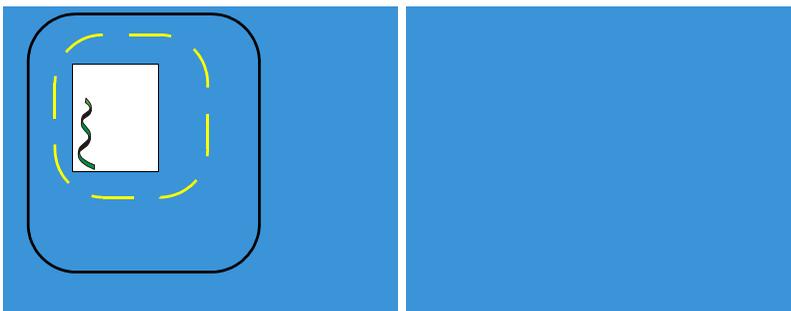


FIGURE 27: **Step 1: A Program Starts**

Control group: A group is created as the program is loaded.

Share group: A group is created as the program begins executing.

Workers group: The thread in the `main()` routine is the workers group.

Lockstep group: There is no lockstep group because the thread is running.

- 2 The program forks a process. (See Figure 28.)

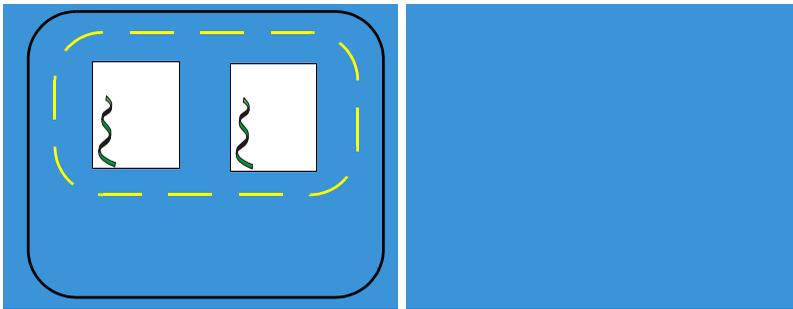


FIGURE 28: Step 2: Forking a Process

Control group: A second process is added to the existing group.

Share group: A second process is added to the existing group.

Workers group: TotalView adds the thread in the second process to the existing group.

Lockstep group: There are no lockstep groups because the threads are running.

- 3 The second process is exec'd. (See Figure 29 on page 31.)

Control group: The group is unchanged.

Share group: TotalView creates a second share group having this exec'd process as a member. TotalView removes this process from the first share group.

Workers group: Both threads are in the workers group.

Lockstep group: There are no lockstep groups because the threads are running.

- 4 The first process hits a break point.

Control group: The group is unchanged.

Share group: The groups are unchanged.

Workers group: The group is unchanged.

Lockstep group: TotalView creates a lockstep group whose member is the thread of the current process. (In this example, each thread is its own lockstep group.)

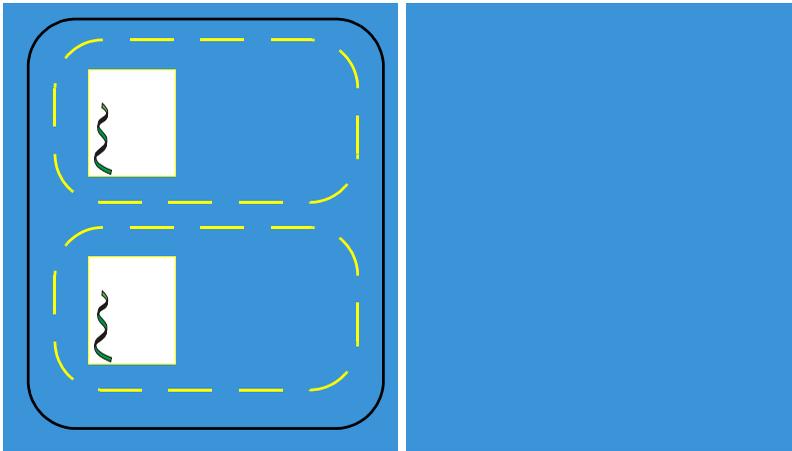


FIGURE 29: **Step 3: Exec'ing a Process**

- 5 The program is continued and TotalView starts a second version of your program from the shell. You attach to it within TotalView and put it in the same control group as your first process. (See Figure 30.)

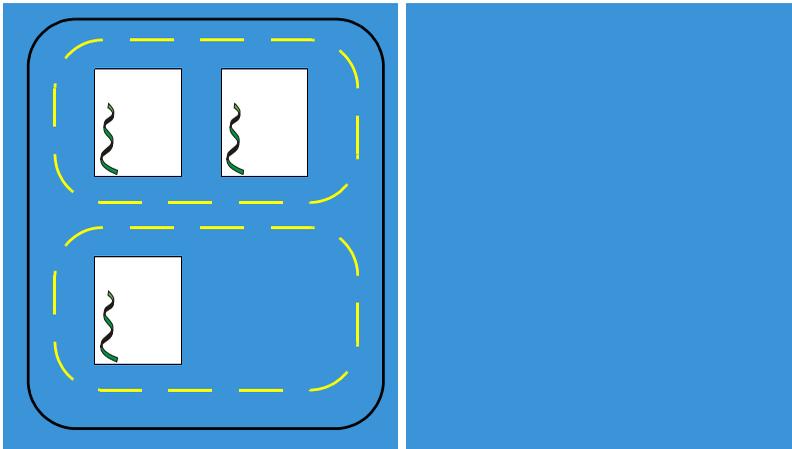


FIGURE 30: **Step 5: Creating a Second Version**

Control group: TotalView adds a third process.

Share group: TotalView adds this third process to the first share group.

Workers group: TotalView adds the thread in this third process to the group.

Lockstep group: There are no lockstep groups because the threads are running.

- 6 Your program creates a process on another computer. (See Figure 31.)

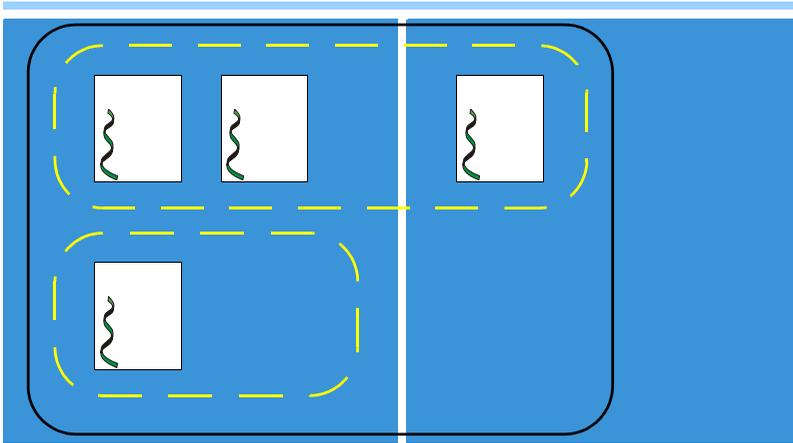


FIGURE 31: Step 6: Creating a Remote Process

Control group: TotalView extends the control group so that it contains the fourth process running on the second computer.

Share group: The first share group now contains this newly created process even though it is running on the second computer.

Workers group: TotalView adds the thread within this fourth process to the workers group.

Lockstep group: There are no lockstep groups because the threads are running.

- 7 A process within control group 1 creates a thread. This adds a second thread to one of the processes. (See Figure 32 on page 33.)

Control group: The group is unchanged.

Share group: The group is unchanged.

Workers group: TotalView adds a fifth thread to this group.

Lockstep group: There are no lockstep groups because the threads are running.

- 8 A breakpoint is set on a line in a process executing in the first share group, and the breakpoint is shared. The process executes until all three processes are at the breakpoint. (See Figure 33 on page 33.)

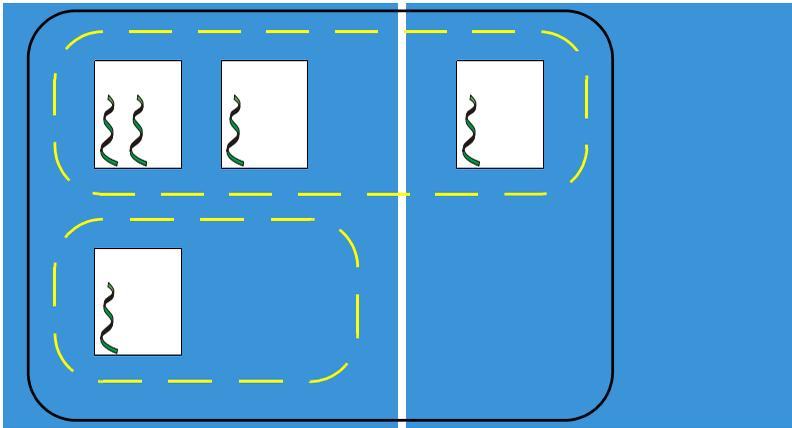


FIGURE 32: Step 7: A Thread Is Created

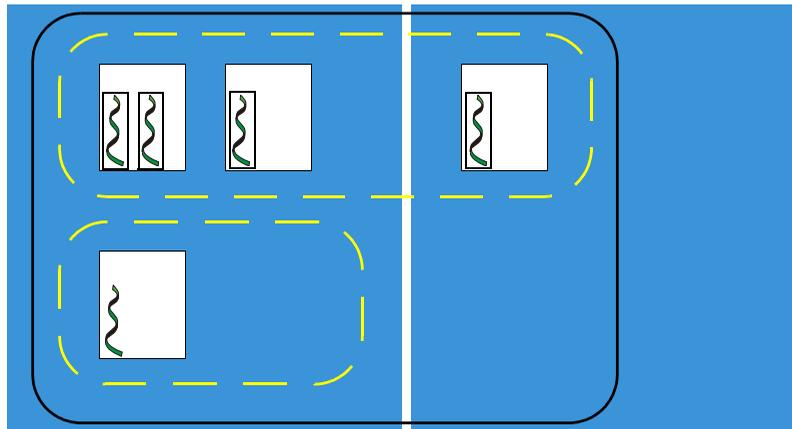


FIGURE 33: Step 8: Hitting a Breakpoint

Control group: The group is unchanged.

Share group: The groups are unchanged.

Workers group: The group is unchanged.

Lockstep groups: TotalView creates a lockstep group whose members are the four threads in the first share group.

9 You tell TotalView to step the lockstep group. (See Figure 34.)

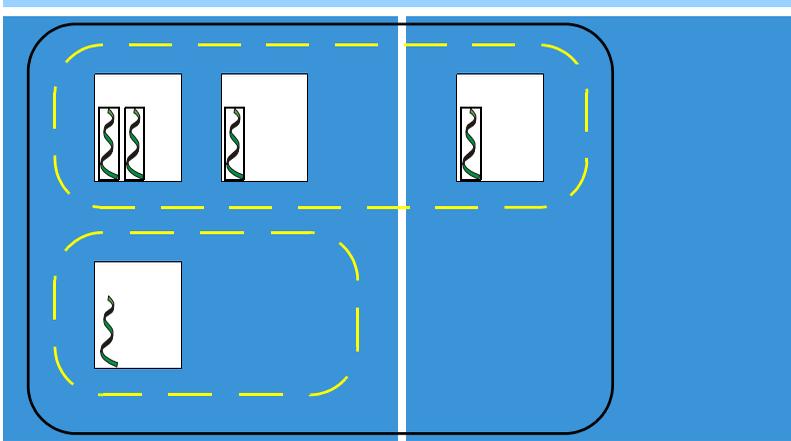


FIGURE 34: Step 9: Stepping the Lockstep Group

Control group: The group is unchanged.

Share group: The groups are unchanged.

Workers group: The group is unchanged.

Lockstep group: The lockstep groups are unchanged. (Note that there are other lockstep groups. This will be explained in Chapter 11.)

Clearly, this example could keep on going until a much more complicated system of processes and threads was created. However, adding more processes and threads won't do anything much different than what's been discussed.

Simplifying What You're Debugging

The reason you're using a debugger is because your program isn't operating correctly and the way you think you're going to solve the problem (unless it is a $\&\%\$ \#$ operating system problem, which, of course, it usually is) is by stopping your program's threads, examining the values assigned to variables, and stepping your program so you can see what's happening as it executes.

Unfortunately, your multiprocess, multithreaded program and the computers upon which it is executing have lots of things executing that you want TotalView to ignore. For example,

you don't want to be examining manager and service threads that the operating system, your programming environment, and your program create.

Also, most of us are incapable of understanding exactly how a program is acting when perhaps thousands of processes are executing asynchronously. Fortunately, there are only a few problems that require full asynchronous behavior at all times.

One of the first simplifications you can make is to change the number of processes. For example, suppose you have a buggy MPI program running on 128 processors. Your first step might be to have it execute in an 8-processor environment.

After you get the program running under TotalView's control, you will want to run the process being debugged to an action point so you can inspect the program's state at that place. In many cases, because your program has places where processes are forced to wait for an interaction with other processes, you can ignore what they are doing.

NOTE TotalView lets you control as many groups, processes, or threads as you need to control. While each can be controlled individually, you will probably have problems remembering what you're doing if you're controlling large numbers of these things independently. The reason that TotalView creates and manages groups is so you can focus on portions of your program.

In most cases, you don't need to interact with everything that is executing. Instead, you want to focus on one process and the data that this process is manipulating. Things get complicated when the process being investigated is using data created by other processes, and these processes might be dependent on other processes.

This means that there is a rather typical pattern to the way you use TotalView to locate problems:

- 1 At some point, you should make sure that the groups you are manipulating do not contain service or manager threads. (You can remove processes and threads from a group with the `dgroups -remove` command or the **Groups > Edit Group** command.)
- 2 Place a breakpoint in a process or thread and begin investigating the problem. In many cases, you are setting a breakpoint at a place where you hope the program is still executing correctly. Because you are debugging a multiprocess, multithreaded program, you will want to set a *barrier point*—this is a special kind of breakpoint—so that all threads and process will stop at the same place.

NOTE Don't step your program except where you need to individually look at what occurs. Using barrier points is much more efficient.

- 3 After execution stops at a barrier point, look at the contents of your variables. Verify that your program state is actually correct.
- 4 Begin stepping your program through its code. In most cases, step your program synchronously or set barriers so that everything isn't running freely.

Here's where things begin to get complicated. You've been focusing on one process or thread. If another process or thread is modifying the data and you become convinced that this is the problem, you'll want to go off to it and see what's going on.

At this point, you need to keep your focus narrow so that you're only investigating a limited number of behaviors. This is where debugging becomes an art. A multiprocess, multi-threaded program can be doing a great number of things. Understanding where to look when problems occur is the "art."

For example, you'll most often execute commands at the default focus. Only when you think that the problem is occurring in another process will you change to that process. You'll still be executing in a default focus, but this time the default focus is concentrated on other process.

While it will often seem like you need to do a lot of shifting to another focus, what you will probably do is:

- Modify the focus so that it affects just the next command. If you are using the GUI, you might select this process and thread from the list displayed in the Root Window. If you are using the CLI, you would use the **dfocus** command to limit the scope of a future command. For example, here's the CLI command that steps thread 7 in process 3:
`dfocus t3.7 dstep`
- Use the **dfocus** command to change focus temporarily, execute a few commands, and then return to the original focus.

What you've been reading is just an overview of the threads, processes, and groups. You'll find a lot more information in Chapter 11, "Using Groups, Processes, and Threads" on page 239.



Part II: Setting Up

This section of the TotalView Users Guide contains information about running TotalView in the different kinds of environments you execute your program in.

Chapter 3: Setting Up a Debugging Session

The way you configure your personal environment is the same on all operating systems and in all environments. This chapter tells you what you need to know to start TotalView and tailor how it works.

You should, at a minimum, glance at this chapter to see what's here so you can come back at a later time, if necessary.

Chapter 4: Setting Up Remote Debugging Sessions

When you are debugging a program that has processes executing on a remote computer, TotalView launches server processes when your program launches remote processes. The primary focus of this chapter is what to do when there are problems.

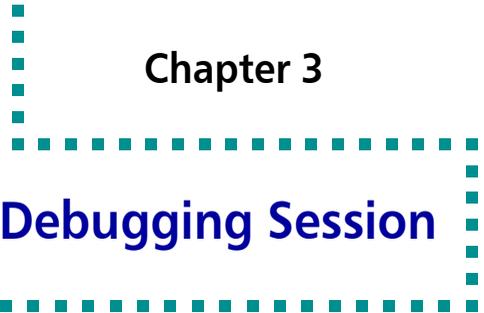
If you aren't having problems, you probably won't need the information in this chapter.

Chapter 5: Setting Up Parallel Debugging Sessions

TotalView lets you debug programs created using many different parallel environments such as OpenMP, MPI, MPICH, UPC, and the like. This chapter discusses each environment.

Because this chapter has individual presentations for the different environments, locate what you need and skip the rest.





Chapter 3

Setting Up a Debugging Session

This chapter explains how to set up a TotalView session. It also describes some common commands and procedures. For information on setting up remote debugging sessions, see Chapter 4, *"Setting Up Remote Debugging Sessions"* on page 73. For information on setting up parallel debugging sessions, see Chapter 5, *"Setting Up Parallel Debugging Sessions"* on page 91.

In this chapter, you will learn about:

- *"Compiling Programs"* on page 40
- *"Exiting from TotalView"* on page 46
- *"Exiting from TotalView"* on page 46
- *"Loading Executables"* on page 46
- *"Attaching to Processes"* on page 49
- *"Detaching from Processes"* on page 52
- *"Examining Core Files"* on page 53
- *"Viewing Process and Thread State"* on page 54
- *"Handling Signals"* on page 56
- *"Setting Search Paths"* on page 59
- *"Setting Command Arguments"* on page 61
- *"Setting Input and Output Files"* on page 62
- *"Setting Preferences"* on page 64
- *"Setting Environment Variables"* on page 70
- *"Monitoring TotalView Sessions"* on page 71

Compiling Programs

Before you start debugging a program, you must compile it. The only thing you do differently when debugging with TotalView is to compile using the `-g` option. This option tells your compiler to generate symbol table debugging information. For example:

```
cc -g -o executable source_program
```

You can also debug programs that you did not compile using the `-g` option or programs for which you don't have source code. For more information, refer to "Viewing the Assembler Version of Your Code" on page 216.

Table 3 presents some general considerations, but you should check "Compilers and Platforms", in the TOTALVIEW REFERENCE GUIDE to determine the exact syntax and any other considerations.

Table 3: Compiler Considerations

Compiler Option or Library	What It Does	When to Use It
Debugging symbols option (usually <code>-g</code>)	Generates debugging information in the symbol table.	Before debugging <i>any</i> program with TotalView.
Optimization option (usually <code>-O</code>)	Rearranges code to optimize your program's execution. Some compilers won't let you use the <code>-O</code> and the <code>-g</code> option at the same time. Even if your compiler lets you use the <code>-O</code> option, don't do it when debugging your program as strange results often occur.	After you finish debugging your program with TotalView.

Table 3: Compiler Considerations

Compiler Option or Library	What It Does	When to Use It
Multiprocess programming library (usually dbfork)	<p>Uses special versions of the fork() and execve() system calls.</p> <p>In some cases, you will need to use -lpthread.</p> <p>Using dbfork is discussed in "Linking with the dbfork Library" contained in the "Compilers and Platforms" Chapter of the TOTALVIEW REFERENCE GUIDE.</p>	<p>Before debugging a multiprocess program that explicitly calls fork() or execve().</p> <p>Refer to "Processes That Call <i>fork()</i>" on page 348 and "Processes That Call <i>execve()</i>" on page 348.</p>

File Extensions

When TotalView reads a file, it uses the file's extension to determine the programming language. If TotalView does the wrong thing, you can have it do the right thing by setting the **TV::suffixes** variable in a startup file. For more information, see the "TotalView Variables" chapter in the TOTALVIEW REFERENCE GUIDE.

Starting TotalView

TotalView can debug programs that run in many different computing environments and which use a variety of parallel processing modes. This section looks at few of the ways you can start TotalView. The "TotalView Command Syntax" chapter in the TOTALVIEW REFERENCE GUIDE contains more detailed information.

In most cases, the command for starting TotalView looks like:

```
totalview [ executable [ corefile ] ] [ options ]
```

where *executable* is the name of the executable file you will be debugging and *corefile* is the name of the core file being examined.

```
CLI EQUIVALENT: totalviewcli [ executable [ corefile ] ] [ options ]
```

In some cases, you may need to do something different. For example, if you are debugging an MPI program, you must invoke TotalView on `mpirun`. You'll find details in Chapter 5, "Setting Up Parallel Debugging Sessions" on page 91.

If the GUI is executing, you can invoke the CLI by selecting the **Tools > Command Line** command. In this case, both the GUI and the CLI will be executing.

Here are some examples that show how to start TotalView:

Just start TotalView

`totalview`

Starts TotalView without loading a program or core file. After TotalView starts, you can load a program by using the **File > New Program** command.

CLI EQUIVALENT: `totalviewcli then dload executable`

Debugging a program

`totalview executable`

Starts TotalView and loads the *executable* program.

CLI EQUIVALENT: `totalviewcli executable`

Debugging a core file

`totalview executable corefile`

Starts TotalView and loads the *executable* program and the *corefile* core file.

CLI EQUIVALENT: `dattach -c corefile -e executable`

Passing arguments to the program being debugged

`totalview executable -a args`

Starts TotalView and passes all the arguments following `-a` to the *executable* program. When you use the `-a` option, you must enter it as the last TotalView option on the command line.

CLI EQUIVALENT: `totalviewcli executable -a args`

If you don't use `-a` and want to add arguments after TotalView loads your program, use the **Process > Startup** command.

CLI EQUIVALENT: `dset ARGS_DEFAULT {value}`

Debugging a program that runs on another computer

`totalview executable --remote hostname_or_address[:portnumber]`

Starts TotalView on your local host and the TotalView Debugger Server (**tvdsvr**) on a remote host. After TotalView begins executing, it loads the program specified by *executable* for remote debugging. You can specify a host name or a TCP/IP address. If you need to, you can also enter the TCP/IP port number.

CLI EQUIVALENT: `totalviewcli executable
-r hostname_or_address[:portnumber]`

For more information on:

- Debugging parallel programs such as MPI, PVM, or UPC, refer to Chapter 5, "Setting Up Parallel Debugging Sessions" on page 91.
- The **totalview** command, refer to "TotalView Command Syntax" in the TOTALVIEW REFERENCE GUIDE.
- Remote debugging, refer to "Starting the TotalView Debugger Server" on page 73 and "TotalView Debugger Server (tvdsvr) Command Syntax" in the TOTALVIEW REFERENCE GUIDE.

Initializing TotalView

When TotalView begins executing, it can grab initialization from many places. The two most commonly used are initialization files that you create and preference files that TotalView creates.

An initialization file is a place where you can store CLI functions, set variables, and execute actions. TotalView will execute this file whenever it starts executing. This file, which you must name **tvdrc**, resides in a **.totalview** subdirectory contained in your home directory.

TotalView can actually read more than one initialization file. You can place these files in your installation directory, the **.totalview** subdirectory, or the directory in

which you invoke TotalView. If the file is present in one or all of these places, TotalView reads and executes its contents. Only the initialization file within your `.totalview` directory has the name `tvdrvc`. The other initialization files have the name `.tvdrvc`. Notice the dot preceding the file name.

NOTE Before Version 6.0, you would place your personal `.tvdrvc` file in your home directory. If you do not move this file into the `.totalview` directory, TotalView will still find it. However, if you also have a `tvdrvc` file in the `.totalview` directory, TotalView will ignore the `.tvdrvc` in your home directory.

TotalView writes your preferences file into your `.totalview` subdirectory. It's name is `preferences6.tvd`. Do not modify this file as TotalView will overwrite it whenever it saves your preferences.

If you add the `-s filename` option to either the `totalview` or `totalviewcli` shell command, TotalView executes the CLI commands contained in `filename`. This startup file will execute after `.tvdrvc` files execute. The `-s` option lets you, for example, initialize the debugging state of your program, run the program you're debugging until it reaches some point where you're ready to begin debugging, and even lets you create a shell command that starts the CLI.

Figure 35 shows the order in which TotalView executes initialization and startup files.

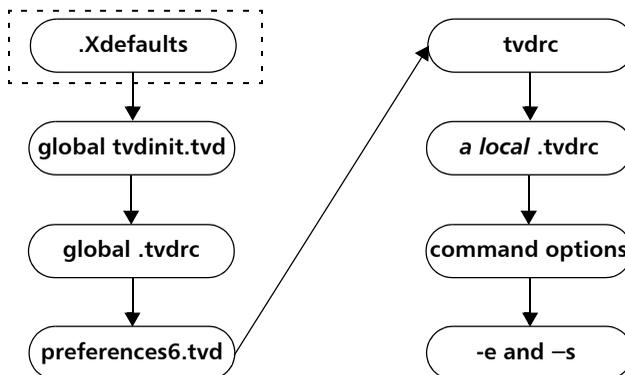


FIGURE 35: **Startup and Initialization Sequence**

The `.Xdefaults` file, which is actually read by the server when you start X Windows, is only used by the GUI. The CLI ignores it. Prior to TotalView release 6.0, the `.Xdefaults` file was extensively used. Beginning at TotalView 6.0, its use is negligible.

NOTE If you have an X resources file, TotalView will read it the first time Release 6.0 starts executing. It will then write any TotalView resources it finds to your `preferences6.tvd` file. If you change a value after this file is written, TotalView will ignore your change. The only exceptions are Visualizer X resources. For information on these resources, go to www.et-nus.com/Support/docs/xresources/XResources.html. You can force TotalView to reread your X resources by deleting your preferences file.

As part of the initialization process, TotalView exports three environment variables into your environment: `LM_LICENSE_FILE`, `TVROOT`, and either `SHLIB_PATH` or `LD_LIBRARY_PATH`.

If you have saved a action point file into the same subdirectory as your program, TotalView automatically reads the information in this file when it loads your program.

NOTE The format of a Release 6.0 action point file differs from that used in earlier releases. While TotalView 6.0 can read action point files created by earlier versions, earlier versions cannot read a Release 6.0 action point file.

You can also invoke scripts by naming them in the `TV::process_load_callbacks` list. For information, see “*Initializing TotalView After Loading an Image*” in the “*Type Transformations*” chapter of the TOTALVIEW REFERENCE GUIDE.

If you are debugging multiprocess programs that run on more than one computer, TotalView caches library information in the `.totalview` subdirectory. If you wish to move this cache to another location, set the `TV::library_cache_directory` to this location. The files within this cache directory can be shared among users.

Exiting from TotalView

You can exit from TotalView by selecting the **File > Exit** command. You can select this command in the Root, Process, and Variable Windows. (See Figure 36.)



FIGURE 36: **File > Exit** Dialog Box

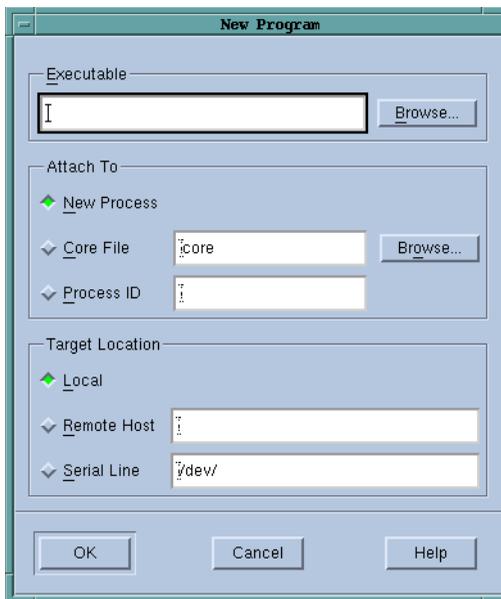
Select **Yes** to exit. Otherwise, select **No**. As TotalView exits, it kills all programs and processes that it started. However, programs and processes that TotalView did not start continue to execute.

CLI EQUIVALENT: `exit`

NOTE If you have a CLI window open, TotalView also closes this window. Similarly, if you type “exit” in the CLI, the CLI will close GUI windows.

Loading Executables

TotalView can debug programs on local and remote hosts and programs that you access over networks and serial lines. The **File > New Program** command, which is located in the Root and Process Windows, loads local and remote programs, core files, and processes that are already running. (See Figure 37.)

FIGURE 37: **File > New Program Dialog Box**

The controls within this dialog box lets you:

- **Load a new executable**

Type the path name into the **Executable** field.

CLI EQUIVALENT: `dload -e executable`

- **Load a core file**

Type the name into the **Core File** field. You must also type the path name of the executable associated with this core file in the **Executable** field.

CLI EQUIVALENT: `dattach -c corefile -e executable`

- **Load a program using process ID**

Type a process ID into the **Process ID** field *and* type the associated executable's path name into the **Executable** field.

CLI EQUIVALENT: `dattach executable pid`

If you need to debug a program on a remote machine, type the machine's host name or IP address in the **Remote Host** field. If the program is local, make sure that you have selected the **Local** button.

CLI EQUIVALENT: `dload executable -r hostname`

You can use a full or relative path name in the **Executable** and **Core File** fields. If you enter a file name, TotalView searches for it in the list of directories named using the **File > Search Path** command or listed in your **PATH** environment variable.

CLI EQUIVALENT: `dset EXECUTABLE_PATH`

If you select **New Process**, TotalView always loads a new copy of the program you named in the **Executable** field. Even if the program is already loaded, TotalView loads another copy.

Debugging over a serial line is discussed in "Debugging Over a Serial Line" on page 86.

Loading Remote Executables

If TotalView fails to automatically load a remote executable, you may need to disable *autolaunching* for this connection and manually start the TotalView Debugger Server (**tvdsvr**) manually. (*Autolaunching* is the process of automatically launching **tvdsvr** processes.) You can disable autolaunching by adding the *hostname:portnumber* suffix to the name you type in the **Remote Host** field. As always, the *portnumber* is the TCP/IP port number on which the debugger server is communicating with TotalView. Refer to "Starting the TotalView Debugger Server" on page 73 for more information.

NOTE You cannot examine core files on remote systems.

You can connect to a remote machine in three ways:

- Using the `-remote` command-line option when you start TotalView. For details on the syntax for the `-remote` command-line option, see "Starting TotalView" on page 41.

- With the **File > New Program** command after you start TotalView.

CLI EQUIVALENT: `dload executable -r hostname`

- By first connecting to a remote host using the **File > New Program** command and then displaying the **Unattached** Page of the Root Window. You can now attach to these programs by diving into them.

CLI EQUIVALENT: **If you're using the CLI, you will need to know the file's name so that you can use the `dattach` command to attach to the program.**

NOTE If TotalView supports your program's parallel process runtime library (for example, MPI, PVM, or UPC), it automatically connects to remote hosts. For more information, see Chapter 5, "Setting Up Parallel Debugging Sessions" on page 91.

Attaching to Processes

If a program you're testing is hung or looping (or misbehaving in some other way), you can attach to it while it is running. You can attach to single processes, multiprocess programs, and these programs can be running remotely.

To attach to a process, either use the **Unattached** Page in the Root Window or use the **File > New Program** command located on the Root and Process Windows. (Using the **Unattached** Page is easier if the process is listed. However, if it's not there, you must use the **File > New Program** command.)

CLI EQUIVALENT: `dattach executable pid`

If the process or any of its children calls the `execve()` routine, you may need to attach to it by creating a new Process Window. This is because TotalView relies on the `ps` command to obtain the process name, and it can make mistakes.

NOTE When you exit from TotalView, TotalView kills all programs and processes that it started. However, programs and processes that were executing before you brought them under TotalView's control continue to execute.

Attaching Using the Unattached Page

Here's the procedure for using the **Unattached** Page to attach a process:

- 1 Go to the Root Window and select the **Unattached** Page tab.

This page lists the process ID, status, and name of each process associated with your username. The processes that appear dimmed are those that are being debugged or those that TotalView won't allow you to debug. For example, you can't debug the TotalView process itself. (See Figure 38.) The processes at the top of this figure are local. The remaining processes are remote.

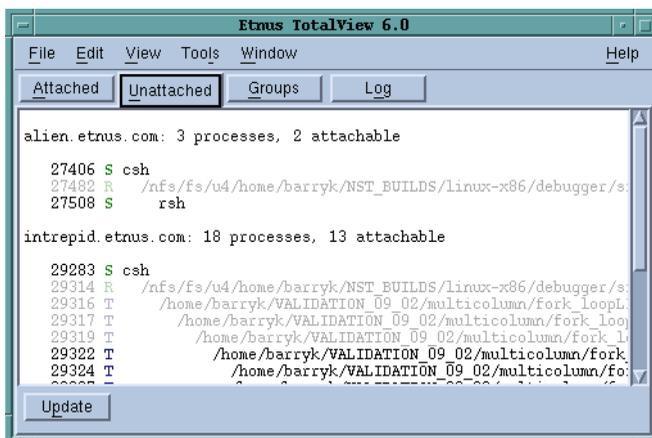


FIGURE 38: **Unattached Page**

If you're debugging a remote process, the **Unattached** Page also shows processes running under your username on each remote host name. You can attach to any of these remote processes. For more information on remote debugging, refer to "Starting the TotalView Debugger Server" on page 73 and "TotalView Debugger Server (tvdsvr) Command Syntax" in the TOTALVIEW REFERENCE GUIDE.

- 2 Dive into the process you wish to debug by double-clicking on it.

A Process Window appears. The right arrow points to the current program counter (PC), indicating where the program was executing when TotalView attached to it.

Attaching Using File > New Program and dattach

Here's the procedure for using the Root Window's **File > New Program** command to attach to a process:

- 1 Use the **ps** shell command to obtain the process ID (PID) of the process.
- 2 Select the **File > New Program** command. TotalView displays the dialog box shown in Figure 39.

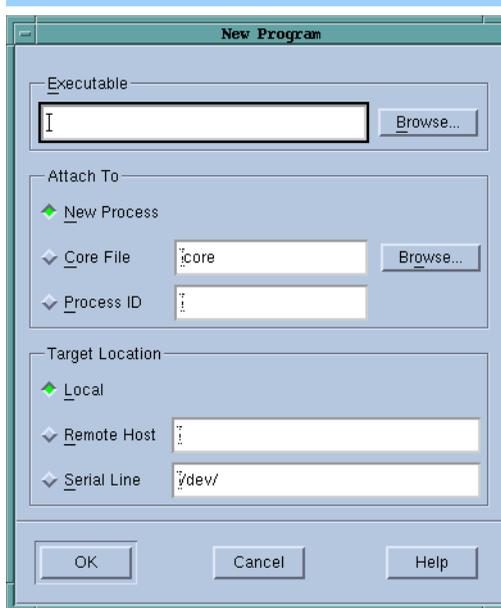


FIGURE 39: **File > New Program Dialog Box**

Enter a file name in the **Executable** field. This name can be a full or relative path name. If you supply a simple file name, TotalView searches for it in the directories named using the **File > Search Path** command or listed in your **PATH** environment variable.

Enter the process ID (PID) of the *unattached* process into the **Process ID** field.

CLI EQUIVALENT: **dattach pid**
dset EXECUTABLE_PATH

- 3 Select **OK**.

If the executable is a multiprocess program, TotalView asks if you want to attach to all relatives of the process. To examine all processes, select **Yes**.

If the process has children that call `execve()`, TotalView tries to determine each child's executable. If TotalView can't figure it out, you must delete (*kill*) the parent process and start it again using TotalView.

A Process Window will appear. In this window, the right arrow points to the current program counter (PC), which is where the program was executing when TotalView attached to it.

Detaching from Processes

You can use the following procedure to detach from processes that TotalView did not create:

- 1 After opening a Process Window on the process, select the **Thread > Continuation Signal** command. Choose the signal that TotalView should send to the process when it detaches from the process. For example, to detach from a process and leave it stopped, set the continuation signal to **SIGSTOP**. (See Figure 40.)

CLI EQUIVALENT: **No equivalent to Thread > Continuation exists.**



FIGURE 40: **Thread > Continuation Signal Dialog Box**

- 2 Select the **Process > Detach** command.

CLI EQUIVALENT: **ddetach**

When you detach from a process, TotalView removes all breakpoints that you have set within it.

Examining Core Files

If a process encounters a serious error and dumps a core file, you can look at it using one of the following methods:

- Start TotalView as follows:

totalview *filename corefile* [*options*]

CLI EQUIVALENT: **totalviewcli** *filename corefile* [*options*]

- Select the **File > New Program** command from the Root Window. In the middle section of the dialog box, type the name of the core file in the **Core File** field, and then select **OK**. In the top portion, enter the *executable*'s name.

CLI EQUIVALENT: **dattach -c corefile -e executable**

NOTE You can only debug local core files. You can, however, debug core files at a remote location if you log on to the remote machine and then start TotalView on the now local core file. In this case, TotalView is running *locally* on the remote machine (that is, TotalView is now local to the machine upon which the application and core file reside).

The Process Window displays the core file, with the Stack Trace, Stack Frame, and Source Panes showing the state of the process when it dumped core. The title bar of the Process Window names the signal that caused the core dump. The right arrow in the line number area of the Source Pane indicates the value of the program counter (PC) when the process encountered the error.

You can examine the state of all variables at the time the error occurred. "Examining and Changing Data" on page 277 contains more information.

If you start a process while you're examining a core file, TotalView stops using the core file and switches to this new process.

Viewing Process and Thread State

Process and thread state is displayed in:

- The **Attached** Page of the Root Window, for processes and threads.
- The **Unattached** Page of the Root Window, for processes.
- The process and thread status bars of the Process Window.
- The Threads Pane of the Process Window.
- The P/T Set Browser.

Figure 41 shows TotalView displaying process state information in the **Attached** Page.

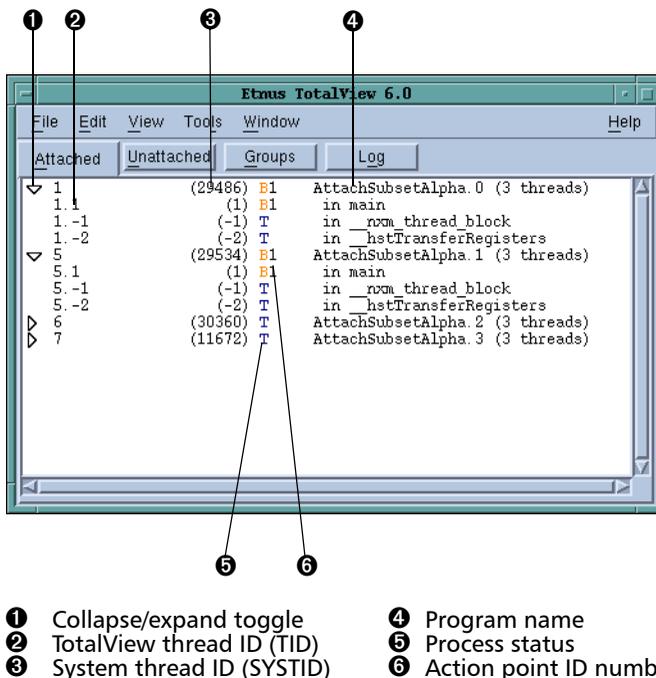


FIGURE 41: **Attached Page Showing Process and Thread Status**

The status of a process includes the process location, the process ID, and the state of the process. (These characters are explained in “*Attached Process States*” on page 55.)

CLI EQUIVALENT: **dstatus and dptsets**

When you use either of these commands, TotalView also displays state information.

The **Unattached** Page lists all processes associated with your username. The information in this page is similar to the information in the **Attached** Page, differing only in that TotalView dims out the processes being debugged. The status bars in the Process Window display similar information. (See Figure 42.)



FIGURE 42: **Process and Thread Labels in the Process Window**

NOTE If the TotalView-assigned thread ID and the system-assigned thread ID are the same, TotalView displays only one ID value.

Attached Process States

TotalView uses the following letters to indicate process and thread state. (The position of these letters in the **Attached** Page is indicated by ⑤ in Figure 41.):

Table 4: Attached Process and Thread States

State Code	State Name
blank	Exited or never created
B	At breakpoint
E	Error reason
K	In kernel
M	Mixed
R	Running
T	Stopped reason
W	At watchpoint

The **Error** state usually indicates that your program received a fatal signal such as **SIGSEGV**, **SIGBUS**, or **SIGFPE** from the operating system. See “*Handling Signals*” on

page 56 for information on controlling how TotalView handles signals that your program receives.

CLI EQUIVALENT: **The CLI prints out a word indicating the state; for example, "breakpoint."**

Unattached Process States

TotalView derives the state information for a process displayed in the **Unattached** Page from the operating system. The state characters TotalView uses to summarize the state of an unattached process do not necessarily match those used by the operating system. Here are the state indicators that TotalView displays:

Table 5: Summary of Unattached Process States

State Code	State Description
I	Idle
R	Running
S	Sleeping
T	Stopped
Z	Zombie

Handling Signals

If your program contains a signal handler routine, you may need to adjust the way TotalView handles signals. You can do this using:

- A dialog box (described in this section)
- The `--signal_handling_mode` command-line option to the **totalview** and **totalviewcli** commands (refer to "TotalView Command Syntax" in the TOTALVIEW REFERENCE GUIDE)

Unless you tell TotalView otherwise, here is how it handles UNIX signals:

Table 6: Default Signal Handling Behavior

Signals that TotalView Passes Back to Your Program		Signals that TotalView Treats as Errors	
SIGHUP	SIGIO	SIGILL	SIGPIPE
SIGINT	SIGIO	SIGTRAP	SIGTERM
SIGQUIT	SIGPROF	SIGIOT	SIGTSTP
SIGKILL	SIGWINCH	SIGEMT	SIGTTIN
SIGALRM	SIGLOST	SIGFPE	SIGTTOU
SIGURG	SIGUSR1	SIGBUS	SIGXCPU
SIGCONT	SIGUSR2	SIGSEGV	SIGXFSZ
SIGCHLD		SIGSYS	

NOTE TotalView uses the SIGTRAP and SIGSTOP signals internally. If a process receives either of these signals, TotalView neither stops the process with an error nor passes the signal back to your program. You cannot alter the way TotalView uses these signals.

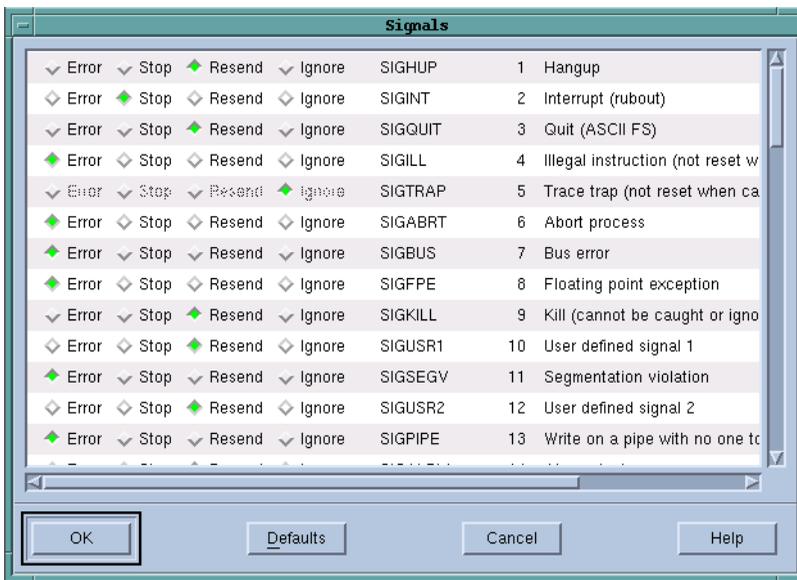
On some systems, hardware registers affect how TotalView and your program handle signals such as SIGFPE. For more information, refer to “*Interpreting Status and Control Registers*” on page 237 and “*Architectures*” in the TOTALVIEW REFERENCE GUIDE.

NOTE If you are using an SGI computer, setting the TRAP_FPE environment variable to any value indicates that your program will trap underflow errors. If you set this variable, however, you will also need to use the controls in the File > Signals Dialog Box to indicate what TotalView should do with SIGFPE errors. (In most cases, you will set SIGFPE to *Resend*.) As an alternative, you can use the `-signal_handling_mode "action_list"` option.

You can change the signal handling mode using the File > Signals command. (See Figure 43 on page 58.)

NOTE The signal names and numbers that TotalView displays are platform specific.

When your program receives a signal, TotalView stops all related processes. If you don't want this behavior, clear the **Stop control group on error signal** button

FIGURE 43: **File > Signals** Dialog Box

(which is found in the **Options** Page of the **File > Preferences** Dialog Box. (See Figure 48.)

CLI EQUIVALENT: `dset TV::warn_step_throw`

When your program encounters an error signal, TotalView opens or raises the Process Window. Clearing the **Open process window on error signal** check box, also found on the **Options** Page in the **File > Preferences** Dialog Box, tells TotalView that it should not open or raise windows.

CLI EQUIVALENT: `dset TV::GUI::pop_on_error`

If processes in a multiprocess program encounter an error, TotalView only opens a Process Window for the first process that encounters an error. (If it did it for all of them, TotalView would be filling up the screen with Process Windows.)

If you select the **Open process window at breakpoint** check box, which is found in the **File > Preferences' Action Points** Page, TotalView opens or raises the Process Window when your program reaches a breakpoint.

CLI EQUIVALENT: **TV::GUI::pop_at_breakpoint**

Make your changes by selecting one of the following radio buttons:

- | | |
|---------------|--|
| Error | Stops the process, places it in the <i>error</i> state, and displays an error in the title bar of the Process Window. If you have also selected the Stop control group on error signal check box, TotalView will also stop all related processes. Select this button for severe error conditions such as SIGSEGV and SIGBUS . |
| Stop | Stops the process and places it in the <i>stopped</i> state. Select this button if you want TotalView to handle this signal as it would a SIGSTOP signal. |
| Resend | Sends the signal back to the process. This setting lets you test your program's signal handling routines. TotalView sets the SIGKILL and SIGHUP signals to Resend as most programs have handlers to handle program termination. |
| Ignore | Discards the signal and continues the process. The process will not know that something had raised a signal. |

NOTE Do not use Ignore mode for fatal signals such as **SIGSEGV** and **SIGBUS**. If you do, TotalView can get caught in a signal/resignal loop with your program; the signal will immediately recur because the failing instruction repeatedly reexecutes.

Setting Search Paths

If your source code, executable, and object files reside in different directories, set search paths for these directories with the **File > Search Path** command. You do not need to use this command if these directories are already named in your environment's **PATH** variable.

These search paths apply to *all* processes that you're debugging. (See Figure 44 on page 60.)

TotalView searches the following directories (in order):

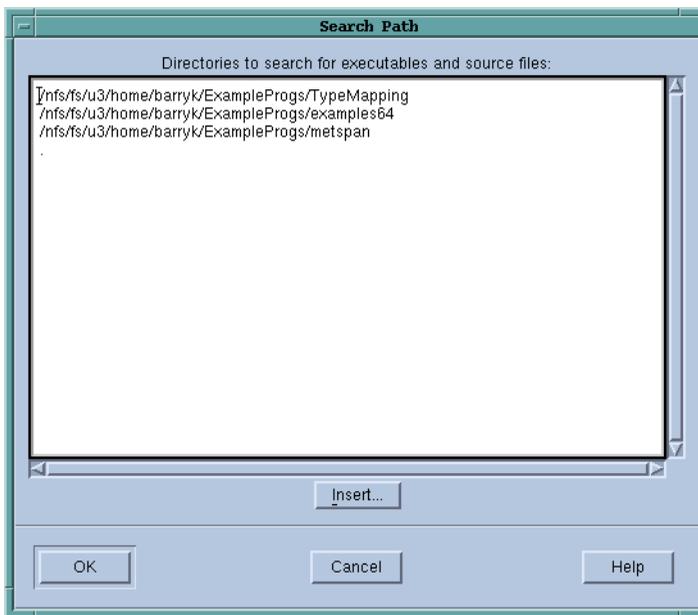


FIGURE 44: **File > Search Path Dialog Box**

- 1 The current working directory (.).
- 2 The directories you specify by using the **File > Search Path** command in the exact order you enter them.
- 3 If you entered a full path name for the executable when you started TotalView, TotalView searches this directory.
- 4 If your executable is a symbolic link, TotalView will look in the directory in which your executable actually resides for the new file.

As you can have multiple levels of symbolic links, TotalView keeps on following links until it finds the actual file. After it has found the executable, it will look in the executable's directory for your file. If it isn't there, it'll back up the chain of links until either it finds the file or determines that the file can't be found.

- 5 The directories specified in your **PATH** environment variable.

When entering directories into this dialog box, you must enter them in the order you want them searched, and you must enter each on its own line.

- You can type path names directly.
- You can cut and paste directory information.

- You can use the **Insert** button to tell TotalView to display the **Select Directory dialog box** that lets you browse through the file system, interactively selecting directories. (See Figure 45 on page 61.)

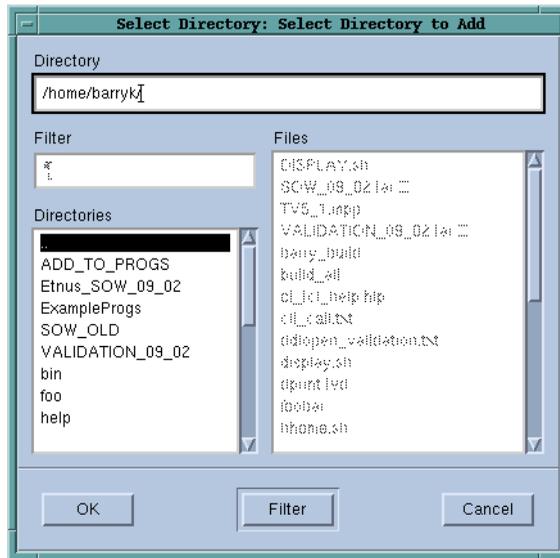


FIGURE 45: **Select Directory Dialog Box**

The current working directory (.) in the **File > Search Path** Dialog Box is the first directory listed in the window. TotalView interprets relative path names as being *relative* to the current working directory.

If you remove the current working directory, TotalView reinserts it at the top of the list.

After you change this list of directories, TotalView again searches for the source file of the routine being displayed in the Process Window.

You can also specify search directories using the `TV::search_path` variable.

Setting Command Arguments

When TotalView creates a process, it uses the name of the file containing the executable code for the process's program name. If your program requires command-

line arguments and you hadn't entered them using TotalView's `-a` command-line option, here's how you can set these arguments *before* you start the process:

- 1 Select the **Arguments** Tab within the **Process > Startup Parameters** Dialog Box. (See Figure 46 on page 62.)

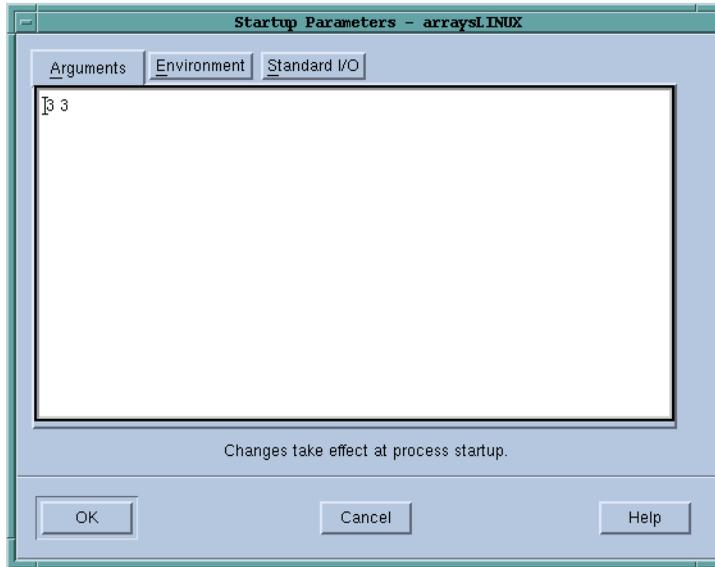


FIGURE 46: **Process > Startup Parameters** Dialog Box: **Arguments** Page

- 2 Type the arguments you want TotalView to pass to your program. Either separate each argument with a space or place each one on a separate line. If an argument contains spaces, enclose the entire argument in double quotes. When you're done, select **OK**.

CLI EQUIVALENT: `dset ARGS_DEFAULT {value}`

Setting Input and Output Files

Before your program begins executing, TotalView defines how it will manage standard input (**stdin**) and standard output (**stdout**). Unless you tell it otherwise, **stdin** and **stdout** use the shell window from which you invoked TotalView.

The **Process > Startup** command lets you redirect **stdin** or **stdout**. You can only do this before your program begins executing. Here's how:

- 1 Select the **Standard I/O** Tab from the dialog box displayed when you invoke the **Process > Startup Parameters** command. (See Figure 47 on page 63.)

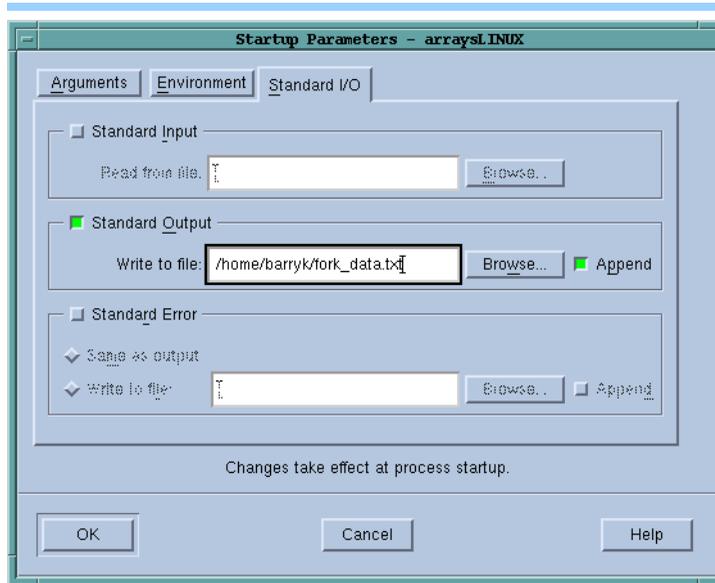


FIGURE 47: **Process > Startup Parameters Dialog Box: Standard I/O Page**

- 2 Type the name of the file, relative to your current working directory. Entering names in these text boxes is equivalent to using `<`, `>`, or `>&` symbols in most shells.
- 3 Select **OK**.

If you select the **Append** check box, TotalView appends new information to the end of the file if the file already exists. If it isn't checked, TotalView overwrites the file's contents.

If you select the **Same as output** check box, TotalView writes **stderr** information to the same place indicated in the **Standard Output** field.

CLI EQUIVALENT: **`drun` and `drerun` have arguments that let you reset `stdin`, `stdout`, and `stderr`.**

Setting Preferences

The **File > Preferences** command lets you tailor many of TotalView's behaviors. This section contains an overview of these preferences. Detailed explanations are in the online Help.

Some settings such as the prefixes and suffixes looked at when loading dynamic libraries can be different from operating system to operating system. If these settings can differ, TotalView will let you set values for each operating system. This is done transparently, which means that you only see an operating system's values when you are running TotalView on that operating system. In general, this applies to the server launch strings and dynamic library paths.

NOTE Every preference has a variable that can be set using the CLI. These variables are described in the TotalView Reference Guide.

- **Options.** This page contains check boxes that are either general in nature or that influence different parts of the system. (See Figure 48.)

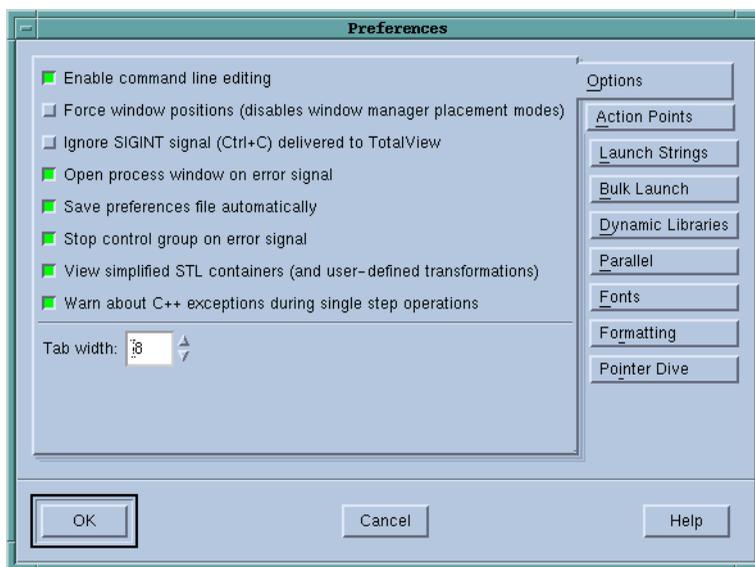


FIGURE 48: **File > Preferences** Dialog Box: Options Page

- **Action Points.** The commands on this page indicate if TotalView should stop anything else when it encounters an action point, the scope of the action point, au-

automatic saving and loading of action points, and if TotalView should open a Process Window for the process encountering a breakpoint. (See Figure 49.)

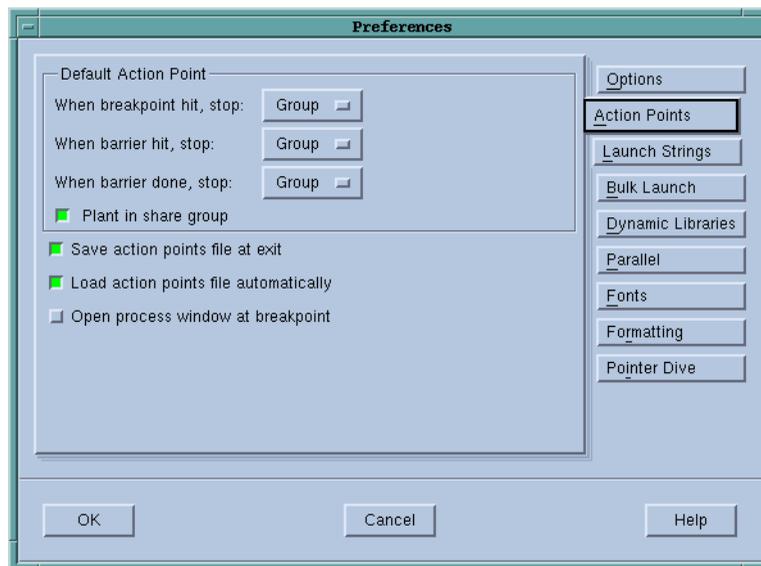


FIGURE 49: **File > Preferences Dialog Box: Action Points Page**

- **Launch Strings.** The three areas of this page let you set the launch string that TotalView uses when it launches the `tvdsrv` remote debugging server, the Visualizer, and a source code editor. Notice that default values exist for these launch strings. (See Figure 50 on page 66.)
- **Bulk Launch.** The fields and commands on this page configure TotalView's bulk launch system. See Chapter 4 for more information. (See Figure 51 on page 66.)
- **Dynamic Libraries.** This page lets you control which symbols are added to TotalView when it loads a dynamic library. (See Figure 52 on page 67.)
- **Parallel.** This page lets you define what will occur when your program goes parallel. (See Figure 53 on page 67.)
- **Fonts.** Use this page to specify the fonts used in the user interface and when TotalView displays your code. (See Figure 54 on page 68.)
- **Formatting.** Use this page to control how your program's variables are displayed. (See Figure 55 on page 68.)
- **Pointer Dive.** Use this page to control how pointers are dereferenced and how pointers to arrays are cast. (See Figure 56 on page 69.)

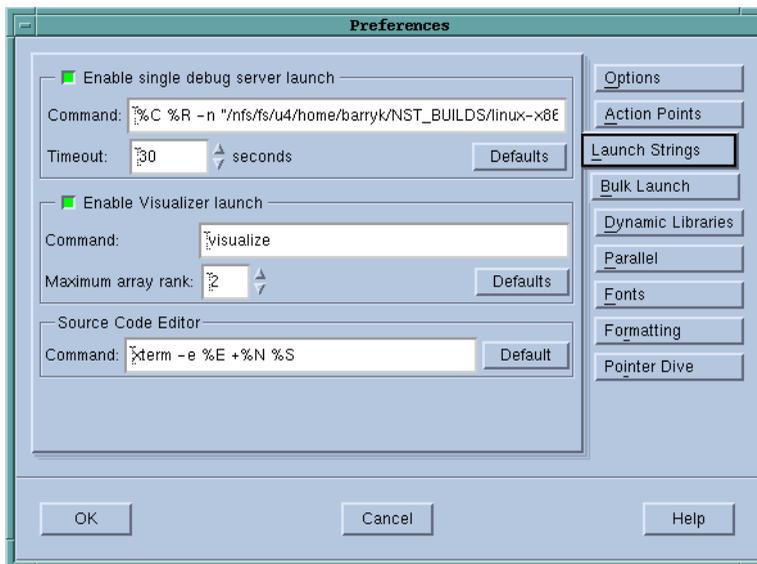


FIGURE 50: File > Preferences Dialog Box: Launch Strings Page

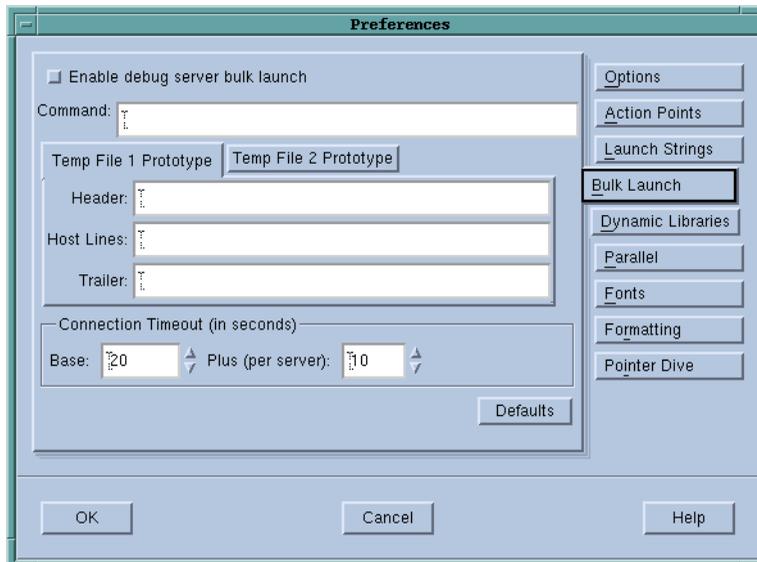


FIGURE 51: File > Preferences Dialog Box: Bulk Launch Page

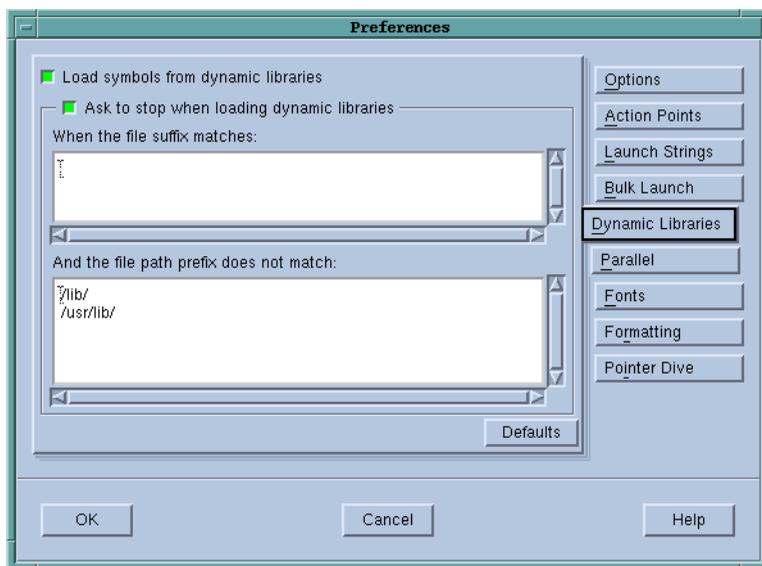


FIGURE 52: File > Preferences Dialog Box: Dynamic Libraries Page

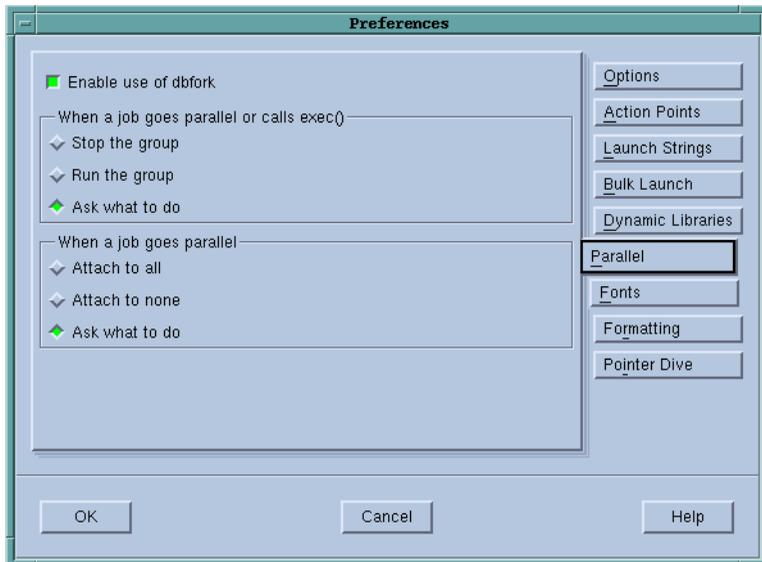


FIGURE 53: File > Preferences Dialog Box: Parallel Page

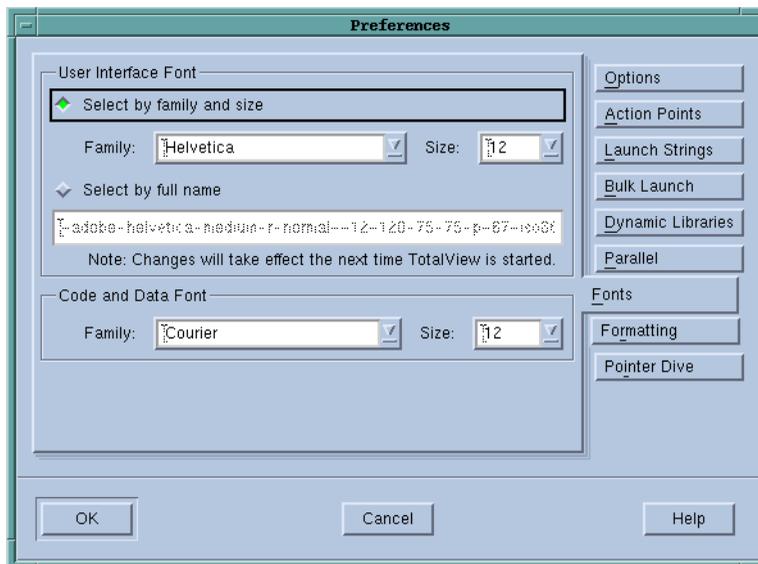


FIGURE 54: File > Preferences Dialog Box: Fonts Page

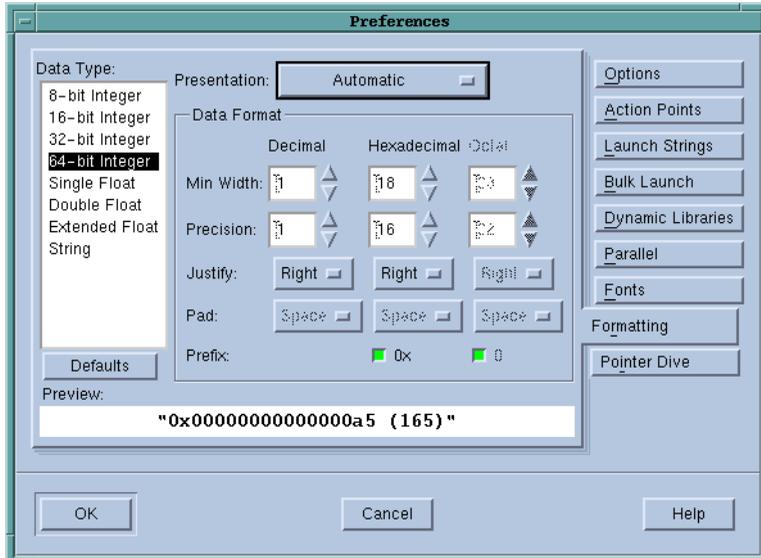


FIGURE 55: File > Preferences Dialog Box: Formatting Page

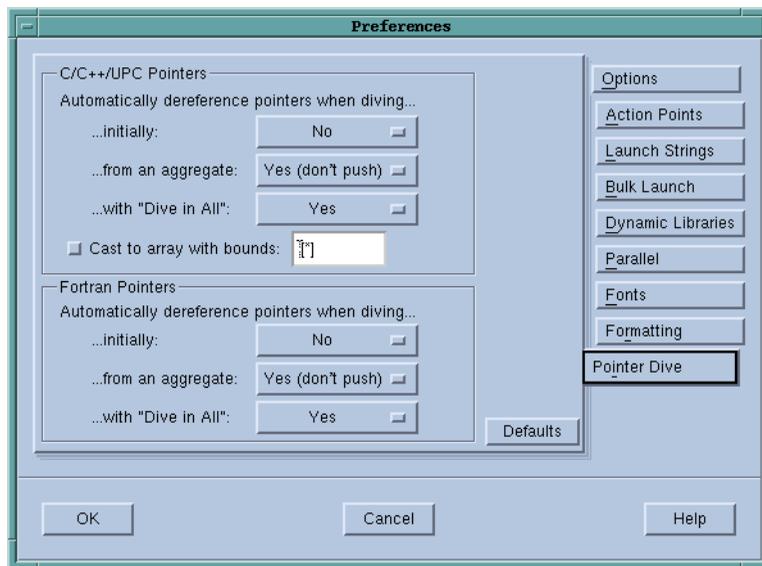


FIGURE 56: **File > Preferences Dialog Box: Pointer Dive Page**

Setting Preferences, Options, and X Resources

While preferences are the best way to set many of TotalView's features and characteristics, TotalView also lets you use variables and command-line options to set features and characteristics.

Older versions of TotalView did not have a preference system. Instead, you needed to set values in your `.Xdefaults` file or in a command-line option. For example, setting `totalview*autoLoadBreakpoints` to true tells TotalView that it should automatically load an executable's breakpoint file when it loads an executable. Because you can also set this option as a preference and set it using the CLI's `dset` command, this X resource has been *deprecated*.

NOTE "Deprecated" means that the feature is still available. While the feature may exist for a while, there's no guarantee that it will continue to work. All "totalview" options have been deprecated. Those used for setting the Visualizer are still supported.

Similarly, documentation for earlier releases told you how to use a command-line option to tell TotalView to automatically load breakpoints, and there were two dif-

ferent command-line options to perform this action. While these methods still work, they are also deprecated.

Visualizer X resources are still supported. Documentation for them can be found at www.etnus.com/Support/docs/xresources/XResources.html.

In some cases, you may want to set a state for one session or you may want to override one of your preferences. (A preference indicates a behavior that you want to occur in all of your TotalView sessions.) This is the function of the command-line options described in “*TotalView Command Syntax*” in the TOTALVIEW REFERENCE GUIDE.

For example, you can use `-bg` to set the background color for TotalView windows in the TotalView session just being invoked. TotalView does not remember changes to its default behavior that you make using command-line options. You have to set them again when you start a new session.

Setting Environment Variables



You can set and edit the environment variables that TotalView passes to processes. When TotalView creates a new process, it passes a list of environment variables to the process. You can add to this list by using the **Environment** Page in the **Process > Startup Parameters** Dialog Box.

NOTE TotalView does not display the variables that it was passed. Instead, it just displays the variables you have added to the environment using this command.

The format for specifying an environment variable is *name=value*. For example, the following definition creates an environment variable named **DISPLAY** whose value is **enterprise:0.0**:

```
DISPLAY=enterprise:0.0
```

To add, delete, or modify environment variables that you enter, select the **Environment** Tab from the dialog box displayed when you invoke the **Process > Startup Parameters** command. See Figure 57.

When entering environment variables, place each on a separate line. The actions you can now perform are:

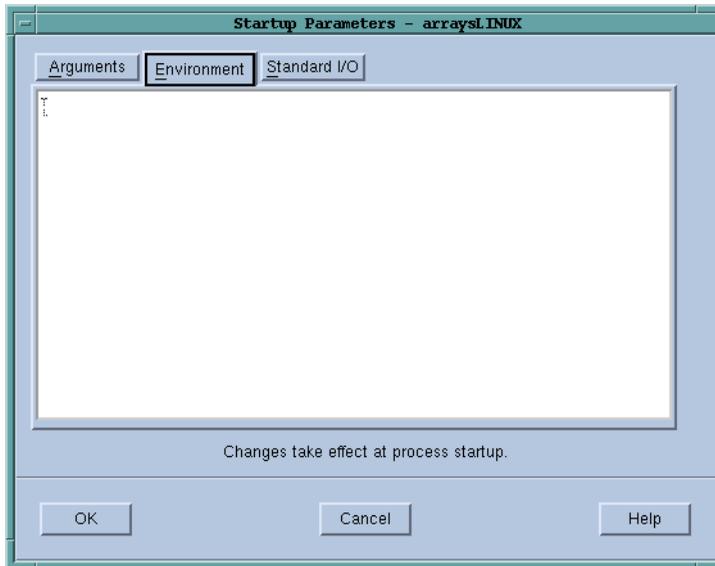


FIGURE 57: **Process > Startup Parameters Dialog Box: Environment Page**

- Changing the name or value of an environment variable by editing a line.
- Adding a new environment variable by inserting a new line and specifying a name and value.
- Deleting an environment variable that you added by deleting a line. If you delete all lines, the process again inherits the environment used by TotalView or **tvdsvr**.

Monitoring TotalView Sessions

TotalView logs all significant events occurring for all processes being debugged. To view the event log, select the Root Window's Log Tab. This page displays a sequential list of these events. See Figure 58 on page 72 for an example.

You can set the amount of information TotalView writes to this window by using the CLI's **dset** command to set the **VERBOSE** variable. If you always want it set to a value, you can set it in your **.tvdr** file. For example:

```
dset VERBOSE WARNING
```

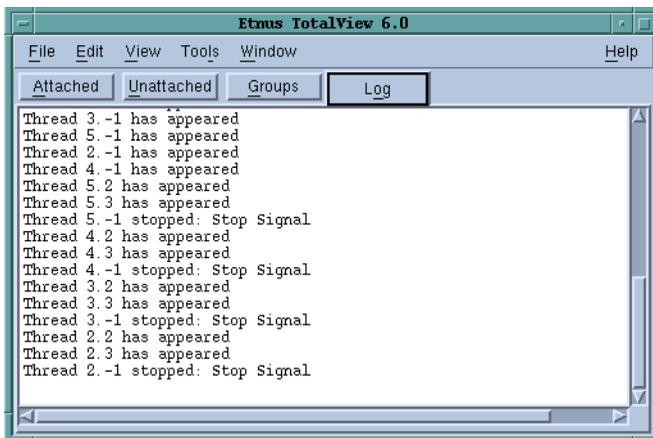


FIGURE 58: Root Window Log Page

Chapter 4

Setting Up Remote Debugging Sessions

This chapter explains how to set up TotalView remote debugging sessions. This chapter discusses:

- “Starting the TotalView Debugger Server” on page 73
- “Debugging Over a Serial Line” on page 86

Starting the TotalView Debugger Server

Debugging a remote process with TotalView is only slightly different than debugging a native process in that:

- TotalView needs to work with a TotalView processes that will be running on remote machines. This remote TotalView process, which TotalView usually automatically launches, is called the TotalView Debugger Server (**tvdsvr**).
- TotalView’s performance depends on your network’s performance. If the network is overloaded, debugging can be slow.

Unless you tell it otherwise, TotalView automatically launches **tvdsvr** in one of the following ways:

- It can independently launch a **tvdsvr** on each remote host. This is called *single-process server launch*.
- It can launch all remote processes at the same time. This is called *bulk server launch*.

Because TotalView can automatically launch **tvdsvr**, there’s nothing you need to do when you’re debugging remote processes. It shouldn’t matter to you if a process is local or remote.

NOTE If the default single-process server launch procedure meets your needs and you're not experiencing any problems accessing remote processes from within TotalView, you can safely ignore the information in this chapter. If you do experience a problem launching the server, you should check that the `tvdsvr` process is in your path.

Topics in this section are:

- "Setting Single-Process Server Launch Options" on page 74
- "Setting Bulk Launch Window Options" on page 76
- "Starting the Debugger Server Manually" on page 79
- "Using the Single-Process Server Launch Command" on page 80
- "Bulk Server Launch on an SGI MIPS Machine" on page 81
- "Bulk Server Launch on an HP Alpha Machine" on page 83
- "Bulk Server Launch on an IBM RS/6000 AIX Machine" on page 83
- "Disabling Autolaunch" on page 84
- "Changing the Remote Shell Command" on page 84
- "Changing the Arguments" on page 85
- "Autolaunch Sequence" on page 86

Setting Single-Process Server Launch Options

The **Enable single debug server launch** preferences in the **Launch Strings** Page of the **File > Preferences** Dialog Box lets you disable autolaunch, change the command that TotalView uses when it launches remote servers, and alters the amount of time TotalView waits when establishing connections to a `tvdsvr` process. (See Figure 59 on page 75.)

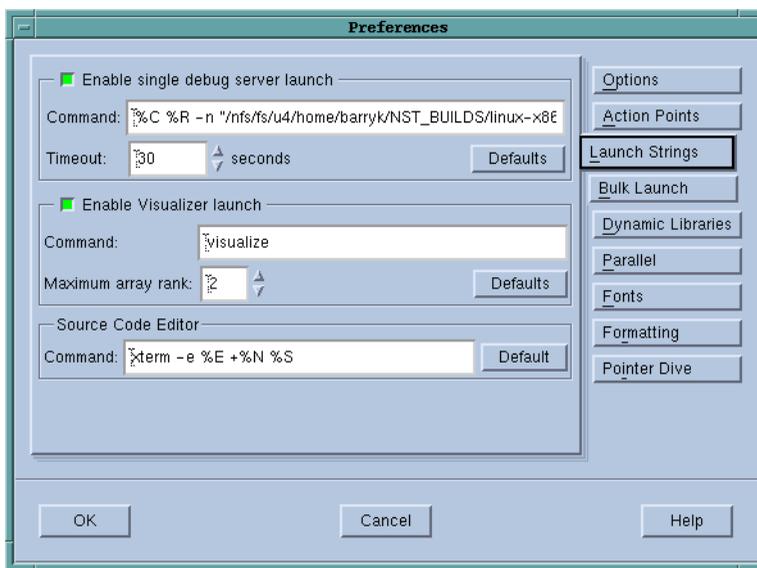


FIGURE 59: **File > Preferences: Server Launch Strings Page**

Enable single debug server launch

If you select this check box, TotalView will independently launch the TotalView Debugger Server (**tvdsvr**) on each remote system.

CLI EQUIVALENT: **dset TV::server_launch_enabled**

Note Even if you have enabled bulk server launch, you probably also want this option to be enabled. TotalView uses this launch string when you start TotalView upon a file when you have named a host within the File > New Program dialog box or have used the `–remote` command line option. You only want to disable single server launch when it can't work.

Command

Enter the command that TotalView will use when it independently launches **tvdsvr**. For information on this command and its options, see "Using the Single-Process Server Launch Command" on page 80.

CLI EQUIVALENT: **dset TV::server_launch_string**

Timeout

After TotalView automatically launches **tvdsvr**, it waits 30 seconds for **tvdsvr** to respond. If the connection isn't made in this time, TotalView times out. You can change the amount of time by entering a value from 1 to 3600 seconds (1 hour).

CLI EQUIVALENT: **dset TV::server_launch_timeout**

You can also preset the timeout value by using a TotalView preference. See the online Help for more information.

If you notice that TotalView fails to launch **tvdsvr** (as shown in the **xterm** window from which you started the debugger) before the timeout expires, select **Yes** in the **Question** Dialog Box that will appear.

Defaults

If you make a mistake or decide you want to go back to TotalView's default settings, select the **Defaults** button.

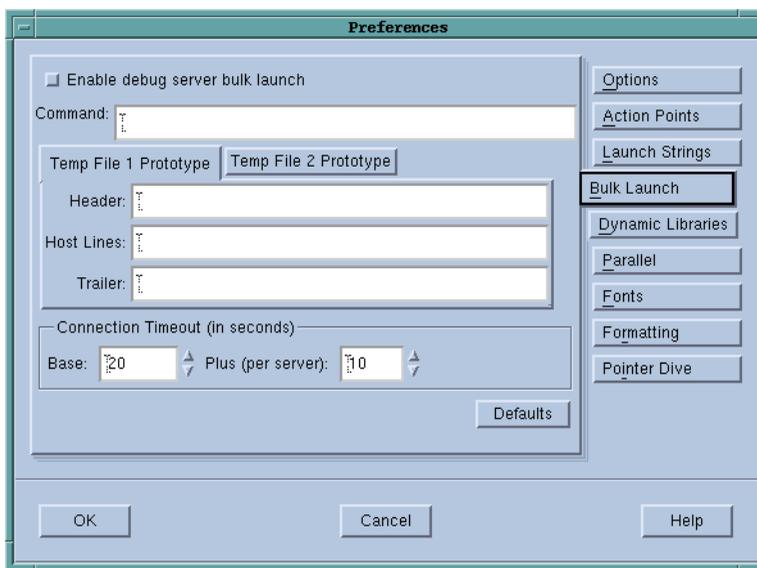
Selecting **Defaults** also throws away any changes you made using a CLI variable. TotalView doesn't immediately change settings after you click the **Defaults** button; instead, it waits until you select the **OK** button.

Setting Bulk Launch Window Options

The fields in the **File > Preferences' Bulk Launch** Page lets you change the bulk launch command, disable bulk launch, and alter connection timeouts that TotalView uses when it launches **tvdsvr** programs. Figure 60 shows this page.

Enable debug server bulk launch

If you select this check box, TotalView uses its bulk launch procedure when launching the TotalView Debugger Server (**tvdsvr**).

FIGURE 60: **File > Preferences: Bulk Launch Page**

By default, bulk launch is disabled; that is, TotalView uses its single-server launch procedure.

CLI EQUIVALENT: **dset TV::bulk_launch_enabled**

Command

If you have enabled bulk launch, TotalView will use this command to launch **tvdsvr**. For information on this command and its options, see “*Bulk Server Launch on an SGI MIPs Machine*” on page 81 and “*Bulk Server Launch on an IBM RS/6000 AIX Machine*” on page 83.

CLI EQUIVALENT: **dset TV::bulk_launch_string**

Temp File 1 Prototype

Temp File 2 Prototype

Both tab pages have three fields. These fields let you specify the contents of temporary files that the bulk launch operation will use. For information on these fields, see “*TotalView Debugger*”

Server (*tvdsvr*) *Command Syntax*" in the TOTALVIEW REFERENCE GUIDE.

```
CLI EQUIVALENT:  dset TV::bulk_launch_tmpfile1_header_line
                  dset TV::bulk_launch_tmpfile1_host_lines
                  dset TV::bulk_launch_tmpfile1_trailer_line
                  dset TV::bulk_launch_tmpfile2_header_line
                  dset TV::bulk_launch_tmpfile2_host_lines
                  dset TV::bulk_launch_tmpefile2_trailer_line
```

Connection Timeout (in seconds)

After TotalView launches **tvdsvr** processes, it waits 20 seconds (the **Base** time) plus 10 seconds for each server that it will launch for responses from successfully connected processes. If connections are not made in this time, TotalView times out.

The **Base** timeout value can be from 1 to 3600 seconds (1 hour). The incremental **Plus** value is from 1 to 360 seconds (6 minutes). See the online Help for information on presetting these values.

```
CLI EQUIVALENT:  dset TV::bulk_launch_base_timeout
                  dset TV::bulk_incr_timeout
```

If you notice that TotalView fails to launch **tvdsvr** (as shown in the **xterm** window from which you started the debugger) before the timeout expires, select **Yes** in the **Question** Dialog Box that will appear.

Defaults

If you make a mistake or decide you want to go back to TotalView's default settings, select the **Defaults** button.

Selecting **Defaults** also throws away any changes you made using a CLI variable. TotalView doesn't immediately change settings after you click the **Defaults** button; instead, it waits until you select the **OK** button.

Starting the Debugger Server Manually

If TotalView can't automatically launch `tvdsvr`, you can start it manually. Unfortunately, this method isn't completely secure: other users can connect to your instance of `tvdsvr` and begin using your UNIX UID.

NOTE If you specify `hostname:portnumber` when opening a remote process, TotalView will not launch a debugger server.

Here is how you manually start `tvdsvr`:

- 1 Begin by insuring that both the bulk launch and single server launch or disabled. To disable the bulk launch, select the **Bulk Launch** Tab within the **File > Preferences** Dialog Box. (You can select this command from the Root Window or the Process Window.) The dialog box shown in Figure 60 on page 77 appears. Next, clear the **Enable debug server bulk launch** check box within the **Bulk Launch** Tab to disable the autolaunch feature and then select **OK**.

CLI EQUIVALENT: `dset TV::bulk_launch_enabled`

Similarly, select the **Server Launch** Tab and clear the **Enable single debug server launch** button.

CLI EQUIVALENT: `dset TV::server_launch_enabled`

- 2 Log in to the remote machine and start `tvdsvr`:

`tvdsvr -server`

If you don't (or can't) use the default port number (4142), you will need to use the `-port` or `-search_port` options. For details, refer to "TotalView Debugger Server (`tvdsvr`) Command Syntax" in the TOTALVIEW REFERENCE GUIDE.

After printing out the port number and the assigned password, the server begins listening for connections. Be sure to remember the password; you'll need to enter it in step 4.

NOTE Because the `-server` option is not secure, it must be explicitly enabled. (This is usually done by your system administrator.) For details, see "`-server`" in the "TotalView Command Syntax" chapter of the *TotalView Reference Guide*.

- 3 From the Root Window, select the **File > New Program** command. Type the program's name in the **Executable** field and the *hostname:portnumber* in the **Remote Host** field and then select **OK**.

CLI EQUIVALENT: **dload executable -r hostname**

- 4 TotalView now tries to connect to **tvdsvr**.

When TotalView prompts you for the password, enter the password that **tvdsvr** displayed in step 2.

Figure 61 on page 80 summarizes the steps used when you start **tvdsvr** manually.

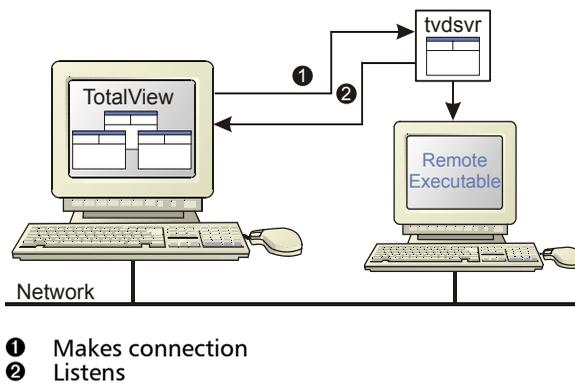


FIGURE 61: Manual Launching of Debugger Server

Using the Single-Process Server Launch Command

Here is the default command string that TotalView uses when it automatically launches the debugger server for a single process:

```
%C %R -n "tvdsvr -working_directory %D -callback %L \  
-set_pw %P -verbosity %V"
```

where:

%C Expands to the name of the server launch command being used. On most platforms, this is **rsh**. On HP machines, this command is **remsh**. If the **TVDSVRLAUNCHCMD** environment variable exists, TotalView uses its value instead of its platform-specific default value.

- %R** Expands to the host name of the remote machine that you specified in the **File > New Program** or **dload** commands.
- n** Tells the remote shell to read standard input from **/dev/null**; that is, the process will immediately receive an EOF (End-Of-File) signal.
- working_directory %D** Makes **%D** the directory to which TotalView will be connected. **%D** expands to the absolute path name of the directory.
- Using this option assumes that the host machine and the target machine are mounting identical file systems. That is, the path name of the directory to which TotalView is connected must be identical on host and target machines.
- After changing to this directory, the shell will invoke the **tvdsvr** command.
- You must make sure the **tvdsvr** directory is in your path on the remote machine.
- callback %L** Establishes a connection from **tvdsvr** to TotalView. **%L** expands to the host name and TCP/IP port number (*hostname:port*) upon which TotalView is listening for connections from **tvdsvr**.
- set_pw %P** Sets a 64-bit password. TotalView must supply this password when **tvdsvr** establishes a connection with it. **%P** expands to the password that TotalView automatically generates. For more information on this password, see “*TotalView Debugger Server (tvdsvr) Command Syntax*” in the TOTALVIEW REFERENCE GUIDE.
- verbosity %V** Sets the verbosity level of the TotalView Debugger Server. **%V** expands to the current TotalView verbosity setting.

You can also use the **%H** option with this command. This option is discussed in “*Bulk Server Launch on an SGI MIPs Machine*” on page 81.

For information on the complete syntax of the **tvdsvr** command, refer to “*TotalView Debugger Server (tvdsvr) Command Syntax*” in the TOTALVIEW REFERENCE GUIDE.

Bulk Server Launch on an SGI MIPs Machine

On an SGI machine, the bulk server launch string is similar to the single-process server launch and is:

```
array tvdsvr -working_directory %D -callback_host %H \
             -callback_ports %L -set_pws %P -verbosity %V
```

where:

-working_directory %D

Makes %D the directory to which TotalView will be connected. %D expands to this directory's absolute path name.

TotalView assumes that the host machine and the target machine mount identical file systems. That is, the path name of the directory to which TotalView is connected must be identical on both host and target machines.

After performing this operation, **tvdsvr** starts executing.

-callback_host %H

Names the host upon which TotalView makes this callback. %H expands to the host name of the machine TotalView is running on.

-callback_ports %L

Names the ports on the host machines that TotalView uses for callbacks. %L expands to a comma-separated list of host names and TCP/IP port numbers (*hostname:port,hostname:port...*). TotalView on which TotalView is listening for connections on a port from the instance of **tvdsvr** on the host.

-set_pws %P

Sets 64-bit passwords. TotalView must supply these passwords when **tvdsvr** establishes the connection with it. %P expands to a comma-separated list of 64-bit passwords that TotalView automatically generates. For more information, see "*TotalView Debugger Server (tvdsvr) Command Syntax*" in the TOTALVIEW REFERENCE GUIDE.

-verbosity %V

Sets **tvdsvr**'s verbosity level. %V expands to the current TotalView verbosity setting.

You must enable **tvdsvr**'s use of the **array** command by adding the following information to the `/usr/lib/array/arrayd.conf` file:

```
#
# Command that allow invocation of the TotalView Debugger
# server when performing a Bulk Server Launch.
#
```

```

command tvdsvr
  invoke /opt/totalview/bin/tvdsvr %ALLARGS
  user %USER
  group %GROUP
  project %PROJECT

```

This assumes that the location of **tvdsvr** is **/opt/totalview/bin**. For information on the syntax of the **tvdsvr** command, refer to “*TotalView Debugger Server (tvdsvr) Command Syntax*” in the TOTALVIEW REFERENCE GUIDE.

Bulk Server Launch on an IBM RS/6000 AIX Machine

On an IBM RS/6000 AIX machine, the bulk server launch string is:

```
%C %H -n "poe -pgmmodel mpmd -resd no -tasks_per_node 1
          -procs %N -hostfile %t1 -cmdfile %t2"
```

where the options unique to this command are:

%N	The number of servers that TotalView will launch.
%t1	A temporary file created by TotalView that contains a list of the hosts tvdsvr will run on. This is the information you enter in the Temp File 1 Prototype field in the Bulk Launch Page . TotalView generates this information by expanding the %R symbol. This is the information you enter in the Temp File 2 Prototype field in the Bulk Launch Page .
%t2	A file that contains the commands to start the tvdsvr processes on each machine. TotalView creates these lines by expanding the following template: <pre>tvdsvr -working_directory %D \ -callback %L -set_pw %P -verbosity %V</pre>

Information on the options and expansion symbols is in the “*TotalView Debugger Server (tvdsvr) Syntax*” chapter of the TOTALVIEW REFERENCE GUIDE.

Bulk Server Launch on an HP Alpha Machine

On an HP Alpha machine, the bulk server launch string is:

```
prun -T -1 tvdsvr -callback_host %H -callback_ports %L
      -set_pws %P -verbosity %V -working_directory %D
```

Information on the options and expansion symbols is in the “*TotalView Debugger Server (tvdsvr) Syntax*” chapter of the TOTALVIEW REFERENCE GUIDE.

Disabling Autolaunch

If after changing the autolaunch options, TotalView still can't automatically start **tvdsvr**, you must disable autolaunching and start **tvdsvr** manually. Here are two ways to do this:

- Clear the **Enable single debug server launch** check box in the **Launch Strings** Page of the **File > Preferences** Dialog Box.

CLI EQUIVALENT: `dset TV::server_launch_enabled`

- When you debug the remote process, as described in “*Starting the TotalView Debugger Server*” on page 73, enter a host name and port number in the bottom section of the **File > New Program** Dialog Box. This disables autolaunching for the current connection.

NOTE If you disable autolaunching, you must start **tvdsvr** before you load a remote executable or attach to a remote process.

Changing the Remote Shell Command

Some environments require that you create your own autolaunch command. You might do this, for example, if your remote shell command doesn't provide the security that your site requires.

If you create your own autolaunch command, you must use the **tvdsvr** command's **-callback** and **-set_pw** arguments.

If you're not sure whether **rsh** (or **remsh** on HP machines) works at your site, try typing “**rsh** *hostname*” (or “**remsh** *hostname*”) from an **xterm** window, where *hostname* is the name of the host upon which you want to invoke the remote process. If this command prompts you for a password, you must add the host name of the host machine to your **.rhosts** file on the target machine.

For example, you could use the following combination of the **echo** and **telnet** commands:

```
echo %D %L %P %V; telnet %R
```

Once **telnet** establishes a connection to the remote host, you could use the **cd** and **tvdsrv** commands directly, using the values of **%D**, **%L**, **%P**, and **%V** that were displayed by the **echo** command. For example:

```
cd directory
tvdsrv -callback hostname:portnumber -set_pw password
```

If your machine doesn't have a command for invoking a remote process, TotalView can't autolaunch the **tvdsrv** and you must disable both single server and bulk server launches.

For information on the **rsh** and **remsh** commands, refer to the manual page supplied with your operating system.

Changing the Arguments

You can also change the command-line arguments passed to **rsh** (or whatever command you use to invoke the remote process).

For example, if the host machine doesn't mount the same file systems as your target machine, the debugger server may need to use a different path to access the executable being debugged. If this is the case, you could change **%D** to the directory used on the target machine.

If the remote executable reads from standard input, you cannot use the **-n** option with your remote shell command because the remote executable will receive an EOF immediately on standard input. If you omit **-n**, the remote executable reads standard input from the **xterm** in which you started TotalView. This means that you should invoke **tvdsrv** from another **xterm** window if your remote program reads from standard input. Here's an example:

```
%C %R "xterm -display hostname:0 -e tvdsrv -callback %L \
      -working_directory %D -set_pw %P -verbosity %V"
```

Now, each time TotalView launches **tvdsrv**, a new **xterm** appears on your screen to handle standard input and output for the remote program.

Autolaunch Sequence

If you want to know more about autolaunch, here is the sequence of actions carried out by you, TotalView, and **tvdsvr**:

- 1 With the **File > New Program** or **dload** commands, you specify the host name of the machine on which you want to debug a remote process, as described in "Starting the TotalView Debugger Server" on page 73.
- 2 TotalView begins listening for incoming connections.
- 3 TotalView launches the **tvdsvr** process with the server launch command. ("Using the Single-Process Server Launch Command" on page 80 describes this command.)
- 4 The **tvdsvr** process starts on the remote machine.
- 5 The **tvdsvr** process establishes a connection with TotalView.

Figure 62 on page 86 summarizes these actions.

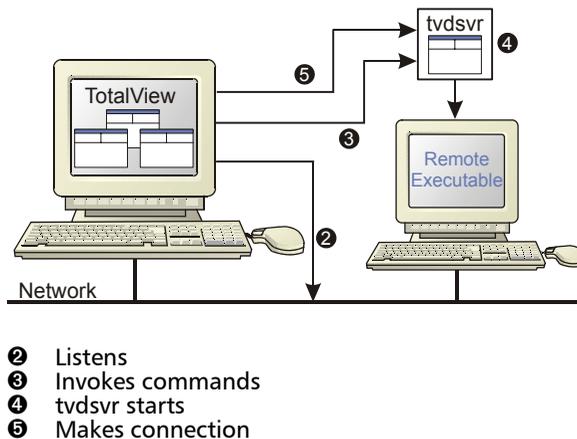


FIGURE 62: Root Window Showing Process and Thread Status

Debugging Over a Serial Line

TotalView allows you to debug programs over a serial line as well as TCP/IP sockets. However, if a network connection exists, you will probably want to use it because performance will be much better.

You will need to have two connections to the target machine: one for the console and the other for TotalView. Do not try to use one serial line as TotalView cannot share a serial line with the console.

Figure 63 on page 87 illustrates a TotalView debugging session using a serial line. In this example, TotalView is communicating over a dedicated serial line with a TotalView Debugger Server running on the target host. A VT100 terminal is connected to the target host's console line, allowing you to type commands on the target host.

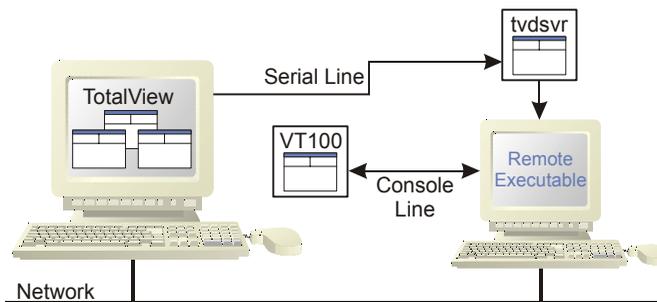


FIGURE 63: TotalView Debugging Session Over a Serial Line

Topics in this section are:

- "Starting the TotalView Debugger Server" on page 87
- "Starting TotalView on a Serial Line" on page 88
- "Using the New Program Window" on page 88

Starting the TotalView Debugger Server

To start a TotalView debugging session over a serial line from the command line, you must first start the TotalView Debugger Server (**tvdsrv**).

Using the console connected to the target machine, start **tvdsrv** and enter the name of the serial port device on the target machine. Here is the syntax of the command you would use:

```
tvdsrv -serial device[:baud=num]
```

where:

<i>device</i>	The name of the serial line device.
<i>num</i>	The serial line's baud rate; if you omit the baud rate, TotalView uses a default value of 38400 .

For example:

```
tvdsvr -serial /dev/com1:baud=38400
```

After it starts, **tvdsvr** waits for TotalView to establish a connection.

Starting TotalView on a Serial Line

Start TotalView on the host machine and include the name of the serial line device.

The syntax of this command is:

```
totalview -serial device[:baud=num] filename
```

or

```
totalviewcli -serial device[:baud=num] filename
```

where:

<i>device</i>	The name of the serial line device on the host machine.
<i>num</i>	The serial line's baud rate. If you omit the baud rate, TotalView uses a default value of 38400 .
<i>filename</i>	The name of the executable file.

For example:

```
totalview -serial /dev/term/a test_pthreads
```

Using the New Program Window



Here is the procedure for starting a TotalView debugging session over a serial line when you're already in TotalView:

- 1 Start the TotalView Debugger Server. (This is discussed in "Starting the TotalView Debugger Server" on page 87).
- 2 Select the **File > New Program** command. TotalView responds by displaying the dialog box shown in Figure 64.

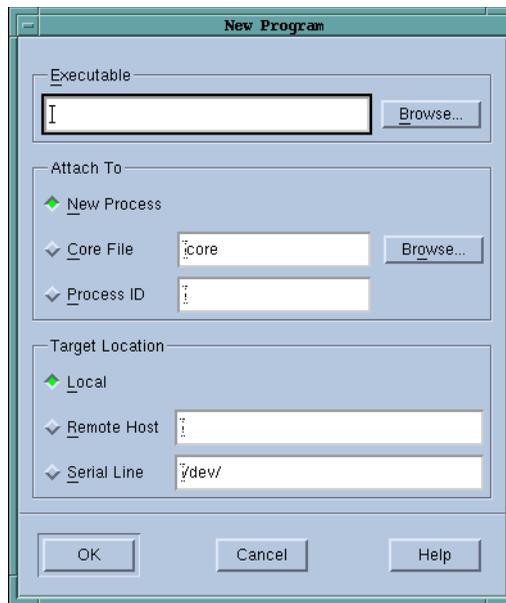


FIGURE 64: **File > New Program Dialog Box**

Type the name of the executable file in the **Executable** field.

Type the name of the serial line device in the **Serial Line** field.

- 3 Select **OK**.

Chapter 5

Setting Up Parallel Debugging Sessions

This chapter explains how to set up TotalView parallel debugging sessions for applications that use the following parallel execution models.

The information in this chapter describes running many different environments on many different architectures. While there is a lot of information in this chapter you do need, you probably don't need the information on many of the environments and architectures. This means that you shouldn't just read this chapter. Instead, go to this book's table of contents and decide what's important to you.

This chapter discusses:

- *"Debugging MPICH Applications"* on page 92
- *"Debugging HP Alpha MPI Applications"* on page 96
- *"Debugging HP MPI Applications"* on page 97
- *"Debugging IBM MPI (PE) Applications"* on page 99
- *"Debugging QSW RMS Applications"* on page 103
- *"Debugging SGI MPI Applications"* on page 104
- *"Debugging Sun MPI Applications"* on page 105
- *"Debugging OpenMP Applications"* on page 113
- *"Debugging PVM and DPVM Applications"* on page 121
- *"Debugging Shared Memory (SHMEM) Code"* on page 128
- *"Debugging UPC Programs"* on page 129
- *"Parallel Debugging Tips"* on page 134

There are a few things that are of general interest:

- TotalView lets you decide which process you want it to attach. You will find information in *"Attaching to Processes"* on page 134.
- If you're using a messaging system, TotalView displays this information visually as a message queue graph and textually in a message queue window.

- The end of this chapter has some hints on how you can approach debugging parallel programs.

Debugging MPICH Applications

To debug Message Passing Interface/Chameleon Standard (MPICH) applications, you must use MPICH version 1.2.3 or later on a homogenous collection of machines. If you need a copy of MPICH, you can obtain it at no cost from Argonne National Laboratory at www.mcs.anl.gov/mpi. (You are strongly urged to use a later version of MPICH. Information on versions that work with TotalView can be found in the TOTALVIEW PLATFORMS document.)

The MPICH library should use the `ch_p4`, `ch_p4mpd`, `ch_shmem`, `ch_lfshmem`, or `ch_mpl` devices. For networks of workstations, `ch_p4` is the default. For shared-memory SMP machines, use `ch_shmem`. On an IBM SP machine, use the `ch_mpl` device. The MPICH source distribution includes all of these devices and you can choose one when you configure and build MPICH.

NOTE When configuring MPICH, you must ensure that the MPICH library maintains all of the information required by TotalView, which means that you must use the `--enable-debug` option with the MPICH `configure` command. (Versions earlier than 1.2 used the `--debug` option.) In addition, the TotalView Release Notes contains information on patching your MPICH 1.2.3 distribution.

Topics in this section are:

- “Starting TotalView on an MPICH Job” on page 92
- “Attaching to an MPICH Job” on page 94
- “MPICH P4 *procgrou*p Files” on page 96

Starting TotalView on an MPICH Job

Before you can bring an MPICH job under TotalView’s control, both TotalView and the TotalView server must be in your path. You can set this up in either a login or shell startup script.

At Version 1.1.2, here’s the command line that starts a job under TotalView’s control:

```
mpirun [ MPICH-arguments ] -tv program [ program-arguments ]
```

For example:

```
mpirun -np 4 -tv sendrecv
```

The MPICH **mpirun** command obtains information from the **TOTALVIEW** environment variable and then uses this information when it starts the first process in the parallel job.

At Version 1.2.4, this changes to:

```
mpirun -dbg=totalview [ other_mpich-arguments ] program [ program-arguments ]
```

For example:

```
mpirun -dbg=totalview -np 4 sendrecv
```

In this case, **mpirun** obtains the information it needs from the **-dbg** command-line option.

NOTE In other contexts, setting this environment variable means that you can use different versions of TotalView or pass command-line options to TotalView.

For example, here is the C shell command that sets the **TOTALVIEW** environment variable so that **mpirun** passes the **-no_stop_all** option to TotalView:

```
setenv TOTALVIEW "totalview -no_stop_all"
```

TotalView begins by starting the first process of your job, the master process, under its control. You can then set breakpoints and begin debugging your code.

On the IBM SP machine with the **ch_mpl** device, the **mpirun** command uses the **po**e command to start an MPI job. While you still must use the MPICH **mpirun** (and its **-tv** option) command to start an MPICH job, the way you start MPICH differs. For details on using TotalView with **po**e, see "Starting TotalView on a PE Job" on page 100.

Starting TotalView using **ch_p4mpd** is similar to starting TotalView using **po**e on IBM or other methods you might use on Sun and HP platforms. In general, you start TotalView using the **totalview** command. Here's the syntax;

```
totalview mpirun [ totalview_arguments ] \
                -a [ mpich-arguments ] program [ program-arguments ]
```

```
CLI EQUIVALENT:  totalviewcli mpirun [ totalview_arguments ] \
                  -a [ mpich-arguments ] program [ program-arguments ]
```

As your program executes, TotalView automatically acquires the processes that are part of your parallel job as your program creates them. Before TotalView begins to acquire them, it asks if you want to stop the spawned processes. If your answer is **Yes**, you can stop processes as they are initialized. This lets you check their states or set breakpoints that are unique to the process. TotalView automatically copies breakpoints from the master process to the slave processes as it acquires them. Consequently, you don't have to stop them just to set these breakpoints.

If you're using the GUI, TotalView updates the Root Window's **Attached** Page to show these newly acquired processes. For more information, see "Attaching to Processes" on page 134.

Attaching to an MPICH Job

TotalView allows you to attach to an MPICH application even if it was not started under TotalView's control. Here is the procedure:

- 1 Start TotalView.
- 2 The Root Window's **Unattached** Page displays the processes that are not yet owned.

```
CLI EQUIVALENT:  dattach executable pid
```

- 3 Attach to the first MPICH process in your workstation cluster by diving into it.

On an IBM SP with the **ch_mpi** device, attach to the **poe** process that started your job. For details, see "Starting TotalView on a PE Job" on page 100. Figure 65 on page 95 shows the Unattached window after some attaching has occurred.

Normally, the first MPICH process is the highest process with the correct image name in the process list. Other instances of the same executable can be:

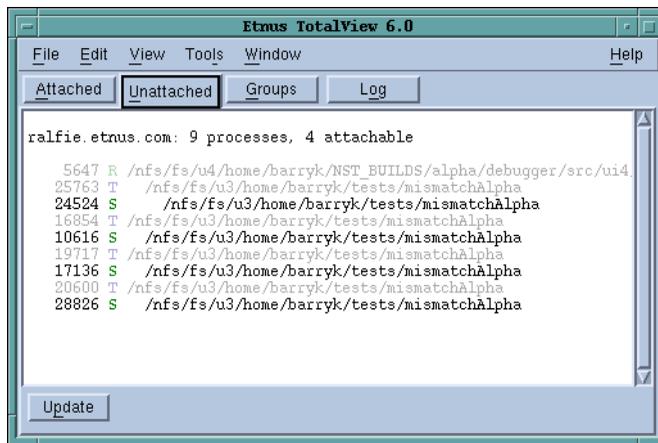


FIGURE 65: Root Window: Unattached Page

- The **p4** listener processes if MPICH was configured with **ch_p4**.
 - Additional slave processes if MPICH was configured with **ch_shmem** or **ch_lfshmem**.
 - Additional slave processes if MPICH was configured with **ch_p4** and have a machine file that places multiple processes on the same machine.
- 4** After you attach to your program's processes, TotalView asks if you also wish to attach to slave MPICH processes. If you do, press **Return** or choose **Yes**. If you do not, select **No**.

If you choose **Yes**, TotalView starts the server processes and acquires all MPICH processes.

As an alternative, you can use the **Group > Attach Subsets** command to pre-define what TotalView should do. For more information, see "Attaching to Processes" on page 134.

In some situations, the processes you expect to see may not exist (for example, they may have crashed or exited). TotalView acquires all the processes it can and then warns you if it could not attach to some of them. If you attempt to dive into a process that no longer exists (for example, using a message queue display), TotalView tells you that the process no longer exists.

MPICH P4 procgroup Files

If you're using MPICH with a P4 **procgroup** file (by using the **-p4pg** option), you must use the *same* absolute path name in your **procgroup** file and on the **mpirun** command line. For example, if your **procgroup** file contains a different path name than that used in the **mpirun** command, even though this name resolves to the same executable, TotalView treats it as different executable, which causes debugging problems.

The following example uses the *same* absolute path name on TotalView's command line and in the **procgroup** file:

```
% cat p4group
local 1 /users/smith/mympichexe
bigiron 2 /users/smith/mympichexe
% mpirun -p4pg p4group -tv /users/smith/mympichexe
```

In this example, TotalView:

- 1 Reads the symbols from **mympichexe** only once.
- 2 Places MPICH processes in the same TotalView share group.
- 3 Names the processes **mympichexe.0**, **mympichexe.1**, **mympichexe.2**, and **mympichexe.3**.

If TotalView assigns names such as **mympichexe<mympichexe>.0**, a problem occurred and you should compare the contents of your **procgroup** file and **mpirun** command line.

Debugging HP Alpha MPI Applications

You can debug HP Alpha MPI applications on the HP Alpha platform. To use TotalView with HP Alpha MPI, you must use HP Alpha MPI version 1.7 or later.

Starting TotalView on a HP Alpha MPI Job

HP Alpha MPI programs are most often started with the **dmpirun** command. You would use very similar command when starting an MPI program under TotalView's control:

```
{ totalview | totalviewcli } dmpirun -a dmpirun-command-line
```

This invokes TotalView and tells it to show you the code for the main program in **dmpirun**. Since you're not usually interested in debugging this code, you can use the **Process > Go** command to let the program run.

CLI EQUIVALENT: **dfocus p dgo**

The **dmpirun** command runs and starts all MPI processes. TotalView will also acquire them and ask if you want to stop them.

NOTE Problems can occur if you rerun HP Alpha MPI programs that are under TotalView control due to resource allocation issues within HP Alpha MPI. Consult the HP Alpha MPI manuals and release notes for information on using `mpiclean` to clean up the MPI system state.

Attaching to a HP Alpha MPI Job

To attach to a running HP Alpha MPI job, attach to the **dmpirun** process that started the job. The procedure for attaching to a **dmpirun** process is the same as the procedure for attaching to other processes. For details, see "Attaching to Processes" on page 49. You can also use the **Group > Attach Subset** command with is discussed in "Attaching to Processes" on page 134.

After you attach to the **dmpirun** process, TotalView asks if you also wish to attach to slave MPICH processes. If you do, press **Return** or choose **Yes**. If you do not, select **No**.

If you choose **Yes**, TotalView starts the server processes and acquires all MPICH processes.

Debugging HP MPI Applications

You can debug HP MPI applications on a PA-RISC 1.1 or 2.0 processor. To use TotalView with HP MPI, you must use HP MPI versions 1.6 or 1.7.

Starting TotalView on an HP MPI Job

TotalView lets you start an MPI program in three ways:

`{ totalview | totalviewcli } program -a mpi-arguments`

This command tells TotalView to start the MPI process. TotalView will then show you the machine code for the HP MPI `mpirun` executable.

CLI EQUIVALENT: `dfocus p dgo`

`mpirun mpi-arguments -tv -f startup_file`

This command tells MPI that it should start TotalView and then start the MPI processes as they are defined within the `startup_file` script. This file names the processes that MPICH will start. Typically, this file has contents that are similar to:

```
-h localhost -np 1 sendrecv
-h localhost -np 1 sendrecvva
```

In this example, `sendrecv` and `sendrecvva` are two different executable programs.

Your HP MPI documentation describes the contents of this startup file.

`mpirun mpi-arguments -tv program`

This command tells MPI that it should start TotalView.

Just before `mpirun` starts the MPI processes, TotalView acquires them and asks if you want to stop the processes before they start executing. If your answer is **yes**, TotalView halts them before they enter the `main()` routine. You can then create breakpoints.

Attaching to an HP MPI Job

To attach to a running HP MPI job, attach to the HP MPI `mpirun` process that started the job. The procedure for attaching to an `mpirun` process is the same as the procedure for attaching to any other process. For details, see "Attaching to Processes" on page 49.

After TotalView attaches to the HP MPI `mpirun` process, it displays the same dialog as it does with MPICH. (See step 4 on page 95 of "Attaching to an MPICH Job" on page 94.)

Debugging IBM MPI (PE) Applications

You can debug IBM MPI Parallel Environment (PE) applications on the IBM RS/6000 and SP platforms.

To take advantage of TotalView's ability to automatically acquire processes, you must be running release 3,1 or later of the Parallel Environment for AIX.

See "Displaying the Message Queue Graph" on page 107 for message queue display information.

Topics in this section are:

- "Preparing to Debug a PE Application" on page 99
- "Starting TotalView on a PE Job" on page 100
- "Setting Breakpoints" on page 101
- "Starting Parallel Tasks" on page 101
- "Attaching to a PE Job" on page 102

Preparing to Debug a PE Application

The following sections describe what you must do before TotalView can display a PE application.

Using Switch-Based Communication

If you're using switch-based communications (either "IP over the switch" or "user space") on an SP machine, you must configure your PE debugging session so that TotalView can use "IP over the switch" for communicating with the TotalView Debugger Server (`tvdsvr`). Do this by setting `adapter_use` to `shared` and `cpu_use` to `multiple`, as follows:

- If you're using a PE host file, add `shared multiple` after all host names or pool IDs in the host file.
- Always use the following arguments on the `poe` command line:
`-adapter_use shared -cpu_use multiple`

If you don't want to set these arguments in the `poe` command line, set the following environment variables before starting `poe`:

```
setenv MP_ADAPTER_USE shared
setenv MP_CPU_USE multiple
```

When using "IP over the switch," the default is usually **shared adapter use** and **multiple cpu use**; to be safe, set them explicitly by using one of these techniques.

When you're using switch-based communications, you must run TotalView on one of the SP or SP2 nodes. Since TotalView will be using "IP over the switch" in this case, you cannot run TotalView on an RS/6000 workstation.

Performing Remote Logins

You must be able to perform a remote login using the **rsh** command. You will also need to enable remote logins by adding the host name of the remote node to the `/etc/hosts.equiv` file or to your `.rhosts` file.

When the program is using switch-based communications, TotalView tries to start the TotalView Debugger Server by using the **rsh** command with the switch host name of the node.

Setting Timeouts

If you receive communications timeouts, you can set the value of the `MP_TIMEOUT` environment variable. For example:

```
setenv MP_TIMEOUT 1200
```

If this variable isn't set, TotalView uses a **timeout** value of 600 seconds.

Starting TotalView on a PE Job

Here is the syntax for running Parallel Environment (PE) programs from the command line:

```
program [ arguments ] [ pe_arguments ]
```

You can use the **poe** command to run programs:

```
poe program [ arguments ] [ pe_arguments ]
```

If, however, you start TotalView on a PE application, you must start **poe** as TotalView's target. The syntax for this is:

```
{ totalview | totalviewcli } poe -a program [ arguments ] [ PE_arguments ]
```

For example:

```
totalview poe -a sendrecv 500 -rmpool 1
```

Setting Breakpoints

After TotalView is running, you can start the **poe** process; this process will then start the program's parallel processes with the **Process > Go** command.

CLI EQUIVALENT: **dfocus p dgo**

TotalView responds by displaying a dialog box—in the CLI, it prints a question—that asks if you want to stop the parallel tasks.

If you want to set breakpoints in your code before they begin executing, answer **Yes**. TotalView initially stops the parallel tasks, which also allows you to set breakpoints. You can set breakpoints and control parallel tasks in the same way as any process controlled by TotalView.

If you have already set and saved breakpoints with the **Action Points > Save All** command and want to reload the file, answer **No**. After TotalView loads these saved breakpoints, the parallel tasks begin executing.

CLI EQUIVALENT: **dactions -save filename**
dactions -load filename

Starting Parallel Tasks

After you set breakpoints, you can start all of the parallel tasks with the Process Window's **Group > Go** command.

CLI EQUIVALENT: **dfocus G dgo**
Abbreviation: G

NOTE No parallel tasks will reach the first line of code in your main routine until all parallel tasks start.

You should be very cautious in placing breakpoints at or before a line that calls **MPI_Init()** or **MPL_Init()** because timeouts can occur while your program is being initialized. Once you allow the parallel processes to proceed into the **MPI_Init()** or **MPL_Init()** call, you should allow all of the parallel processes to proceed through it

within a short time. For more information on this, see “*Avoid unwanted timeouts*” on page 140.

Attaching to a PE Job

To take full advantage of TotalView’s **poe**-specific automation, you need to attach to **poe** itself, and let TotalView automatically acquire the **poe** processes on its various nodes. This set of acquired processes will include the processes you want to debug.

Attaching from a Node Running poe

Here’s the procedure for attaching TotalView to **poe** from the node running **poe**.

- 1 Start TotalView in the directory of the debug target.

If you can’t start TotalView in the debug target directory, you can start TotalView by editing the TotalView Debugger Server (**tvdsvr**) command line before attaching to **poe**. See “*Using the Single-Process Server Launch Command*” on page 80.

- 2 In the Root Window’s **Unattached** Page, find the **poe** process list, and attach to it by diving into it. When necessary, TotalView launches TotalView Debugger Servers. TotalView will also update the Root Window’s **Attached** Page and open a Process Window for the **poe** process.

CLI EQUIVALENT: **dattach poe pid**

- 3 Locate the process you want to debug and dive on it. TotalView responds by opening a Process Window for it.



If your source code files are not displayed in the Source Pane, you may not have told TotalView where these files reside. You can fix this by invoking the **File > Search Path** command to add directories to your search path.

Attaching from a Node Not Running poe

The procedure for attaching TotalView to **poe** from a node not running **poe** is essentially the same as the procedure for attaching from a node running **poe**. Since you did not run TotalView from the node running **poe** (the startup node), you won’t

be able to see **poe** on the process list in your Root Window's **Attached** Page and you won't be able to start it by diving into it.

The procedure for placing **poe** within this list is:



- 1 Connect TotalView to the startup node. For details, see "Starting the TotalView Debugger Server" on page 73 and "Attaching to Processes" on page 49.
- 2 Select the Root Window's **Unattached** Page, and then invoke the **Window > Update** command.
- 3 Look for the process named **poe** and continue as if attaching from a node running **poe**.

CLI EQUIVALENT: `dattach -r hostname poe poe-pid`

Debugging QSW RMS Applications

TotalView supports automatic process acquisition on AlphaServer SC systems and 32-bit Red Hat Linux systems that use Quadrics's RMS resource management system with the QSW switch technology.

NOTE Message queue display is only supported if you are running version 1, patch 2 or later, of AlphaServer SC.

Starting TotalView on an RMS Job

To start a parallel job under TotalView's control, use TotalView as though you were debugging **prun**:

```
{ totalview | totalviewcli } prun -a prun-command-line
```

TotalView starts up and shows you the machine code for RMS **prun**. Since you're not usually interested in debugging this code, use the **Process > Go** command to let the program run.

CLI EQUIVALENT: `dfocus p dgo`

The RMS **prun** command executes and starts all MPI processes. After TotalView acquires them, it asks if you want to stop them at startup. If you answer **yes**,

TotalView halts them before they enter the main program. You can then create breakpoints.

Attaching to an RMS Job

To attach to a running RMS job, attach to the RMS **prun** process that started the job.

You attach to the **prun** processes the same way you attach to other processes. For details on attaching to processes, see "Attaching to Processes" on page 49.

After you attach to the RMS **prun** process, TotalView asks if you also wish to attach to slave MPICH processes. If you do, press **Return** or choose **Yes**. If you do not, select **No**.

If you choose **Yes**, TotalView starts the server processes and acquires all MPICH processes.

As an alternative, you can use the **Group > Attach Subsets** command to predefine what TotalView should do. For more information, see "Attaching to Processes" on page 134.

Debugging SGI MPI Applications

TotalView can acquire processes started by SGI MPI, which is part of the Message Passing Toolkit (MPT) 1.3 and 1.4 packages.

Message queue display is supported by release 1.3 and 1.4 of the Message Passing Toolkit. See "Displaying the Message Queue Graph" on page 107 for message queue display.

Starting TotalView on a SGI MPI Job

SGI MPI programs are normally started by using the **mpirun** command. You would use a similar command to start an MPI program under TotalView's control:

```
{ totalview | totalviewcli } mpirun -a mpirun-command-line
```

This invokes TotalView and tells it to show you the machine code for `mpirun`. Since you're not usually interested in debugging this code, use the **Process > Go** command to let the program run.

CLI EQUIVALENT: `dfocus p dgo`

The SGI MPI `mpirun` command runs and starts all MPI processes. After TotalView acquires them, it asks if you want to stop them at startup. If you answer **Yes**, TotalView halts them before they enter the main program. You can then create breakpoints.

If you set a verbosity level that allows informational messages, TotalView also prints a message showing the name of the array and the value of the array services handle (**ash**) to which it is attaching.

Attaching to an SGI MPI Job

To attach to a running SGI MPI job, attach to the SGI MPI `mpirun` process that started the job. The procedure for attaching to an `mpirun` process is the same as the procedure for attaching to any other process. For details, see "Attaching to Processes" on page 49.

After you attach to the `mpirun` process, TotalView asks if you also wish to attach to slave MPICH processes. If you do, press **Return** or choose **Yes**. If you do not, select **No**.

If you choose **Yes**, TotalView starts the server processes and acquires all MPICH processes.

As an alternative, you can use the **Group > Attach Subsets** command to predefine what TotalView should do. For more information, see "Attaching to Processes" on page 134.

Debugging Sun MPI Applications

TotalView can debug a Sun MPI program and can display Sun MPI message queues. This section describes how to perform *job startup* and *job attach*.

- 1 Type the following command

```
totalview mprun [ totalview_args ] -a [ mpi_args ]
```

For example:

```
totalview mprun -g blue -a -np 4 /usr/bin/mpi/conn.x
```

CLI EQUIVALENT: **totalviewcli mprun [totalview_args] -a [mpi_args]**

When the TotalView Process Window appears, select the **Go** button.

CLI EQUIVALENT: **dfocus p dgo**

TotalView may display a dialog box that says:

Process mprun is a parallel job. Do you want to stop the job now?

- 2 If you had compiled using the **-g** option, clicking **Yes** tells TotalView to open a Process Window showing your source. All processes will be halted.

Attaching to a Sun MPI Job

This section describes how to attach to an already running **mprun** job.

- 1 Find the host name and process identifier (PID) of the **mprun** job by typing **mpps -b**. For more information, refer to the **mpps(1M)** manual page.

Here is sample output from this command:

JOBNAME	MPRUN_PID	MPRUN_HOST
cre.99	12345	hpc-u2-9
cre.100	12601	hpc-u2-8

- 2 After selecting **File > New Program**, type **mprun** in the **Executable** field and type the PID in the **Process ID** field.

CLI EQUIVALENT: **dattach mprun mprun-pid**
 For example:
dattach mprun 12601

- If TotalView is running on a different node than the `mprun` job, enter the host name in the **Remote Host** field.

CLI EQUIVALENT: `dattach -r host-name mprun mprun-pid`

Displaying the Message Queue Graph

TotalView can graphically display your MPI program's message queue state. If you select the Process Window's **Tools > Message Queue Graph** command, TotalView displays a window with a large empty area. After you select the ranks to be monitored, the kind of messages, and message states, TotalView updates this window to show the current queue state. Figure 66 shows a sample window.

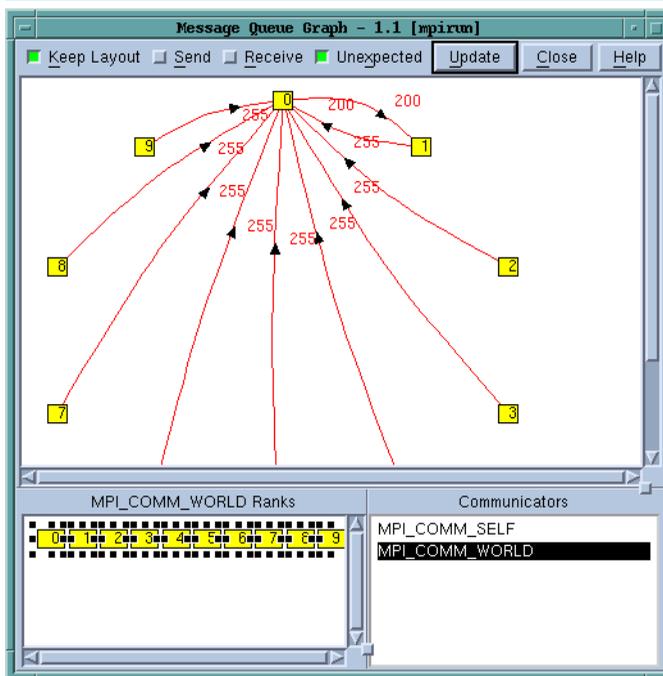


FIGURE 66: **Tools > Message Queue Graph** Window

The numbers in the boxes indicate the MPI message tag number. Diving on a box tells TotalView to open a Process Window for that process.

The numbers next to the arrows indicate the number of messages that existed when TotalView created the graph. Diving on an arrow tells TotalView that it should display its **Tools > Message Queue** Window, which will have detailed information about the messages. A grey box indicates a process to which TotalView is not attached.

The colors used to draw the lines and arrows have the following meaning:

- **Green:** sent messages
- **Blue:** receive messages
- **Red:** unexpected messages

This graph shows you the state of your program at a particular instant. Selecting the **Update** button tells TotalView that it should update the display.

While you can use this window in many ways, here are some to consider:

- Pending messages often indicate that a process can't keep up with the amount of work it is expected to perform. These messages indicate places where you may be able to improve your program's efficiency.
- Unexpected messages can indicate that something is wrong with your program because the receiving process doesn't know how to process the message. The red lines indicated unexpected messages.
- After a while, the shape of the graph tends to tell you something about how your program is executing. If something doesn't look right, you might want to determine why it looks wrong.
- You can change the shape of the graph by dragging either nodes or the arrows. This is often useful when you're comparing sets of nodes and their messages with one another. TotalView doesn't remember the places to which you have dragged the nodes and arrows. This means that if you select the **Display** button after you arrange the graph, your changes are lost.

Topics related to this one are:

- "*Message Queue Display Overview*" on page 109
- "*Using Message Operations*" on page 110
- "*OpenMP Stack Parent Token Line*" on page 120

Displaying the Message Queue

The **Tools > Message Queue** Dialog Box displays your MPI program's message queue state textually. This can be useful when you need to find out why a deadlock occurred.

To use the message queue display feature, you must be using one of the following versions of MPI:

- MPICH version 1.2.3 or later.
- HP Alpha MPI (DMPI) version 1.8, 1.9, and 1.96.
- HP HP-UX version 1.6 and 1.7.
- IBM MPI Parallel Environment (PE) version 3.1 or 3.2, but only programs using the threaded IBM MPI libraries. MOD is not available with earlier releases, or with the non-thread-safe version of the IBM MPI library. Therefore, to use TotalView MOD with IBM MPI applications, you must use the `mpicc_r`, `mpxlf_r`, or `mpxlf90_r` compilers to compile and link your code.
- For the SGI MPI TotalView message queue display, you must obtain the Message Passing Toolkit (MPT) release 1.3 and 1.4. Check with SGI for availability.

Message Queue Display Overview

After an MPI process returns from the call to `MPI_Init()`, you can display the internal state of the MPI library by selecting the **Tools > Message Queue** command. The information is shown in Figure 67 on page 110.

This window displays the state of the process's MPI communicators. If user-visible communicators are implemented as two internal communicator structures, TotalView displays both of them. One will be used for point-to-point operations and the other for collective operations.

NOTE You cannot edit any of the fields in the Message Queue Window.

The contents of the Message Queue Window are only valid when a process is stopped.

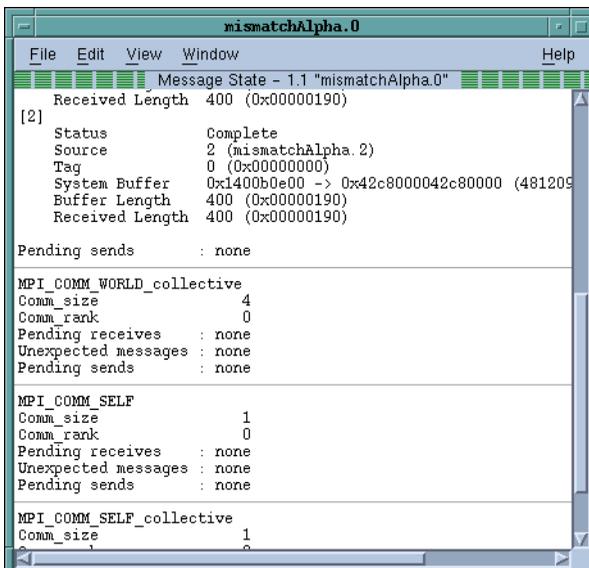


FIGURE 67: Message Queue Window

Using Message Operations

For each communicator, TotalView displays a list of pending receive operations, pending unexpected messages, and pending send operations. Each operation has an index value displayed in brackets ([*n*]). The online Help for this window contains a description of the fields that can be displayed.

Topics in this section are:

- "Diving on MPI Processes" on page 111
- "Diving on MPI Buffers" on page 111
- "Pending Receive Operations" on page 111
- "Unexpected Messages" on page 112
- "Pending Send Operations" on page 112

Diving on MPI Processes

To display more detail, you can dive into fields in the Message Queue Window. When you dive into a process field, TotalView does one of the following:

- Raises its Process Window if it exists.
- Sets the focus to an existing Process Window on the requested process.
- If a Process Window doesn't exist, TotalView creates a new one for the process.

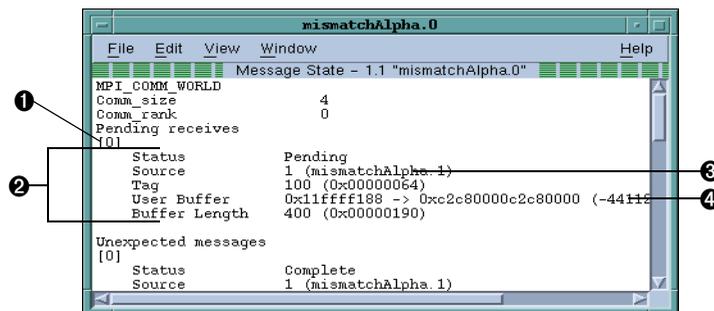
Diving on MPI Buffers

When you dive into the buffer fields, TotalView opens a Variable Window. It also guesses what the correct format for the data should be based on the buffer's length and the data's alignment. If TotalView guesses incorrectly, you can edit the type field in the Variable Window.

NOTE TotalView doesn't use the MPI data type to set the buffer type.

Pending Receive Operations

TotalView displays each pending receive operation in the **Pending receives** list. Figure 68 shows an example of an MPICH pending receive operation.



- ❶ Operation index
- ❷ One receive operation
- ❸ Diving here displays a Process Window
- ❹ Diving here displays a Variable Window

FIGURE 68: Message Queue Window Showing Pending Receive Operation

NOTE TotalView displays all receive operations maintained by the IBM MPI library. You should set the environment variable `MP_EUIDEVELOP` to the value `DEBUG` if you want blocking operations to be visible; otherwise, the library only maintains non-blocking operations. For more details on the `MP_EUIDEVELOP` environment variable, consult the IBM *Parallel Environment Operations and Use* manual.

Unexpected Messages

The **Unexpected messages** portion of the **Message Queue** Window shows information for retrieved and enqueued messages that are not yet matched with a receive operation.

Some MPI libraries such as MPICH only retrieve already received messages as a side effect of calls to functions such as `MPI_Recv()` or `MPI_Iprobe()`. (In other words, while some versions of MPI may know about the message, the message may not yet be in a queue.) This means that TotalView can't list a message until after the destination process makes a call that retrieves it.

Pending Send Operations

TotalView displays each pending send operation in the **Pending sends** list.

MPICH does not normally keep information about pending send operations. However, when you configure MPICH, you can tell it to maintain a list of them. Start your program under TotalView's control and use `mpirun -ksq`, or `-KeepSendQueue` to see these messages.

Depending on the device for which MPICH was configured, blocking send operations may or may not be visible. However, if TotalView doesn't display them, you can see that these operations occurred because the call is in the stack backtrace.

If you attach to an MPI program that isn't maintaining send queue information, TotalView displays the following message:

```
Pending sends : no information available
```

MPI Debugging Troubleshooting

If you can't successfully start TotalView on MPI programs, check the following:

- Can you successfully start MPICH programs without TotalView?

The MPICH code contains some useful scripts that let you verify that you can start remote processes on all of the machines in your machines file. (See `tstmachines` in `mpich/util`.)

- You won't get a message queue display if you get the following warning:

The symbols and types in the MPICH library used by TotalView to extract the message queues are not as expected in the image *<your image name>*. This is probably an MPICH version or configuration problem.

You need to check that you are using MPICH Version 1.1.0 or later and that you have configured it with the `-debug` option. (You can check this by looking in the `config.status` file at the root of the MPICH directory tree).

- Does the TotalView Debugger Server (`tvdsrv`) fail to start?

`tvdsrv` must be in your `PATH` when you log in. Remember that TotalView uses `rsh` to start the server, and that this command doesn't pass your current environment to remotely started processes.

- Make sure you have the correct MPI version and have applied all required patches. See the TOTALVIEW RELEASE NOTES for up-to-date information.
- Under some circumstances, MPICH kills TotalView with the `SIGINT` signal. You can see this behavior when you use the `Group > Delete` command to restart an MPICH job.

CLI EQUIVALENT: `dfocus g ddelete`

If TotalView exits and terminates abnormally with a **Killed** message, try setting the `TV::ignore_control_c` variable to true.

Debugging OpenMP Applications

TotalView supports many OpenMP C and Fortran compilers. Supported compilers and architectures are listed in the TOTALVIEW PLATFORMS document and on our Web site.

Here are some of the features that TotalView supports:

- Source-level debugging of the original OpenMP code.
- The ability to plant breakpoints throughout the OpenMP code, including lines that are executed in parallel.
- Visibility of OpenMP worker threads.
- Access to **SHARED** and **PRIVATE** variables in OpenMP **PARALLEL** code.
- A stack-back link token in worker threads' stacks so that you can find their master stack.
- Access to OMP **THREADPRIVATE** data in code compiled by the IBM and Guide, SGI IRIX, and HP Alpha compilers.

The code examples used in this section are included in the TotalView distribution in the `examples/omp_simplef` file.

NOTE On the SGI IRIX platform, you must use the MIPSpro 7.3 compiler or later to debug OpenMP.

Topics in this section are:

- "Debugging OpenMP Programs" on page 114
- "OpenMP Private and Shared Variables" on page 117
- "OpenMP **THREADPRIVATE** Common Blocks" on page 118
- "OpenMP Stack Parent Token Line" on page 120

Debugging OpenMP Programs

Debugging OpenMP code is very similar to debugging multithreaded code, differing only in that the OpenMP compiler makes the following special code transformations:

- The most visible transformation is *outlining*. The compiler pulls the body of a parallel region out of the original routine and places it into an *outlined routine*. In some cases, the compiler will generate multiple outlined routines from a single parallel region. This allows multiple threads to execute the parallel region. The outlined routine's name is based on the original routine's name.
- The compiler inserts calls to the OpenMP runtime library.

- The compiler splits variables between the original routine and the outlined routine. Normally, shared variables are maintained in the master thread's original routine, and private variables are maintained in the outlined routine.
- The master thread creates threads to share the workload. As the master thread begins to execute a parallel region in the OpenMP code, it creates the worker threads, dispatches them to the outlined routine, and then calls the outlined routine itself.

TotalView OpenMP Features

TotalView makes these transformations visible in the debugging session. Here are some things you should know:

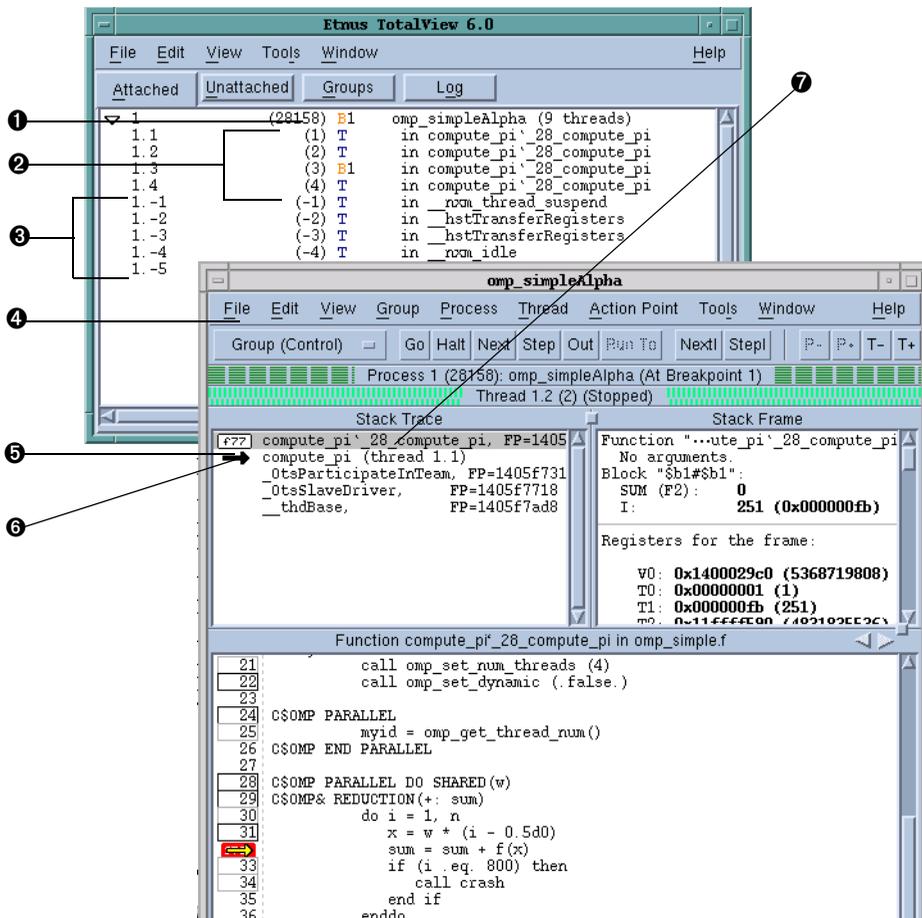
- The compiler may generate multiple outlined routines from a single parallel region. This means that a single line of source code can generate multiple blocks of machine code inside different functions.
- You can't single step into or out of a parallel region. Instead, set a breakpoint inside the parallel region and allow the process to run to it. Once inside a parallel region, you can single step within it.
- OpenMP programs are multithreaded programs, so the rules for debugging multithreaded programs apply.

Figure 69 on page 116 shows a sample OpenMP debugging session.

OpenMP Platform Differences

The following list contains information on platform differences:

- On HP Alpha Tru64 UNIX and on the Guide compilers, the OpenMP threads are implemented by the compiler as **pthreads**, and on SGI IRIX as **sprocs**. TotalView shows the threads' logical and/or system thread ID, not the OpenMP thread number.
- The OpenMP master thread has logical thread ID number 1. The OpenMP worker threads have a logical thread ID number greater than 1.
- In HP Alpha Tru64 UNIX, the system manager threads have a negative thread ID; as they do not take part in your OpenMP program, you should never manipulate them.
- SGI OpenMP uses the **SIGTERM** signal to terminate threads. Because TotalView stops a process when the process receives a **SIGTERM**, the OpenMP process



- | | | | |
|---|-------------------------------|---|--|
| ① | OpenMP master thread | ⑤ | "Original" routine name |
| ② | OpenMP worker threads | ⑥ | Stack parent token (select or dive to view master) |
| ③ | Manager threads (don't touch) | ⑦ | "Outlined" routine name |
| ④ | Slave Thread Window | | |

FIGURE 69: Sample OpenMP Debugging Session

doesn't terminate. If you want the OpenMP process to terminate instead of stop, set the default action for the **SIGTERM** signal to *Resend*.

- When you stop the OpenMP master thread in a **PARALLEL DO** outlined routine, the stack backtrace shows the following call sequence:

- The outlined routine called from.
- The OpenMP runtime library called from.
- The original routine (containing the parallel region).
- When you stop the OpenMP worker threads in a **PARALLEL DO** outlined routine, the stack backtrace shows the following call sequence:
 - Outlined routine called from the special stack parent token line.
 - The OpenMP runtime library called from.
- Select or dive on the stack parent token line to view the original routine's stack frame in the OpenMP master thread.

OpenMP Private and Shared Variables

TotalView allows you to view both OpenMP private and shared variables.

The compiler maintains OpenMP private variables in the outlined routine, and treats them like local variables. See "*Displaying Local Variables and Registers*" on page 284. In contrast, the compiler maintains OpenMP shared variables in the master thread's original routine stack frame. However, Guide compilers pass shared variables to the outlined routine as parameter references.

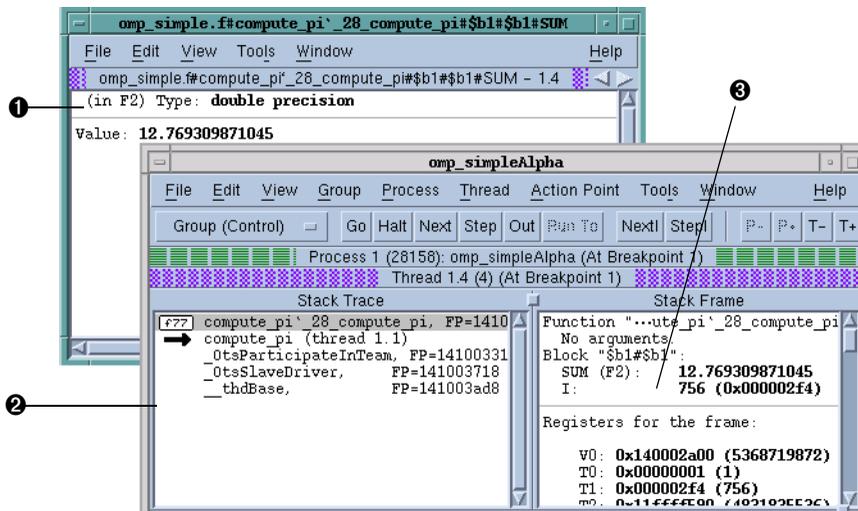
TotalView lets you display shared variables through a Process Window focused on the OpenMP master thread or through one of the OpenMP worker threads, as follows:

- 1 Select the outlined routine in the Stack Trace Pane; or select the original routine stack frame in the OpenMP master thread.
- 2 Dive on the variable name, or select the **View > Lookup Variable** command. When prompted, enter the variable name.

CLI EQUIVALENT: `dprint`
You will need to set your focus to the OpenMP master thread first.

TotalView will open a Variable Window displaying the value of the OpenMP shared variable, as shown in Figure 70 on page 118.

Shared variables are stored on the OpenMP master thread's stack. When displaying shared variables in OpenMP worker threads, TotalView uses the stack context of the



- ❶ OpenMP shared variables have master thread's context
- ❷ Original routine's stack frame selected
- ❸ Stack Frame Pane includes shared variables

FIGURE 70: OpenMP Shared Variable

OpenMP master thread to find the shared variable. TotalView uses the OpenMP master thread's context when displaying the shared variable in a Variable Window.



You can also view OpenMP shared variables in the Stack Frame by selecting the original routine stack frame in the OpenMP master thread, or by selecting the stack parent token line in the Stack Trace Pane of OpenMP worker threads, as shown in Figure 70.

OpenMP THREADPRIVATE Common Blocks

The HP Alpha Tru64 UNIX OpenMP and SGI IRIX compilers implement OpenMP **THREADPRIVATE** common blocks by using the thread local storage system facility. This facility stores a variable declared in OpenMP **THREADPRIVATE** common blocks at different memory locations in each thread in an OpenMP process. This allows the variable to have different values in each thread. In contrast, the IBM and Guide compilers use the pthread key facility.

On SGI, the `THREADPRIVATE` variables are mapped to the same virtual address. However, they have different physical addresses.

Here's how you can view a variable in an OpenMP `THREADPRIVATE` common block, or the OpenMP `THREADPRIVATE` common block itself:

- 1 In the Threads Pane of the Process Window, select the thread containing the private copy of the variable or common block you would like to view.
- 2 In the Stack Trace Pane of the Process Window, select the stack frame that will allow you to access the OpenMP `THREADPRIVATE` common block variable. You can select either the outlined routine or the original routine for an OpenMP master thread. You must, however, select the outlined routine for an OpenMP worker thread.
- 3 From the Process Window, dive on the variable name or common block name. Or select the **View > Lookup Variable** command. When prompted, enter the name of the variable or common block. You may need to append an underscore (`_`) after the common block name.

CLI EQUIVALENT: `dprint`

TotalView opens a Variable Window displaying the value of the variable or common block for the selected thread.

See "Displaying Variables" on page 281 for more information on displaying variables.



- 4 To view OpenMP `THREADPRIVATE` common blocks or variables across all threads, you can use the Variable Window's **View > Laminate Threads** command. See "Displaying a Variable in All Processes or Threads" on page 333.

Figure 71 on page 120 shows Variable Windows displaying OpenMP `THREADPRIVATE` common blocks. Because the Variable Window has the same thread context as the Process Window from which it was created, the title bar patterns for the same thread match. In the laminated views, the values of the common block across all threads are displayed.

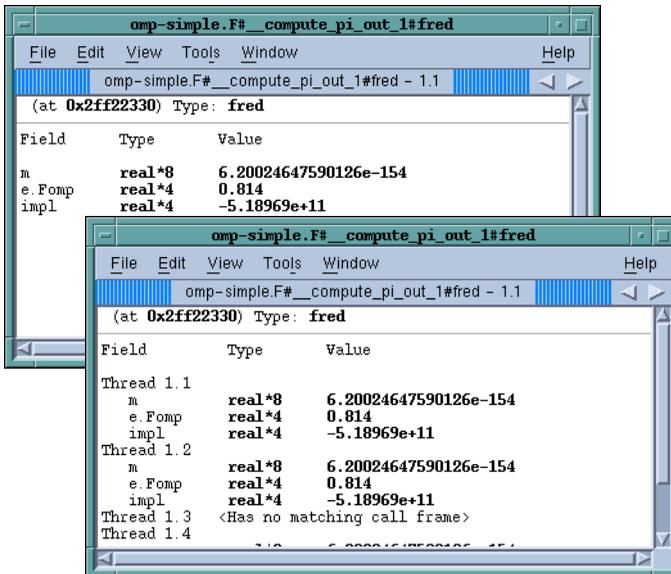


FIGURE 71: OpenMP THREADPRIVATE Common Block Variables

OpenMP Stack Parent Token Line

TotalView inserts a special stack parent token line in the Stack Trace Pane of OpenMP worker threads when they are stopped in an outlined routine.

When you select or dive on the stack parent token line, the Process Window switches to the OpenMP master thread, allowing you to see the stack context of the OpenMP worker thread's routine. This context includes the OpenMP shared variables. (See Figure 72.)

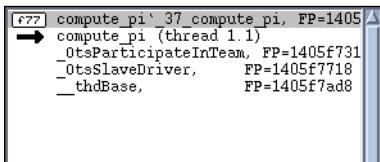


FIGURE 72: OpenMP Stack Parent Token Line

Debugging PVM and DPVM Applications

You can debug applications that use the Parallel Virtual Machine (PVM) library or the HP Alpha Tru64 UNIX Parallel Virtual Machine (DPVM) library with TotalView on some platforms. TotalView supports ORNL PVM Version 3.4.4 on all platforms and DPVM Version 1.9 or later on the HP Alpha platform.

NOTE See the TotalView *Platforms* document for the most up-to-date information regarding your PVM or DPVM software.

For tips on debugging parallel applications, see “*Parallel Debugging Tips*” on page 134.

Topics in this section are:

- “*Supporting Multiple Sessions*” on page 121
- “*Setting Up ORNL PVM Debugging*” on page 122
- “*Starting an ORNL PVM Session*” on page 122
- “*Starting a DPVM Session*” on page 123
- “*Automatically Acquiring PVM/DPVM Processes*” on page 124
- “*Attaching to PVM/DPVM Tasks*” on page 126

Supporting Multiple Sessions

When you debug a PVM or DPVM application, TotalView becomes a PVM tasker. This lets it establish a debugging context for your session. You can run:

- One TotalView PVM or DPVM debugging session for a user and for an architecture; that is, different users can’t interfere with each other on the same machine or same machine architecture.

One user can start TotalView to debug the same PVM or DPVM application on different machine architectures. However, a single user can’t have multiple instances of TotalView debugging the same PVM or DPVM session on a single machine architecture.

For example, suppose you start a PVM session on Sun 5 and HP Alpha machines. You must start two TotalView sessions: one on the Sun 5 machine to debug the Sun 5 portion of the PVM session, and one on the HP Alpha machine to debug the HP Alpha portion of the PVM session. These two TotalView sessions are separate and don’t interfere with one another.

- Similarly, in one TotalView session, you can run either a PVM application or a DPVM application, but not both. However, if you run TotalView on a HP Alpha, you can have two TotalView sessions: one debugging PVM and one debugging DPVM.

Setting Up ORNL PVM Debugging

To enable PVM, create a symbolic link from the PVM **bin** directory (which is `$HOME/pvm3/bin/$PVM_ARCH/tvdsvr`) to the TotalView Debugger Server (**tvdsvr**). With this link in place, TotalView invokes **pvm_spawn()** to spawn the **tvdsvr** tasks.

For example, if **tvdsvr** is installed in the `/opt/totalview/bin` directory, enter the following command:

```
ln -s /opt/totalview/bin/tvdsvr \
    $HOME/pvm3/bin/$PVM_ARCH/tvdsvr
```

If the symbolic link doesn't exist, TotalView can't spawn **tvdsvr**. When TotalView can't spawn **tvdsvr**, it displays the following error:

```
Error spawning TotalView Debugger Server: No such file
```

Starting an ORNL PVM Session

Start the ORNL PVM daemon process before you start TotalView. See the ORNL PVM documentation for information about the PVM daemon process and console program. The following steps outline this procedure.

- 1 Use the **pvm** command to start a PVM console session—this command starts the PVM daemon. If PVM isn't running when you start TotalView (with PVM support enabled), TotalView exits with the following message:
Fatal error: Error enrolling as PVM task: pvm error
- 2 If your application uses groups, start the **pvmgs** process before starting TotalView. PVM groups are unrelated to TotalView process groups. For information about TotalView process groups, refer to "Examining Groups" on page 223.
- 3 You can use the `-pvm` command-line option to the **totalview** command. As an alternative, you can set the `TV::pvm` variable in a startup file. The command-line options override the a CLI variable. For more information, refer to "TotalView Command Syntax" in the TOTALVIEW REFERENCE GUIDE.

- 4 Set the TotalView directory search path to include the PVM directories. This directory list must include those needed to find both executable and source files. The directories you use will vary, but should always contain the current directory and your home directory.



You can set the directory search path by setting the `TV::search_path` variable or you can use the **File > Search Directory** command. Refer to “*Setting Search Paths*” on page 59 for more information.

For example, to debug the PVM examples, you can place the following directories in your search path:

```
$HOME
$PVM_ROOT/xep
$PVM_ROOT/xep/$PVM_ARCH
$PVM_ROOT/src
$PVM_ROOT/src/$PVM_ARCH
$PVM_ROOT/bin/$PVM_ARCH
$PVM_ROOT/examples
$PVM_ROOT/examples/$PVM_ARCH
$PVM_ROOT/gexamples
$PVM_ROOT/gexamples/$PVM_ARCH
```



- 5 Verify that the action taken by TotalView for the **SIGTERM** signal is appropriate. (You can examine the current action by using the Process Window’s **File > Signals** command. Refer to “*Handling Signals*” on page 56 for more information.)

PVM uses the **SIGTERM** signal to terminate processes. Because TotalView stops a process when the process receives a **SIGTERM**, the process is not terminated. If you want the PVM process to terminate, set the action for the **SIGTERM** signal to **Resend**.

Continue with “*Automatically Acquiring PVM/DPVM Processes*” on page 124.

Starting a DPVM Session

Starting a DPVM debugging session is similar to starting any other TotalView debugging session. The only additional requirement is that you must start the DPVM daemon before you start TotalView. See the DPVM documentation for information about the DPVM daemon and its console program.

- 1 Use the **dpvm** command to start a DPVM console session; starting the session also starts the DPVM daemon. If DPVM isn't running when you start TotalView (with DPVM support enabled), TotalView displays the following error message before it exits:

Fatal error: Error enrolling as DPVM task: dpvm error

- 2 You can enable DPVM support in two ways. The first uses the **TV::dpvm** CLI variable. As an alternative, you can add the **-dpvm** command-line option to the **totalview** command. This option enables DPVM support.

The command-line options override the **TV:dpvm** command variable. For more information on the **totalview** command, refer to "TotalView Command Syntax" in the TOTALVIEW REFERENCE GUIDE.



- 3 Verify that the default action taken by TotalView for the **SIGTERM** signal is appropriate. You can examine the default actions with the Process Window's **File > Signals** command in TotalView. Refer to "Handling Signals" on page 56 for more information.

DPVM uses the **SIGTERM** signal to terminate processes. Because TotalView stops a process when the process receives a **SIGTERM**, the process is not terminated. If you want the DPVM process to terminate, set the action for the **SIGTERM** signal to **Resend**.

If you enable PVM support using the **TV::pvm** variable and you need to use DPVM, you must use both **-no_pvm** and **-dpvm** command-line options when you start TotalView. Similarly, when enabling DPVM support us the **TV::dpvm** variable, you can must use the **-no_dpvm** and **-pvm** command-line options.

NOTE You cannot use CLI variables to start both PVM and DPVM.

Automatically Acquiring PVM/DPVM Processes

This section describes how TotalView automatically acquires PVM and DPVM processes in a PVM or DPVM debugging session. Specifically, TotalView uses the PVM tasker to intercept **pvm_spawn()** calls.

When you start TotalView as part of a PVM or DPVM debugging session, it takes the following actions:

- TotalView makes sure that no other PVM or DPVM taskers are running. If TotalView finds a tasker on a host that it is debugging, it displays the following message and then exits:
Fatal error: A PVM tasker is already running on host '*host*'
- TotalView finds all the hosts in the PVM or DPVM configuration. Using the `pvm_spawn()` call, TotalView starts a TotalView Debugger Server (`tvdsvr`) on each remote host that has the same architecture type as the host TotalView is running on. It tells you it has started a debugger server by displaying:
Spawning TotalView Debugger Server onto PVM host '*host*'

If you add a host with a compatible machine architecture to your PVM or DPVM debugging session after you start TotalView, TotalView automatically starts a debugger server on that host.

After all debugger servers are running, TotalView will intercept every PVM or DPVM task created with the `pvm_spawn()` call on hosts that are part of the debugging session. If a PVM or DPVM task is created on a host with a different machine architecture, TotalView ignores that task.

When TotalView receives a PVM or DPVM tasker event, it takes the following actions:

- 1 TotalView reads the symbol table of the spawned executable.
- 2 If a saved breakpoint file for the executable exists and you have enabled automatic loading of breakpoints, TotalView loads breakpoints for the process.
- 3 TotalView asks if you want to stop the process before it enters the `main()` routine.

If you answer **Yes**, TotalView stops the process before it enters `main()` (that is, before it executes any user code). This allows you to set breakpoints in the spawned process before any user code executes. On most machines, TotalView stops a process in the `start()` routine of the `cr0.o` module if it is statically linked. If the process is dynamically linked, TotalView stops it just after it finishes running the dynamic linker. Because the Process Window displays assembler instructions, you will need to use the **View > Lookup Function** command to display the source code for the `main()` routine.

CLI EQUIVALENT: `dlist function-name`

For more information on this command, refer to “*Finding the Source Code for Functions*” on page 213.

Attaching to PVM/DPVM Tasks

You can attach to a PVM or DPVM task if the task meets the following criteria:

- The machine architecture on which the task is running is the same as the machine architecture on which TotalView is running.
- The task must be created. (This is indicated when flag 4 is set in the PVM Tasks and Configuration Window.)
- The task must not be a PVM tasker. If flag 400 is clear in the PVM Tasks and Configuration Window, the process is a tasker.
- The executable name must be known. If the executable name is listed as a dash (–), TotalView cannot determine the name of the executable. (This can occur if a task was not created with the `pvm_spawn()` call.)

To attach to a PVM or DPVM task, complete the following steps:



- 1 Select **Tools > PVM Tasks** command from TotalView’s Root Window.

The PVM Tasks is displayed, as shown in Figure 73. This window displays current information about PVM tasks and hosts—TotalView automatically updates this information as it receives events from PVM.

Since PVM doesn’t always generate an event that allows TotalView to update this window, you should use the **Windows > Update** command to ensure that you are seeing the most current information.

For example, you can attach to the tasks named `xep` and `mtile` in Figure 73 because flag 4 is set. In contrast, you can’t attach to the `tvdsrv` and `–` (dash) executables because flag 400 is set.

- 2 Dive on a task entry that meets the criteria for attaching to tasks. TotalView attaches to the task.
- 3 If the task to which you attached has related tasks that can be debugged, TotalView asks if you want to attach to these related tasks. If you answer **Yes**, TotalView attaches to them. If you answer **No**, it only attaches to the task you dove on.

After attaching to a task, TotalView looks for attached tasks that are related to this task; if there are related tasks, TotalView places them in the same control group. If

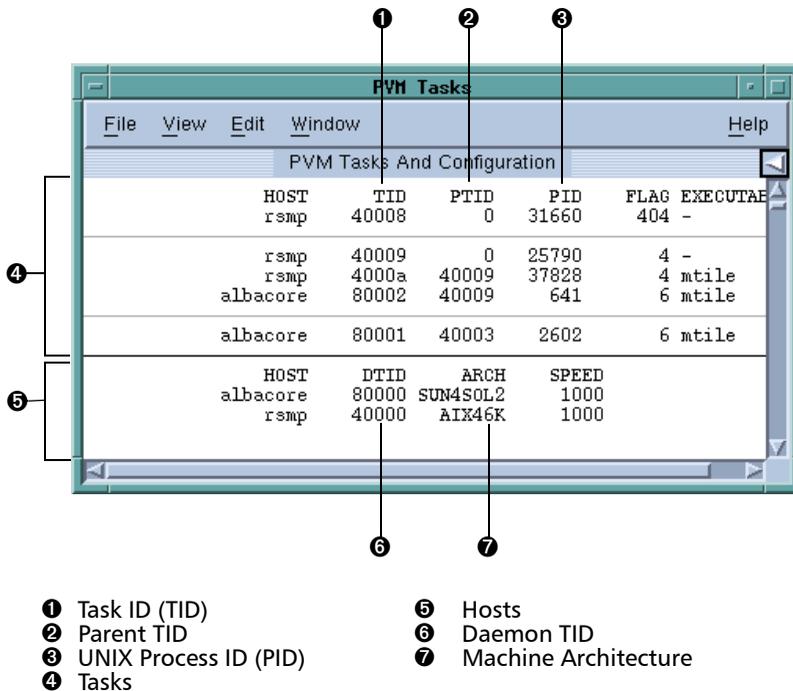


FIGURE 73: PVM Tasks and Configuration Window

TotalView is already attached to a task you dove on, it simply opens and raises the Process Window for the task.

Reserved Message Tags

TotalView uses PVM message tags in the range 0xDEB0 through 0xDEBF to communicate with PVM daemons and the TotalView Debugger Server. Avoid sending messages that use these reserved tags.

Cleanup of Processes

The `pvmgs` process registers its task ID in the PVM database. If the `pvmgs` process is terminated, the `pvm_joingroup()` routine hangs because PVM won't clean up the database. If this happens, you must terminate and then restart the PVM daemon.

TotalView attempts to clean up the TotalView Debugger Server daemons (`tvdsrvr`), that also act as taskers. If some of these processes do not terminate, you must manually terminate them.

Debugging Shared Memory (SHMEM) Code

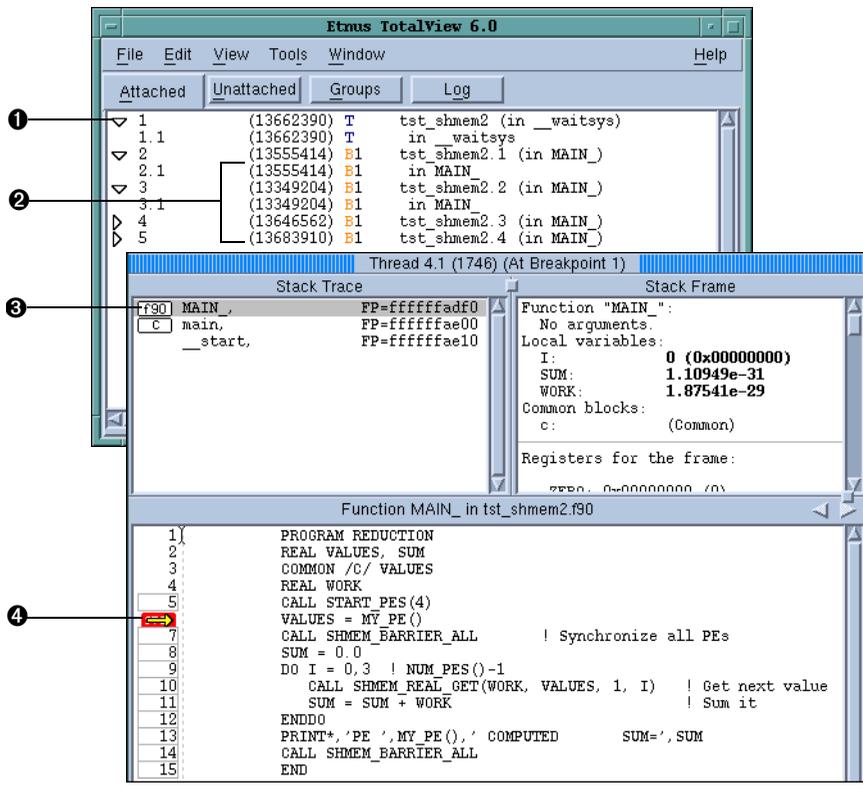
TotalView supports the SGI IRIX logically shared, distributed memory access (SHMEM) library.

To debug a SHMEM program, follow these steps:

- 1 Link it with the **dbfork** library. See “*Linking with the dbfork Library*” in the “*Compilers and Platforms*” chapter of the TOTALVIEW REFERENCE GUIDE.
- 2 Start TotalView on your program. See Chapter 3, “*Setting Up a Debugging Session*” on page 39.
- 3 Set at least one breakpoint after the call to the **start_pes()** SHMEM routine. (This is illustrated in Figure 74.)

NOTE You cannot single-step over the call to **start_pes()**.

The call to **start_pes()** creates new worker processes that return from the **start_pes()** call and execute the remainder of your program. The original process never returns from **start_pes()**, but instead stays in that routine, waiting for the worker processes it created to terminate.



- ① SHMEM starter process
- ② SHMEM worker processes
- ③ Select a worker process in the Root Window
- ④ Set breakpoint *after* the call to start_pes()

FIGURE 74: SHMEM Sample Session

Debugging UPC Programs

TotalView lets you debug UPC programs that compiled using the HP Compaq Alpha UPC 2.0 and the Intrepid (SGI gcc UPC) compilers. This section only discusses UPC-specific features of TotalView. It not an introduction to the UPC language. If you're looking for an introduction, you'll find information at <http://www.gwu.edu/~upc>.

NOTE When debugging UPC code, TotalView requires help from a UPC assistant library that your compiler vendor provides. You may need to include the location of this library in your `LD_LIBRARY_PATH` variable. Etnus also provides assistants that you can use. You can find these assistants at <http://www.etnus.com/Products/TotalView/developers/index.html>

Topics in this section are:

- “*Viewing Shared Objects*” on page 130
- “*Pointer to Shared*” on page 132

Invoking TotalView

Here’s how to invoke TotalView upon UPC programs:

- When running on an SGI system using the gcc UPC compiler, invoke TotalView upon your UPC program in the same way as any other program. For example:

```
totalview prog_upc -a args_to_foo_upc
```
- When running on HP Compaq SC machines, debug your UPC code in the same way that you would debug other kinds of parallel code. That is, invoke TotalView upon `prun`. For example:

```
totalview prun -a -n <node_count> prog_upc args_to_prog_upc
```

Viewing Shared Objects

Totalview displays UPC shared objects, and will fetch data from the UPC thread with which it has an affinity. For example, TotalView always fetches shared scalar variables from thread 0.

The upper-left figure in Figure 75 displays elements of a large shared array. You can manipulate an examine shared arrays the same as any other array. For example, you can slice, filter, obtain statistical information on, and so forth. (For more information on displaying array data, see Chapter 13, “*Examining Arrays*” on page 319) The bottom-right illustration shows a 10-element slice of this array.

In this illustration, the **Shared Address** area TotalView tells you that it is displaying the address of the array.

As the array is shared, it has an additional property: the element’s affinity. You can display this information if you select the Variable Window’s **View > Node Display**

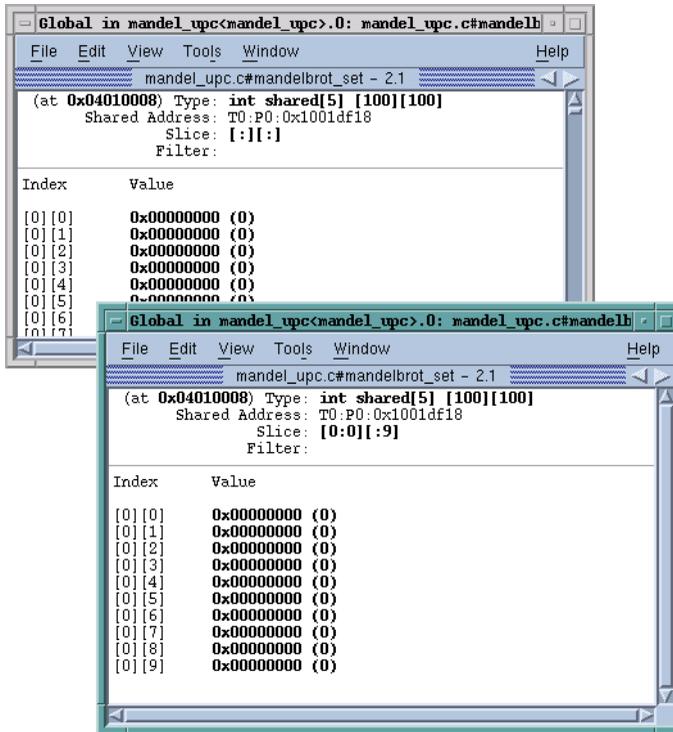


FIGURE 75: A Sliced UPC Array

command. This command tells TotalView to add a column that indicates the node associated with the value. This is shown in the Figure 76 on page 131.

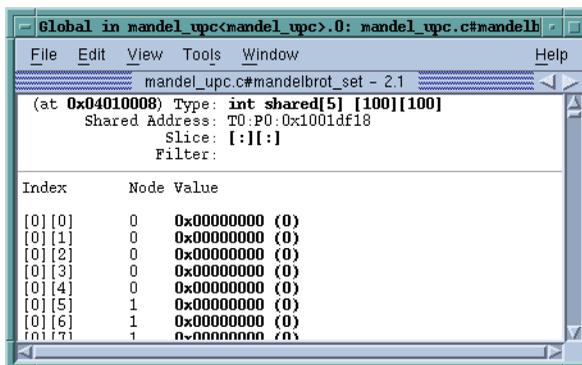


FIGURE 76: UPC Variable Window Showing Nodes

You can also use the **Tools > Visualize Distribution** to visualize this array. For more information on visualization, see "Using the Visualizer to Display Array Data" on page 163.

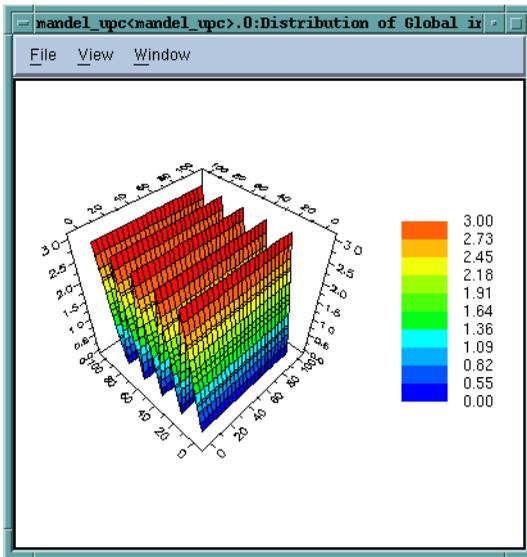


FIGURE 77: Laminated UPC Variable Window

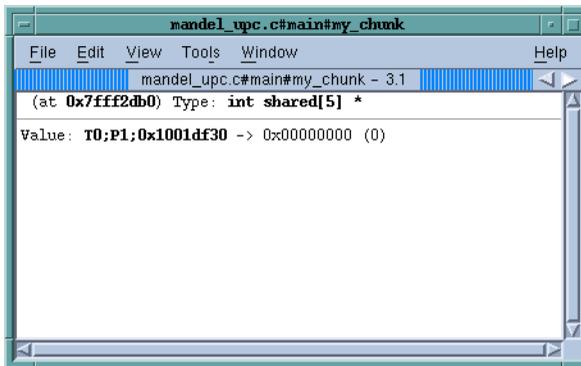
Pointer to Shared

TotalView understands pointer-to-shared data and displays the components of the data, as well as the target of the pointer to shared. For example, Figure 78 shows what is displayed:

Because the **Type** field shows the full type name, TotalView is telling you that this is a pointer to a shared **int** with a block size of 5.

In this figure, TotalView also displays the **upc_threadof ("T0")**, the **upc_phaseof ("P1")**, and the **upc_addrfield (0x0x1001df30)** components of this variable.

In the same way that TotalView normally shows the target of a pointer variable, it also shows the target of a UPC pointer variable. TotalView will fetch the target of the pointer from the UPC thread with which the pointer has affinity.

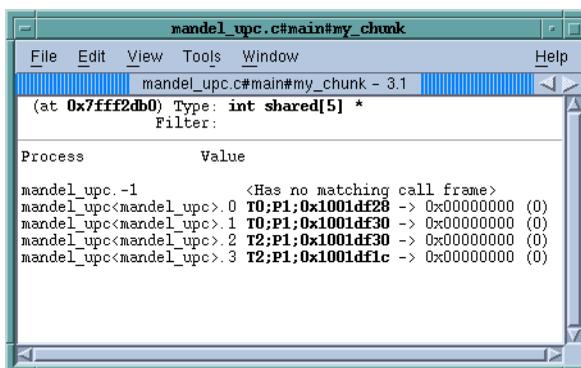
FIGURE 78: **Pointer to a Shared Variable**

You can update the pointer by selecting the pointer value and editing the thread, phase, or address values. If the phase is corrupt, you'll see something like the following in the **Value** field:

```
Value: T0;P6;0x3ffc0003b00 <Bad phase [max 4]> ->
        0xc0003c80 (-1073726336)
```

This example is indicating that the pointer is invalid because the phase is outside the legal range. It shows a similar message if the thread is invalid.

Since the pointer is not shared, you can use the **Tools > Laminate** command to display the value from each of the UPC threads.

FIGURE 79: **UPC Laminated Variable**

Parallel Debugging Tips

This section contains some information that you may find useful when debugging parallel programs. The topics in this section are:

- "Attaching to Processes" on page 134
- "General Parallel Debugging Tips" on page 137
- "MPICH Debugging Tips" on page 139
- "IBM PE Debugging Tips" on page 140

Attaching to Processes

In a typical multiprocess job, you're interested in what's occurring in some of your processes and not as much interested in others. By default, TotalView tries to attach to all the processes that you program starts. If there are a lot of processes, there may be considerable overhead involved in opening and communicating with the jobs. You can minimize this overhead by using the **Group > Attach Subsets** command, which displays the dialog box shown in Figure 80.

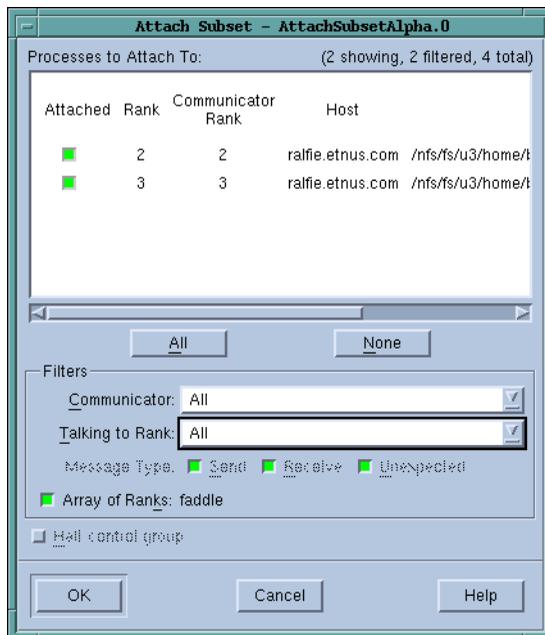


FIGURE 80: **Group > Attach Subset** Dialog Box

By selecting the boxes at the left side of the list, you tell TotalView which processes it should attach to. So, while your program will launch all of these processes, TotalView will only attach to the processes that you have selected here.

The four controls underneath the All and the None buttons let you limit what TotalView automatically attaches to.

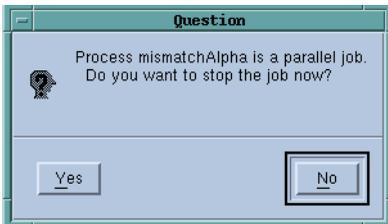
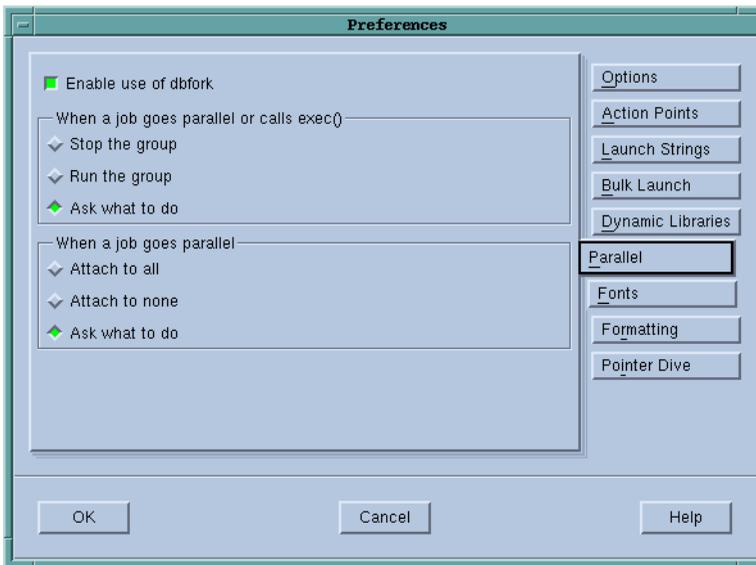
- The **Communicator** control specifies that the processes must be involved with the communicators that you select. For example, if something goes wrong that involves a communicator, selecting it from the list tells TotalView that it should only attach to the processes that use that communicator.
- The **Talking to Rank** control further limits the processes to those that you name here. Most of the entries in this list are just the process numbers. Two other entries are useful: **All** and **MPI_ANY_SOURCE**.
- The three checkboxes in the **Message Type** area add yet another qualifier. Checking a box tells TotalView that it should only display communicators that are involved with a **Send**, **Receive**, or **Unexpected** messages.

After you've found the problem, you can detach from these nodes by selecting **None**. In most cases, you would use the **All** button to set all the check boxes, then clear the ones that you're not interested in.

Many applications place the ranks numbers in a variable so they can be referred to easily. If you do this, you can display the variable in a Variable Window and then select the **Tools > Attach Subset (Array of Ranks)** command to display this dialog box

While you can use the **Group > Attach** command at any time, you would probably use it immediately before TotalView launches processes. Unless you have set preferences otherwise, TotalView will stop and ask if you want it stop your processes. When selected, the **Halt control group** check box also tells TotalView that it stop a process just before it begins executing. (See Figure 81 on page 136.)

The commands on the **Parallel** Page with the **File > Preferences** Dialog Box let you control what TotalView will do when your program goes parallel. (See Figure 82 on page 136.)

FIGURE 81: **Stop Before Going Parallel Question Dialog Box**FIGURE 82: **File > Preferences: Parallel Page**

The **When a job goes parallel or calls exec()** radio buttons have the following meanings:

- **Stop the group:** Stops the control group immediately after the processes are created.
- **Run the group:** Allows all newly created processes in the control group to run freely.
- **Ask what to do:** Asks what should occur. If you select this option, TotalView will ask if it should start the created processes.

CLI EQUIVALENT: **dset TV::parallel_stop**

The **When a job goes parallel** radio buttons have the following meaning:

- **Attach to all:** TotalView automatically attaches to all processes when they begin executing.
- **Attach to none:** TotalView will not attach to any created process when it begins executing.
- **Ask what to do:** Asks what should occur. If you select this option, TotalView opens the same dialog box that is displayed when you select **Group > Attach Subsets**. TotalView will then attach to the processes that you have selected. Note that this dialog box isn't displayed when you set the preference. Instead, it controls what will happen when your program creates parallel processes.

CLI EQUIVALENT: `dset TV::parallel_attach`

General Parallel Debugging Tips

Here are some tips that are useful for debugging most parallel programs:

■ Breakpoint behavior

When you're debugging message-passing and other multiprocess programs, it is usually easier to understand the program's behavior if you change the default stopping action of breakpoints and barrier breakpoints. By default, when one process in a multiprocess program hits a breakpoint, TotalView will stop all the other processes.

To change the default stopping action of breakpoints and barrier breakpoints, you can set TotalView preferences. Information on these preferences can be found in the online Help. These preferences tell TotalView if it should allow other processes and threads to continue to run when a process or thread hits the breakpoint.

These options only affect the default behavior. As usual, you can choose a behavior for a breakpoint by setting the breakpoint properties in the **File > Preferences's Action Points** Page. See "*Setting Breakpoints for Multiple Processes*" on page 346.

■ Process synchronization

TotalView has two features that make it easier to get all of the processes in a multiprocess program synchronized and executing a line of code.

Process barrier breakpoints and the process hold/release features work together to help you control the execution of your processes. See “*Barrier Points*” on page 350.

The Process Window’s **Group > Run To** command is a special kind of stepping command. It allows you to run a group of processes to a selected source line or instruction. See “*Stepping (Part I)*” on page 241.

■ Using group commands

Group commands are often more useful than process commands.

It is often more useful to use the **Group > Go** command to restart the whole application instead of the **Process > Go** command.

CLI EQUIVALENT: **dfocus g dgo**
Abbreviation: **G**

You would then use the **Group > Halt** command instead of **Process > Halt**.

CLI EQUIVALENT: **dfocus g dhalt**
Abbreviation: **H**

The group-level single-stepping commands such as **Group > Step** and **Group > Next** allow you to single-step a group of processes in a parallel. See “*Stepping (Part I)*” on page 241.

CLI EQUIVALENT: **dfocus g dstep**
Abbreviation: **S**
dfocus g dnext
Abbreviation: **N**

■ Process-level stepping

If you use a process-level single-stepping command in a multiprocess program, TotalView may appear to be hung (it continuously displays the watch cursor). If you single-step a process over a statement that can’t complete without allowing another process to run and that process is stopped, the stepping process appears to hang. This can occur, for example, when you try to single-step a process over a communication operation that cannot complete without the participation of another process. When this happens, you can abort the single-step operation by selecting **Cancel** in the **Waiting for Command to Complete** Window that

TotalView will display. As an alternative, consider using a group-level single-step command.

CLI EQUIVALENT: **Type Ctrl+C**

NOTE Etnus receives many bug reports about processes being hung. In almost all cases, the reason is that one process is waiting for another. Using the Group debugging commands almost always solves this problem.

■ Determining which processes and threads are executing

The TotalView Root Window helps you determine where various processes and threads are executing. When you select a line of code in the Process Window, the Root Window's **Attached** Page is updated to show which processes and threads are executing that line. See "Displaying Thread and Process Locations" on page 232.

■ Viewing variable values

You can view (laminare) the value of a variable that is replicated across multiple processes or multiple threads in a single Variable Window. See "Displaying a Variable in All Processes or Threads" on page 333.

■ Restarting from within TotalView

You can restart a parallel program at any time. If your program runs too far, you can kill the program by selecting the **Group > Delete** command. This command kills the master process and all the slave processes. Restarting the master process (for example, **mpirun** or **poe**) recreates all of the slave processes. Startup is faster when you do this because TotalView doesn't need to reread the symbol tables or restart its server processes since they are already running.

CLI EQUIVALENT: **dfocus g dkill**

MPICH Debugging Tips

Here are some debugging tips that apply only to MPICH:

■ Passing options to mpirun

You can pass options to TotalView through the MPICH **mpirun** command.

To pass options to TotalView when running **mpirun**, you can use the **TOTALVIEW** environment variable. For example, you can cause **mpirun** to invoke TotalView with the **-no_stop_all** option as in the following C shell, example:

```
setenv TOTALVIEW "totalview -no_stop_all"
```

■ Using `ch_p4`

If you start remote processes with `MPICH/ch_p4`, you may need to change the way TotalView starts its servers.

By default, TotalView uses `rsh` to start its remote server processes. This is the same behavior as `ch_p4`. If you configure `MPICH/ch_p4` to use a different start-up mechanism from another process, you will probably also need to change the way that TotalView starts the servers.

For more information about `tvdsrv` and `rsh`, see “*Setting Single-Process Server Launch Options*” on page 74. For more information about `rsh`, see “*Using the Single-Process Server Launch Command*” on page 80.

IBM PE Debugging Tips

Here are some debugging tips that apply only to IBM MPI (PE):

■ Avoid unwanted timeouts

Timeouts can occur if you place breakpoints that stop other processes too soon after calling `MPI_Init()` or `MPL_Init()`. If you create “stop all” breakpoints, the first process that gets to the breakpoint stops all the other parallel processes that have not yet arrived at the breakpoint. This may cause a timeout.

To turn the option off, select the Process Window’s **Action Point > Properties** command while the line with the stop symbol is selected. After the **Properties Dialog Box** appears, you should select the **Process** button in the **When Hit, Stop** area and also select **Plant in share group**.

CLI EQUIVALENT: `dbarrier location -stop_when_hit process`

■ Control the `poe` process

Even though the `poe` process continues under TotalView control, you should not attempt to start, stop, or otherwise interact with it. Your parallel tasks require that `poe` continue to run. For this reason, if `poe` is stopped, TotalView automatically continues it when you continue any parallel task.

■ Avoid slow processes due to node saturation

If you try to debug a PE program in which more than three parallel tasks run on a single node, the parallel tasks on each node may run noticeably slower than they would run if you were not debugging them.

In general, the number of processes you are running on a node should be the same as the number of processors in the node.

This becomes more noticeable as the number of tasks increases, and, in some cases, the parallel tasks may make hardly any progress. This is because PE uses the **SIGALRM** signal to implement communications operations, and AIX requires that debuggers must intercept all signals. As the number of parallel tasks on a node increases, TotalView becomes saturated and can't keep up with the **SIGALRM**s being sent, thus slowing down the tasks.



Part III: Using the GUI

The two chapters in this part of the Users Guide only contain information that you'll need if you're using TotalView's GUI.

Chapter 6: Using TotalView's Windows

Describes using the mouse and the more important Windows.

Chapter 7: Visualizing Programs and Data

Some of TotalView's commands and tools are only useful if you're using the GUI. For example, the Visualizer graphically displays an array's data.



Chapter 6

Using TotalView's Windows

This chapter introduces you to the TotalView interface and describes:

- *"Using the Mouse Buttons"* on page 145
- *"Using the Root Window"* on page 146
- *"Using the Process Window"* on page 150
- *"Diving into Objects"* on page 152
- *"Resizing and Positioning Windows and Dialog Boxes"* on page 155
- *"Editing Text"* on page 156
- *"Saving the Contents of Windows"* on page 157

Using the Mouse Buttons

TotalView uses the buttons on your three-button mouse as follows:

Table 7: Mouse Button Functions

Button	Action	Purpose	How to Use It
Left	Select	Selects or edits object, scrolls in windows and panes	Move the cursor over the object and click the button.
Middle	Paste	Writes information previously copied or cut into the clipboard	Move the cursor to where you will be inserting the information and click the button; not all windows support pasting.
	Dive	Displays more information or replaces window contents	Move the cursor over an object, then click the middle mouse button.

Table 7: Mouse Button Functions (cont.)

Button	Action	Purpose	How to Use It
Right	Context menu	Displays a menu with commonly used commands	<p>Move the cursor over an object and click the button.</p> <p>Most windows and panes have context menus; dialog boxes do not have context menus.</p>

In most cases, a single-click selects what's under the cursor and a double-click dives on the object. However, if the field is editable, TotalView goes into its edit mode where you can alter the selected item's value.

In some places such as the Stack Trace Pane, selecting a line tells TotalView that it should perform an action. In this pane, TotalView dives on the selected routine. (In this case, *diving* means that TotalView finds the selected routine and show it in the Source Pane.)

In the line number area of the Source Pane, a left mouse click sets a breakpoint at that line. TotalView shows you that it has set a breakpoint by displaying a **STOP** icon instead of a line number.

Selecting the **STOP** icon a second time deletes the breakpoint. If you change any of the breakpoint's properties are if you had created an evaluation point—this is indicated by an **EVAL** icon—selecting the icon disables it. For more information on breakpoints and evaluation points, refer to Chapter 14, "Setting Action Points" on page 337.

Using the Root Window

The Root Window appears when you start TotalView. If you do not enter a program name when starting TotalView, it is the only window that appears. If you indicate a program name, TotalView also open a Process Window containing the program's source code.

The *Root Window* contains the following four tabbed pages:

- **Attached:** Displays a list of all the processes and threads being debugged. Initially—that is, before your program begins executing—the Root Window just

contains the name of the program being debugged. As processes and threads are created, TotalView adds them to this list. Associated with each is a name, location (if a remote process), process ID, status, and a list of executing threads for each process. It also shows the thread ID, status, and the routine being executed in each thread.

Figure 83 shows the **Attached** Page for an executing multithreaded multiprocess program.

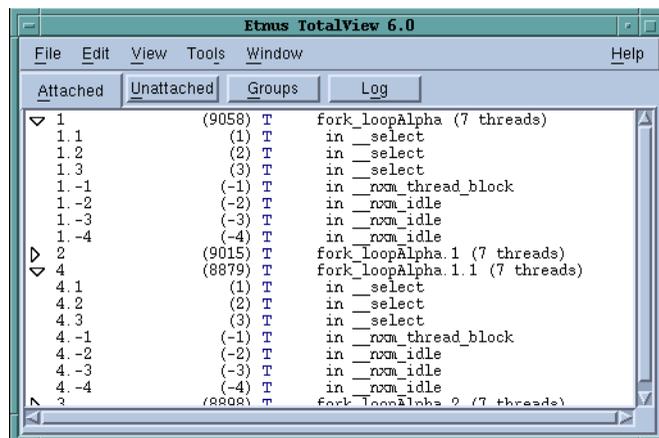


FIGURE 83: **Root Window Attached Page**

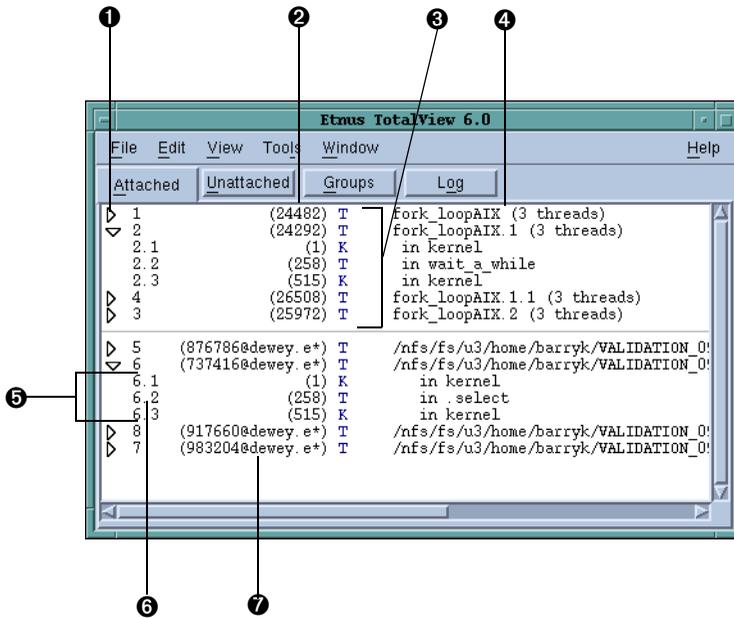
Notice the triangles on the left. If a triangle is pointing right, you can click on it to display the process's threads. If it is pointing down, clicking on it conceals this thread information.

When debugging a remote process, TotalView displays an abbreviated version of the host name on which the process is running in brackets ([]). The full host name appears in brackets in the title bar of the Process Window. In Figure 84 on page 148, the process is running on the machine **dewey.etnus.com**. This name is abbreviated in the Root Window. This figure also describes the contents of the columns in this window.

When you dive on a line in this window, TotalView displays the source for the process or thread in a Process Window.

- **Unattached:** Displays processes over which you have control. If you can't attach to one of these processes, TotalView displays it in gray. Figure 85 on page 148 shows the **Unattached** Page.

Diving on processes in this pane tells TotalView to attach to them.



- ❶ Collapse/expand toggle
- ❷ Process ID (PID)
- ❸ Thread status
- ❹ Program name
- ❺ Thread list
- ❻ Thread ID (TID/SYSTID)
- ❼ Remote process location

FIGURE 84: Root Window Showing Remote

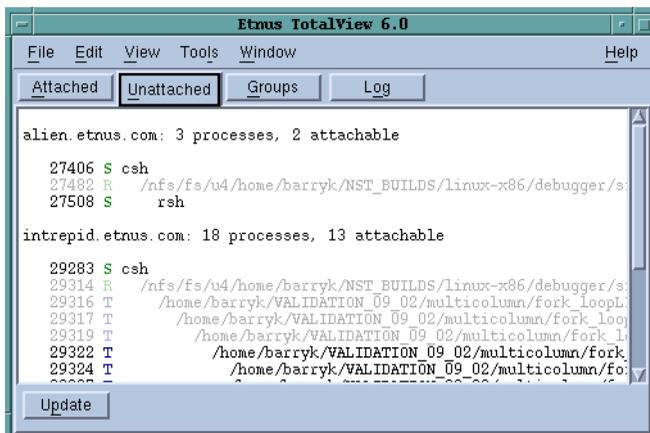


FIGURE 85: Root Window Unattached Page

- Groups:** Lists the groups used by your program. The top pane in Figure 86 lists all of your program's groups. This list includes all the groups that TotalView creates and all that you create using the CLI. When you select a group in the top pane, the group's members are displayed in the bottom pane.

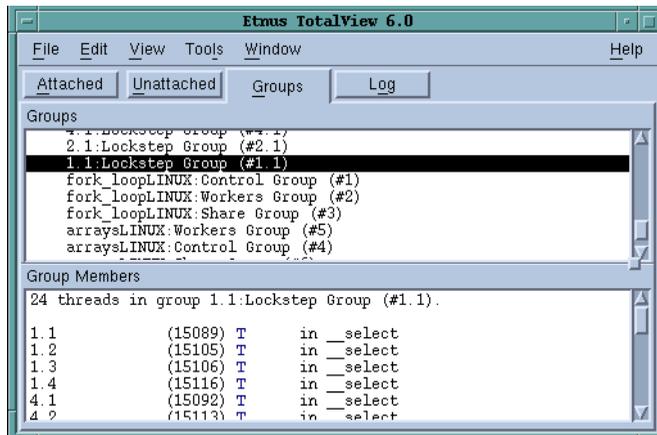


FIGURE 86: Root Window Groups Page

- Log:** Contains a log of the debugging actions. (See Figure 87.) This information is sometimes useful when analyzing the behavior of misbehaving multiprocess/multithreaded programs.

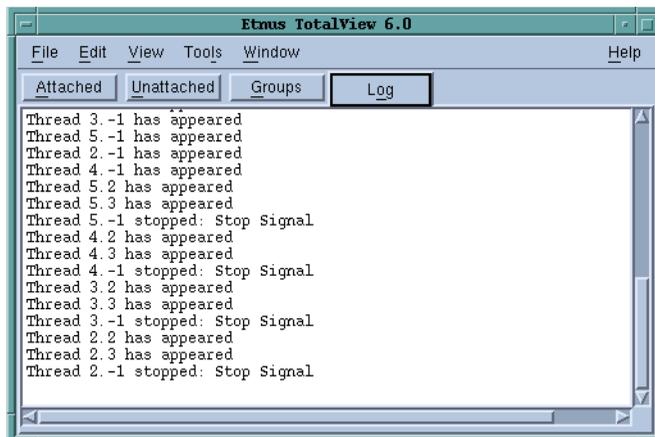


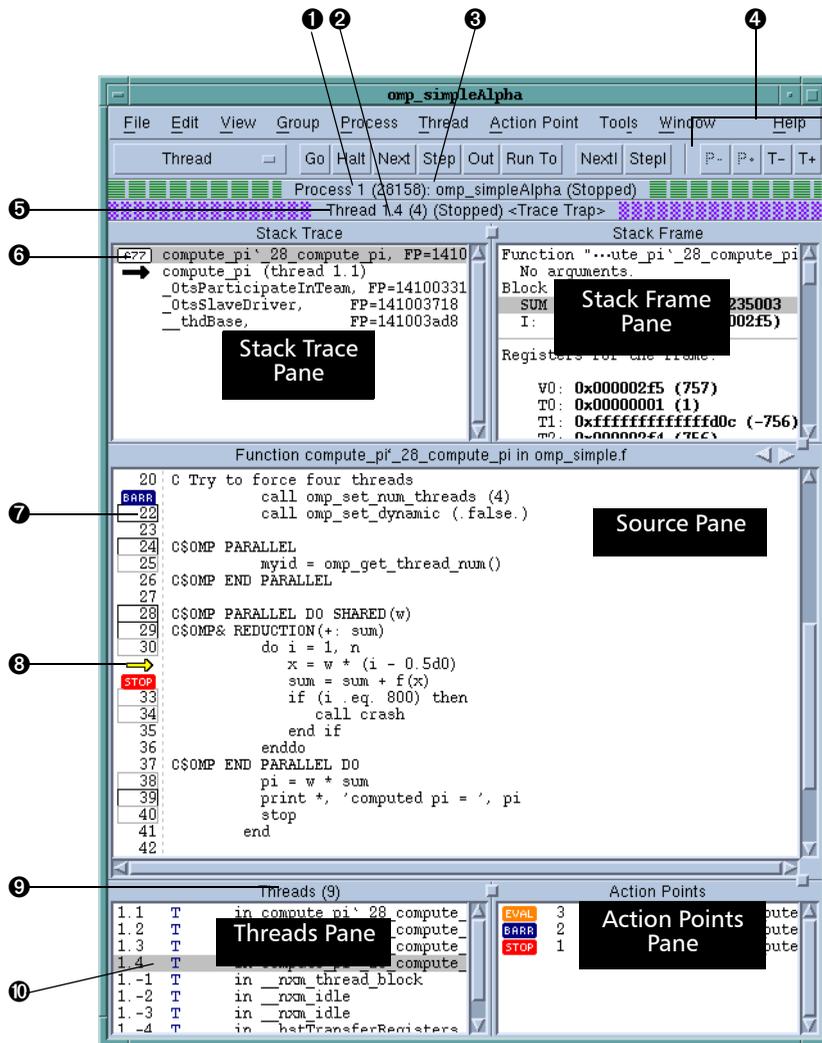
FIGURE 87: Root Window Log Page

Using the Process Window

The *Process Window*, which is shown in Figure 88, contains the code for the process or thread you're debugging, as well as other related information. This window contains five *panes* of information. The large scrolling list in the middle of the Process Window is the Source Pane. (The contents of these panes are discussed later in this section.)

As you examine the Process Window, notice the following:

- The thread ID shown in the Root Window and in the process's Threads Pane is the TotalView-assigned logical thread ID (TID) and system-assigned thread ID (SYSTID). On systems such as HP Alpha Tru64 UNIX where the TID and SYSTID values are the same, TotalView displays only the TID value.
In other windows, TotalView uses the value *pid.tid* to identify a process's threads. The *Threads Pane* shows the list of threads that currently exist in the process. When you select a different thread in this list, TotalView updates the Stack Trace Pane, Stack Frame Pane, and Source Pane to show the information for that thread. When you dive on a different thread in the thread list, TotalView finds or opens a new window displaying information for that thread.
The number in the Threads Pane title (🍎 in Figure 88 on page 151) is the number of threads that currently exist in the process.
- The *Stack Trace Pane* shows the call stack of routines that the selected thread is executing. You can move up and down the call stack by clicking on the routine's name (stack frame). When you select a different stack frame, TotalView updates the Stack Frame and Source Panes to show the information about the routine you just selected.
- The *Stack Frame Pane* displays all of a routine's parameters, its local variables, and the registers for the selected stack frame.
- The information displayed in the Stack Trace and Stack Frame Panes reflects the state of the process when it was last stopped. This means that the information they are displaying is not up-to-date while the thread is running.
- The left margin of the *Source Pane* displays line numbers and action point icons. You can place a breakpoint at any line whose line number is contained within a box. (See Figure 89.) The box indicates that executable code was created by the source code.



- | | |
|-----------------------|-----------------------|
| ① Process status | ⑥ Language of routine |
| ② Thread ID (TID) | ⑦ Line number area |
| ③ Process ID (PID) | ⑧ Current PC |
| ④ Navigation controls | ⑨ Thread count |
| ⑤ Thread status | ⑩ Selected thread |

FIGURE 88: Process Window

When you place a breakpoint on a line, TotalView places a STOP icon over the line number. An arrow over the line number shows the current location of the program counter (PC) within the selected stack frame. See Figure 89.

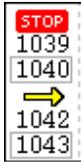


FIGURE 89: **Line Numbers, with Stop Icon and PC Arrow**

Each thread has its own unique program counter (PC). When you stop a multiprocess or multithreaded program, the routine displayed in the Stack Trace Pane for a thread depends on the thread's PC. Because threads execute asynchronously, you'll usually find that threads are stopped at different places. (When your thread hits a breakpoint, the TotalView default is to stop all the other threads in the process as well.)

- The *Action Points Pane* shows the list of breakpoints, evaluation points, and watchpoints for the process.

Diving into Objects

Diving, which is clicking your middle mouse button on something in a TotalView window, is one of TotalView's more distinguishing features.

NOTE In some cases, single-clicking preforms a dive. For example, single-clicking on a function name in the Stack Trace Pane tells TotalView to dive into the function. In other cases, double-clicking does the same thing. While this may sound confusing, it's pretty intuitive and you'll be diving without thinking almost instantaneously.

For example, diving on processes and threads in the Root Window is the quickest way to display a Process Window that contains information about what you're diving on. The procedure is simple: dive on a process or thread and TotalView takes care of the rest. Here's another example: diving on variables in the Process Window tells TotalView to display information about the variable in a Variable Window.

Table 8 describes typical diving operations.

Table 8: Diving

Dive on:	Information Displayed by Diving:
Process or thread	When you dive on a thread in the Root Window, TotalView finds or opens a Process Window for that process. If it doesn't find a matching window, TotalView replaces the contents of an existing window and shows you the selected process.
Subroutine	The source code for the routine replaces the current contents of the Source Pane—this is called a nested dive. When this occurs TotalView places a right angle bracket (>) in the process's title. Every time it dives, it adds another angle bracket. See Figure 90, which follows this table. A routine must be compiled with source-line information (usually, with the -g option) for you to dive into it and see source code. If the subroutine wasn't compiled with this information, TotalView displays the routine's assembler code.
Variable	The contents of the variable appear in a separate Variable Window.
Pointer	TotalView dereferences the pointer and shows the result in a separate Variable Window. Given the nature of pointers, you may need to cast the result into something that is more to your liking.
Array element, structure element, or referenced memory area	The contents of the element or memory area replaces the contents that were in the Variable Window—this is known as a <i>nested</i> dive.
Routine in the Stack Trace Pane	The stack frame and source code for the routine appear in a Process Window.

Function >>>pthread_mutex_lock

FIGURE 90: **Nested Dive**

TotalView tries to reuse windows whenever possible. For example, if you dive on a variable and that variable is already being displayed in a window, TotalView pops the window to the top of the display. If you want the information to appear in a separate window, use the **View > Dive Anew** command.

NOTE Using on a process or a thread may not create a new window if TotalView determines that it can reuse a Process Window. If you really want to see information in two windows, use the Process Window's **Window > Duplicate** command.

When you dive into functions in the Process Window or when you are chasing pointers or following structure elements in the Variable Window, you can move back and forth between your selections by using the *forward* and *backward* icons. The location of the two controls is shown in the boxed area in Figure 91.

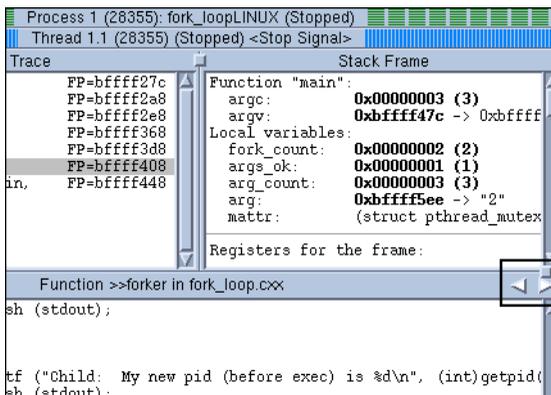


FIGURE 91: **Backward and Forward Buttons**

For additional information about displaying variable contents, refer to “*Diving in Variable Windows*” on page 291.

Other windowing commands that you can use are:

- **Window > Duplicate:** (Variable Window) Creates a duplicate copy of the current Variable Window.
- **Window > Duplicate Base:** (Variable Window) Creates a copy of the current Variable Window. Unlike what happens when you use the **Window > Duplicate** command, this command retains the dive stack.
- **File > Close:** Closes an open window.
- **File > Close Relatives:** Closes windows that are related to the current window. The current window isn't closed.
- **File > Close Similar:** Closes the currently open window and all windows similar to it. When you have lots of similar windows, this is a great time-saver.

Resizing and Positioning Windows and Dialog Boxes

You can resize most of TotalView's windows and dialog boxes. While TotalView tries to do the right thing, you can push things to the point where shrinking doesn't work very well. Figure 92 shows a before and after look where a dialog box was made too small.

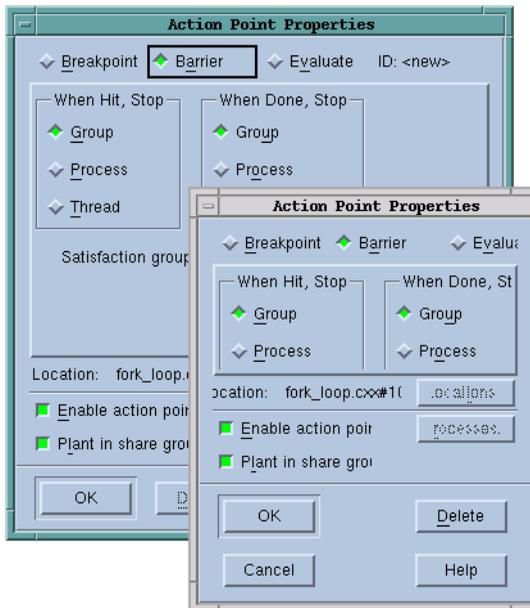


FIGURE 92: **Resizing (and Sometimes Its Consequences)**

Many programmers like to have their windows always appear in the same position in each session. TotalView has two commands that can help:

- **Window > Memorize:** Tells TotalView it should remember the position of the current window. The next time you bring up this window, it'll be in this position.
- **Window > Memorize All:** Tells TotalView it should remember the positions of just about all of its windows. The next time you bring up any of the windows displayed when you had used this command, it will be in the same position.

Most modern window managers such as KDE or Gnome do an excellent job managing window position. If you are using an older window manager such as **twm** or **mwm**, you may want to select the **Force window positions (disables window**

manager placement modes) check box option located on the **Options** Page of the **File > Preferences** Dialog Box. This tells TotalView to manage a window's position and size. If it isn't selected, TotalView only manages a window's size.

Editing Text

The TotalView field editor lets you change the values of fields in windows or change text fields in dialog boxes. To edit text:

- 1 Click the left mouse button to select the text you wish to change. If you can edit the selected text, it appears within a highlighted rectangle, and you will see an editing cursor. (See Figure 93.)

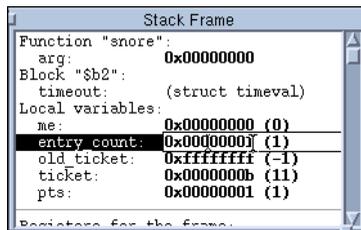


FIGURE 93: Editing Cursor

- 2 Edit the text and press Return.

Like other Motif-based applications, you can use your mouse to copy and paste text within TotalView and to other X Windows applications by using your mouse buttons.

You can also manipulate text by using **Edit > Copy**, **Edit > Cut**, **Edit > Paste**, and **Edit > Delete**.

In most cases, clicking your middle mouse button tells TotalView to dive. However, if TotalView is displaying an editing cursor, clicking your middle mouse button tells TotalView to paste information.

Saving the Contents of Windows

You can write an ASCII equivalent to most pages and panes by using the **File > Save Pane** command. This command also lets you pipe data to UNIX shell commands. (See Figure 94.)

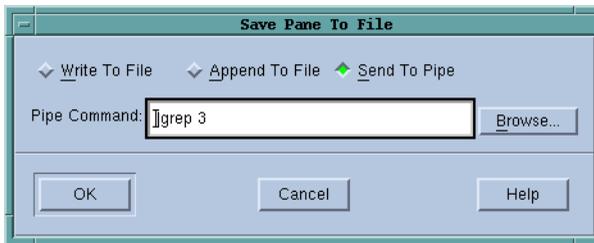


FIGURE 94: **File > Save Pane Dialog Box**

When piping information, TotalView sends what you've typed to `/bin/sh`. This means that you can enter a series of shell commands. For example, here is a command that ignores the top five lines of output, compares the current ASCII text to an existing file, and writes the differences to another file:

```
| tail +5 | diff - file > file.diff
```


Chapter 7

Visualizing Programs and Data

TotalView provides a set of tools that allow you to visualize how your program is performing and the values of variables. This chapter describes:

- *"Displaying Your Program's Call Tree"* on page 159
- *"Displaying Memory Statistics"* on page 161
- *"Using the Visualizer to Display Array Data"* on page 163

Other visualization tools are described in the following sections:

- *"Using the P/T Set Browser"* on page 271
- *"Displaying the Message Queue Graph"* on page 107

Displaying Your Program's Call Tree

Debugging is an art, not a science. Debugging often means having the "intuition" to know what a problem means and where to look for it. Locating a problem is often 90% or more of the effort. TotalView's call tree is one tool that helps you get an understanding of what your program is doing so that you can begin to understand how your program is executing.

Use the **Tools > Call Tree** command in the Process Window to tell TotalView to display a Call Tree Window. (See Figure 95 on page 160.)

The call tree is a diagram showing all the currently active routines. These routines are linked by arrows indicating that one routine is called by another. TotalView's call tree is a *dynamic* call tree in that it displays the call tree at the time when TotalView creates it. The **Update** button tells TotalView to recreate this display.

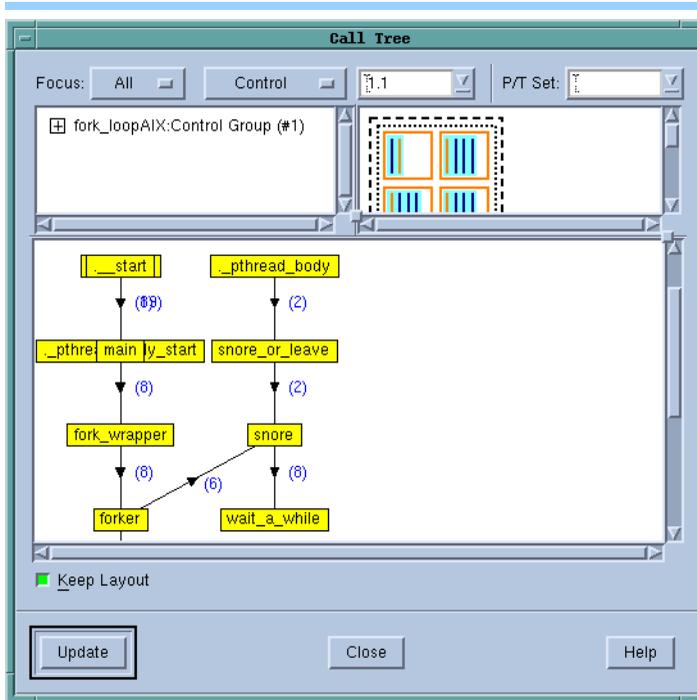


FIGURE 95: **Tools > Call Tree Dialog Box**

NOTE You'll find information on using the P/T Set Controls in the top portion of this window in Chapter 11, *"Using Groups, Processes, and Threads"* on page 239.

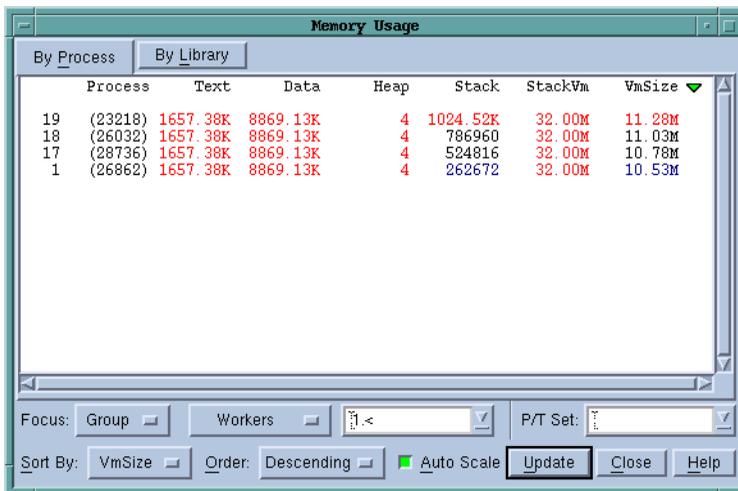
You can tell TotalView to display a call tree for the processes and threads specified with the controls at the top of this window. If you don't touch these controls, TotalView displays a call tree for the group defined in the icon bar of your Process Window. If TotalView is displaying the call tree for a multiprocess or multithreaded program, numbers next to the arrows indicate how many times a routine is on the call stack.

As you begin to understand your program, you will see that it has a rhythm and a dynamic that is reflected in this diagram. As you examine and understand this structure, you will sometimes see things that don't look right—which is a subjective response to how your program is operating. These places are often where you want to begin looking for problems.

Looking at the call tree can also tell you where bottlenecks are occurring. For example, if one routine is used by many other routines and that routine controls a shared resource, this thread may be negatively affecting performance. For example, in Figure 95, the **snore** routine might be a bottleneck. Giving good names to routines helps. For example, if you see lots of routines in a routine named **snore**, you've probably designed things to that routines will be waiting there, so this wouldn't represent a problem.

Displaying Memory Statistics

The **Tools > Memory Usage** Window tells you how your program is using memory and where this memory is being used. The best way to use this window is to compare memory use over time so that you can tell if your program is leaking memory. If a program is leaking memory, you'll see that the amount of memory being used steadily increases over time. Figure 96 shows the By Process Pane of the Memory Usage window.



The screenshot shows the 'Memory Usage' window with the 'By Process' pane selected. The table displays memory usage for four processes, with columns for Process, Text, Data, Heap, Stack, StackVm, and VmSize. The VmSize column shows values in red (11.28M, 11.03M, 10.78M) and blue (10.53M).

Process	Text	Data	Heap	Stack	StackVm	VmSize
19 (23218)	1657.38K	8869.13K	4	1024.52K	32.00M	11.28M
18 (26032)	1657.38K	8869.13K	4	786960	32.00M	11.03M
17 (28736)	1657.38K	8869.13K	4	524816	32.00M	10.78M
1 (26862)	1657.38K	8869.13K	4	262672	32.00M	10.53M

FIGURE 96: **Tools > Memory Usage** Window

When you see this information within TotalView, the maximum value of an item is displayed in **red** and the minimum value is displayed in **blue**. This information spec-

ifies the amount of memory used by the executable's text and data segments, and the TotalView process IDs.

You can change the sort order by clicking on a column's header or by using the controls at the bottom of the window.

Notice that if you add the memory values of all columns but the last, the sum doesn't equal this last column's value. This is because most operating systems divide segments into pages, and information in a segment doesn't cross page boundaries.

CLI EQUIVALENT: **dmstat**

Here's the definition for most of these columns:

Text	The amount of memory needed to store your program's machine code instructions.
Data	The amount of memory required to store initialized data.
Heap	The amount of memory currently being used for data created at runtime.
Stack	The amount of memory used by the currently executing routine and all the routines that have invoked it.
StackVm	The logical size of the stack is the difference between the current value of the stack pointer and address from which the stack originally grew. This value can differ from the size of the virtual memory mapping in which the stack resides. This value is that size difference.
VmSize	The sum of the sizes of the mappings in the process's address space.

NOTE The online Help has more information.

Using the P/T Set controls at the top of the window is discussed in Chapter 11, "Using Groups, Processes, and Threads" on page 271.

The **By Library** Pane (shown in Figure 97 on page 163) shows which library files are contained within your executable.

Library	Text	Data	Processes
...barryk/tests/forked_memAIX	4098	8389170	1 17 18 19
/usr/ccs/bin/usla	40104	0	1 17 18 19
/usr/lib/libcrypt.a	800	440	1 17 18 19
/usr/lib/libc.a	1652160	692384	1 17 18 19

FIGURE 97: **Tools > Memory Usage Window: By Library Pane**

Clicking on any of the columns tells TotalView to sort the values in that column. You can change the order from ascending to descending (or descending to ascending) by clicking a second time on the column heading. If you're so inclined, there are commands for sorting these columns in the menubar.

Using the Visualizer to Display Array Data

The TotalView Visualizer creates graphic images of your program's array data.

NOTE The Visualizer isn't available on Linux Alpha and 32-bit SGI Irix.

Topics in this section are:

- "How the Visualizer Works" on page 164
- "Configuring TotalView to Launch the Visualizer" on page 165
- "Visualizing Data Manually" on page 168
- "Visualizing Data Programmatically" on page 169
- "Using the Visualizer" on page 170
- "Using the Graph Window" on page 173

- “Using the Surface Window” on page 177
- “Launching the Visualizer from the Command Line” on page 180

How the Visualizer Works

The Visualizer is a stand-alone program that is integrated with TotalView. This relationship gives you a lot of flexibility:

- If you launch the Visualizer from within TotalView, you can visualize your program’s data as you are debugging your program.
- You can save the data that would be sent to the Visualizer, and then invoke the Visualizer from the command line and have it read this previously written data. (See Figure 98.)

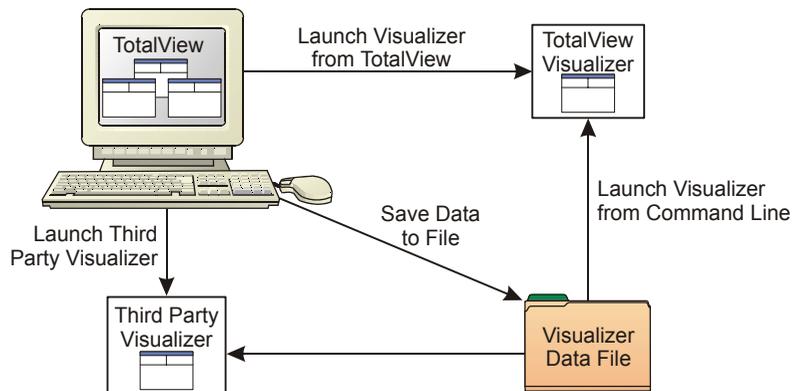


FIGURE 98: **TotalView Visualizer Relationships**

- Because TotalView is sending a data stream to the Visualizer, you can even replace our Visualizer with any tool that can read this data.

NOTE The online Help contains information on adapting a third-party visualizer so that it can be used with TotalView.

Visualizing your program’s data is a two step process:

- 1 You select the data that you want visualized.
- 2 You tell the Visualizer how it should display this data.

TotalView marshals the program’s data and pipes it to the Visualizer. The Visualizer reads this data and displays it for analysis. (See Figure 99 on page 165.)

TotalView: Extracts data from an array

TotalView Visualizer: Displays the array data graphically

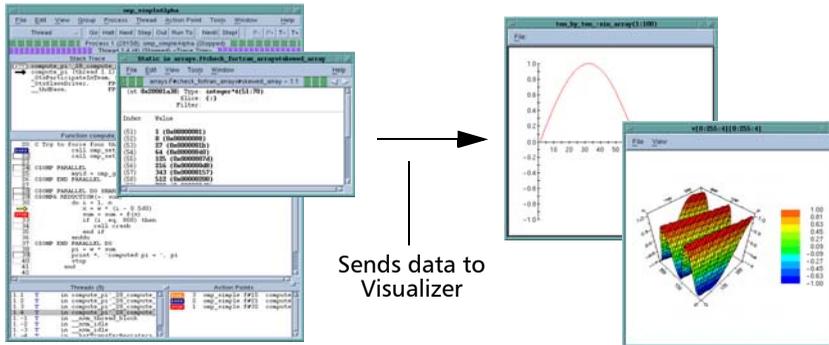


FIGURE 99: TotalView Visualizer Connection

Configuring TotalView to Launch the Visualizer

TotalView launches the Visualizer when you select the **Tools > Visualize** command from the Variable Window. It will also launch it if or when you use a `$visualize` function within an evaluation point and the **Tools > Evaluate** Dialog Box.

TotalView lets you set a preference that disables visualization. This lets you turn off visualization when your program executes code containing evaluation points, without having to individually disable all the evaluation points.

To change the Visualizer launch options interactively, select **File > Preferences**, and then select the **Launch Strings** Tab. (See Figure 100 on page 166.)

Using the commands in this page, you can:

- Customize the command TotalView uses to start a visualizer by entering the visualizer's startup command in the **Command** edit box. Entering information in this field is discussed a little later in this section.
- Change the autolaunch option. If you want to disable visualization, clear the **Enable Visualizer Launch** check box.
- Change the maximum permissible rank. Edit the value in the **Maximum array rank** edit field to save the data exported from the debugger or display it in a different visualizer. A rank's value can range from 1 to 16.

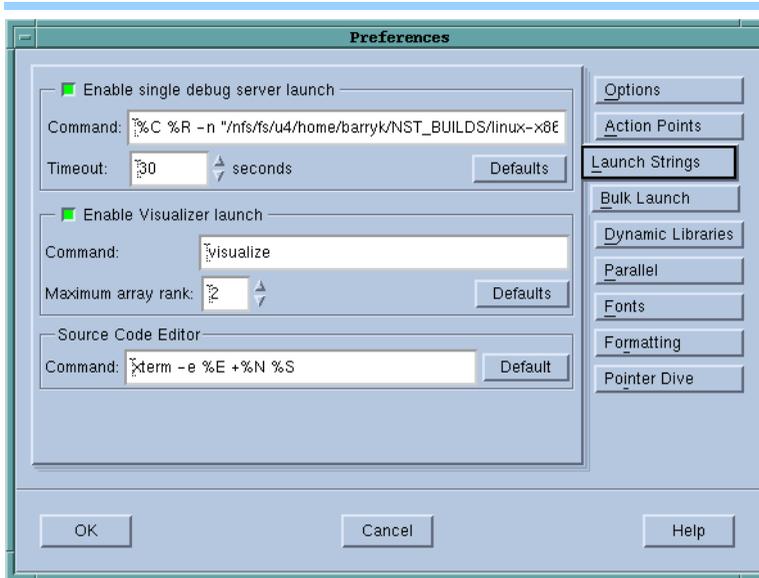


FIGURE 100: **File > Preferences Launch Strings Page**

Setting the maximum permissible rank to either 1 or 2 (the default is 2) ensures that the TotalView Visualizer can use your data—the Visualizer displays only two dimensions of data. This limit doesn't apply to data saved in files or to third-party visualizers that can display more than two dimensions of data.

- Clicking on the **Defaults** button returns all values to their defaults. This reverts options to their defaults even if you have used X resources to change them.

If you disable visualization while the Visualizer is running, TotalView closes its connection to the Visualizer. If you reenables visualization, TotalView launches a new Visualizer process the next time you visualize something.

Visualizer Launch Command

You can change the shell command that TotalView uses to launch the Visualizer by editing the Visualizer launch command. (In most cases, the only reasons you'd do this is if you're having path problems or you're running a different visualizer.) You could also change what's entered here so that you can view this information at another time. Here's an example:

```
cat > your_file
```

Later, you can visualize this information using either of the following commands:

```
visualize -persist < your_file
visualize -file your_file
```

You can preset the Visualizer launch options by setting X resources. These resources are described on our Web site. For more information, go to www.etnus.com/Support/docs/.

Data Types That TotalView Can Visualize

The data selected for visualization is called a *dataset*. Each dataset is tagged with a TotalView-generated numeric identifier that lets the Visualizer know whether it is seeing a new dataset or an update to an existing dataset. TotalView treats stack variables at different recursion levels or call paths as different datasets.

TotalView can visualize one- and two-dimensional arrays of character, integer, or floating-point data. If an array has more than two dimensions, you can visualize part of it using an array slice that creates a subarray having fewer dimensions.

Figure 101 shows a three-dimensional variable sliced into two dimensions by selecting a single index in the middle dimension.

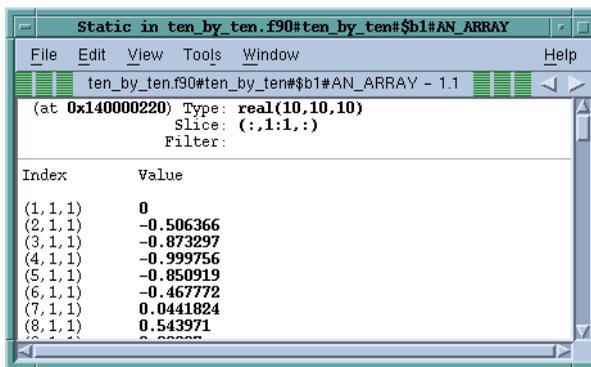


FIGURE 101: A Three-Dimensional Array Sliced into Two Dimensions

Viewing Data

Different datasets can require different views to display their data. For example, a graph is more suitable for displaying one-dimensional datasets or two-dimensional

datasets if one of the dimensions has a small extent; however, a surface view is better for displaying a two-dimensional dataset.

When TotalView launches the Visualizer, one of the following actions will occur:

- If a Data Window is currently displaying the dataset, the Visualizer raises it to the top of the desktop. If the window was minimized, the Visualizer restores it.
- If you haven't visualized the dataset in this session, the Visualizer chooses a method based on how well your dataset matches what is best shown for each kind of visualization method. You can enable and disable this feature from the **Options** menu in the Visualizer's Directory Window.
- If you've previously visualized a dataset but you've killed its window, the Visualizer creates a new Data Window by using the most recent visualization method.

Visualizing Data Manually

Before you can visualize an array, you must:

- Open a Variable Window for the array's data.
- Stop program execution where the array's values are set to what you want them to be when they are visualized.

Figure 102 shows an Variable Window containing an array.

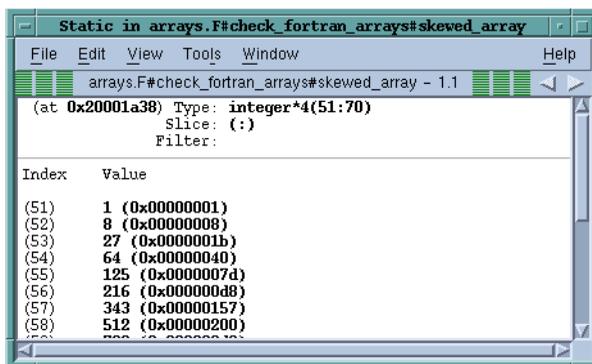


FIGURE 102: Variable Window

You can restrict the data being visualized by editing the **Type** and **Slice** fields. For example, editing the **Slice** fields limits the amount of data being visualized. (See

"*Displaying Array Slices*" on page 319.) Limiting the amount increases the Visualizer's speed.

After selecting the Variable Window's **Tools > Visualize** command, the Visualizer begins executing and then creates its window. The data sent to the Visualizer isn't automatically updated as you step through your program. Instead, you must explicitly update the display by reentering the **Tools > Visualize** command.

TotalView can visualize laminated variables. (See "*Visualizing a Laminated Variable Window*" on page 336.) The process or thread index will be one of the visualized data's dimensions. This means that you can only visualize scalar or vector information. If you don't want the process or thread index to be a dimension, use a nonlaminated display.

Visualizing Data Programmatically

TotalView's **\$visualize** function allows you add visualization to expressions in evaluation action points or with expressions entered in the **Tools > Evaluate** Window. If you enter this function within an expression, TotalView will interpret rather than compile the expression, which can greatly decrease performance. See "*Defining Evaluation Points and Conditional Breakpoints*" on page 354 for information about compiled and interpreted expressions. Adding this function also lets you visualize several different variables from a single expression or evaluation point.

Using **\$visualize** in an evaluation point lets you animate the changes that occur in your data because the Visualizer will update the array's display every time TotalView reaches the evaluation point. Here's this function's syntax:

```
$visualize ( array [, slice_string ])
```

The *array* argument names the dataset being visualized. The optional *slice_string* argument is a quoted string defining a constant slice expression that modifies the *array* parameter's dataset.

Here are six examples showing how you can use this function. Notice that the array's dimension ordering differs in C and in Fortran.

```

C          $visualize (my_array);
          $visualize (my_array, "[::2][10:15]");
          $visualize (my_array, "[12][:]");

Fortran   $visualize (my_array)
          $visualize (my_array, '(11:16,:::2)')
          $visualize (my_array, '(:,13)')

```

The first example in each programming language group visualizes the entire array. The second example selects every second element in the array's major dimension; it also clips the minor dimension to all elements in the range. The third example reduces the dataset to a single dimension by selecting one subarray.

You may need to cast your data so that TotalView will know what the array's dimensions are. Here's a C function declaration that passes a two-dimensional array parameter. Notice that it does not specify the major dimension's extent.

```

void my_procedure (double my_array[][32])
{ /* procedure body */ }

```

Here's how you can cast the array so that TotalView can visualize it. For example:

```

$visualize (*(double[32][32]*)my_array);

```

Sometimes, it's hard to know what to specify. You can quickly refine array and slice arguments, for example, by entering **\$visualize** into the **Tools > Evaluate** Dialog Box. When you select the **Evaluate** button, you'll quickly see the result. You can even use this technique to display several arrays simultaneously.

Using the Visualizer

The Visualizer uses two types of windows:

■ Data Windows

These are the windows that display your data. The commands in a Data Window let you set viewing options and change the way the Visualizer displays your data.

■ A Directory Window

This window lists the datasets that you can visualize. Use this window to set global options and to create views of your datasets. Commands in this window let you obtain different views of the same data by opening more than one Data Window.

The top window in Figure 103 on page 171 is a Directory Window. The two remaining windows show a surface and a graph view.

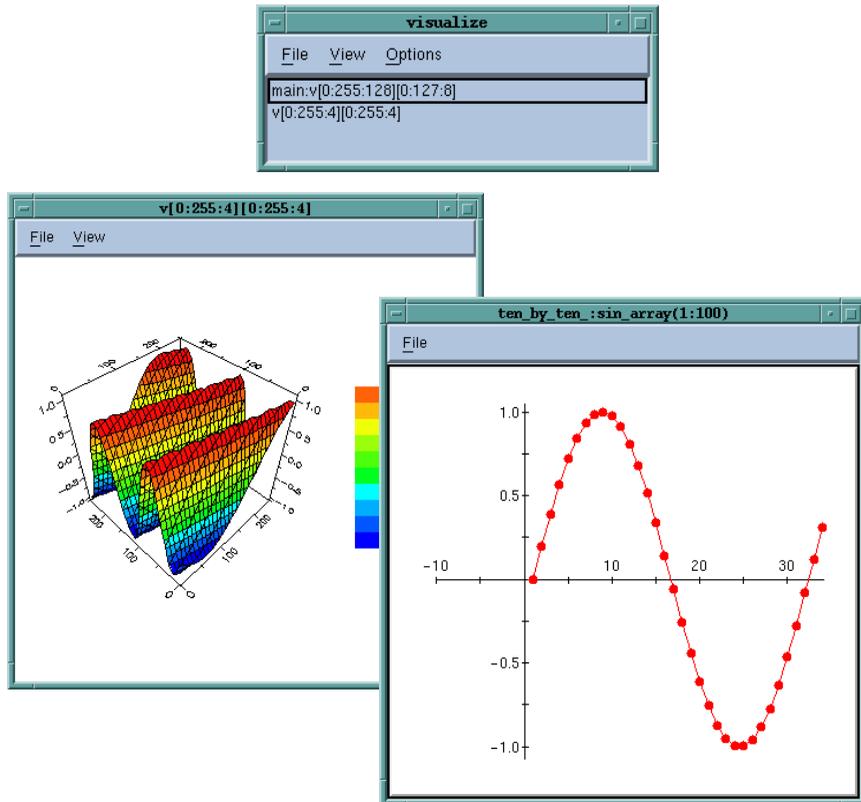


FIGURE 103: Sample Visualizer Windows

Directory Window

The Directory Window contains a list of the datasets you can display. You can select a dataset by clicking on it. Double-clicking on the dataset tells the Visualizer to display it. While you can display multiple datasets, you can only select one dataset at a time.

The **View** menu lets you select **Graph** or **Surface** visualization. Whenever TotalView sends a new dataset to the Visualizer, the Visualizer updates its dataset list. To de-

lete a dataset from the list, click on it, display the **File** menu, and then select **Delete**. (It's usually easier to just close the Visualizer.)

Here are the commands contained in the Directory Window's menu bar:

- File > Delete** Deletes the currently selected dataset. It removes the dataset from the dataset list and destroys the Data Windows displaying it.
- File > Exit** Closes all windows and exits the **Visualizer**.
- View > Graph** Creates a new Graph Window; see "Using the Graph Window" on page 173 for more detail.
- View > Surface** Creates a new Surface Window; see "Using the Surface Window" on page 177 for more detail.
- Options > Auto Visualize**
This item is a toggle; when enabled, the **Visualizer** automatically visualizes new datasets as they are read. Typically, this option is left on. If, however, you have large datasets and need to configure how the Visualizer displays it, you may want to disable this option.

Data Windows

Data Windows display graphic images of your data. Figure 104 on page 173 shows a surface view and a graph view. Every Data Window contains a menu bar and a drawing area. The Data Window title is its dataset identification.

The Data Window menu commands are as follows:

- File > Close** Closes the Data Window.
- File > Delete** Deletes the Data Window's dataset from the dataset list. This also destroys other Data Windows viewing the dataset.
- File > Directory** Raises the Directory Window to the front of the desktop. If you have minimized the Directory Window, the Visualizer restores it.
- File > New Base Window**
Creates a new Data Window having the same visualization method and dataset as the current Data Window.
- File > Options** Pops up a window of viewing options.

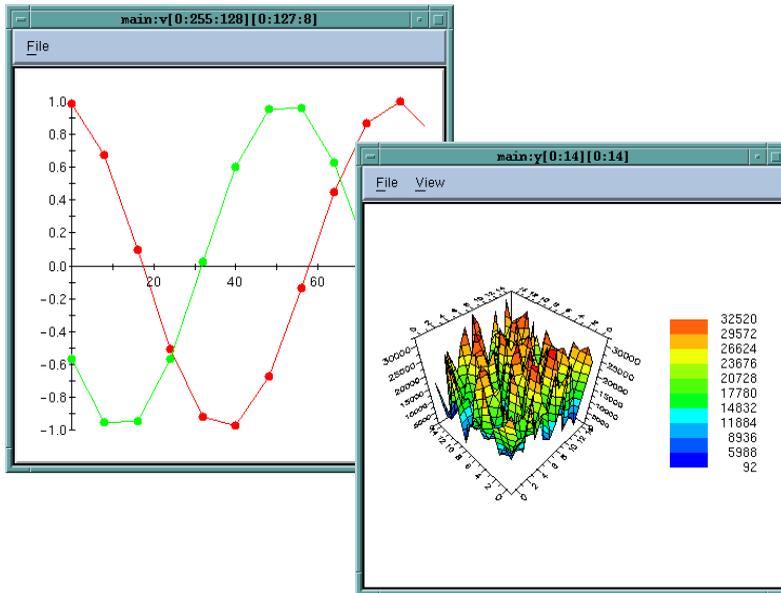


FIGURE 104: Sample Visualizer Data Windows

The drawing area displays the image of your data. You can interact with the drawing area to alter the view of your data. For example, if the Visualizer is showing a surface, you can rotate the graph to view it from different angles. You can also get the value and indices of the dataset element nearest the cursor by clicking on it. A pop-up window displays the information. (See Figure 105 on page 174.)

Using the Graph Window

The Graph Window displays a two-dimensional graph of one- or two-dimensional datasets. If the dataset is two-dimensional, the Visualizer displays multiple graphs. When you first create a Graph Window on a two-dimensional dataset, the Visualizer uses the dimension with the larger number of elements for the X axis. It then draws a separate graph for each subarray having the smaller number of elements. If you don't like this choice, you can transpose the data.

NOTE You probably don't want to use a graph to visualize two-dimensional datasets with large extents in both dimensions, as the display will be very cluttered.

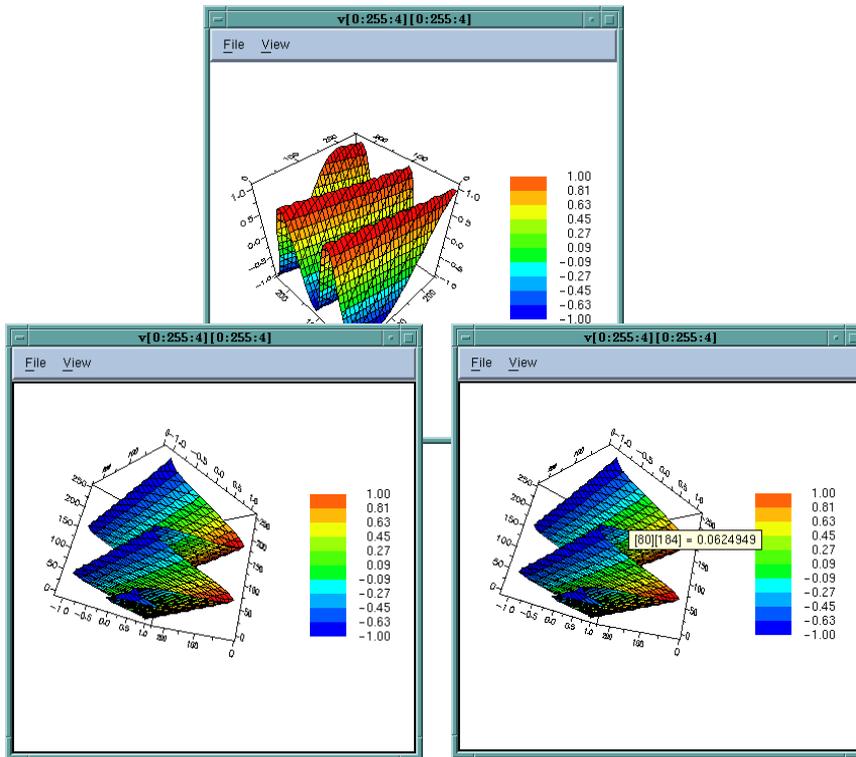


FIGURE 105: **Rotating and Querying**

You can display graphs with markers for each element of the dataset, with lines connecting dataset elements, or with both lines and markers as shown in Figure 106 on page 175. See “*Displaying Graphs*” on page 174 for more details. Multiple graphs are displayed in different colors. The **X** axis of the graph is annotated with the indices of the long dimension. The **Y** axis shows you the data value.

You can scale and translate the graph, or pop up a window displaying the indices and values for individual dataset elements. See “*Manipulating Graphs*” on page 176 for details.

Displaying Graphs

The **File > Options** Dialog Box lets you control how the Visualizer displays the graph. (See Figure 107 on page 175.)

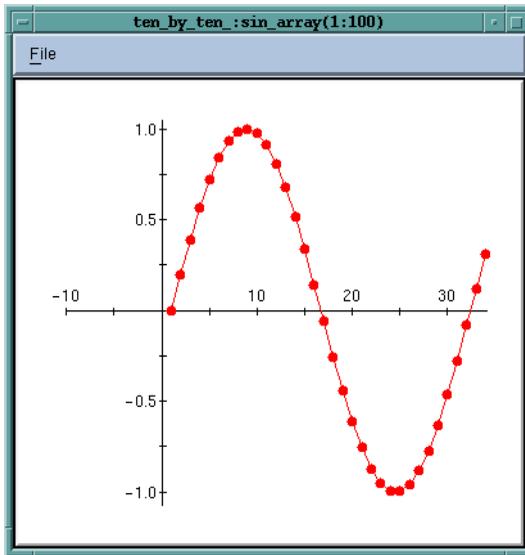


FIGURE 106: Visualizer Graph Data Window

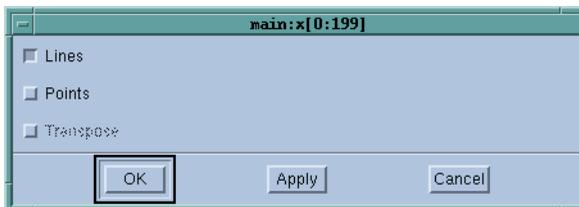


FIGURE 107: Graph Options Dialog Box

Here's what the check boxes in this dialog box mean.

- | | |
|------------------|---|
| Lines | If this is set, the Visualizer displays lines connecting dataset elements. |
| Points | If this is set, the Visualizer displays markers for dataset elements. |
| Transpose | If this is set, the Visualizer inverts the X and Y axis of the displayed graph. |

Manipulating Graphs

You can manipulate the way the Visualizer displays a graph by using the following actions:

- | | |
|-------------------|--|
| Scale | Press the Control key and hold down the middle mouse button. Move the mouse down to zoom in on the center of the drawing area, or up to zoom out. |
| Translate | Press the Shift key and hold down the middle mouse button. Moving the mouse drags the graph. |
| Zoom | Press the Control key and hold down the left mouse button. Drag the mouse button to create a rectangle that encloses an area. The Visualizer scales the graph to fit the drawing area. |
| Reset View | Select View > Reset to reset the display to its initial state. |
| Query | Hold down the left mouse button near a graph marker. A window pops up displaying the dataset element's indices and value. |

Figure 108 shows a graph view of two-dimensional random data created by selecting **Points** and clearing **Lines** in the Data Window's **Graph Options** Dialog Box.

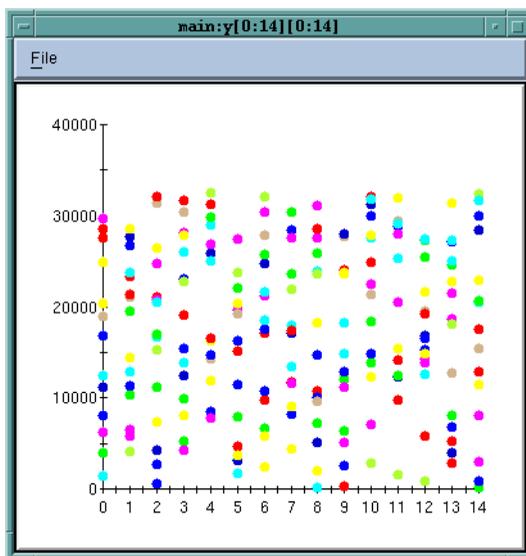


FIGURE 108: **Display of Random Data**

Using the Surface Window

The Surface Window displays two-dimensional datasets as a surface in two or three dimensions. The dataset's array indices map to the first two dimensions (X and Y axes) of the display. Figure 109 shows a two-dimensional map, where the dataset values are shown using only the **Zone** option. (This demarcates ranges of element values.) For a zone map with contour lines, turn the **Zone** and **Contour** settings on and **Mesh** and **Shade** off.

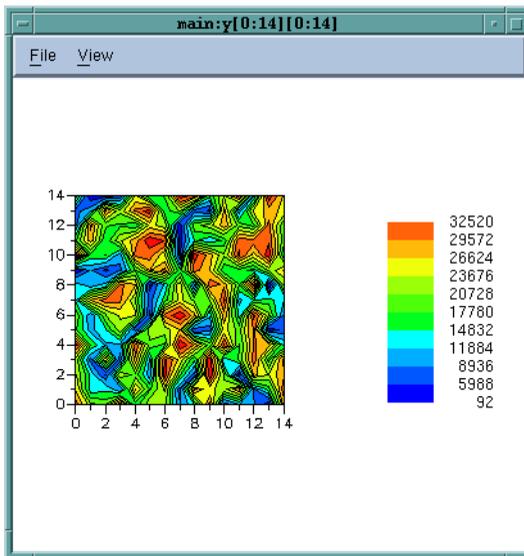


FIGURE 109: Two-Dimensional Surface Visualizer Data Display

You can display random data by selecting only the **Zone** setting and turning **Mesh**, **Shade**, and **Contour** off. The display shows where the data is located, and you can click on the display to get the values of the data points.

Figure 110 on page 178 shows a three-dimensional surface that maps element values to the height (Z axis).

Displaying Surface Data

The Surface Window's **File > Options** command lets you control how the Visualizer displays the graph. (See Figure 111 on page 178.)

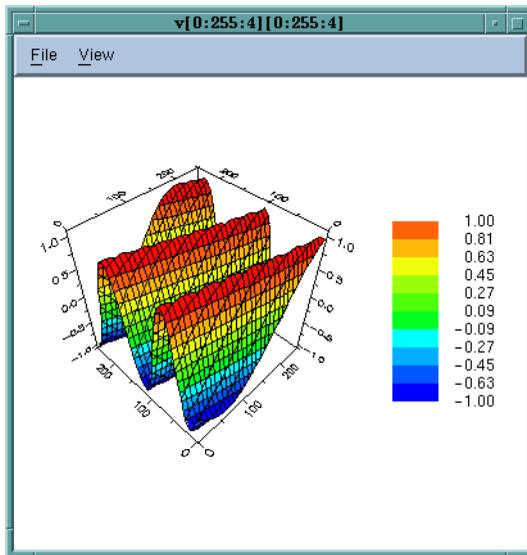


FIGURE 110: Three-Dimensional Surface Visualizer Data Display

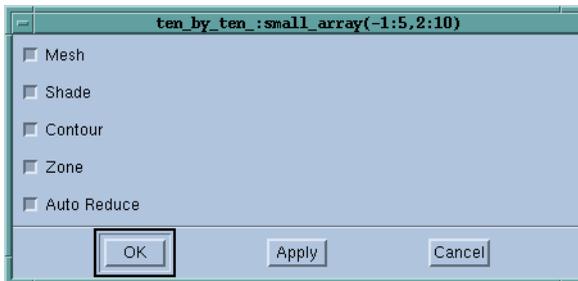


FIGURE 111: Surface Options Dialog Box

This dialog box has the following choices:

- Mesh** If this option is set, the Visualizer displays the surface as a three dimensional mesh, with the X-Y grid projected onto the surface. If you don't set this or the **Shade** option, the Visualizer displays the surface in two dimensions. (See Figure 109.)
- Shade** If this option is set, the Visualizer displays the surface in three dimensions and shaded either in a "flat" color to differentiate the top and bottom sides of the surface, or in colors corresponding to the value if the **Zone** option is also set. When nei-

ther this nor the **Mesh** option are set, the Visualizer displays the surface in two dimensions. (See Figure 109.)

- Contour** If this option is set, the Visualizer displays contour lines indicating ranges of element values.
- Zone** If this option is set, the Visualizer displays the surface in colors showing ranges of element values.
- Auto Reduce** If this option is set, the Visualizer derives the displayed surface by averaging over neighboring elements in the original dataset. This speeds up visualization by reducing the resolution of the surface. Clear this option if you want to accurately visualize all dataset elements.

The **Auto Reduce** option allows you to choose between viewing all your data points—which takes longer to appear in the display—or viewing the averaging of data over a number of nearby points.

You can reset the viewing parameters to those used when the Visualizer first came up by selecting the **View > Reset** command, which restores all translation, rotation, and scaling to its initial state and enlarges the display slightly.

Manipulating Surface Data

The following commands change the display or give you information about it:

- Query** Hold down the **left mouse** button near the surface. A window pops up displaying the nearest dataset element’s indices and value.
- Rotate** Hold down the **middle mouse** button and **drag** the mouse to freely rotate the surface. You can also press the **X**, **Y**, or **Z** keys to select a single axis of rotation. The Visualizer lets you rotate the surface in two dimensions simultaneously.

While you’re rotating the surface, the Visualizer displays a wire-frame bounding box of the surface and moves it as your mouse moves.
- Scale** Press the **Control** key and hold down the **middle mouse** button. Move the mouse down to zoom in on the center of the drawing area, or up to zoom out.

- | | |
|------------------|--|
| Translate | Press the Shift key and hold down the middle mouse button. Moving the mouse drags the surface. |
| Zoom | Press the Control key and hold down the left mouse button. Drag the mouse button to create a rectangle that encloses the area of interest. The Visualizer then translates and scales the area to fit the drawing area. See Figure 112 on page 181. |

Launching the Visualizer from the Command Line

To start the Visualizer from the shell, use the following syntax:

```
visualize [ -file filename | -persist ]
```

where:

- | | |
|-----------------------|---|
| -file filename | Reads data from <i>filename</i> instead of reading from standard input. |
| -persist | Continues to run after encountering an EOF on standard input. If you don't use this option, the Visualizer exits as soon as it reads all of the data. |

By default, the Visualizer reads its datasets from standard input and exits when it reads an EOF. When started by TotalView, the Visualizer reads its data from a pipe, ensuring that the Visualizer exits when TotalView does. If you want the Visualizer to continue to run after it exhausts all input, invoke it by using the **-persist** option.

If you want to read data from a file, invoke the Visualizer with the **-file** option:

```
visualize -file my_data_set_file
```

The Visualizer reads all the datasets in the file. This means that the images you see represent the last versions of the datasets in the file.

The Visualizer supports the generic X toolkit command-line options. For example, you can start the Visualizer with the Directory Window minimized by using the **-iconic** option. Your system manual page for the X server or the X WINDOW SYSTEM USER'S GUIDE by O'Reilly & Associates lists the generic X command-line options in detail.

You can also customize the Visualizer by setting X resources in your resource files or on the command line with the **-xrm resource_setting** option. The available resources are described in "TotalView Command Syntax" in the TOTALVIEW REFERENCE

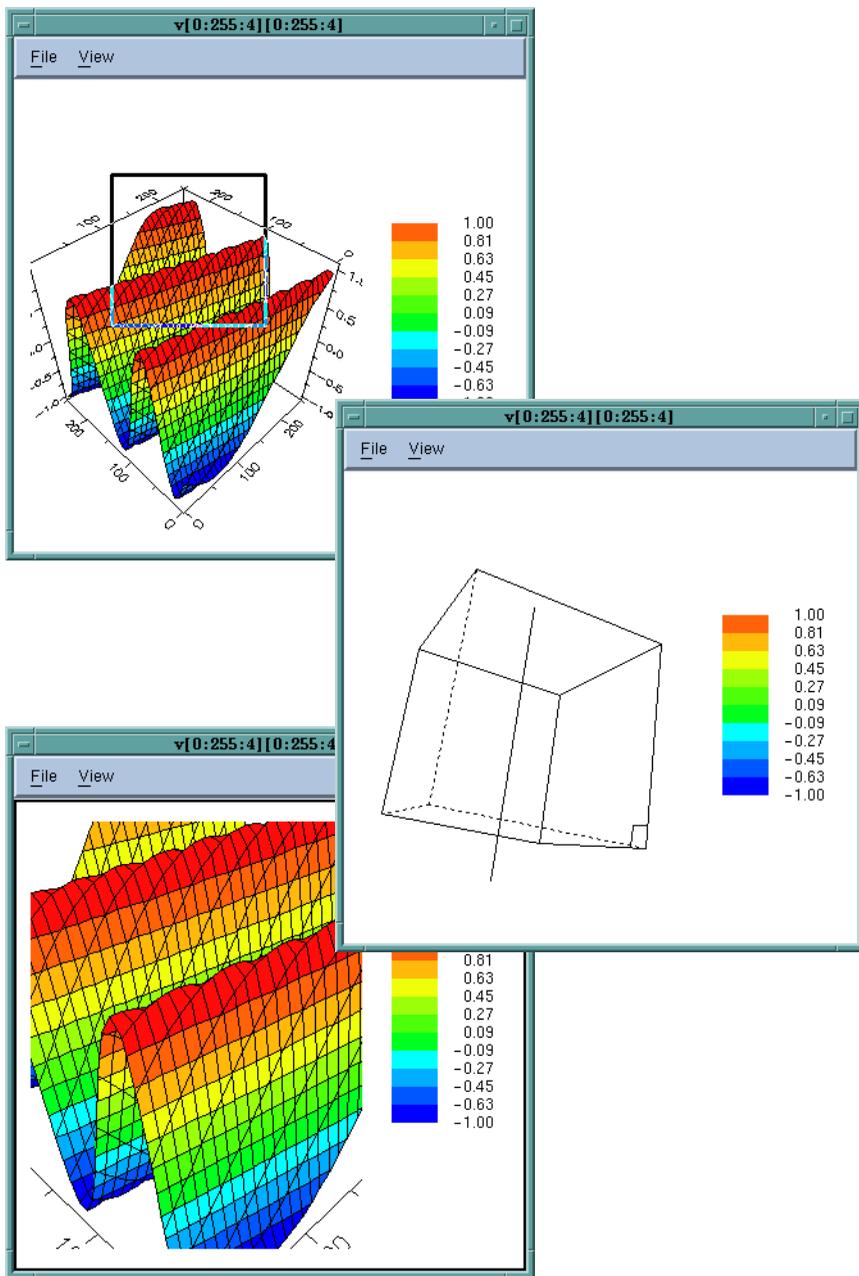


FIGURE 112: **Zooming, Rotating, About an Axis**

GUIDE. Use of X resources to modify the default behavior of TotalView or the TotalView Visualizer is described in greater detail on our Web site at www.et-nus.com/Support/docs/xresources/XResources.html.



Part IV: Using the CLI

While other parts of this book deal with both the GUI and the CLI or with just the GUI, the chapters in this part deal exclusively with the CLI. Most CLI commands must have a process/thread focus for what they will be doing. See Chapter 11: “Using Groups, Processes, and Threads” on page 239 for more information.

Chapter 8: Seeing the CLI at Work

While you can use the CLI as a stand-alone debugger, using the GUI is usually easier. Where the CLI shines is in creating debugging functions that are unique to your program or in automating repetitive tasks. This chapter presents a few Tcl macros in which TotalView CLI commands are embedded.

While most of these examples are simple, you are urged to, at a minimum, skim over this information so you get a feel for what can be done.

Chapter 9: Using the CLI

You can use TotalView’s CLI commands without knowing much about Tcl, which is the approach taken in this chapter. Here you will read about how to enter CLI commands and how the CLI and TotalView interact with one another when used in a nongraphical way.



Chapter 8

Seeing the CLI at Work

The CLI is a command-line debugger that is completely integrated with TotalView. You can use it and never use the TotalView GUI or you can use it and the GUI simultaneously. Because the CLI is embedded within a Tcl interpreter, you can also create debugging functions that exactly meet your needs. When you do this, you can use these functions in the same way that you use TotalView's built-in CLI commands.

This chapter contains a few macros that show how the CLI programmatically interacts with your program and with TotalView. Reading a few examples without bothering too much with details will give you an appreciation for what the CLI can do and how you can use it. As you will see, you really need to have a basic knowledge of Tcl before you can make full use of all CLI features.

The chapter presents a few macros. In each macro, all Tcl commands that are unique to the CLI are displayed in bold. The macros in this chapter are for:

- *"Setting the EXECUTABLE_PATH State Variable"* on page 185
- *"Initializing an Array Slice"* on page 187
- *"Printing an Array Slice"* on page 187
- *"Writing an Array Variable to a File"* on page 189
- *"Automatically Setting Breakpoints"* on page 189

Setting the EXECUTABLE_PATH State Variable

The following macro recursively descends through all directories starting at a location that you enter. (This is indicated by the *root* argument.) The macro will ignore directories named in the *filter* argument. The result is then set as the value of the CLI EXECUTABLE_PATH state variable.

```

# Usage:
#
# rpath [root] [filter]
#
# If root is not specified, start at the current
# directory. filter is a regular expression that removes
# unwanted entries. If it is not specified, the macro
# automatically filters out CVS/RCS/SCCS directories.
#
# The TotalView search path is set to the result.

proc rpath {{root "."} {filter "/(CVS|RCS|SCCS) (/|$)"} } {

    # Invoke the UNIX find command to recursively obtain a
    # list of all directory names below "root".
    set find [split [exec find $root -type d -print] \n]

    set npath ""

    # Filter out unwanted directories.
    foreach path $find {
        if {![regexp $filter $path]} {
            append npath ":"
            append npath $path
        }
    }

    # Tell TotalView to use it.
    dset EXECUTABLE_PATH $npath
}

```

In this macro, the last statement sets the `EXECUTABLE_PATH` state variable. This is the only statement that is unique to the CLI. All other statements are standard Tcl.

The `dset` command, like most interactive CLI commands, begins with the letter `d`. (The `dset` command is only used in assigning values to CLI state variables. In contrast, values are assigned to Tcl variables by using the standard Tcl `set` command.)

Initializing an Array Slice

The following macro initializes an array slice to a constant value:

```
array_set (var lower_bound upper_bound val) {
    for {set i $lower_bound} {$i <= $upper_bound} {incr i} {
        dassign $var\($i) $val
    }
}
```

The CLI **dassign** command assigns a value to a variable. In this case, it is setting the value of an array element. Here is how you use this function:

```
d1.<> dprint list3
list3 = {
    (1) = 1 (0x00000001)
    (2) = 2 (0x00000001)
    (3) = 3 (0x00000001)
}
d1.<> array_set list 2 3 99
d1.<> dprint list3
list3 = {
    (1) = 1 (0x00000001)
    (2) = 99 (0x00000063)
    (3) = 99 (0x00000063)
}
```

Printing an Array Slice

The following macro prints a Fortran array slice. This macro, like other ones shown in this chapter, relies heavily on Tcl and uses unique CLI commands sparingly.

```
proc pf2Dslice {anArray i1 i2 j1 j2 {i3 1} {j3 1} \
    {width 20}} {
    for {set i $i1} {$i <= $i2} {incr i $i3} {
        set row_out ""
        for {set j $j1} {$j <= $j2} {incr j $j3} {
            set ij [capture dprint $anArray\($i,$j\)]
            set ij [string range $ij \
                [expr [string first "=" $ij] + 1] end]
            set ij [string trimright $ij]
            if {[string first "-" $ij] == 1} {
                set ij [string range $ij 1 end]}
        }
    }
}
```

```

        append ij "
        append row_out " " \
            [string range $ij 0 $width] " "
    }
    puts $row_out
}
}

```

NOTE The CLI's `dprint` command lets you specify a slice. For example: `dprint a(1:4,1:4)`.

After invoking this macro, the CLI prints a two-dimensional slice (`i1:i2:i3, j1:j2:j3`) of a Fortran array to a numeric field whose width is specified by the `width` argument. This width doesn't include a leading minus (-) sign.

All but one line is standard Tcl. This line uses the `dprint` command to obtain the value of one array element. This element's value is then captured into a variable. The CLI `capture` command allows a value that is normally printed to be sent to a variable. For information on the difference between values being displayed and values being returned, see "CLI Output" on page 200.

Here are several examples:

```

d1.<> pf2Dslice a 1 4 1 4
0.841470956802 0.909297406673 0.141120001673-0.756802499294
0.909297406673-0.756802499294-0.279415488243 0.989358246326
0.141120001673-0.279415488243 0.412118494510-0.536572933197
-0.756802499294 0.989358246326-0.536572933197-0.287903308868
d1.<> pf2Dslice a 1 4 1 4 1 1 17
0.841470956802 0.909297406673 0.141120001673-0.756802499294
0.909297406673-0.756802499294-0.279415488243 0.989358246326
0.141120001673-0.279415488243 0.412118494510-0.536572933197
-0.756802499294 0.989358246326-0.536572933197-0.287903308868
d1.<> pf2Dslice a 1 4 1 4 2 2 10
0.84147095 0.14112000
0.14112000 0.41211849

```

```
d1.<> pf2Dslice a 2 4 2 4 2 2 10
-0.75680249 0.98935824
 0.98935824-0.28790330
d1.<>
```

Writing an Array Variable to a File

There are many times when you would like to save the value of an array so that you can analyze its results at a later time. The following macro writes array values to a file.

```
proc save_to_file {var fname} {
    set values [capture dprint $var]
    set f [open $fname w]

    puts $f $values
    close $f
}
```

The following shows how you might use this macro. Notice that using the **exec** command lets **cat** display the file that was just written.

```
d1.<> dprint list3
list3 = {
  (1) = 1 (0x00000001)
  (2) = 2 (0x00000002)
  (3) = 3 (0x00000003)
}
d1.<> save_to_file list3 foo
d1.<> exec cat foo
list3 = {
  (1) = 1 (0x00000001)
  (2) = 2 (0x00000002)
  (3) = 3 (0x00000003)
}
d1.<>
```

Automatically Setting Breakpoints

In many cases, your knowledge of what a program is doing lets you make predictions as to where problems will occur. The following CLI macro parses comments

that you can include within a source file and, depending on the comment's text, sets a breakpoint or an evaluation point.

Immediately following this listing is an excerpt from a program that uses this macro.

```
# make_actions: Parse a source file, and insert
# evaluation and breakpoints according to comments.
#
proc make_actions {{filename ""}} {

    if {$filename == ""} {
        puts "You need to specify a filename"
        error "No filename"
    }

    # Open the program's source file and initialize a
    # few variables.
    set fname [set filename]
    set fsource [open $fname r]
    set lineno 0
    set incomment 0

    # Look for "signals" that indicate the kind of
    # action point; they are buried in the comments.
    while {[gets $fsource line] !=-1} {
        incr lineno
        set bpline $lineno

        # Look for a one-line evaluation point. The
        # format is ... /* EVAL: some_text */.
        # The text after EVAL and before the "*/" in
        # the comment is assigned to "code".
        if [regexp "/\\* EVAL: *(.*)\\*/" $line all code] {
            dbbreak $fname\\#$bpline -e $code
            continue
        }

        # Look for a multiline evaluation point.
        if [regexp "/\\* EVAL: *(.*)" $line all code] {
            # Append lines to "code".
            while {[gets $fsource interiorline] !=-1} {
                incr lineno
            }
        }
    }
}
```

```

# Tabs will confuse dbreak.
regsub -all \t $interiorline \
    " " interiorline

# If "*" is found, add the text to "code",
# then leave the loop. Otherwise, add the
# text, and continue looping.
if [regexp "(.*)\\*/" $interiorline \
    all interiorcode]{
    append code \n $interiorcode
    break
} else {
    append code \n $interiorline
}
}
dbreak $fname\#$bpline -e $code
continue
}

# Look for a breakpoint.
if [regexp "/\\* STOP: .*" $line] {
dbreak $fname\#$bpline
continue
}

# Look for a command to be executed by Tcl.
if [regexp "/\\* *CMD: *(.*)\\*/" $line all cmd] {
    puts "CMD: [set cmd]"
eval $cmd
}
}
close $fsource
}

```

The only similarity between this example and the previous three is that almost all of the statements are Tcl. The only purely CLI commands are the instances of the **dbreak** command that sets evaluation points and breakpoints.

The following excerpt from a larger program shows how you would embed comments within a source file that would be read by the next macro.

```

...
struct struct_bit_fields_only {
    unsigned f3 : 3;
    unsigned f4 : 4;
    unsigned f5 : 5;
    unsigned f20 : 20;
    unsigned f32 : 32;
} sbfo, *sbfo = &sbfo;
...
int main()
{
    struct struct_bit_fields_only *lbfo = &sbfo;
    ...
    int i;
    int j;
    sbfo.f3 = 3;
    sbfo.f4 = 4;
    sbfo.f5 = 5;
    sbfo.f20 = 20;
    sbfo.f32 = 32;
    ...
    /* TEST: Check to see if we can access all the
       values */
    i=i;    /* STOP: */
    i=1;    /* EVAL: if (sbfo.f3 != 3) $stop; */
    i=2;    /* EVAL: if (sbfo.f4 != 4) $stop; */
    i=3;    /* EVAL: if (sbfo.f5 != 5) $stop; */
    ...
    return 0;
}

```

The `make_actions` macro reads a source file one line at a time. As it reads these lines, the regular expressions look for comments that begin with `/* STOP`, `/* EVAL`, and `/* CMD`. After parsing the comment, it sets a breakpoint at a *stop* line, an evaluation point at an *eval* line, or executes a command at a *cmd* line.

Using evaluation points can be confusing because evaluation point syntax differs from that of Tcl. In this example, the `$stop` command is a command contained in TotalView (and the CLI). It is not a Tcl variable. In other cases, the evaluation statements will be in the C or Fortran programming languages.

Chapter 9

Using the CLI

The two components of the Command Line Interface (CLI) are the Tcl-based programming environment and the commands added to the Tcl interpreter that allow you to debug your program. This chapter looks at how these components interact and describes how you specify processes, groups, and threads.

This chapter tends to emphasize interactive use of the CLI rather than using the CLI as a programming language because many of the concepts that will be discussed are easier to understand in an interactive framework. However, everything in this chapter can be used in both environments.

Topics discussed in this chapter are:

- *"Tcl and the CLI"* on page 193
- *"Starting the CLI"* on page 196
- *"CLI Output"* on page 200
- *"Command Arguments"* on page 201
- *"Using Namespaces"* on page 202
- *"Command and Prompt Formats"* on page 203
- *"Built-In Aliases and Group Aliases"* on page 203
- *"Effects of Parallelism on TotalView and CLI Behavior"* on page 204
- *"Controlling Program Execution"* on page 206

Tcl and the CLI

The TotalView CLI is built within version 8.0 of Tcl, so TotalView's CLI commands are built into Tcl. This means that the CLI is not a library of commands that you can bring into other implementations of Tcl. Because the Tcl you are running is the

standard 8.0 version, the TotalView CLI supports all libraries and operations that run using version 8.0 of Tcl.

Integrating CLI commands into Tcl makes them intrinsic Tcl commands. This lets you enter and execute all CLI commands in exactly the same way as you enter and execute built-in Tcl commands. As CLI commands are also Tcl commands, you can embed Tcl primitives and functions within CLI commands and embed CLI commands within sequences of Tcl commands.

For example, you can create a Tcl list that contains a list of threads, use Tcl commands to manipulate that list, and then use a CLI command that operates on the elements of this list. Or you create a Tcl function that dynamically builds the arguments that a process will use when it begins executing.

The CLI and TotalView

The following figure illustrates the relationship between the CLI, the TotalView GUI, the TotalView core, and your program:

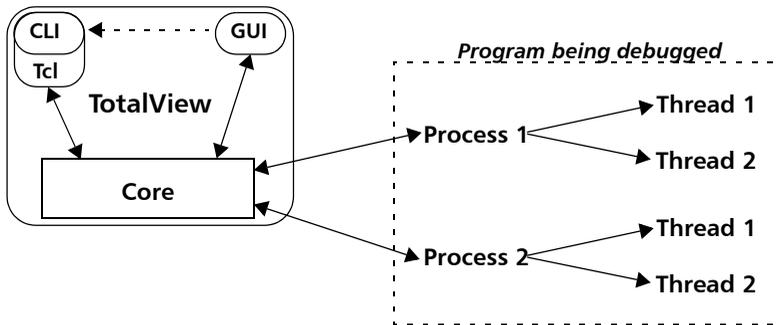


FIGURE 113: The CLI and TotalView

The CLI and the GUI are components that communicate with the TotalView core, which is what actually does the work. In this figure, the dotted arrow between the GUI and the CLI indicates that you can invoke the CLI from the GUI. The reverse isn't true: you can't invoke the GUI from the CLI.

In turn, the TotalView core communicates with the processes that make up your program and receives information back from these processes, and passes them

back to the component that sent the request. If the GUI is also active, the core also updates the GUI's windows. For example, stepping your program from within the CLI changes the PC in the Process Window, updates data values, and so on.

The CLI Interface

The way in which you interact with the CLI is by entering a CLI or Tcl command. (As entering a Tcl command does exactly the same thing in the CLI as it does when interacting with a Tcl interpreter, entering commands and the command environment won't be discussed here.) Typically, the effect of executing a CLI command is one or more of the following:

- The CLI displays information about your program.
- A change takes place in your program's state.
- A change takes place in the information that the CLI maintains about your program.

After the CLI executes your command, it displays a prompt. Although CLI commands are executed sequentially, commands executed by your program may not be. For example, the CLI doesn't require that your program be stopped when it prompts for and performs commands. It only requires that the last CLI command be complete before it can begin executing the next one. In many cases, the processes and threads being debugged continue to execute after the CLI finished doing what you've asked it to do.

Because actions are occurring constantly, state information and other kinds of messages that the CLI displays are usually mixed in with the commands that you type. You may want to limit the amount of information TotalView displays by setting the **VERBOSE** variable to **WARNING** or **ERROR**. (For more information, see the "Variables" chapter in the TOTALVIEW REFERENCE GUIDE.)

Pressing Ctrl+C while a CLI command is executing interrupts that CLI command or executing Tcl macro. If the CLI is displaying its prompt, typing Ctrl+C stops executing processes.

Starting the CLI

You can start the CLI in two ways:

- You can start the CLI from within the TotalView window by selecting the **Tools > Command Line** command in the Root and Process Windows. After selecting this command, TotalView opens a window into which you can enter CLI commands.
- You can start the CLI directly from a shell prompt by typing **totalviewcli**. (This assumes that the TotalView binary directory is in your path.)

Figure 114 is a snapshot of a CLI window that shows part of a program being debugged.

```

TotalView Command Line Input
d1.< s
81 >          denorms(i) = x'00000001'
d1.< s
820> 40  continue
d1.< dlist -n 5
79
800          do 40 i = 1, 500
81          denorms(i) = x'00000001'
820> 40  continue
83          do 42 i = 500, 1000
84          denorms(i) = x'80000001'
d1.< dstatus
1          (4656)      Breakpoint [arraysLINUX]
1,1        (4656/4656)  Breakpoint PC=0x08048fa8, [arrays.F#82]
d1.< dwhere
> 0 MAIN_          PC=0x08048fa8, FP=0xbfffdac8 [arrays.F#82]
1 main           PC=0x0804909e, FP=0xbfffdac8 [/nfs/fs/u3/home/barryk/Examp
leProgs/arraysLINUX]
2 _libc_start_main PC=0x40065647, FP=0xbfffd08 [.../sysdeps/generic/libc-sta
rt.c#129]
d1.< dup
1 main           PC=0x0804909e, FP=0xbfffdac8 [/nfs/fs/u3/home/barryk/Examp
leProgs/arraysLINUX]
d1.<
  
```

FIGURE 114: CLI xterm Window

If you have problems entering and editing commands, it could be because you invoked the CLI from a shell or process that manipulates your **stty** settings. You can eliminate these problems if you use the **stty sane** CLI command. (If the **sane** option isn't available, you will have to change values individually.)

If you start the CLI with the **totalviewcli** command, you can use all of the command-line options that you can use when starting TotalView except those that have to do with the GUI. (In some cases, TotalView displays an error message if you try. In others, it just ignores what you've done)

Startup Example

Here is a very small CLI script:

```
#
source make_actions.tcl
#
dload fork_loop
dset ARGS_DEFAULT {0 4 -wp}
dstep
catch {make_actions fork_loop.cxx} msg
puts $msg
```

This script begins by loading and interpreting the **make_actions.tcl** file, which was described in Chapter 8, “*Seeing the CLI at Work*” on page 185. It then loads the **fork_loop** executable, sets its default startup arguments, and then steps one source-level statement.

If you stored this in a file named **fork_loop.tvd**, here is how you would tell TotalView to start the CLI and execute this file:

```
totalviewcli -s fork_loop.tvd
```

Information on TotalView’s command-line options is in the “*TotalView Command Syntax*” chapter of the TOTALVIEW REFERENCE GUIDE.

The following example places a similar set of commands in a file that you would invoke from the shell:

```
#!/bin/sh
# Next line exec. by shell, but ignored by Tcl because: \
exec totalviewcli -s "$0" "$@"

#
source make_actions.tcl
#
dload fork_loop
dset ARGS_DEFAULT {0 4 -wp}
dstep
catch {make_actions fork_loop.cxx} msg
puts $msg
```

Notice that the only difference is the first few lines in the file. In the second line, the shell ignores the backslash continuation character while Tcl processes it. This means that the shell will execute the **exec** command while Tcl will ignore it.

Starting Your Program

The CLI lets you start debugging operations in several ways. To execute your program from within the CLI, enter a **dload** command followed by the **drun** command. The following example uses the **totalviewcli** command to start the CLI. This is followed by **dload** and **drun** commands. As this was not the first time the file was run, breakpoints exist from a previous session.

NOTE In this listing, the CLI prompt is "d1.<>". The information preceding the ">" symbol indicates the processes and threads upon which the current command acts. The prompt is discussed in "Command and Prompt Formats" on page 203.

```
% totalviewcli
Copyright 1999-2002 by Etnus, LLC. ALL RIGHTS RESERVED.
Copyright 1999 by Etnus, Inc.
Copyright 1989-1996 by BBN Inc.
d1.<> dload arraysAlpha      #load the arraysAlpha program
1
d1.<> dactions                # Show the action points
No matching breakpoints were found
d1.<> dlist -n 10 75
    75          real16_array (i, j) = 4.093215 * j+2
    76 #endif
    77 26      continue
    78 27      continue
    79
    80 do 40 i = 1, 500
    81     denorms(i) = x'00000001'
    82 40      continue
    83 do 42 i = 500, 1000
    84     denorms(i) = x'80000001'
d1.<> dbreak 80                # Add two action points
1
d1.<> dbreak 83
2
d1.<> drun                    # Run the program to the action point
```

This two-step operation of loading and then running allows you to set action points before execution begins. It also means that you can execute a program more than once. At a later time, you can use the **drerun** command to restart your program, perhaps sending it new command-line arguments. In contrast, reentering the **dload** command tells the CLI to reload the program into memory (for example, after edit-

ing and recompiling the program). The **dload** command always creates new processes. This means that you'll get a new process each time and the old one will still be around.

The **dkill** command terminates one or more processes of a program started by using **dload**, **drun**, or **drerun**. The following example continues where the previous example left off:

```
d1.<> dkill                # kill process
d1.<> drun                  # runs arraysLINUX from start
d1.<> dlist -e -n 3        # show lines about current spot
79
80@>      do 40 i = 1, 500
81          denorms(i) = x'00000001'
d1.<> dwhat master_array   # Tell me about master_array
In thread 1.1:
Name: master_array; Type: integer(100); Size: 400 bytes;
      Addr: 0x140821310
      Scope: ##arraysAlpha#arrays.F#check_fortran_arrays
      (Scope class: Any)
      Address class: proc_static_var (Routine static
      variable)
d1.<> dgo                    # Start program running
d1.<> dwhat denorms         # Tell me about denorms
In thread 1.1:
Name: denorms; Type: <void>; Size: 8 bytes; Addr:
      0x1408214b8
      Scope: ##arraysAlpha#arrays.F#check_fortran_arrays
      (Scope class: Any)
      Address class: proc_static_var (Routine static
      variable)
d1.<> dprint denorms(0)    # Show me what's stored
      denorms(0) = 0x0000000000000001 (1)
d1.<>
```

Because information is interleaved, you may not realize that the prompt has appeared. It is always safe to use the Enter key to have the CLI redisplay its prompt. If a prompt isn't displayed after you press Enter, then you know that the CLI is still executing.

CLI Output

A CLI command can either print its output to a window or return the output as a character string. If the CLI executes a command that returns a string value, it also prints the returned string. Most of the time, you won't care about the difference between *printing* and *returning-and-printing*. Either way, the CLI displays information in your window. And, in both cases, printed output is fed through a simple *more* processor. (This is discussed in more detail in the next section.)

Here are two cases where it matters whether TotalView directly prints output or returns and then prints it:

- When the Tcl interpreter executes a list of commands, TotalView only prints the information returned from the last command. It doesn't show information returned by other commands.
- You can only assign the output of a command to a variable if the CLI returns a command's output. You can't assign output that the interpreter prints directly to a variable or otherwise manipulate it unless you save it using the **capture** command.

For example, the **dload** command returns the ID of the process object that was just created. The ID is normally printed—unless, of course, the **dload** command appears in the middle of a list of commands. For example:

```
{dload test_program; dstatus}
```

In this example, the CLI doesn't display the ID of the loaded program since **dload** was not the last command.

When information is returned, you can assign it to a variable. For example, the next command assigns the ID of a newly created process to a variable:

```
set pid [dload test_program]
```

Because you can't assign the output of the **help** command to a variable, the following doesn't work:

```
set htext [help]
```

This statement assigns an empty string to **htext** because **help** doesn't return text; it just prints it.

To capture the output of a command that prints its output, use the **capture** command. For example, here's how to place **help** command output into a variable:

```
set htext [capture help]
```

NOTE You can only capture the output from commands. You can't capture the informational messages displayed by the CLI that describe process state. If you are using the GUI, TotalView also writes this information to the Root Window's Log Pane. If it is being written there, you can use the File > Save Pane command to write this information to a file.

"more" Processing

When the CLI displays output, it sends data through a simple *more*-like process. This prevents data from scrolling off the screen before you view it. After you see the **MORE** prompt, press Enter to see the next screen of data. If you type **q** (followed by pressing the Enter key), the CLI discard any data it hasn't yet displayed.

You can control the number of lines displayed between prompts by using the **dset** command to set the **LINES_PER_SCREEN** CLI variable. (For more information, see the TOTALVIEW REFERENCE GUIDE.)

Command Arguments

The default command arguments for a process are stored in the **ARGS(*num*)** variable, where *num* is the CLI ID for the process. If you don't set the **ARGS(*num*)** variable for a process, the CLI uses the value stored in the **ARGS_DEFAULT** variable. TotalView sets the **ARGS_DEFAULT** when you use the **-a** option when starting the CLI or the GUI.

NOTE The **-a** option tells TotalView to pass everything that follows on the command line to the program.

For example:

```
totalviewcli -a argument-1, argument-2, . . .
```

To set (or clear) the default arguments for a process, you can use **dset** to modify the **ARGS()** variables directly, or you can start the process with the **drun** command. For example, here is how you can clear the default argument list for process 2:

```
dunset ARGS (2)
```

The next time process 2 is started, the CLI uses the arguments contained in **ARGS_DEFAULT**.

You can also use the **dunset** command to clear the **ARGS_DEFAULT** variable. For example:

```
dunset ARGS_DEFAULT
```

All commands (except **drun**) that can create a process—including **dgo**, **drerun**, **dcont**, **dstep**, and **dnext**—pass the default arguments to the new process. The **drun** command differs in that it replaces the default arguments for the process with the arguments that are passed to it.

Using Namespaces

CLI interactive commands exist within the primary Tcl namespace (**::**). Some of the TotalView state variables also reside in this namespace. Seldom used functions and functions that are not primarily used interactively reside in other namespaces. These namespaces also contain most TotalView state variables. (The variables that appear in other namespaces are usually related to TotalView preferences. The namespaces that TotalView uses are:

- TV::** Contains commands and variables that you will use when creating functions. While they can be used interactively, this is not their primary role.
- TV::GUI::** Contains state variables that define and describe properties of the user interface such as window placement, color, and the like.

If you discover other namespaces beginning with **TV**, you have found a place containing internal functions and variables. These objects can (and will) change and disappear, so don't use them. Also, don't create namespaces that begin with **TV**, as you could cause problems by interfering with built-in functions and variables.

The CLI's **dset** command lets you set the value of these variables. You can have the CLI display a list of these variables by specifying the namespace. For example:

```
dset TV::
```

Command and Prompt Formats

The appearance of the CLI prompt lets you know that the CLI is ready to accept a command. This prompt lists the current focus, and then displays a greater-than symbol (>) and a blank space. (The *current focus* is the processes and threads to which the next command applies.) For example:

<code>d1.<></code>	The current focus is the default set for each command, focusing on the first user thread in process 1.
<code>g2.3></code>	The current focus is process 2, thread 3; commands act on the entire group.
<code>t1.7></code>	The current focus is thread 7 of process 1.
<code>gW3.></code>	All worker threads in the control group containing process 3.
<code>p3/3</code>	All processes in process 3, group 3.

You can change the prompt's appearance by using the `dset` command to set the `PROMPT` state variable. For example:

```
dset PROMPT "Kill this bug! > "
```

Built-In Aliases and Group Aliases

Many CLI commands have an alias that allows you to abbreviate the command's name. (An alias is one or more characters that Tcl interprets as a command or command argument.)

NOTE The "alias" command, which is described in the *TotalView Reference Guide*, lets you create your own aliases.

After a few minutes of entering CLI commands, you will quickly come to the conclusion that it is much more convenient to use the command abbreviation. For example, you could type:

```
dfocus g dhalt
```

(This command tells the CLI to halt the current group.) It is much easier to type:

```
f g h
```

While less-used commands are often typed in full, a few commands are almost always abbreviated. These command include **dbreak (b)**, **ddown (d)**, **dfocus (f)**, **dgo (g)**, **dlist (l)**, **dnext (n)**, **dprint (p)**, **dstep (s)**, and **dup (u)**.

The CLI also includes uppercase “group” versions of aliases for a number of commands, including all stepping commands. For example, the alias for **dstep** is “**s**”; in contrast, “**S**” is the alias for “**dfocus g dstep**”. (The first command tells the CLI to step the process. The second steps the control group.)

There are two ways in which group aliases differ from the kind of group-level command that you would type:

- They do not work if the current focus is a list. The **g** focus specifier modifies the current focus, and it can only be applied if the focus contains just one term.
- They always act on the group, no matter what width is specified in the current focus. Therefore, **dfocus t S** does a step-group command.

Effects of Parallelism on TotalView and CLI Behavior

A parallel program consists of some number of processes, each involving some number of threads. Processes fall into two categories, depending on when they are created:

■ Initial process

A preexisting process from the normal run-time environment (that is, created outside TotalView) or one that was created as TotalView loaded the program.

■ Spawned process

A new process created by a process executing under the CLI’s control.

TotalView assigns an integer value to each individual process and thread under its control. This *process/thread identifier* can be the system identifier associated with the process or thread. However, it can be an arbitrary value created by the CLI. Process numbers are unique over the lifetime of a debugging session; in contrast, thread numbers are only unique while the process exists.

Process/thread notation lets you identify the component that a command targets. For example, if your program has two processes, and each has two threads, four threads exist:

Thread 1 of process 1
Thread 2 of process 1
Thread 1 of process 2
Thread 2 of process 2

You would identify the four threads as follows:

- 1.1—Thread 1 of process 1
- 1.2—Thread 2 of process 1
- 2.1—Thread 1 of process 2
- 2.2—Thread 2 of process 2

Kinds of IDs

Multithreaded, multiprocess, and distributed programs contain a variety of IDs. Here is some background on the kinds used in the CLI and TotalView:

- | | |
|----------------------------|---|
| System PID | This is the process ID and is generally called the PID. |
| Debugger PID | This is an ID created by TotalView that lets it identify processes. It is a sequentially numbered value beginning at 1 that is incremented for each new process. Note that if the target process is killed and restarted (that is, you use the dkill and drun commands), the debugger PID doesn't change. The system PID, however, changes since the operating system has created a new target process. |
| System TID | This is the ID of the system kernel or user thread. On some systems (for example, AIX), the TIDs have no obvious meaning. On other systems, they start at 1 and are incremented by 1 for each thread. |
| TotalView thread ID | This is usually identical to the system TID. On some systems (such as AIX where the threads have no obvious meaning), TotalView uses its own IDs. |
| pthread ID | This is the ID assigned by the Posix pthreads package. If this differs from the system TID, it is a pointer value that points to the pthread ID. |

Controlling Program Execution

Knowing what's going on and where your program is executing is simple in a serial debugging environment. Your program is either stopped or running. When it is running, an event such as arriving at a breakpoint can occur. This event tells the debugger to stop the program. Sometime later, you will tell the serial program to continue executing. Multiprocess and multithreaded programs are more complicated. Each thread and each process has its own execution state. When a thread (or set of threads) triggers a breakpoint, TotalView must decide what it should do about the other threads and processes. Some may stop; some may continue to run.

Advancing Program Execution

Debugging begins by entering a **dload** or **dattach** command. If you use the **dload** command, you must use the **drun** command to start the program executing. These three commands work at process level and you can't use them to start an individual threads. (This is also true for the **dkill** command.)

To advance program execution, you enter a command that causes one or more threads to execute instructions. The commands are applied to a P/T set. (P/T sets are discussed in Chapters 2 and 11.) Because the set doesn't have to include all processes and threads, you can cause some processes to be executed while holding others back. You can also advance program execution by increments, *stepping* the program forward, and you can define the size of the increment. For example, "**dnext 3**" executes the next three statements and then pauses what you've been stepping.

Typically, debugging a program means that you have the program run, and then you stop it and examine its state. In this sense, a debugger can be thought of as tool that allows you to alter a program's state in a controlled way. And debugging is the process of stopping the process to examine its state. However, the term "stop" has a slightly different meaning in a multiprocess, multithreaded program; in these programs, *stopping* means that the CLI holds one or more threads at a location until you enter a command that tells them to start executing again.

Action Points

Action points tell the CLI that it should stop a program's execution. You can specify four different kinds of action points:

- A *breakpoint* (see **dbreak** in the TOTALVIEW REFERENCE GUIDE) stops the process when the program reaches a location in the source code.
- A *watchpoint* (see **dwatch** in the TOTALVIEW REFERENCE GUIDE) stops the process when the value of a variable is changed.
- A *barrier point* (see **dbarrier** in the TOTALVIEW REFERENCE GUIDE), as its name suggests, effectively prevents processes from proceeding beyond a point until all other related processes arrive. This gives you a method for synchronizing the activities of processes. (Note that you can only set a barrier on processes; you can't set them on individual threads.)
- An *evaluation point* (see **dbreak** in the TOTALVIEW REFERENCE GUIDE) lets you programmatically evaluate the state of the process or variable when execution reaches a location in the source code. Evaluation points typically do not stop the process; instead, they perform an action. In most cases, an evaluation point stops the process when some condition that you specify is met.

NOTE Extensive information on action points can be found in Chapter 14, "Setting Action Points" on page 337.

Each action point is associated with an *action point identifier*. You use these identifiers when you need to refer to the action point. Like process and thread identifiers, action point identifiers are assigned numbers as they are created. The ID of the first action point created is 1. The second ID is 2, and so on. These numbers are never reused during a debugging session.

The CLI and TotalView only let you assign one action point to a source code line, but you can make this action point as complex as you need it to be. (Setting multiple action points is required by debuggers that limit what you can do.)



Part V: Debugging

The chapters in this part of the TotalView Users Guide describe how you actually go about debugging your programs. The preceding sections describe, for the most part, what you need to do before you get started with TotalView. In contrast, the chapters in this section are what TotalView is really about.

Chapter 10: Debugging Programs

Reading this chapter will help you find your way around your program. It also tells you how to start it under TotalView's control, and the ways to step your program's execution. Of course, it also tells you how to halt, terminate, and restart your program.

Chapter 11: Using Groups, Processes, and Threads

The stepping information in Chapter 10 describes the commands and the different kinds of stepping. In a multiprocess, multithreaded program, you may need to finely control what is executing. This chapter tells you how to do this.

Chapter 12: Examining and Changing Data

As your program executes, you will want to examine what the value stored in a variable is. This chapter tells you how.

Chapter 13: Examining Arrays

Displaying the information in arrays presents special problems. This chapter tells how TotalView solves these problems.

Chapter 14: Setting Action Points

TotalView's action points let you control how your programs execute and what happens when your program reaches statements that you define as important. Action points also let you monitor changes to a variable's value.



Chapter 10

Debugging Programs

This chapter explains how to perform basic debugging tasks with TotalView. The topics discussed are:

- *"Searching and Looking Up Program Elements"* on page 211
- *"Viewing the Assembler Version of Your Code"* on page 216
- *"Editing Source Text"* on page 218
- *"Manipulating Processes and Threads"* on page 219
- *"Using Stepping Commands"* on page 229
- *"Executing to a Selected Line"* on page 231
- *"Displaying Thread and Process Locations"* on page 232
- *"Continuing with a Specific Signal"* on page 233
- *"Deleting Programs"* on page 234
- *"Restarting Programs"* on page 235
- *"Checkpointing Programs and Processes"* on page 235
- *"Setting the Program Counter"* on page 236
- *"Interpreting Status and Control Registers"* on page 237

Searching and Looking Up Program Elements

TotalView provides several ways for you to navigate and find information in your source file. Topics in this section are:

- *"Searching for Text"* on page 212
- *"Looking for Functions and Variables"* on page 212
- *"Finding the Source Code for Functions"* on page 213

- “Finding the Source Code for Files” on page 215
- “Resetting the Stack Frame” on page 216

Searching for Text

You can search for text strings in most windows with the **Edit > Find** command. After invoking this command, TotalView displays the dialog box shown in Figure 115.



FIGURE 115: **Edit > Find** Dialog Box

Controls in this dialog box let you perform case-sensitive searches, continue searching from the beginning of the file if the string isn’t found in the region beginning at the currently selected line and ending at the last line of the file, and keep the dialog box up between searches. You can also tell TotalView if it should search towards the bottom of the file (**Down**) or the top (**Up**).

After you have found a string, you can find another instance of it by using the **Edit > Find Again** command.

Looking for Functions and Variables

Having TotalView locate a variable or a function is usually easier than scrolling through your sources to look for it. Do this with the **View > Lookup Function** and **View > Lookup Variable** commands. Figure 116 on page 213 shows the dialog box displayed by these commands.

If TotalView doesn’t find the name and it can find something similar, it displays a dialog box containing the names of functions that could match. (See Figure 117.)

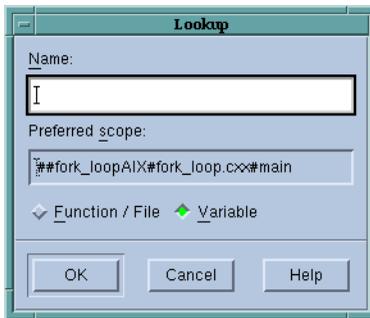


FIGURE 116: View > Lookup Variable Dialog Box

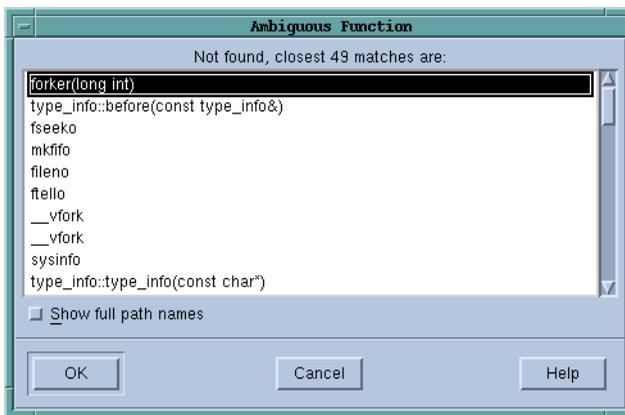


FIGURE 117: Ambiguous Function Dialog Box

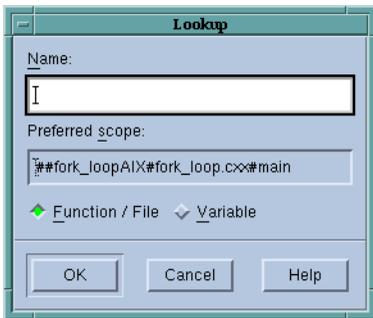
If the one you want is listed, click on its name and then, select **OK** to have it displayed in the Source Pane.

Finding the Source Code for Functions

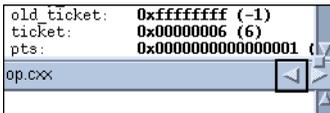
The **View > Lookup Function** command lets you search for a function’s declaration. Here’s the dialog box displayed after you select this command.

CLI EQUIVALENT: `dlist function-name`

After locating your function, TotalView displays it in the Source Pane. If you didn’t compile the function using `-g`, TotalView displays disassembled machine code.

FIGURE 118: **View > Lookup Function Dialog Box**

When you want to return to the previous contents of the Source Pane, use the undive icon located in the upper right corner of the Source Pane and just below the Stack Frame Pane. In the following figure, a square surrounds the undive icon.

FIGURE 119: **Undive/Dive Controls**

You can also use the **View > Reset** command to discard the dive stack so that the Source Pane is displaying the PC it displayed when you last stopped execution.



The **File > Edit Source** command (see “*Editing Source Text*” on page 218 for details) lets you display a file in a text editor. The default editor is **vi**. However, TotalView will use the editor named in an **EDITOR** environment variable or the editor you name in the **Source Code Editor** field of the **File > Preferences’s Launch Strings Page**.

Another method of locating a function’s source code is to dive into a source statement in the Source Pane that shows the function being called. After diving, you’ll see the source.

Resolving Ambiguous Names

Sometimes the function name you specify is ambiguous. For example, you may have specified the name of:

- A static function, and your program contains different versions of it.

- A member function in a C++ program, and multiple classes have a member function with that name.
- An overloaded function or a template function.

Figure 120 on page 215 shows the dialog box that TotalView displays when it encounters an ambiguous function name. You can resolve the ambiguity by clicking on the function you desire.

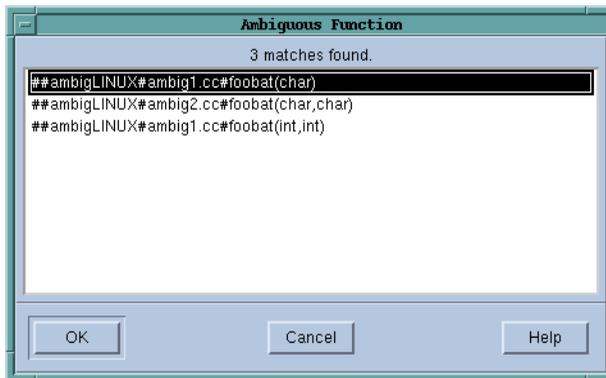


FIGURE 120: Ambiguous Function Dialog Box

Finding the Source Code for Files

You can display a file's source code by selecting the **View > Lookup Function** command and entering the file name in the dialog box shown in Figure 121 on page 215.

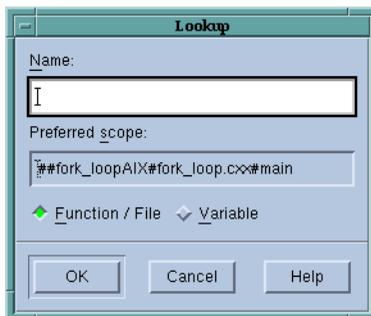


FIGURE 121: View > Lookup Function Dialog Box

If a header file contains source lines that produce executable code, you can enter its name here.

Resetting the Stack Frame

After moving around your source code to look at what's happening in different places, you can return to the executing line of code for the current stack frame by selecting the **View > Reset** command. This command places the PC arrow onto the screen.

This command is also useful when you want to undo the effect of scrolling or when you move to different locations using, for example, the **View > Lookup Function** command.

If the program hasn't started running, the **View > Reset** command displays the first executable line in your main program. This is useful when you are looking at your source code and want to get back to the first statement your program will execute.

Viewing the Assembler Version of Your Code

You can display your program in source form or as assembler. Here are the commands that you can use:

Source code (Default)

Use the **View > Source As > Source** command.

Assembler code

Use the **View > Source As > Assembler** command.

Both Source and assembler

Use the **View > Source As > Both** command.

The Source Pane divides into two parts. The left contains the program's source code and the right contains the assembler version of this code. You can set breakpoints in either of these panes. Note that setting an action point at the first instruction after a source statement, is equivalent to setting it at that source statement.

The commands in the following table tell TotalView to display your assembler code by using symbolic or absolute addresses:

Table 9: Assembler Code Display Styles

Command	TotalView Shows
View > Assembler > By Address	Absolute addresses for locations and references; this is the default
View > Assembler > Symbolically	Symbolic addresses (function names and offsets) for locations and references

NOTE You can also display assembler instructions in a Variable Window. For more information, see “*Displaying Machine Instructions*” on page 290.

The following three figures illustrate the different ways TotalView can display assembler code. In Figure 122 on page 217, the second column (the one to the right of the line numbers) shows the absolute address location. The fourth column shows references using absolute addresses.

```

Function snore in fork_loop.cxx
: 0x080498a3: 0x08
: 0x080498a4: 0x00
: 0x080498a5: 0x74 je 0x80498ac
: 0x080498a6: 0x05
637 0x080498a7: 0xe8 call tight_loop()
: 0x080498a8: 0x38
: 0x080498a9: 0xff
: 0x080498aa: 0xff
639 0x080498ac: 0x83 cmpl $0, do_segvr
: 0x080498ad: 0x3d
: 0x080498ae: 0x3c
: 0x080498af: 0xde
: 0x080498b0: 0x04
: 0x080498b1: 0x08
: 0x080498b2: 0x00
: 0x080498b3: 0x74 je 0x80498ed
: 0x080498b4: 0x38
: 0x080498b5: 0x8b movl -4(%ebp), %eax
: 0x080498b6: 0x45
: 0x080498b7: 0xfc
: 0x080498b8: 0x3b cmpl do_segvr_index, %eax
: 0x080498b9: 0x05

```

FIGURE 122: Address Only (Absolute Addresses)

This next figure shows information symbolically. The second column shows locations using functions and offsets.

The final “assembler” figure (see Figure 124 on page 218) shows the split Source Pane, with one side showing the program’s source code and the other showing the assembler version. In this example, the assembler is shown symbolically. How it is

Symbolic Address	Hex Address	Instruction
snore(void*)+0x6f:	0x08	
snore(void*)+0x70:	0x00	
snore(void*)+0x71:	0x74	je snore(void*)+0x78
snore(void*)+0x72:	0x05	
637: snore(void*)+0x73:	0xe8	call tight_loop()
snore(void*)+0x74:	0x38	
snore(void*)+0x75:	0xff	
snore(void*)+0x76:	0xff	
snore(void*)+0x77:	0xff	
639: snore(void*)+0x78:	0x83	cmpl \$0, do_segvs
snore(void*)+0x79:	0x3d	
snore(void*)+0x7a:	0x3c	
snore(void*)+0x7b:	0xde	
snore(void*)+0x7c:	0x04	
snore(void*)+0x7d:	0x08	
snore(void*)+0x7e:	0x00	
snore(void*)+0x7f:	0x74	je snore(void*)+0xb9
snore(void*)+0x80:	0x38	
snore(void*)+0x81:	0x8b	movl -4(%ebp), %eax
snore(void*)+0x82:	0x45	
snore(void*)+0x83:	0xfc	
snore(void*)+0x84:	0x3b	cmpl do_segvs_index, %eax
snore(void*)+0x85:	0x05	

FIGURE 123: Assembler Only (Symbolic Addresses)

shown depends upon whether you've selected **View > Assembler > By Address** or **View > Assembler > Symbolically**.

Source Line	Source Code	Symbolic Address	Hex Address
655	struct timeval timeout;	snore(void *)+0xfc:	0xe86:
656	long bad_addr;	snore(void *)+0x100:	0xe86:
657	int *foo;	snore(void *)+0x104:	0x906:
658	int bar;	snore(void *)+0x108:	0xe88:
659	wait_a_while (&timeout);	snore(void *)+0x10c:	0xe86:
660	bad_addr = -3;	snore(void *)+0x110:	0x386:
661	foo = (int *)bad_addr;	snore(void *)+0x114:	0x906:
662	bar = *foo;	snore(void *)+0x118:	0x386:
663	*foo = bar + 1;	snore(void *)+0x11c:	0x4bf:
664	}	snore(void *)+0x120:	0xe86:
665		snore(void *)+0x124:	0xe86:
666	for (;;) {	snore(void *)+0x128:	0x2c0:
667	{	snore(void *)+0x12c:	0x418:
668	struct timeval timeout;	snore(void *)+0x130:	0x480:
669	wait_a_while (&timeout);	snore(void *)+0x134:	0xe84:
670	if (verbose)	snore(void *)+0x138:	0x546:
671	printf ("Thread %ld woke u	snore(void *)+0x13c:	0x387:
672	if (use mut)	snore(void *)+0x140:	0x480:

FIGURE 124: Both Source and Assembler (Symbolic Addresses)

Editing Source Text

The **File > Edit Source** command lets you examine the current routine in a text editor. TotalView uses an *editor launch string* to determine how to start your editor.

TotalView expands this string into a command that TotalView sends to the **sh** shell.

The fields within the **Launch Strings** Page of the **File > Preferences** Window let you name the editor and the way TotalView launches the editor. The online help for this page contains information on setting this preference.

Manipulating Processes and Threads

Topics discussed in this section are:

- "Using the Toolbar to Select a Target" on page 219
- "Stopping Processes and Threads" on page 220
- "Updating Process Information" on page 221
- "Holding and Releasing Processes and Threads" on page 221
- "Examining Groups" on page 223
- "Displaying Groups" on page 224
- "Placing Processes into Groups" on page 225
- "Starting Processes and Threads" on page 225
- "Creating a Process Without Starting It" on page 226
- "Creating a Process by Single-Stepping" on page 227
- "Stepping and Setting Breakpoints" on page 227

Using the Toolbar to Select a Target

The Process Window's toolbar has three groups of buttons. The first group, which is a single pulldown list, defines the *focus* of the command selected in the second group of the toolbar. The third group changes the process and thread being displayed. Figure 125 shows this toolbar.



FIGURE 125: **The Toolbar**

When you are doing something to a multi-process, multi-threaded program, TotalView needs to know which processes and threads it should act upon. In the CLI, you specify this target using the **dfocus** command. When using the GUI, you specify the focus using this pulldown. For example, if you select **Thread**, and then select the **Step** button, TotalView steps the current thread. In contrast, if you select **Process Workers** and then select the **Go** button, TotalView tells all the processes that are in the same workers group as the current thread (this thread is called the *thread of interest*).

NOTE Chapter 11, “Using Groups, Processes, and Threads” on page 239 fully describes how TotalView manages processes and threads. While TotalView gives you the ability to control the precision your application requires, most applications do not need this level of interaction. In almost all cases, using the controls in the toolbar gives you all the control you need.

Stopping Processes and Threads

To stop a group, process, or thread, select a **Halt** command from the **Group**, **Process**, or **Thread** pulldown menu on the toolbar.

CLI EQUIVALENT: **dhalt**
Halts a group, process, or thread. Setting the focus changes the scope.

The three **Halt** commands differ in the scope of what they halt. In all cases, TotalView uses the current thread (which is called the thread of interest or TOI) to determine what else it will halt. For example, suppose you select **Process > Halt**. This tells TotalView to determine the process in which the TOI is running. It will then halt this process. Similarly, if you select **Group > Share > Halt**, TotalView determines what processes are in the share group the current thread participates in. It then stops all of these processes.

NOTE For more information on the TOI, see “Defining the GOI, POI, and TOI” on page 239.

When you select the **Halt** button in the toolbar instead of the commands within the menubar, TotalView decides what it should stop based on what is set in the two toolbar pulldown lists.

After entering a **Halt** command, TotalView updates any windows that can be updated. When you restart the process, execution continues from the point where TotalView had stopped the process.

Updating Process Information

Normally, TotalView only updates information when the thread being executed stops executing. You can force TotalView to update a window if you use the **Window > Update** command. You'll need to use this command if you want to see what a variable's value is while your program is executing.

NOTE When you use this command, TotalView momentarily stops execution so that it can obtain the information it needs. It then restarts the thread.

Holding and Releasing Processes and Threads

When you are running a multiprocess or multithreaded program, there will be many times when you will want to synchronize execution to the same statement. You can do this manually using a *hold* command, or you can let TotalView do this by setting a barrier point.

When a process or a thread is *held*, any command that it receives that tells it to execute are ignored. For example, assume that you place a hold on a process in a control group that contains three processes. After you select **Group > Control > Go**, two of the three processes will resume executing. The held process ignores the **Go** command.

At a later time, you will want to run whatever is being held. Do this using a **Release** command. When you release a process or a thread, you are telling it that it can run. But you still need to tell it to execute, which means that it is waiting to receive an execution command such as **Go**, **Out**, or **Step**.

Manually holding and releasing processes and threads is useful in these instances:

- When you need to run a subset of the processes and threads. You can manually hold all but the ones you want to run.

- When a process or thread is held at a barrier point and you want to run it without first running all the other processes or threads in the group to that barrier. In this case, you'd release the process or the thread manually and then run it.

TotalView can also hold a process or thread if it stops at a barrier breakpoint. You can manually release a process or thread being held at a barrier breakpoint. See "Barrier Points" on page 350 for more information on manually holding and releasing barrier breakpoint.

When TotalView is holding a process, the Root and Process Windows display a held indicator, which is the letter **H**. When TotalView is holding a thread, it displays the letter **h**.

Here are four ways to hold or release a thread, process, or group of processes:

- You can hold a group of processes with the **Group > Hold** command.
- You can release a group of processes with the **Group > Release** command.
- You can toggle the hold/release state of a process by selecting and clearing the **Process > Hold** command.
- You can toggle the hold/release state of a thread by selecting and clearing the **Thread > Hold** command.

CLI EQUIVALENT: **dhold and dunhold**
Setting the focus changes the scope.

If a process or a thread is running when you enter a hold or release command, TotalView stops the process or thread, and then holds it.

NOTE TotalView allows you to hold and release processes independently from threads.

Notice that the Process pulldown menu contains a **Hold Threads** and a **Release Threads** command. While they appear to do the same thing, they are used in a slightly different way. If you use Hold Threads on a multi-threaded process, you'll be placing a hold on all threads. This is seldom what you want. If you then uncheck the **Threads > Hold** command, TotalView allows you to release the one you want. This is an easy way to select one or two threads when your program has a lot of

threads. You can verify that you're doing the right thing by looking at the thread's status in the Root Window's Attached pane.

CLI EQUIVALENT: `dhold -thread`
`dhold -process`
`dunhold -thread`

Examining Groups

When you debug a multiprocess program, TotalView adds processes to both a control and a share group as the process starts. (See Chapter 2, "Understanding Threads, Processes, and Groups" on page 17 for information on groups.)

NOTE These groups are not related to either UNIX process groups or PVM groups.

Because a program can have more than one control group and more than one share group, it decides where to place a process based on the type of system call (`fork()` or `execve()`) that created or changed the process. The two types of process groups are:

Control Group Includes the parent process and all related processes. A control group includes children that a process forks (processes that share the same source code as the parent). It also includes forked children that subsequently call a function such as `execve()`. That is, a control group can contain processes that don't share the same source code as the parent.

Control groups also include processes created in parallel programming disciplines like MPI.

Share Group Is the set of processes in a control group that share the same source code. Members of the same share group share action points.

NOTE See Chapter 11, "Using Groups, Processes, and Threads" on page 239 for a complete discussion of groups.

TotalView automatically creates share groups when your processes fork children that call `execve()` or when your program creates processes that use the same code as some parallel programming models such as MPI do.

TotalView names processes according to the name of the source program. Here are the naming rules it uses:

- It names the parent process after the source program.
- The name for forked child processes differs from the parent in that TotalView appends a numeric suffix (*.n*). If you're running an MPI program, the numerical suffix is the process's rank in `COMM_WORLD`.
- If a child process calls `execve()` after it is forked, TotalView places a new executable name within angle brackets (`<>`).

In Figure 126, assume that the **generate** process doesn't fork any children, and that the **filter** process forks two child process. Later, the first child forks another child and then calls `execve()` to execute the **expr** program. In this figure, the middle column shows the names that TotalView will use.

Process Groups	Process Names	Relationship	
Control Group 1	<ul style="list-style-type: none"> └─ Share Group 1 └─ Share Group 2 	<ul style="list-style-type: none"> └─ filter └─ filter.1 └─ filter.2 └─ filter<expr>.1.1 	<ul style="list-style-type: none"> parent process #1 child process #1 child process #2 grandchild process #1
Control Group 2	<ul style="list-style-type: none"> └─ Share Group 3 	<ul style="list-style-type: none"> └─ generate 	<ul style="list-style-type: none"> parent process #2

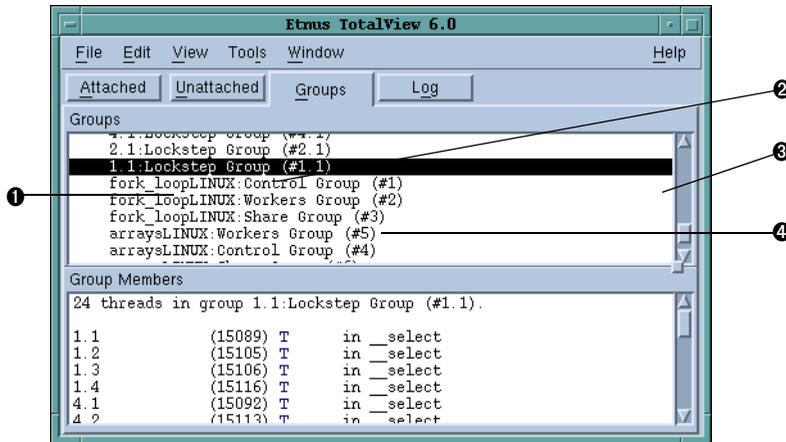
FIGURE 126: Example of Control Groups and Share Groups

Displaying Groups

To display a list of process and thread groups, select the Root Window's Groups Tab. (See Figure 127.)

CLI EQUIVALENT: `dptsets`

When you select a group in the top list pane, TotalView updates the bottom pane to show the group's members. After TotalView updates the bottom pane, you can dive into anything shown there.



- ❶ Name of executable
- ❷ Type of process or thread group
- ❸ Select a group in the top pane to display members in the bottom pane
- ❹ Group number

FIGURE 127: Root Window: Groups Page

Placing Processes into Groups

TotalView uses your executable's name to determine the share group the program belongs to. If the path names are identical, it assumes they are the same program. If the path names differ, TotalView assumes they are different even if the filename within the path name is the same and will place them in different share groups.

Starting Processes and Threads

To start a process, go to the Process Window and select a **Go** command from the **Group**, **Process**, or **Thread** pulldown menus.

After you select a **Go** command, TotalView decides what it will run based on the current thread. It uses this thread, which is called the Thread of Interest (TOI), to decide what other threads it should run. For example, if you enter **Group >**

Workers > Go, TotalView continues all threads in the workers group associated with this thread.

CLI EQUIVALENT: **dfocus g dgo**
Abbreviation: **G**

The commands you will use most often are **Group > Go** and **Process > Go**. The **Group > Go** command creates and starts the current process and all other processes in the multiprocess program. There are some limitations, however. TotalView only resumes a process if it:

- Is not being held.
- Already exists and is stopped.
- Is at a breakpoint.

Using a **Group > Go** command on a process that's already running starts the other members of the process's control group.

CLI EQUIVALENT: **dgo**

If the process hasn't yet been created, a **Go** command creates and starts it. *Starting* a process means that all threads in the process resume executing unless you are individually holding a thread.

NOTE TotalView disables the **Thread > Go** command if asynchronous thread control is not available. If you enter a thread-level command in the CLI when asynchronous thread controls aren't available, TotalView will try to perform an equivalent action. For example, it will continue a process instead of a thread.

For a single-process program, **Process > Go** and **Group > Go** are equivalent. For a single-threaded process, **Process > Go** and **Thread > Go** are equivalent.

Creating a Process Without Starting It

The **Process > Create** command creates a process and stops it before the first statement in your program executes. If you had linked a program with shared libraries, TotalView allows the dynamic loader to map into these libraries. Creating a process without starting it is useful when you need to:

- Create watchpoints or change the values of global variables after a process is created, but before it runs.
- Debug C++ static constructor code.

CLI EQUIVALENT: **dstepi**

While there is no equivalent to the **Process > Create** command, executing **dstepi** produces the same effect.

Creating a Process by Single-Stepping

The TotalView single-stepping commands allow you to create a process and run it to a location in your programs. The single-stepping commands available from the **Process** menu are as shown in the following table:

Table 10: Creating a Process by Stepping

GUI Command	CLI Command	Creates the process and ...
Process > Step	dfocus p dstep	Runs it to the first line of the main() routine.
Process > Next	dfocus p dnest	Runs it to the first line of the main() routine; this is the same as Process > Step .
Process > Step Instruction	dfocus p dstepi	Stops it before any of your program executes.
Process > Next Instruction	dfocus p dnexti	Runs it to the first line of the main() routine. this is the same as Process > Step .
Process > Run To	dfocus p duntil	Runs it to the line or instruction selected in the Process Window.

If a group-level or thread-level stepping command creates a process, the behavior is the same as if it were a process-level command.

NOTE Chapter 11, *"Using Groups, Processes, and Threads"* on page 239 contains a detailed discussion of setting the focus for stepping commands.

Stepping and Setting Breakpoints

Several of the single-stepping commands require that you select a source line or machine instruction in the Source Pane. To select a source line, place the cursor over the line and click your left mouse button. If you select a source line that has

more than one instantiation, TotalView will try to do the right thing. For example, if you select a line within a template so you can set a breakpoint on it, you'll actually set a breakpoint on all instantiations of the template. If this isn't what you want, you can use the **Location** button in the **Action Point > Properties** Dialog Box to change which instantiations will have the breakpoints. (See Figure 128 on page 228.)

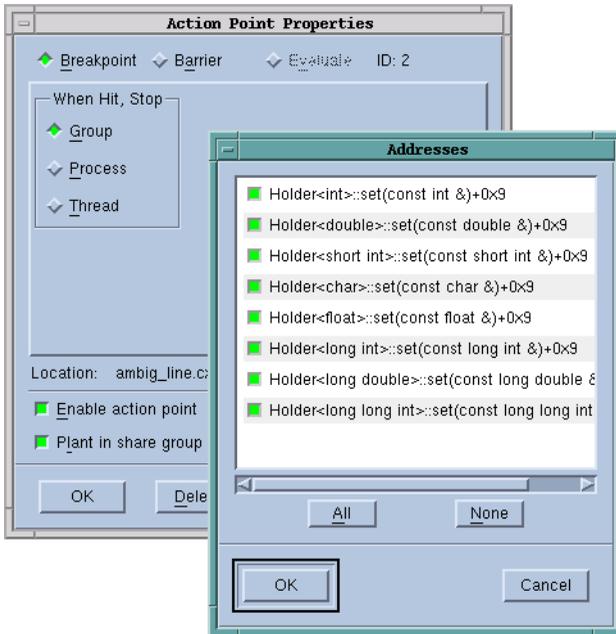


FIGURE 128: Action Point Properties and Address Dialog Boxes

If TotalView can't figure out which instantiation to set a breakpoint at, it will display its Ambiguous Line Dialog Box. (See Figure 129 on page 229.)

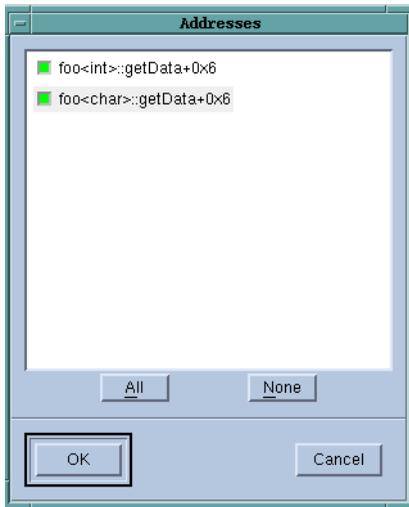


FIGURE 129: Ambiguous Addresses Dialog Box

Using Stepping Commands

While different programs have different requirements, the most common stepping mode is to set group focus to **Control** and the target to **Process** or **Group**. You can now select stepping commands from the **Process** or **Group** menus or use commands in the toolbar.

CLI EQUIVALENT: `dfocus g`
`dfocus p`

Here are some things to remember about single-stepping commands:

- To cancel a single-step command, put the cursor in the Process Window and select **Group > Halt** or **Process > Halt**.

CLI EQUIVALENT: `dhalt`

- If your program reaches a breakpoint while stepping over a function, TotalView cancels the operation and your program stops at the breakpoint.

- If you enter a source-line stepping command and the primary thread is executing in a function that has no source-line information, TotalView performs an assembler-level stepping command.

When TotalView steps through your code, it steps a line at a time. This means that if you have more than one statement on a line, a step instruction executes all of the instructions.

Stepping into Function Calls

The stepping commands execute one line in your program. If you are using the CLI, you can use a numeric argument that indicates how many source lines TotalView should step. For example, here's the CLI instruction for stepping three lines:

```
dstep 3
```

If the source line or instruction contains a function call, TotalView steps into it. If TotalView can't find the source code and the function was compiled with `-g`, it displays the function's machine instructions.

It is possible that you might not realize your program is calling a function. For example, if you've overloaded an operator, you'll step into the code that defines the overloaded operator.

NOTE If the function being stepped into wasn't compiled with `-g`, TotalView will always step over the function.

The TotalView GUI has eight **Step** commands and eight **Step Instruction** commands. These commands are located on the **Group**, **Process**, and **Thread** pull-downs. The difference is the focus.

```
CLI EQUIVALENT:  dfocus ... dstep  
                 dfocus ... dstepi
```

Stepping Over Function Calls

When you step over a function, TotalView stops execution when the primary thread returns from the function and reaches the source line or instruction after the function call.

The TotalView GUI has eight **Next** commands that execute a single source line while stepping over functions, and eight **Next Instruction** commands that execute a single machine instruction while stepping over functions. These commands are on the **Group**, **Process**, and **Thread** menus.

CLI EQUIVALENT: **dfocus ... dnext**
dfocus ... dnexti

Executing to a Selected Line

If you don't need to stop execution every time execution reaches a specific line, you can tell TotalView to run your program to a selected line or machine instruction. After selecting the line on which you want the program to stop, invoke one of the eight **Run To** commands defined within the TotalView GUI. These commands are on the **Group**, **Process**, and **Thread** menus.

CLI EQUIVALENT: **dfocus ... duntil**

Executing to a selected line is discussed in greater depth in Chapter 11, "Using Groups, Processes, and Threads" on page 239.

If your program reaches a breakpoint while running to a selected line, TotalView stops at that breakpoint.

If your program calls recursive functions, you can select a nested stack frame in the Stack Trace Pane. When you do this, TotalView determines where to stop execution by looking at:

- The frame pointer (FP) of the selected stack frame.
- The selected source line or instruction to determine.

CLI EQUIVALENT: **dup and ddown**

Executing to the Completion of a Function

You can step your program out of a function by using the **Out** commands. The eight commands within the TotalView GUI are located on the **Group**, **Process**, and **Thread** menus.

CLI EQUIVALENT: `dfocus ... dout`

If the source line that is the *goal* of the **Out** operation has more than one statement, TotalView will stop execution just after the routine from which it just emerged. For example, suppose this is your source line:

```
routine1; routine2
```

Suppose you step into **routine1**, then use an **Out** command. While the PC arrow hasn't moved, the actual PC is just after **routine1**. This means that if you use a step command, you will step into **routine2**.

TotalView's PC arrow will not move, when the source line only has one statement on it. The internal PC will, of course, have changed.



You can also return out of several functions at once, by selecting the routine in the Stack Trace Pane that you want to go to, and then selecting an **Out** command.

If your program calls recursive functions, you can select a nested stack frame in the Stack Trace Pane to indicate which instance you'll be running out of.

Displaying Thread and Process Locations

You can see which processes and threads in the share group are at a location by selecting a source line or machine instruction in the Source Pane of the Process Window. TotalView dims thread and process information in the Root Window's **Attached** Page for share group members if the thread or process is not at the selected line. TotalView considers a process to be at the selected line if any of the threads in the process are at that line. Selecting a line in the Process Window that is already selected removes the dimming in the **Attached** Page.

CLI EQUIVALENT: `dstatus`

The **Attached** Page reflects the line that you last selected. If you have several Process Windows open, the information in the **Attached** Page will change depending on the line you selected last in each Process Window. The display can also change after an operation that changes the process state or when you issue a **Window > Update** command.

Figure 130 on page 233 shows an **Attached** Page with dimmed process information. In this example, the parallel program was run to a barrier breakpoint, and one process (`dmpirun<cpu>.1`) was stepped to the next source line.

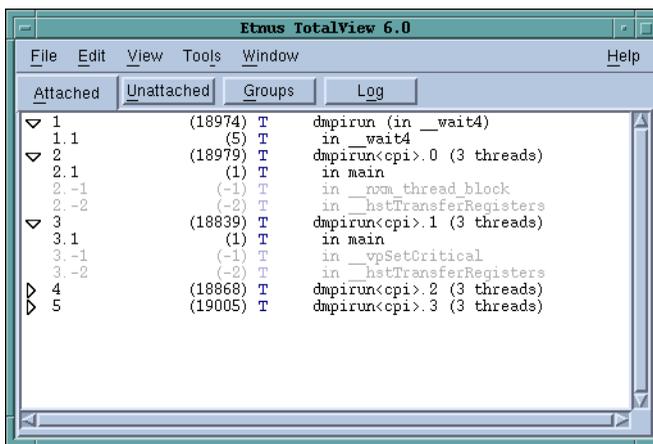


FIGURE 130: Dimmed Process Information in the Root Window

Since the MPI starter process (`dmpirun`) isn't in the same share group as the processes running the `cpu` program, TotalView doesn't dim its process information.

Continuing with a Specific Signal

Letting your program continue after sending it a signal is useful when your program contains a signal handler. Here's how you tell TotalView to do this:

- 1 Select the Process Window's **Thread > Continuation Signal** command. (See Figure 131.)

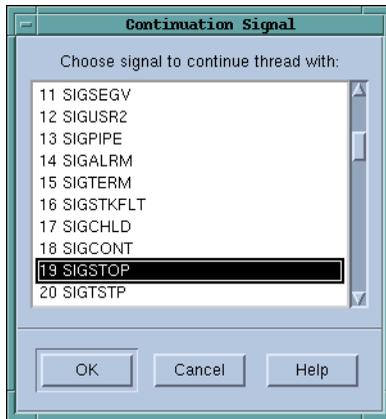


FIGURE 131: Thread > Continuation Signal Dialog Box

- 2 Select the signal to be sent to the thread and then select **OK**.

The continuation signal is set for the thread contained in the current Process Window. If the operating system can deliver multi-threaded signals, you can set a separate continuation signal for each thread. If it can't, this command clears continuation signals set for other threads in the process.

- 3 Continue execution of your program with commands such as **Process > Go**, **Step**, **Next**, or **Detach**.

TotalView continues the threads and sends it the specified signals.

NOTE You can clear the continuation signal by selecting signal 0.

Deleting Programs

To delete all the processes in a control group, use the **Group > Delete** command. The next time you start the program, for example, by using the **Process > Go** command, TotalView creates and starts a fresh master process.

CLI EQUIVALENT: `dfocus g dkill`

Restarting Programs

You can use the **Group > Restart** command to restart a program that is running or one that is stopped but hasn't exited.

CLI EQUIVALENT: **drerun**

If the process is part of a multiprocess program, TotalView deletes all related processes, restarts the master process, and runs the newly created program.

The **Group > Restart** command is equivalent to the **Group > Delete** command followed by the **Process > Go** command.

Checkpointing Programs and Processes

On SGI IRIX platforms, you can save the state of selected processes and then use this saved information to restart the processes from the position where they were saved. For more information, see the Process Window's **Tools > Create Checkpoint** and **Tools > Restart Checkpoint** commands in TotalView's Help information. (See Figure 132.)

CLI EQUIVALENT: **dcheckpoint**
drestart

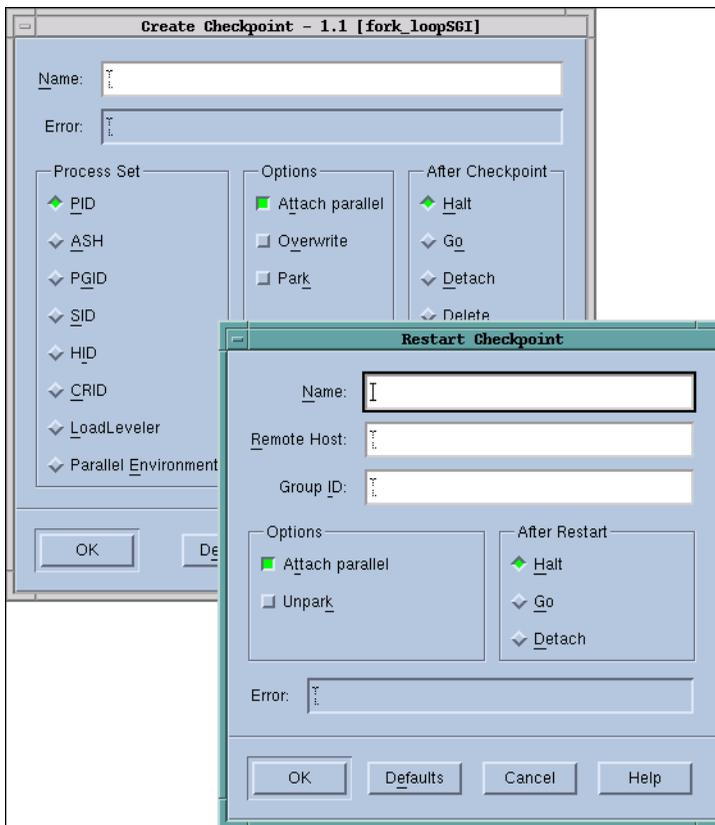


FIGURE 132: Checkpoint and Restart Dialog Boxes

Setting the Program Counter

TotalView lets you resume execution at a different statement than the one at which it stopped execution by resetting the value of the program counter (PC). For example, you might want to skip over some code, execute some code again after changing certain variables, or restart a thread that is in an error state.

Setting the PC can be crucial when you want to restart a thread that is in an error state. Although the PC icon in the line number area points to the source statement that caused the error, the PC actually points to the failed machine instruction in the source statement. You need to explicitly reset the PC to the correct instruction.

(You can verify the actual location of the PC before and after resetting it by display-

ing it in the Stack Frame Pane or displaying both source and assembler code in the Source Pane.)

In TotalView, you can set the PC of a stopped thread to a selected source line, a selected instruction, or an absolute value (in hexadecimal). When you set the PC to a selected line, the PC points to the memory location where the statement begins. For most situations, setting the PC to a selected line of source code is all you need to do.

To set the PC to a selected line:

- 1 If you need to set the PC to a location somewhere in a line of source code, display the **View > Source As > Both** command. TotalView responds by displaying the assembler code.
- 2 Select the source line or instruction in the Source Pane. TotalView highlights the line.
- 3 Select the **Thread > Set PC** command. TotalView asks for confirmation, resets the PC, and moves the PC icon to the selected line.

When you select a line and ask TotalView to set the PC to that line, TotalView attempts to force the thread to continue execution at that statement in the currently selected stack frame. If the currently selected stack frame is not the top stack frame, TotalView asks if it can unwind the stack:

```
This frame is buried. Should we attempt to unwind the
stack?
```

If you select **Yes**, TotalView discards deeper stack frames (that is, all stack frames that are more deeply nested than the selected stack frame) and resets the machine registers to their values for the selected frame. If you select **No**, TotalView sets the PC to the selected line, but it leaves the stack and registers in their current state. Since you can't assume that the stack and registers have the right values, selecting **No** is almost always the wrong thing to do.

Interpreting Status and Control Registers

The Stack Frame Pane in the Process Window lists the contents of CPU registers for the selected frame—you might need to scroll past the stack local variables to see

them. To learn about the meaning of these registers, you need to consult the user's guide for your CPU and "Architectures" in the TOTALVIEW REFERENCE GUIDE.

CLI EQUIVALENT: `dprint register`
You must quote the initial \$ character in the register name; for example, `dprint \"$r1`.

For your convenience, TotalView displays the bit settings of many CPU registers symbolically. For example, TotalView symbolically displays registers that control rounding and exception enable modes. You can edit the values of these registers and continue execution of your program. For example, you might do this to examine the behavior of your program with a different rounding mode.

Since the registers that are displayed vary from platform to platform, see "Architectures" in the TOTALVIEW REFERENCE GUIDE for information on how TotalView displays this information on your CPU. For general information on editing the value of variables (including registers), refer to "Displaying Areas of Memory" on page 289.

Chapter 11

Using Groups, Processes, and Threads

While the specifics of how multiprocess, multithreaded programs execute differ greatly from platform to platform and environment to environment, all share some general characteristics. This chapter discusses TotalView's process/thread model. It also describes the way in which you tell the GUI and the CLI what processes and threads it should direct a command to.

The topics discussed in this chapter are:

- *"Defining the GOI, POI, and TOI"* on page 239
- *"Setting a Breakpoint"* on page 241
- *"Stepping (Part I)"* on page 241
- *"Using P/T Set Controls"* on page 245
- *"Setting Process and Thread Focus"* on page 247
- *"Setting Group Focus"* on page 253
- *"Stepping (Part II): Some Examples"* on page 268
- *"Using P/T Set Operators"* on page 270
- *"Using the P/T Set Browser"* on page 271
- *"Using the Group Editor"* on page 275

Defining the GOI, POI, and TOI

This chapter consistently uses three related acronyms:

- GOI, which means Group of Interest
- POI, which means Process of Interest
- TOI, which means Thread of Interest

These terms are important in TotalView's process/thread model because TotalView must determine the scope of what it will do when executing a command. For example, Chapter 2 introduced the kinds of groups contained within TotalView. For reasons that will become obvious in this chapter, that chapter ignored what happens when you execute a TotalView command upon a group. For example, what does "stepping a group" actually mean? Which processes and threads will TotalView actually step? What happens to processes and threads that aren't in this group?

Associated with these three terms is a fourth: *arena*. The *arena* is the collection of processes, threads, and groups that are affected by a debugging command. This collection is called an *arena list*.

In the GUI, the arena is most often set using the two pulldown menus in the toolbar. If you examine the menubar, you'll see that there are 8 *next* commands. The difference between them is the arena; that is, the difference between the *next* commands is the processes and threads that are the target of what the *next* command runs.

When TotalView executes any action command, the arena decides the scope of what can run. It doesn't, however, determine what will run. Depending on the command, TotalView determines the TOI, POI, or GOI, and then executes the command's action upon that thread, process, or group. For example, assume that you tell TotalView to step the current control group.

- TotalView needs to know what the TOI is so it can determine what threads are in the lockstep group—TotalView only allows you to step a lockstep group.
- The lockstep group is part of a share group.
- This share group is also contained in a control group.

So, by knowing what the TOI is, the TotalView GUI also knows what the GOI is. This is important because, as you will see, while TotalView now knows what it will step (the threads in the lockstep group), it also knows what it will allow to run freely while it is stepping these threads. In the CLI, the P/T set determines the TOI.

Using the GOI, POI, and TOI will become clearer as you read the rest of this chapter.

Setting a Breakpoint

You can set breakpoints in your program by selecting the boxed line numbers in the Source Code pane of a Process window. A boxed line number indicates that the line generates executable code. A **STOP** icon masking a line number indicates that a breakpoint is set on the line. Selecting the **STOP** icon clears the breakpoint.

When a program reaches a breakpoint, it stops. You can let the program resume execution in any of the following ways:

- Use single-step commands described in *"Using Stepping Commands"* on page 229.
- Use the set program counter command to resume program execution at a specific source line, machine instruction, or absolute hexadecimal value. See *"Setting the Program Counter"* on page 236.
- Set breakpoints at lines you choose and allow your program to execute to that breakpoint. *"Setting Breakpoints and Barriers"* on page 339.
- Set conditional breakpoints that cause a program to stop after it evaluates a condition that you define, for example "stop when a value is less than 8. See *"Setting Evaluation Points"* on page 356.

TotalView provides additional features for working with breakpoints, process barrier breakpoints, and evaluation points. For more information, refer to Chapter 14, *"Setting Action Points"* on page 337.

Stepping (Part I)

TotalView's stepping commands allow you to:

- Execute one source line or machine instruction at a time; for example, **Process > Step** in the GUI and **dstep** in the CLI.

CLI Equivalent: **dstep**

- Run to a selected line, which acts like a temporary breakpoint; for example, **Process > Run To**.

CLI Equivalent: **duntil**

- Run until a function call returns. For example, **Process > Out**.

CLI Equivalent: `dout`

In all cases, stepping commands operate on the Thread of Interest (TOI). In the GUI, the TOI is the selected thread in the current Process Window. In the CLI, the TOI is the thread that TotalView uses to determine the scope of the stepping operation.

On all platforms except SPARC Solaris, TotalView uses *smart* single-stepping to speed up stepping of one-line statements containing loops and conditions, such as Fortran 90 array assignment statements. *Smart* stepping occurs when TotalView realizes that it doesn't need to step through an instruction. For example, assume that you have the following statements:

```
integer iarray (1000,1000,1000)
iarray = 0
```

These two statements define one billion scalar assignments. If your machine steps every instruction, you will probably never get past this statement. *Smart* stepping means that TotalView will single-step through the assignment statement at a speed that is very close to your machine's native speed.

Other topics in this section are:

- "Group Width" on page 242
- "Process Width" on page 243
- "Thread Width" on page 243

Group Width

TotalView's behavior when stepping at group width depends on whether the Group of Interest (GOI) is a process group or a thread group. In the following lists, *goal* means the place at which things should stop executing. For example, if you are doing a *step* command, it is the next line. If it is a *run to* command, it is the selected line.

If the GOI is a:

- *Process group*, TotalView examines the group and identifies which of its processes has a thread stopped at the same location as the TOI (a *matching* process). TotalView runs these matching processes until one of its threads arrives at the goal. When this happens, TotalView stops the thread's process. The command finishes when it has stopped all of these "matching" processes.
- *Thread group*, TotalView runs all processes in the control group. However, as each thread arrives at the goal, TotalView just stops that thread; the rest of the threads in the same process continue executing. The command finishes when all threads in the GOI arrive at the goal. When the command finishes, TotalView will stop all processes in the control group.
TotalView doesn't wait for threads that are not in the same share group as the TOI since they are executing different code and can never arrive at the goal.

Process Width

TotalView's behavior when stepping at process width (which is the default) depends on whether the Group of Interest (GOI) is a process group or a thread group. If the GOI is a:

- *Process group*, TotalView runs all threads in the process, and execution continues until the TOI arrives at its goal, which can be the next statement, the next instruction, and so on. Only when the TOI reaches the goal will TotalView stop the other threads in the process.
- *Thread group*, TotalView allows all threads in the GOI and all manager threads to run. As each member of the GOI arrives at the goal, TotalView stops it; the rest of the threads continue executing. The command finishes when all members of the GOI arrive at the goal. At that point, TotalView stops the whole process.

Thread Width

When TotalView performs a stepping command, it decides what it will step based on the *width*. Using the toolbar, you specify width using the left-most pulldown. This pulldown has three items: **Group**, **Process**, and **Thread**.

Stepping at thread width tells TotalView that it should just run that thread. It does not step other threads. In contrast, process width tells TotalView that it should run

all threads in the process that are allowed to run while the TOI is stepped. While TotalView is stepping the thread, manager threads are running freely.

Stepping a thread isn't the same as stepping a thread's process because a process can have more than one thread.

NOTE Thread-stepping is not implemented on Sun platforms. On SGI platforms, thread-stepping is not available with pthread programs. If, however, your program's parallelism is based on SGI's sprocs, thread-stepping is available.

Thread-level single-step operations can fail to complete if the TOI needs to synchronize with a thread that isn't running. For example, if the TOI requires a lock that another held thread owns, and steps over a call that tries to acquire the lock, the primary thread can't continue successfully. You must allow the other thread to run in order to release the lock. In this case, you should instead use process-width stepping.

Using "Run To" and duntil Commands

The **duntil** and "Run To" commands differ from other step commands when you apply them to a process group. (These commands tells TotalView to execute program statements *until* a selected statement is reached.) When applied to a process group, TotalView identifies all processes in the group already having a thread stopped at the goal. These are the *matching* processes. TotalView then runs only the nonmatching processes. Whenever a thread arrives at the goal, TotalView stops its process. The command finishes when it has stopped all members of the group. This lets you *sync up* all the processes in a group in preparation for group-stepping them.

Here is what you should know if you're running at process width:

- Process group** If the Thread of Interest (TOI) is already at the goal location, TotalView steps the TOI past the line before the process is run. This allows you to use the **Run To** command repeatedly within loops.
- Thread group** If any thread in the process is already at the goal, TotalView temporarily holds it while other threads in the process run. After all threads in the thread group reach the goal, TotalView stops the process. This allows you to synchronize the threads in the POI at a source line.

If you're running at group width:

- Process group** TotalView examines each process in the process and share group to determine if at least one thread is already at the goal. If a thread is at the goal, TotalView holds its process. Other processes are allowed to run. When at least one thread from each of these processes is held, the command completes. This lets you synchronize at least one thread in each of these processes at a source line. If you're running a control group, this synchronizes all processes in the share group.
- Thread group** TotalView examines all the threads in the thread group that are in the same share group as the TOI to determine if a thread is already at the goal. If it is, TotalView holds it. Other threads are allowed to run. When all of the threads in the TOI's share group reach the goal, TotalView stops the TOI's *control* group and the command completes. This lets you synchronize thread group members. If you're running a workers group, this synchronizes all worker threads in the share group.

The process stops when the TOI and at least one thread from each process in the group or process reach the command stopping point. This lets you synchronize a group of processes and bring them to one location.

You can also run to a selected line in a nested stack frame, as follows:

- 1 Select a nested frame in the Stack Trace Pane.
- 2 Select a source line or instruction in the function.
- 3 Issue a **Run To** command.

TotalView executes the primary thread until it reaches the selected line in the selected stack frame.

Using P/T Set Controls

A few TotalView windows have P/T set control elements. For example, Figure 133 on page 246 shows the top portion of the Process Window.

This pull-down menu differs from the P/T set controls on other elements. On other windows, there are two pull-downs. However, in the context of the Process Window, elements from the two pull-downs have been combined both to eliminate actions

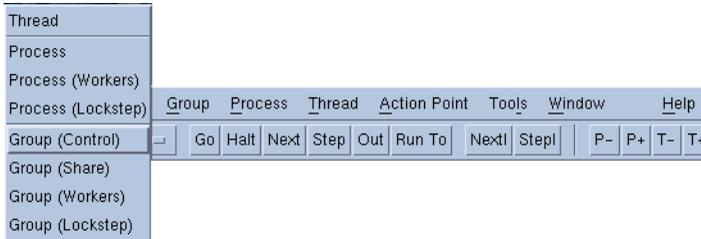


Figure 133: The P/T Set Control in the Process Window

that don't have meaning. When you select a group and a modifier, you are telling TotalView that when you press one of the remaining buttons on the toolbar, this element names the focus upon which TotalView will act. For example, if Thread is selected and you select Step, TotalView steps the current thread. If Process (workers) is selected and you select Halt, TotalView halts all processes associated with the current threads workers group. If you were running a multiprocess program, other processes would continue to execute.

Here's what the controls look like in other windows may have additional P/T set controls.

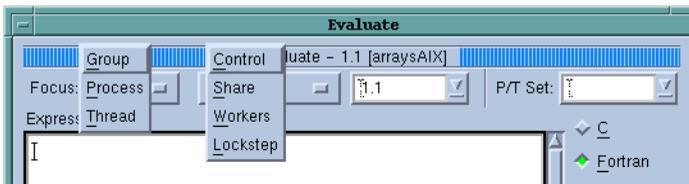


Figure 134: The P/T Set Control in the Tools > Evaluate Window

The first pull-down menu, which is called the *Width Pull-down*, has three elements on it: **Group**, **Process**, and **Thread**. Your choices here indicate the width of the command. For example, if **Group** is selected, a **Go** command continues the group. Which group TotalView will continue is set by the choices on the second pull-down menu. The *Width Pull-down* tells TotalView where it should look when it tries to determine what it will manipulate. The second pull-down, which is called the *Scope Pull-down*, tells TotalView which processes and threads within the scope defined by the *Width Pull-down* it should manipulate. For example, you could tell TotalView to step

the threads defined in the current workers group that are contained in the current process.

Finally, the *P/T Selector* (the third pulldown menu from the left) lets you change the focus of the action from the currently defined process and threads to any other process and thread that TotalView controls. That is, this changes the POI and TOI

The **P/T Set** expression box on the right allows you to directly enter a P/T set expression. The focus of what you enter is modified by the other P/T set controls.

What is selected can be quite complicated when you use the GUI to set these controls, or when you specify a focus using the CLI.

Setting Process and Thread Focus

NOTE While the previous sections have emphasized the GUI, this section and the ones that follow emphasize the CLI. In all cases, the selection of what TotalView runs is based directly or indirectly on P/T set syntax. While the focus is obvious in the CLI, it is often buried within the internals of the GUI. Reading the rest of this chapter is important when you want to have full asynchronous debugging control over your program. Having this level of control, however, is seldom necessary.

When it executes a command, TotalView must decide which processes and threads it should act on. Most commands have a default set of threads and processes and, in most cases, you won't want to change the default. In the GUI, the default is the process and thread in the current Process Window. In the CLI, this default is indicated by the focus, which is shown in the CLI's prompt.

There are times, however, when you'll need to change this default. This section begins a rather intensive look at how you tell TotalView what processes and threads it should use as the target of a command.

Topics in this section are:

- "*Process/Thread Sets*" on page 248
- "*Arenas*" on page 249
- "*Specifying Processes and Threads*" on page 250

Process/Thread Sets

All TotalView commands operate on a set of processes and threads. This set is called a P/T (*Process/Thread*) *set*. The right-hand text box in windows containing P/T set controls lets you construct these sets. In the CLI, you specify a P/T set as an argument to a command such as **dfocus**. If you're using the GUI, TotalView creates this list for you based on which Process Window has focus.

Unlike a serial debugger where each command clearly applies to the only executing thread, TotalView can control and monitor many threads with their PCs at many different locations. The P/T set indicates the groups, processes, and threads that are the target of the CLI command. No limitation exists on the number of groups, processes, and threads in a set.

A P/T set is a Tcl list containing one or more P/T identifiers. (The next section, "Arenas" on page 249, explains what a P/T identifier is.) Tcl lets you create lists in two ways:

- You can enter these identifiers within braces { }.
- You can use Tcl commands that create and manipulate lists.

These lists are then used as arguments to a command. If you're entering one element, you usually do not have to use Tcl's list syntax.

For example, the following list contains specifiers for process 2, thread 1, and process 3, thread 2:

```
{p2.1 p3.2}
```

If you do not explicitly specify a P/T set in the CLI, TotalView defines a target set for you. (In the GUI, the default set is determined by the current Process Window.) This set is displayed as the (default) CLI prompt. (For information on this prompt, see "Command and Prompt Formats" on page 203.)

You can change the focus upon which a command acts by using the **dfocus** command. If the CLI executes **dfocus** as a unique command, it changes the default P/T set. For example, if the default focus is process 1, the following command changes the default focus to process 2:

```
dfocus p2
```

After TotalView executes this command, all commands that follow will focus on process 2.

NOTE In the GUI, you'd set the focus by displaying a Process Window containing this process.

If you begin a command with **dfocus**, TotalView changes the target for just the command that follows. After the command executes, TotalView restores the *old* default. The following example shows both of these ways to use the **dfocus** command. Assume that the current focus is process 1, thread 1. The following commands change the default focus to group 2 and then step the threads in this group twice:

```
dfocus g2
dstep
dstep
```

In contrast, if the current focus is process 1, thread 1, the following commands step group 2 and then step process 1, thread 1:

```
dfocus g2 dstep
dstep
```

Some commands only operate at the process level; that is, you cannot apply them to a single thread (or group of threads) in the process but must apply them to all or to none.

Arenas

A P/T identifier often indicates a number of groups, processes, and threads. For example, assume that two threads executing in process 2 are stopped at the same statement. This means that TotalView places the two stopped threads into lockstep groups. If the default focus is process 2, stepping this process actually steps both of these threads.

TotalView uses the term *arena* to define the processes and threads that are the target of an action. In this case, the arena has two threads. Many CLI commands can act on one or more arenas. For example, here is a command with two arenas:

```
dfocus {p1 p2}
```

The two arenas are process 1 and process 2.

So, what is the GOI, POI, and TOI when there is an arena list? In this case, each arena within the list will have its own GOI, POI, and TOI.

Specifying Processes and Threads

A previous section described a P/T set as being a list, but ignored what the individual elements of the list are. A better definition is that a P/T set is a list of arenas, where an *arena* consists of the processes, threads, and groups that are affected by a debugging command. Each *arena specifier* describes a single arena in which a command will act; the *list* is just a collection of arenas. Most commands iterate over the list, acting individually on an arena. Some CLI output commands, however, will combine arenas and act on them as a single target.

An arena specifier includes a *width* and a TOI. (“Widths” are discussed later in this section.) In the P/T set, the TOI specifies a target thread, while the width specifies how many threads surrounding the thread of interest are affected.

The Thread of Interest (TOI)

The TOI is specified as **p.t**, where **p** is the TotalView process ID (PID) and **t** is the TotalView thread ID (TID). The **p.t** combination identifies the POI (Process of Interest) and TOI. The TOI is the primary thread affected by a command. This means that it is the primary focus for a TotalView command. For example, while the **dstep** command always steps the TOI, it can run the rest of the threads in the POI and step other processes in the group.

In addition to using numerical values, you can also use two special symbols:

- The less-than (<) character indicates the *lowest number worker thread* in a process and is used instead of the TID value. If, however, the arena explicitly names a thread group, < means the lowest numbered member of the thread group. This symbol lets TotalView select the first user thread, which may not be thread 1; for example, the first and only user thread may be thread number 3 on HP Alpha systems.
- A dot (.) indicates the current set. While this is seldom used interactively, it can be useful in scripts.

Process and Thread Widths

You can enter a P/T set in two ways. If you're not manipulating groups, the format is:

```
[width_letter][pid][.tid]
```

NOTE The next section extends this format to include groups.

For example, **p2.3** indicates process 2, thread 3.

While the syntax seems to indicate that you do not need to enter any element, TotalView requires that you enter at least one. Because TotalView will try to determine what it can do based on what you type, it will try to fill in what you omit. The only requirement is that when you use more than one element, you use them in the order shown here.

You can leave out parts of the P/T set if what you do enter is unambiguous. A missing width or PID is filled in from the current focus. A missing TID is always assumed to be `<`. For more information, see "Incomplete Arena Specifiers" on page 267.

The *width_letter* indicates which processes and threads are part of the arena. The letters you can use are:

- | | | |
|----------|----------------------|--|
| t | <i>Thread width</i> | A command's target is the indicated thread. |
| p | <i>Process width</i> | A command's target is the process containing the TOI. |
| g | <i>Group width</i> | A command's target is the group containing the POI. This indicates control and share groups. |
| a | <i>All processes</i> | A command's target is all threads in the GOI that are in the POI. |
| d | <i>Default width</i> | A command's target depends on the default for each command. This is also the width to which the default focus is set. For example, the dstep command defaults to process width (run the process while stepping one thread), and the dwhere command defaults to thread width. Default widths are listed in "Default Arena Widths" in the TOTALVIEW REFERENCE GUIDE. |

You must use lowercase letters to enter these widths.

Figure 135 on page 252 illustrates how these specifiers relate to one another.

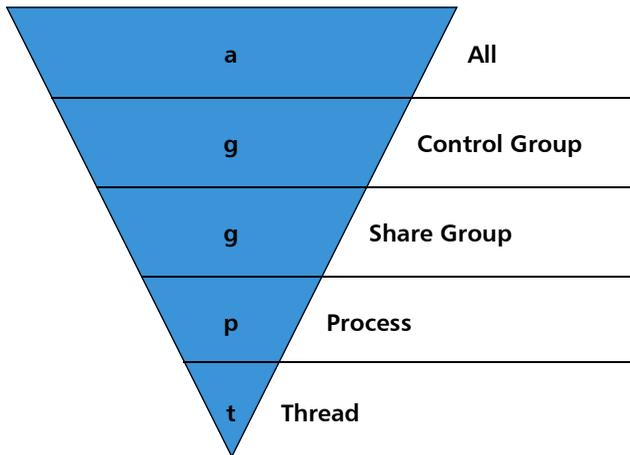


Figure 135: Width Specifiers

Notice that the “g” specifier indicates control and share groups. This inverted triangle is indicating that the arena focuses on a greater number of entities as you move from thread level at the bottom to “all” level at the top.

As mentioned previously, the TOI specifies a target thread, while the width specifies how many threads surrounding the TOI are also affected. For example, the **dstep** command always requires a TOI, but entering this command can:

- Step just the TOI during the step operation (single-thread single-step).
- Step the TOI and step all threads in the process containing the TOI (process-level single-step).
- Step all processes in the group that have threads at the same PC (program counter) as the TOI (group-level single-step).

This list doesn’t indicate what happens to other threads in your program when TotalView steps your thread. For more information, see “*Stepping (Part II): Some Examples*” on page 268.

To save a P/T set definition for later use, assign the specifiers to a Tcl variable. For example:

```
set myset {g2.3 t3.1}
dfocus $myset dgo
```

As the **dfocus** command returns its focus set, you can save this value for later use. For example:

```
set save_set [dfocus]
```

Specifier Examples

Here are some sample specifiers:

- g2.3** Select process 2, thread 3 and set the width to "group".
- t1.7** Commands act only on thread 7 or process 1.
- d1.<** Use the default set for each command, focusing on the first user thread in process 1. The "<" sets the TID to the first user thread.

Setting Group Focus

TotalView has two kinds of groups: process groups and thread groups. Process groups only contain processes, and thread groups only contain threads. The threads in a thread group can be drawn from more than one process.

Topics in this section are:

- *"Specifying Groups in P/T Sets"* on page 255
- *"Arena Specifier Combinations"* on page 257
- *"All' Does Not Always Mean All"* on page 259
- *"Setting Groups"* on page 261
- *"Using the 'g' Specifier: An Extended Example"* on page 263
- *"Focus Merging"* on page 266
- *"Incomplete Arena Specifiers"* on page 267
- *"Lists with Inconsistent Widths"* on page 267

TotalView has four predefined groups. Two of these only contain processes while the other two only contain threads. As you will see, TotalView also allows you to create your own groups, and these groups can have elements that are processes and threads. The predefined process groups are:

■ Control Group

Contains the parent process and all related processes. A control group includes children that were forked (processes that share the same source code as the parent) and children that were forked but which subsequently called `execve()`.

Assigning a new value to the `CGROUP(dpid)` variable for a process changes that process's control group. In addition, the `dgroups -add` command lets you add members to a group in the CLI. In the GUI, you use the `Group > Edit` command.

■ Share Group

Contains all members of a control group that share the same executable image. TotalView automatically places processes in share groups based on their control group and their executable image.

NOTE You can't change a share group's members. In addition, dynamically loaded libraries may vary between share group members.

The predefined thread groups are:

■ Workers Group

Contains all worker threads from all processes in the control group. The only threads not contained in a worker's group are your operating system's manager threads.

■ Lockstep Group

Contains every stopped thread in a share group that has the same PC. TotalView creates one lockstep group for every thread. For example, suppose two threads are stopped at the same PC. TotalView will create two lockstep groups. While each lockstep group has the same two members, they differ in that each has a different TOI. While there are some circumstances where this is important to you, you can ignore this distinction in most cases. That is, while two lockstep groups exist if two threads are stopped at the same PC, ignoring the second lockstep group is almost never harmful.

The group ID's value for a lockstep group differs from the ID of other groups. Instead of having an automatically and transient integer ID, the lockstep group ID is `pid.tid`, where `pid.tid` identifies the thread with which it is associated. For example, the lockstep group for thread 2 in process 1 is `1.2`.

In general, if you're debugging a multiprocess program, the control group and share group differ only when the program has children that it forked with by calling `execve()`.

Specifying Groups in P/T Sets

This section extends the arena specifier syntax to include groups.

If you do not include a group specifier, the default is the control group. For example, the CLI only displays a target group in the focus string if you set it to something other than the default value.

NOTE Target group specifiers are most often used with the stepping commands, as they give these commands more control over what's being stepped.

Here is how you add groups to an arena specifier:

```
[width_letter][group_indicator][pid][.tid]
```

This format adds the *group_indicator* to the what was discussed in "Specifying Processes and Threads" on page 250.

In the description of this syntax, everything appears to be optional. But, while no single element is required, you must enter at least one element. TotalView will determine other values based on the current focus.

TotalView lets you identify a group by using a letter, number, or name.

A Group Letter

You can name one of TotalView's predefined sets. These sets are identified by letters. For example, the following command sets the focus to the workers group:

```
dfocus W
```

The group letter, which is always uppercase, can be:

C *Control group*

All processes in the control group.

D *Default control group*

All processes in the control group. The only difference between this specifier and the **C** specifier is that **D** tells the CLI that it should not display a group letter within the CLI prompt.

S *Share group*

The set of processes in the control group that have the same executable as the arena's TOI.

- W** *Workers group*
The set of all worker threads in the control group.
- L** *Lockstep group*
A set containing all threads in the share group that have the same PC as the arena's TOI. If you step these threads as a group, they will proceed in lockstep.

A Group Number

You can identify a group by the number TotalView assigns to it. For example, here is how you set the focus to group 3:

```
dfocus 3/
```

Notice the trailing slash. This slash lets TotalView know that you're specifying a group number instead of a PID. The slash character is optional if you're using a *group_letter*. However, you must use it as a separator when entering a numeric group ID and a **pid.tid** pair. For example, the following example identifies process 2 in group 3:

```
p3/2
```

A Group Name

You can name a set that you define. You enter this name with slashes. For example, here is how you would set the focus to the set of threads contained in process 3 and that are also contained in a group called **my_group**:

```
dfocus p/my_group/3
```

Arena Specifier Combinations

The following table lists what's selected when you use arena and group specifiers to step your program.

Table 11: Specifier Combinations

Specifier	Specifies
aC	All threads.
aS	All threads.
aW	All threads in all workers groups.
aL	All threads. Every thread is a member of a control group and a member of a share group and a member of a lockstep group. Consequently, three of these definitions mean "all threads."
gC	All threads in the Thread of Interest's (TOI) control group.
gS	All threads in the TOI's share group.
gW	All worker threads in the control group containing the TOI.
gL	All threads in the same share group within the process containing the TOI that have the same PC and the TOI.
pC	All threads in the control group of the Process of Interest (POI). This is the same as gC .
pS	All threads in the process that participate in the same share group as the TOI.
pW	All worker threads in the POI.
pL	All threads in the POI whose PC is the same as the TOI.
tC	Very little. These four combinations, while syntactically correct, are meaningless. The t specifier overrides the group specifier. So, for example, tW and t both name the current thread.
tS	
tW	
tL	

NOTE Stepping commands behave differently if the group being stepped is a process group rather than a thread group. For example, "aC" and "aS" perform the same action while "aL" is different.

If you don't add a PID or TID to your arena specifier, TotalView does it for you, taking the PID and TID from the current or default focus.

Here are some additional examples. These add PIDs and TIDs to the specifier combinations just discussed:

<code>pW3</code>	All worker threads in process 3.
<code>pW3.<</code>	All worker threads in process 3. Notice that the focus of this specifier is the same as the previous example's.
<code>gW3</code>	All worker threads in the control group containing process 3. Notice the difference between this and <code>pW3</code> , which restricts the focus to one of the processes in the control group.
<code>gL3.2</code>	All threads in the same <i>share</i> group as process 3 that are executing at the same PC as thread 2 in process 3. The reason this is a share group and not a control group is that different share groups can reside only in one control group.
<code>/3</code>	Specifies processes and threads in process 3. As the arena width, POI, and TOI are inherited from the existing P/T set, the exact meaning of this specifier depends on the previous context. While the "/" is unnecessary because no group is indicated, it is syntactically correct.
<code>g3.2/3</code>	The 3.2 group ID is the name of the lockstep group for thread 3.2. This group includes all threads in process 3's share group that are executing at the same PC as thread 2.
<code>p3/3</code>	Sets the process to process 3. The Group of Interest (GOI) is set to group 3. If group 3 is a process group, most commands ignore the group setting. If group 3 is a thread group, most commands act on all threads in process 3 that are also in group 3. Setting the process with an explicit group should be done with care, as what you get may not be what you expect given that commands, depending on their scope, must look at the TOI, POI, and GOI.

NOTE Specifying thread width with an explicit group ID probably doesn't mean much.

In the following examples, the first argument to the `dfocus` command defines a temporary P/T set that the CLI command (the last term) will operate on. The `dstatus` command lists information about processes and threads. These examples assume that the global focus was "`d1.<`" initially.

`dfocus g dstatus`

Displays the status of all threads in the control group.

dfocus gW dstatus

Displays the status of all worker threads in the control group.

dfocus p dstatus

Displays the status of all worker threads in the current focus process. The width here, as in the previous example, is process and the (default) group is the control group; the intersection of this width and the default group creates a focus that is the same as in the previous example.

dfocus pW dstatus

Displays the status of all worker threads in the current focus process. The width is process level and the target is the workers group.

The following example shows how the prompt changes as you change the focus. In particular, notice how the prompt changes when you use the **C** and the **D** group specifiers.

```
d1.<> f C
dC1.<
dC1.<> f D
d1.<
d1.<>
```

Two of these lines end with "<". These lines aren't prompts. Instead, they are the value returned by TotalView when it executes the **dfocus** command.

'All' Does Not Always Mean All

When you use stepping commands, TotalView determines the scope of what runs and what stops by looking at the TOI. This section looks at the differences in be-

havior when you use the **a** (all) arena specifier. Here is what runs when you use this arena:

Table 12: a (all) Specifier Combinations

Specifier	Specifies
aC	All threads.
aS	All threads.
aW	All threads in all workers groups.
aL	All threads.
	Every thread is a member of a control group and a member of a share group and a member of a lockstep group. Consequently, three of these definitions mean "all threads."

This is the same information as was presented in "*Arena Specifier Combinations*" on page 257. Here are some combinations and the meaning of these combinations:

- f aC dgo** Runs everything. If you're using the **dgo** command, everything after the **a** is ignored: **a/aPizza/17.2**, **ac**, **aS**, and **aL** do the same thing. TotalView runs everything.
- f aC duntil** While everything runs, TotalView must wait until something reaches a goal. It really isn't obvious what this *thing* is. Since **C** is a process group, you might guess that all processes run until at least one thread in every participating process arrives at a goal. The reality is that since this goal must reside in the current share group, this command completes as soon as all processes in the TOI's share group have at least one thread at the goal. Processes in other control groups run freely until this happens.
Notice that the TOI determines the goal. If there are other control groups, they do not participate in the goal.
- f aS duntil** This command does the same thing as the **f aC until** command because, as was just mentioned, the goals for **f aC until** and **f aS until** are the same, and the processes that are in this scope are identical.
While more than one share group can exist in a control group, these other share groups do not participate in the goal.
- f aL duntil** While everything will run, it is again not clear what should occur. **L** is a thread group, so you might expect that the **duntil**

command will wait until all threads in all lockstep groups arrive at the goal. Instead, TotalView defines the set of threads that it will run to a goal as just those thread in the TOI's lockstep group. While there are other lockstep groups, these lockstep groups do not participate in the goal. So, while the TOI's lockstep threads are progressing towards their goal, all threads that were previously stopped run freely.

f aW duntil While everything will run, TotalView will wait until all members of the TOI's workers group arrive at the goal.

There are two broad distinctions between process and thread group behavior:

- When the focus is on a process group, TotalView waits until just one thread from each participating process arrives at the goal. The other threads just run, and TotalView doesn't care where they end up.

When focus is on a thread group, every participating thread must arrive at the goal.

- When the focus is on a process group, TotalView steps a thread over the goal breakpoint and continues the process if it isn't the "right thread."

When the focus is on a thread group, TotalView holds a thread even if it isn't the right thread. It also continues the rest of the process. Of course, if your system doesn't support asynchronous thread control, TotalView treats thread specifiers as if they were process specifiers.

With this in mind, **f aL dstep** does not step all threads. Instead, it steps only the threads in the TOI's lockstep group. All other threads run freely until the stepping process for these lockstep threads completes.

Setting Groups

This section presents a series of examples that set and create groups. Many of the examples use CLI commands that have not yet been introduced. You will probably need to refer to the command's definition before you can appreciate what's occurring. These commands are described in the TOTALVIEW REFERENCE GUIDE.

NOTE If you will only be using the GUI, there's nothing you need to know in this section.

Here are some methods for indicating that thread 3 in process 2 is a worker thread.

dset WGROUP(2.3) \$WGROUP(2)

Assigns the group ID of the thread group of worker threads associated with process 2 to the **WGROUP** variable. (Assigning a nonzero value to **WGROUP** indicates that this is a worker group.)

dset WGROUP(2.3) 1

This is a simpler way of doing the same thing as the previous example.

dfocus 2.3 dworker 1

Adds the groups in the indicated focus to a workers group.

dset CGROUP(2) \$CGROUP(1)

dgroups --add --g \$CGROUP(1) 2

dfocus 1 dgroups --add 2

These three commands insert process 2 into the same control group as process 1.

dgroups --add --g \$WGROUP(2) 2.3

Adds process 2, thread 3 to the workers group associated with process 2.

dfocus tW2.3 dgroups --add

This is a simpler way of doing the same thing as the previous example.

Here are some additional examples:

dfocus g1 dgroups --add --new thread

Creates a new thread group that contains all the threads in all the processes in the control group associated with process 1.

set mygroup [dgroups --add --new thread \$GROUP(\$SGROUP(2))]

dgroups --remove --g \$mygroup 2.3

dfocus g\$mygroup/2 dgo

These three commands define a new group containing all the threads in process 2's share group except for thread 2.3 and then continues that set of threads. The first command creates a new group containing all the threads from the share group, the second removes thread 2.3, and the third runs the remaining threads.

Using the 'g' Specifier: An Extended Example

The meaning of the **g** width specifier is sometimes not clear when it is coupled with a group scope specifier. Why have a **g** specifier when you have four other group specifiers? Stated in another way, isn't something like **gL** redundant?

The simplest answer, and the reason you'll most often use **g**, is that it forces the group when the default focus indicates something different from what you want it to be.

Here's an example that shows this. The first step is to set a breakpoint in a multi-threaded OMP program and execute the program until it hits the breakpoint:

```
d1.<> dbreak 35
Loaded OpenMP support library libguidedb_3_8.so :
                                                KAP/Pro Toolset 3.8
1
d1.<> dcont
Thread 1.1 has appeared
Created process 1/37258, named "tx_omp_guide_llnl1"
Thread 1.1 has exited
Thread 1.1 has appeared
Thread 1.2 has appeared
Thread 1.3 has appeared
Thread 1.1 hit breakpoint 1 at line 35 in
".breakpoint_here"
```

The default focus is **d1.<**, which means that the CLI is at its default width, The POI is 1, and the TOI is the lowest numbered nonmanager thread. Because the default width for the **dstatus** command is "process," the CLI displays the status of all processes. Typing **dfocus p dstatus** produces the same output.

```
d1.<> dstatus
1:      37258   Breakpoint [tx_omp_guide_llnl1]
  1.1:  37258.1 Breakpoint PC=0x1000acd0,
                        [./tx_omp_llnl1.f#35]
  1.2:  37258.2 Stopped   PC=0xfffffffffffffffffff
  1.3:  37258.3 Stopped   PC=0xd042c944
```

```
d1.<> dfocus p dstatus
1:          37258   Breakpoint [tx_omp_guide_llnl1]
  1.1: 37258.1 Breakpoint PC=0x1000acd0,
        [./tx_omp_llnl1.f#35]
  1.2: 37258.2 Stopped   PC=0xfffffffffffffffffff
  1.3: 37258.3 Stopped   PC=0xd042c944
```

Here's what the CLI displays when you ask for the status of the lockstep group. (The rest of this example will use the **f** abbreviation for **dfocus** and **st** for **dstatus**.)

```
d1.<> f L st
1:          37258   Breakpoint [tx_omp_guide_llnl1]
  1.1: 37258.1 Breakpoint PC=0x1000acd0,
        [./tx_omp_llnl1.f#35]
```

This command tells the CLI to display the status of the threads in thread 1.1's lockstep group as this thread is the TOI. The **f L focus** command narrows the set so that the display only includes the threads in the process that are at the same PC as the TOI.

NOTE By default, the **dstatus** command displays information at "process" width. This means that you don't need to type "f pL dstatus".

The next command runs thread 1.3 to the same line as thread 1.1. The next command then displays the status of all the threads in the process:

```
d1.<> f t1.3 duntil 35
  35@>      write(*,*)"i= ",i,
            "thread= ",omp_get_thread_num()
d1.<> f p dstatus
1:          37258   Breakpoint [tx_omp_guide_llnl1]
  1.1: 37258.1 Breakpoint PC=0x1000acd0,
        [./tx_omp_llnl1.f#35]
  1.2: 37258.2 Stopped   PC=0xfffffffffffffffffff
  1.3: 37258.3 Breakpoint PC=0x1000acd0,
        [./tx_omp_llnl1.f#35]
```

As expected, the CLI has added a thread to the lockstep group:

```
d1.<> f L dstatus
1:          37258   Breakpoint [tx_omp_guide_llnl1]
  1.1: 37258.1 Breakpoint PC=0x1000acd0,
        [./tx_omp_llnl1.f#35]
  1.3: 37258.3 Breakpoint PC=0x1000acd0,
        [./tx_omp_llnl1.f#35]
```

The next set of commands begins by narrowing the width of the default focus to thread width—notice that the prompt changes—and then displays the contents of the lockstep group.

```
d1.<> f t
t1.<> f L dstatus
1:          37258   Breakpoint [tx_omp_guide_llnl1]
    1.1: 37258.1 Breakpoint PC=0x1000acd0,
                        [./tx_omp_llnl1.f#35]
```

While the lockstep group of the TOI has two threads, the current focus has only one thread, and that thread is, of course, part of the lockstep group. Consequently, the lockstep group *in the current focus* is just the one thread even though this thread's lockstep group has two threads.

If you ask for a wider width (**p** or **g**) with **L**, the CLI displays more threads from the lockstep group of thread 1.1.

```
t1.<> f pL dstatus
1:          37258   Breakpoint [tx_omp_guide_llnl1]
    1.1: 37258.1 Breakpoint PC=0x1000acd0,
                        [./tx_omp_llnl1.f#35]
    1.3: 37258.3 Breakpoint PC=0x1000acd0,
                        [./tx_omp_llnl1.f#35]
t1.<> f gL dstatus
1:          37258   Breakpoint [tx_omp_guide_llnl1]
    1.1: 37258.1 Breakpoint PC=0x1000acd0,
                        [./tx_omp_llnl1.f#35]
    1.3: 37258.3 Breakpoint PC=0x1000acd0,
                        [./tx_omp_llnl1.f#35]
t1.<>
```

NOTE If the TOI is 1.1, “L” refers to group number 1.1, which is the lockstep group of thread 1.1.

Because this example only contains one process, the **pL** and **gL** specifiers produce the same result when used with **dstatus**. If, however, there were additional processes in the group, you would only see them when you use the **gL** specifier.

Focus Merging

When you specify more than one focus for a command, the CLI will merge them together. In the following example, the focus indicated by the prompt—this focus is called the *outer* focus—controls the display. Notice what happens when **dfocus** commands are strung together:

```
t1.<> f d
d1.<
d1.<> f tL dstatus
1:          37258      Breakpoint [tx_omp_guide_llnl1]
  1.1: 37258.1 Breakpoint PC=0x1000acd0,
                        [./tx_omp_llnl1.f#35]

d1.<> f tL f p dstatus
1:          37258      Breakpoint [tx_omp_guide_llnl1]
  1.1: 37258.1 Breakpoint PC=0x1000acd0,
                        [./tx_omp_llnl1.f#35]
  1.3: 37258.3 Breakpoint PC=0x1000acd0,
                        [./tx_omp_llnl1.f#35]

d1.<> f tL f p f D dstatus
1:          37258      Breakpoint [tx_omp_guide_llnl1]
  1.1: 37258.1 Breakpoint PC=0x1000acd0,
                        [./tx_omp_llnl1.f#35]
  1.2: 37258.2 Stopped PC=0xfffffffffffffffffff
  1.3: 37258.3 Breakpoint PC=0x1000acd0,
                        [./tx_omp_llnl1.f#35]

d1.<> f tL f p f D f L dstatus
1:          37258      Breakpoint [tx_omp_guide_llnl1]
  1.1: 37258.1 Breakpoint PC=0x1000acd0,
                        [./tx_omp_llnl1.f#35]
  1.3: 37258.3 Breakpoint PC=0x1000acd0,
                        [./tx_omp_llnl1.f#35]

d1.<>
```

Stringing multiple focuses together may not produce the most readable result. In this case, it shows how one **dfocus** command can modify what another sees and will act on. The ultimate result is an arena that a command will act on. In these examples, the **dfocus** command is telling the **dstatus** command what it should be displaying.

Incomplete Arena Specifiers

In general, you do not need to completely specify an arena. TotalView provides values for any missing elements. TotalView either uses built-in default values or obtains them from the current focus. Here is how TotalView fills in missing pieces:

- If you don't use a width, TotalView uses the width from the current focus.
- If you don't use a PID, TotalView uses the PID from the current focus.
- If you set the focus to a list, there is no longer a default arena. This means that you must explicitly name a width and a PID. You can, however, omit the TID. (If you omit the TID, TotalView defaults to <.)

You can type a PID without typing a TID. If you omit the TID, TotalView uses its default of "<", where "<" indicates the lowest numbered worker thread in the process. If, however, the arena explicitly names a thread group, < means the lowest numbered member of the thread group.

TotalView doesn't use the TID from the current focus, since the TID is a process-relative value.

- A dot before or after the number lets TotalView know if you're specifying a process or a thread. For example, "1." is clearly a PID, while ".7" is clearly a TID. If you type a number without typing a dot, the CLI most often interprets the number as being a PID.
- If the width is **t**, you can omit the dot. For instance, **t7** refers to thread 7.
- If you enter a width and don't specify a PID or TID, TotalView uses the PID and TID from the current focus. If you use a letter as a group specifier, TotalView obtains the rest of the arena specifier from the default focus.
- You can use a group ID or tag followed by a "/". TotalView obtains the rest of the arena from the default focus.

Of course, focus merging can also influence how TotalView fills in missing specifiers. For more information, see "Focus Merging" on page 266.

Lists with Inconsistent Widths

TotalView lets you create lists containing more than one width specifier. While this can be very useful, it can be confusing. Consider the following:

```
{p2 t7 g3.4}
```

This list is quite explicit: all of process 2, thread 7, and all processes in the same group as process 3, thread 4. However, how should TotalView use this set of processes, groups, and threads?

In most cases, TotalView does what you would expect it to do: a command iterates over the list and acts on each arena. If TotalView cannot interpret an inconsistent focus, it prints an error message.

Some commands work differently. Some use each arena's width to determine the number of threads on which it will act. This is exactly what the **dgo** command does. In contrast, the **dwhere** command creates a call graph for process-level arenas, and the **dstep** command runs all threads in the arena while stepping the TOI. It may wait for threads in multiple processes for group-level arenas. The command description in the TOTALVIEW REFERENCE GUIDE will point out anything that you need to watch out for.

Stepping (Part II): Some Examples

Here are some examples of things that you'll probably do using the CLI's stepping commands:

■ Step a single thread

While the thread runs, no other thread runs (except kernel manager threads).

Example: `dfocus t dstep`

■ Step a single thread while the process runs

A single thread runs into or through a critical region.

Example: `dfocus p dstep`

■ Step one thread in each process in the group

While one thread in each process in the share group runs to a goal, the rest of the threads run freely.

Example: `dfocus g dstep`

■ Step all worker threads in the process while nonworker threads run

Runs worker threads through a parallel region in lockstep.

Example: `dfocus pW dstep`

- **Step all workers in the share group**

All processes in the share group participate. The nonworker threads run.

Example: `dfocus gW dstep`

- **Step all threads that are at the same PC as the TOI**

TotalView selects threads from one process or from the entire share group. This differs from the previous two bullets in that TotalView uses the set of threads that are in lockstep with the TOI rather than using the workers group.

Example: `dfocus L dstep`

In the following examples, the default focus is set to `d1.<`.

- | | |
|----------------------------|--|
| dstep | Steps the TOI while running all other threads in the process. |
| dfocus W dnext | Runs the TOI and all other worker threads in the process to the next statement. Other threads in the process run freely. |
| dfocus W duntil 37 | Runs all worker threads in the process to line 37. |
| dfocus L dnext | Runs the TOI and all other stopped threads at the same PC to the next statement. Other threads in the process run freely. Threads that encounter a temporary breakpoint in the course of running to the next statement usually join the lockstep group. |
| dfocus gW duntil 37 | Runs all worker threads in the share group to line 37. Other threads in the control group run freely. |
| UNW 37 | Performs the same action as the previous command: runs all worker threads in the share group to line 37. This example uses the predefined UNW alias instead of the individual commands. That is, UNW is an alias for dfocus gW duntil . |
| SL | Finds all threads in the share group that are at the same PC as the TOI and steps them all one statement. This command is the built-in alias for dfocus gL dstep . |
| sl | Finds all threads in the current process that are at the same PC as the TOI, and steps them all one statement. This command is the built-in alias for dfocus L dstep . |

Using P/T Set Operators

At times, you do not want all of one kind of group or process to be in the focus set. TotalView lets you use the following three operators to manage your P/T sets:

- | Creates a union; that is, all members of the sets.
- Creates a difference; that is, all members of the first set that are not also members of a second set.
- & Creates an intersection; that is, all members of the first set that are also members of the second set.

For example, here is how you would create a union of two P/T sets:

```
p3 | L2
```

A set operator only operates on two sets. You can, however, apply these operations repeatedly. For example:

```
p2 | p3 & L2
```

This statement creates a union between **p2** and **p3**, and then creates an intersection between the union and **L2**. As this example suggests, TotalView associates sets from left to right. You can change this order by using parentheses. For example:

```
p2 | (p3 & pL2)
```

Typically, these three operators are used with the following P/T set functions:

- breakpoint()** Returns a list of all threads that are stopped at a breakpoint.
- error()** Returns a list of all threads stopped due to an error.
- existent()** Returns a list of all threads.
- held()** Returns a list of all threads that are held.
- nonexistent()** Returns a list of all processes that have exited or which, while loaded, have not yet been created.
- running()** Returns a list of all running threads.
- stopped()** Returns a list of all stopped threads.
- unheld()** Returns a list of all threads that are not held.
- watchpoint()** Returns a list of all threads that are stopped at a watchpoint.

The argument that all of these operators use is a P/T set. You specify this set in the same way that a P/T set is specified for the **dfocus** command. For example, **watchpoint(L)** returns all threads in the current lockstep group.

The dot (.) operator, which indicates the current set, can be helpful when you are editing an existing set.

The following examples should clarify how you use these operators and functions. The P/T set that is the argument to these operators is **a** (all).

f {breakpoint(a) | watchpoint(a)} dstatus

Shows information about all threads that stopped at breakpoints and watchpoints. The **a** argument is the standard P/T set indicator for "all".

f {stopped(a) - breakpoint(a)} dstatus

Shows information about all stopped threads that are not stopped at breakpoints.

f {. | breakpoint(a)} dstatus

Shows information about all threads in the current set as well as all threads stopped at a breakpoint.

f {g.3 - p6} duntil 577

Runs thread 3 along with all other processes in the group to line 577. However, do not run anything in process 6.

f {(\$PTSET) & p123}

Uses just process 123 within the current P/T set.

Using the P/T Set Browser

There's no question that specifying P/T sets can be confusing. As has been mentioned, there are few programs that need all the power that TotalView's P/T set syntax provides. In all cases, however, the ability to previsualize what the contents of a P/T set will be before you execute the command is essential. This is what the P/T Set Browser is designed to do. The browser, which is accessed from the Root Window's **Tool** menu, shows the current state of processes and threads as well as show what is or will be selected when you specify a P/T set. Figure 136 on page 272 shows a P/T browser displaying information about a multiprocess, multithreaded program.

The top part of this window contains the standard P/T set controls. (See "Using P/T Set Controls" on page 245 for more information.) The large area on the left is a "tree"

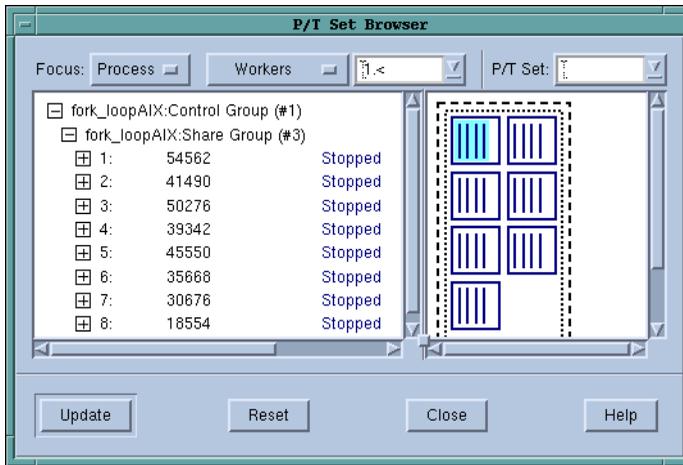


Figure 136: A P/T Set Browser Window

control where clicking on the “+” shows more information, and clicking on a “–” (not shown in this figure) condenses the information. Here you will find a list of all your program’s processes and threads. The information is organized in a hierarchy, with the outermost level being your program’s control groups. In a control group, information is further organized by share group, where you are shown the processes contained in a share group. Finally, if the innermost “+” symbols were clicked, the browser would show information on the threads within a process.

The control and share group numbers displayed in this window are the same as those that are displayed in the **Groups** Page in the Root Window.

The right-hand side contains a graphical depiction of your program’s threads. In the preceding figure, notice that TotalView has highlighted some of the threads. These are the threads in the current focus, which in this case is “1.<”. As you make changes to the P/T set, the threads highlighted in the right-hand side change, showing you what the scope of a P/T set definition is. The next figure contains a variety of P/T set examples.

- ❶ This P/T set displayed differs from the one in Figure 136 in that the **Focus** pulldown menu is now set to **All**. TotalView responds by highlighting all threads on the right-hand side.
- ❷ The **Focus** pulldown menu was changed back to **Process** but the number of processes was limited to 1, 2, and 3. Before these changes were

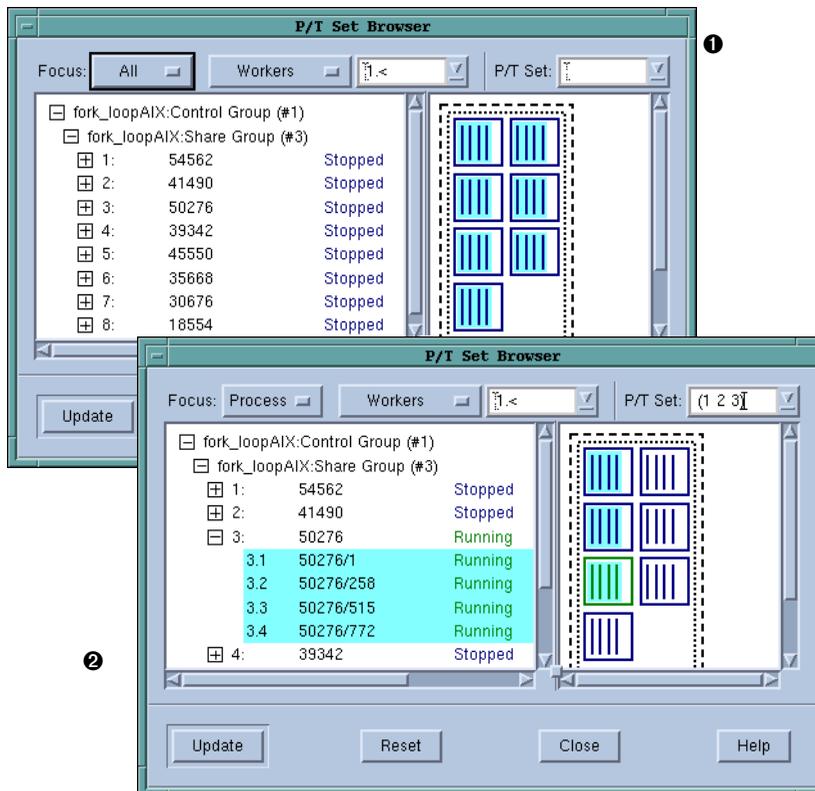


Figure 137: P/T Set Browser Windows (Part 1)

made, process 3 was told to go. As you can see, the browser shows those processes as running.

- ③ Thread 3.1 was halted.
 - ④ Thread 2.4 was selected with the mouse. It doesn't matter if it was selected in the left or right-hand sides, as selecting it causes it to be highlighted in both. After selecting a thread, you can extend the selection by clicking your mouse's left button on another thread while holding down the Shift key. You can select noncontiguous threads by holding down the Control key while clicking your mouse's left button.
- If you are seeing this document online, you'll notice that your selection is in gray while the selection indicating the P/T set is in blue.

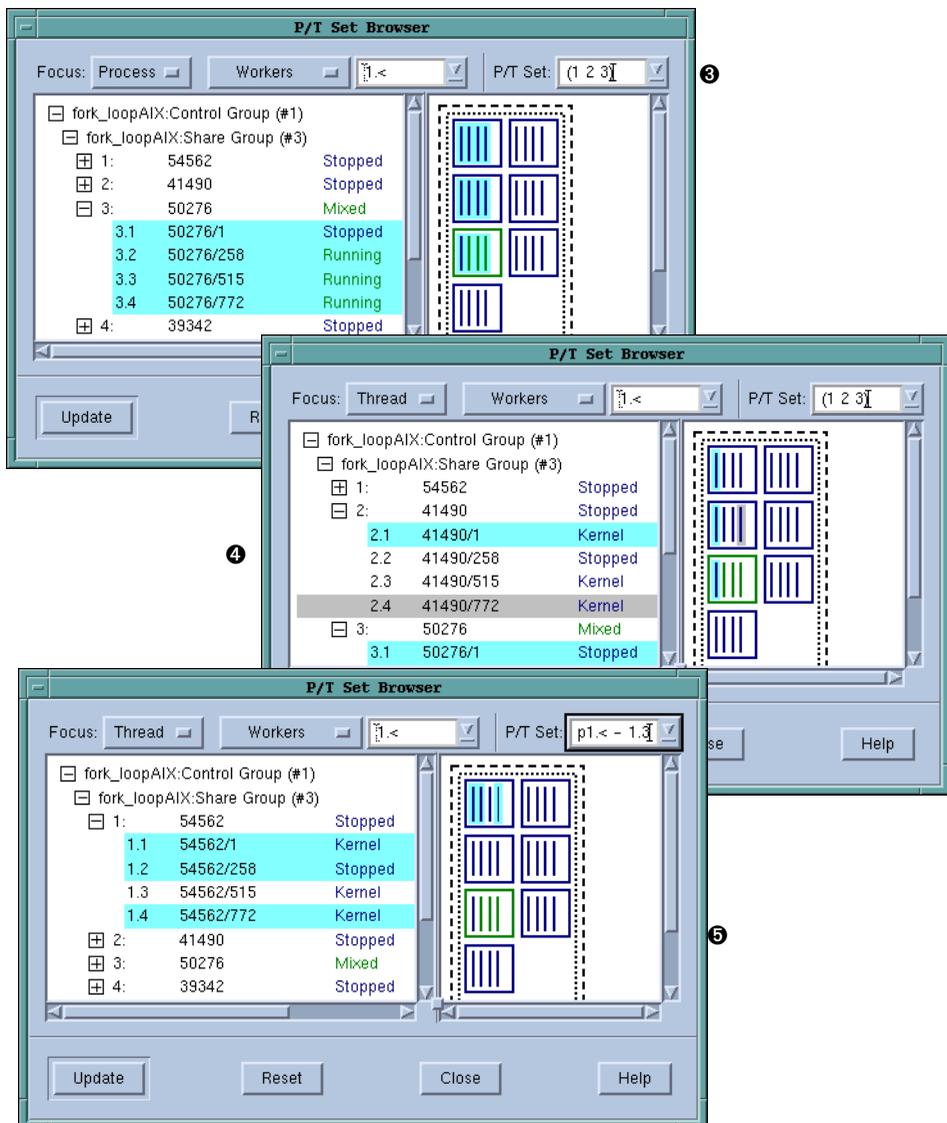


Figure 138: P/T Set Browser Windows (Part 2)

5

The P/T set information was modified to show a difference expression; in this case, thread 1.3 was eliminated from the set of threads named by "p1.<".

The elements on the right side are drawn within two boxes. These boxes represent the control and share groups. Clicking on them tells the browser to select that group.

Using the Group Editor

The visual group editor, which is displayed after you select the **Group > Edit Group** command, can simplify the way in which you create named groups.

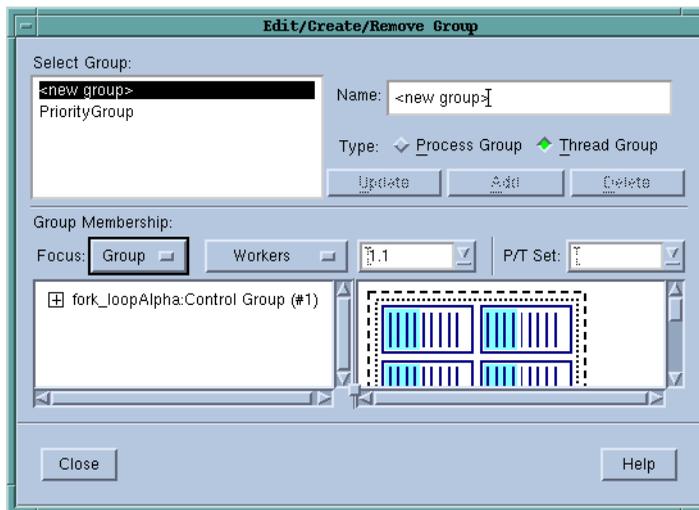


Figure 139: Group > Edit Group

This dialog box can be divided into two halves. The top half allows you to add, update, and delete named sets. The bottom half contains controls that allow you to specify which processes and threads will become part of the group. These controls are discussed in the previous section.

The controls in the upper portion work generally as you'd expect them to. The only thing to be careful about is that you must define the group, and be sure to give the group a name, before you click on the **Add** button. Details on using the controls in this dialog box are contained in the online Help.

Chapter 12

Examining and Changing Data

This chapter explains how to examine and change data as you debug your program. The topics discussed in this chapter are:

- *"Changing How Data Is Displayed"* on page 277
- *"Displaying Variables"* on page 281
- *"Diving in Variable Windows"* on page 291
- *"Scoping and Symbol Names"* on page 294
- *"Changing the Values of Variables"* on page 296
- *"Changing the Data Type of Variables"* on page 297
- *"Working with Opaque Data"* on page 306
- *"Changing Types to Display Machine Instructions"* on page 307
- *"Changing Types to Display Machine Instructions"* on page 307
- *"Displaying C++ Types"* on page 307
- *"Displaying Fortran Types"* on page 310
- *"Displaying Thread Objects"* on page 317

Changing How Data Is Displayed

One of the problems you'll face is that TotalView, like all debuggers, displays data in the way that your compiler stored it. The following two sections let you change the the way TotalView displays this information. These sections are:

- *"Displaying STL Variables"* on page 278
- *"Changing Size and Precision"* on page 280

Displaying STL Variables

The C++ STL (Standard Template Library) greatly simplifies the way in which you can access data. By offering standard and prepackaged ways to organize data, you do not have to be concerned with the mechanics of the access method. The only real downside to using the STL is while debugging. This is because the information you are shown is the compiler's view of your data rather than the STL's logical view. For example, here is how your compiler sees a map compiled using GNU C++:

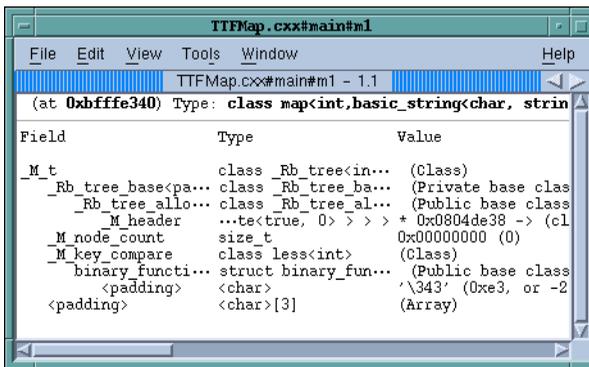


FIGURE 140: An Untransformed Map

TotalView comes with a set of transforms that changes how it displays some STL data. For example, Figure 141 shows the transformed map.

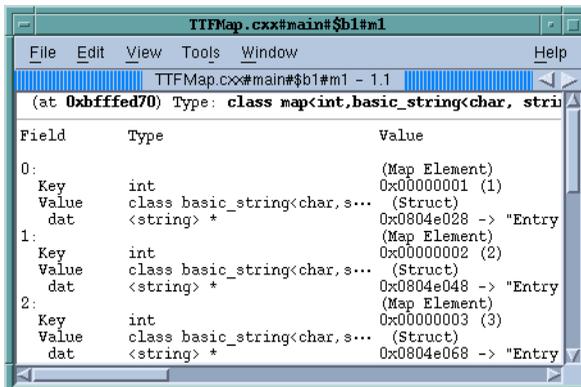


FIGURE 141: A Transformed Map

TotalView can transform STL vectors, lists, and maps using native and GCC compilers on IBM AIX, IRIX/MIPS, and HP Tru64 Alpha. It also supports GCC and Intel's Version 7 C++ 32-bit compiler running on Red Hat x86 platform.

Figure 142 on page 279 shows an untransformed and transformed list and vector.

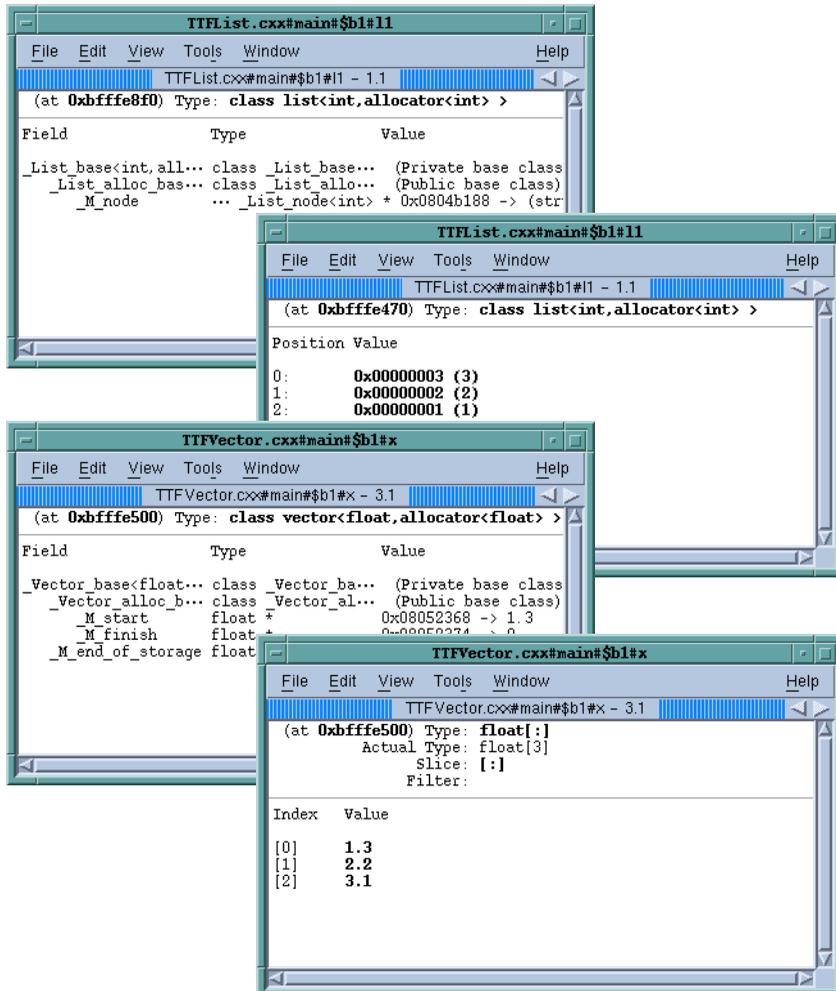


FIGURE 142: List and Vector Transformations

NOTE You can create your own transformations. The process for doing this is described in the "Creating Type Transformations Guide".

By default, TotalView transforms these data structures. If you do not want them transformed, uncheck the **View simplified STL containers (and user-defined transformations)** checkbox within the **File > Preference's Options Page**.

```
dset TV::tff { true | false }
```

Changing Size and Precision

In most cases, TotalView does a reasonable job of displaying a variable's value. If TotalView's defaults don't meet your needs, you can indicate the precision at which to display simple data types by using the **Formatting Page** of the **File > Preferences** Dialog Box. (See Figure 143 on page 280.)

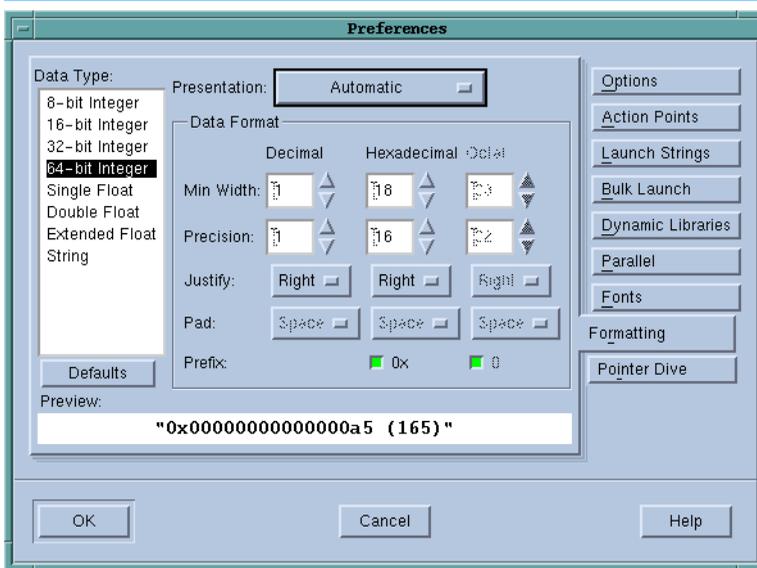


FIGURE 143: **File > Preferences: Formatting Page**

After selecting one of the data types listed on the left, you can set how many character positions a value will use when TotalView displays it (**Min Width**) and how many numbers it should display to the right of the decimal place (the **Precision**). You can also tell TotalView how it should align the value in the **Min Width** area and if it should pad numbers with zeros or spaces.

While the way in which these controls relate and interrelate may appear to be complex, the **Preview** area shows you exactly the result of a change. After you play with the controls for a minute or so, what each control does will be clear. You will probably need to set the **Min Width** value to a larger number than you need it to be to see the results of a change. For example, if the **Min Width** doesn't allow a number to be justified, it could appear that nothing is happening.

CLI EQUIVALENT: You can set these properties from within the CLI. To obtain a list of variables that you can set, type:
`dset TV::data_format*`

Displaying Variables

TotalView displays variables that are local to the current stack frame in the Process Window's Stack Frame Pane. For non-simple variables—for example, pointers, arrays, and structs—this pane doesn't show the data; instead, you need to dive on the variable to bring up a Variable Window that contains the variable's information. For example, diving on an array variable tells TotalView to display the entire contents of the array.

NOTE Dive on a variable by clicking your middle mouse button on it.

If you dive on simple variables or registers, TotalView still brings up a Variable Window. In this case, you'll see some additional information about the variable.

Topics in this section are:

- "Displaying Program Variables" on page 282
- "Browsing for Variables" on page 283
- "Displaying Local Variables and Registers" on page 284
- "Displaying Long Variable Names" on page 286
- "Automatic Dereferencing" on page 287
- "Displaying Areas of Memory" on page 289
- "Displaying Machine Instructions" on page 290
- "Closing Variable Windows" on page 290

Displaying Program Variables

You can display local and global variables by:

- Diving into the variable in either the Source or Stack Panes.
- Selecting the **View > Lookup Variable** command. When prompted, enter the name of the variable.

CLI EQUIVALENT: `dprint variable`

A Variable Window appears for the global variable. (See Figure 144 on page 282.)

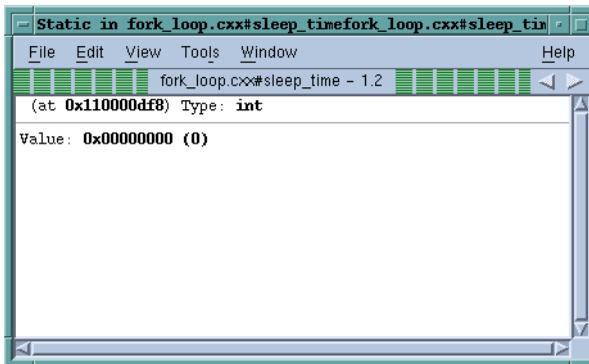


FIGURE 144: Variable Window for a Global Variable

Displaying Variables in the Current Block

In many cases, you're not really interested in just seeing one variable. Instead, you want to see all of the variables in the current block. If you dive on a block label within the Stack Frame Pane, TotalView opens a Variable Window containing just those variables. See Figure 145.

You can dive on any variable in this window to see more information about it.

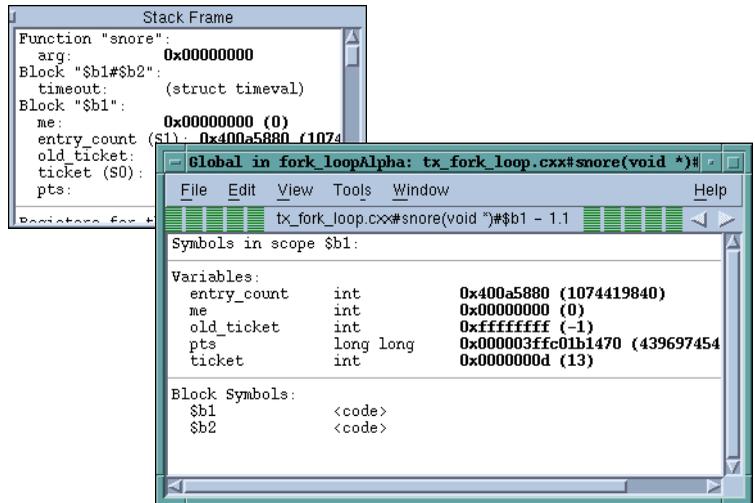


FIGURE 145: Displaying Scoped Variables

Browsing for Variables

The Process Window's **Tools > Program Browser** command displays a window containing all your executable's components. By clicking on a library or program name, you can access all of the variables contained within it. (See Figure 146.)

The window in the upper left corner shows programs and libraries that are loaded. If you have loaded more than one program with the **File > New Program** command, only the currently selected process list will appear. The center window contains a list of files that make up the program as well as other related information. Diving again on a line displays a Variable Window that contains variables and other information related to the file. Figure 147 shows three more diving operations.

The screen in the upper-left corner shows a Variable Window created by diving on one of the files in Figure 147. The center screen dives on a block in that subroutine. Finally, the screen in the lower-right corner shows a variable. (These screens were created using the **View > Dive Anew** command.) If you just dived on a line in a Variable Window, the new contents replace the old contents, and you can use the dive/undive  icons to move back and forth.

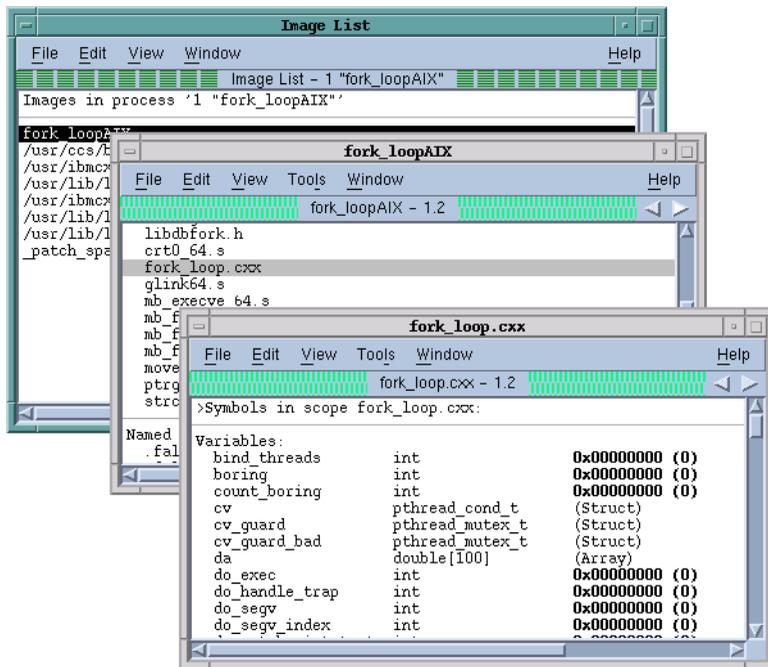


FIGURE 146: Program Browser and Variables Window (Part 1)

Displaying Local Variables and Registers

In the Stack Frame Pane, diving into a formal parameter, local variable, or register tells TotalView to display a Variable Window. You can also dive into parameters and local variables in the Source Pane. The displayed Variable Window shows the name, address, data type, and value for the object. (See Figure 148 on page 285.)

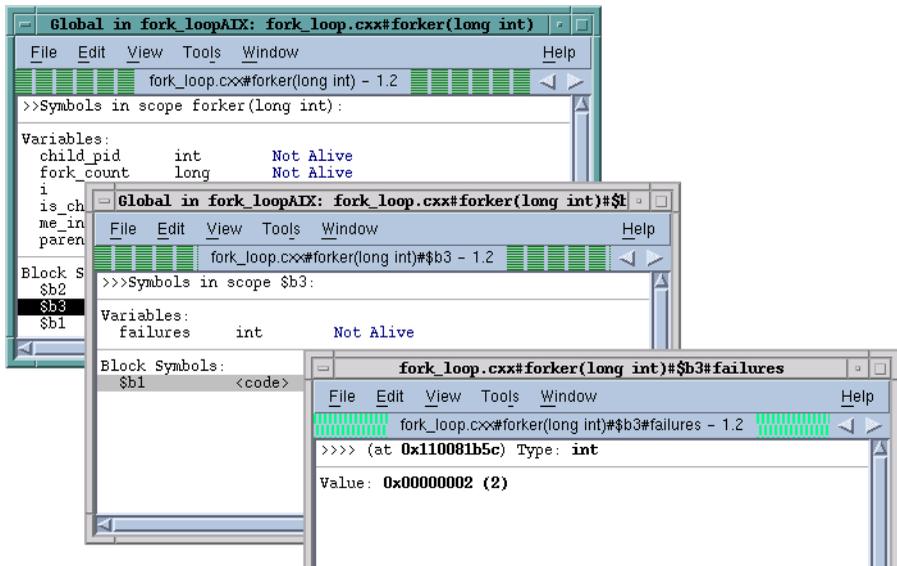


FIGURE 147: Program Browser and Variables Window (Part 2)

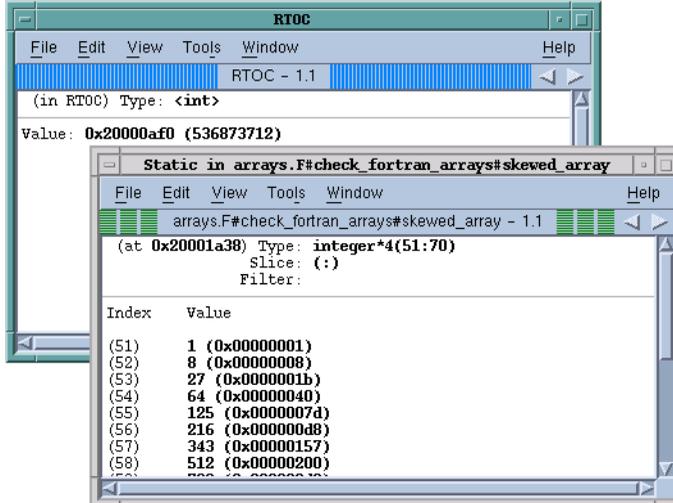


FIGURE 148: Diving into Local Variables and Registers

The top-left window shows the result of diving on a register, while the bottom-right window shows the results of diving on an array variable.

CLI EQUIVALENT: `dprint variable`

This command lets you view variables and expressions without having to select or find them.



You can also display a local variable by using the **View > Lookup Variable** command. When prompted, enter the name of the variable in the dialog box that appears.

If Variable Windows remain open while a process or thread runs, TotalView updates the information in these windows when the process or thread stops. If TotalView can't find a stack frame for a displayed local variable, it displays **Stale** in the pane's header to warn you that you can't trust the data, since the variable no longer exists.

When you debug recursive code, TotalView doesn't automatically refocus a Variable Window onto different instances of a recursive function. If you have a breakpoint in a recursive function, you'll need to explicitly open a new Variable Window to see the value of a local variable in that stack frame.

CLI EQUIVALENT: `dwhere, dup, and dprint`

You'll locate the stack frame using `dwhere`, move to it using `dup`, and then display the value using `dprint`.

Displaying Long Variable Names



If TotalView doesn't have enough space to display all the characters in a variable name, it inserts ellipses (...) to indicate that it has truncated the name. Typically, this occurs when it is displaying demangled C++ names or STL variables.

Figure 149 shows three windows. The top-left Variable Window contains a series of STL names. The other two windows show what TotalView displays when you click on the ellipses. Notice that one of the windows has an **Apply** button. This indicates that the field is editable.

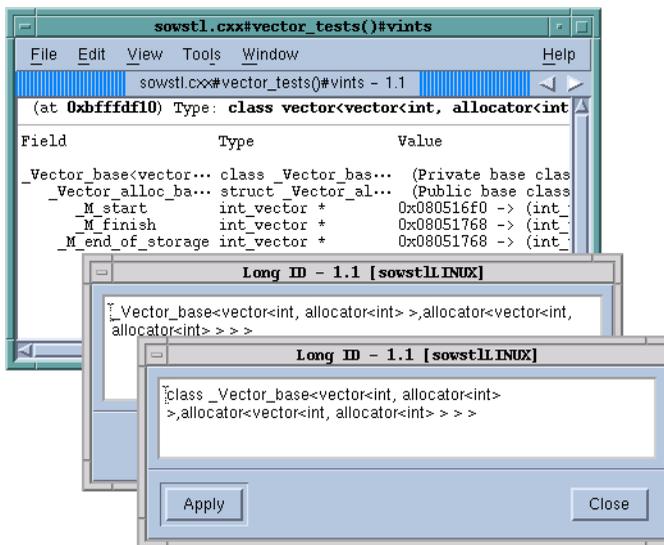


FIGURE 149: Displaying Long STL Names

Automatic Dereferencing

In most cases, you aren't interested in the value contained in a pointer variable. Instead, you want to see what the pointer is pointing to. Using the controls contained in the **File > Preferences's Pointer Dive Page**, you can tell TotalView if it should automatically dereference pointers. (See Figure 150.)

This preference is especially useful when you want to visualize data that is linked together with pointers, as it can present the data as a unified array. Because the data appears to be a unified array, you can use TotalView's array manipulation commands and the Visualizer to view this data.

Each pulldown list has three settings: **No**, **Yes**, and **Yes (don't push)**. The meaning for **No** is obvious: automatic dereferencing will not occur. Both of the remaining values tell TotalView that it should automatically dereference pointers. The difference between the two is based on whether you can use the **Back** command to see the undereferenced pointer's value. If you set this to **Yes**, you can see the value. Setting

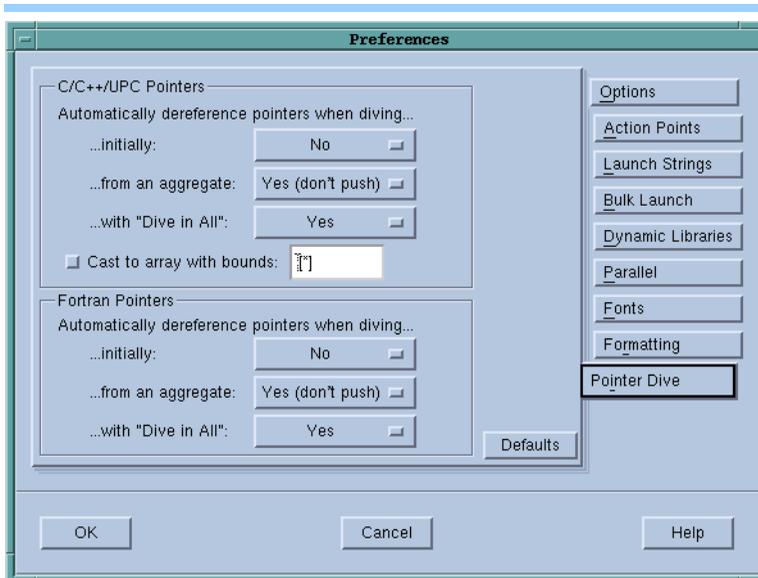


FIGURE 150: **File > Preferences: Pointer Dive Page**

it to **Yes (don't push)** means you can't use the **Back** command to see the pointer's value.

CLI EQUIVALENT: `TV::auto_array_cast_bounds`
`TV::auto_deref_in_all_c`
`TV::auto_deref_in_all_fortran`
`TV::auto_deref_initial_c`
`TV::auto_deref_initial_fortran`
`TV::auto_deref_nested_c`
`TV::auto_deref_nested_fortran`

The three situations in which automatic dereferencing can occur are:

- When TotalView *initially* displays a value.
- When you dive on a value within an *aggregate* or structure.
- When you use the **Dive in All** command.

Displaying Areas of Memory

You can display areas of memory in hexadecimal and decimal values. Do this by selecting the **View > Lookup Variable** command and then entering one of the following in the dialog box that appears:

- **A hexadecimal address**

When you enter a single address, TotalView displays the word of data stored at that address.

CLI EQUIVALENT: `dprint address`

- **A pair of hexadecimal addresses**

When you enter a pair of addresses, TotalView displays the data (in word increments) from the first to the last address. To enter a pair of addresses, enter the first address, a comma, and the last address.

CLI EQUIVALENT: `dprint address,address`

NOTE All hexadecimal constants must have a “0x” prefix. You can use an expression to enter these addresses.

The Variable Window for an area of memory displays the address and contents of each word. (See Figure 151.)

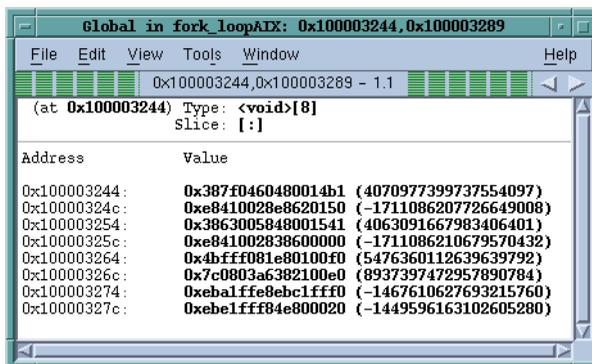


FIGURE 151: Variable Window for Area of Memory

TotalView displays the memory area's starting location at the top of the window's data area. In the window, TotalView displays information in hexadecimal and decimal.

Displaying Machine Instructions

You can display the machine instructions for entire routines as follows:

- Dive into the address of an assembler instruction in the Source Pane (such as **main+0x10** or **0x60**). A Variable Window displays the instructions for the entire function and highlights the instruction you dived into.
- Dive into the PC in the Stack Frame Pane. A Variable Window lists the instructions for the entire function containing the PC, and highlights the instruction the PC points to. (See Figure 152 on page 290.)

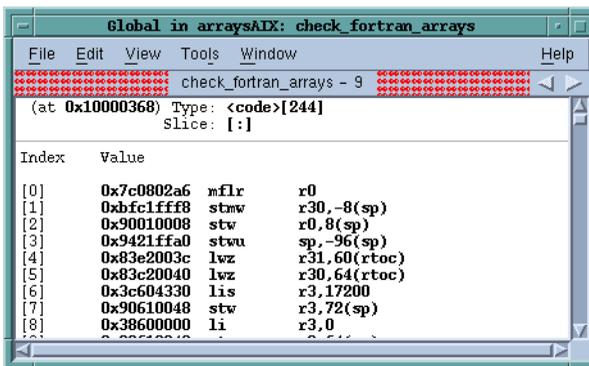


FIGURE 152: Variable Window with Machine Instructions

- Cast a variable to type `<code>` or array of `<code>`, as described in "Changing Types to Display Machine Instructions" on page 307. (See Figure 153 on page 291.)

Closing Variable Windows

When you're finished analyzing the information in a Variable Window, use the **File > Close** command to close the window. You can also use the **File > Close Similar** command to close all Variable Windows.

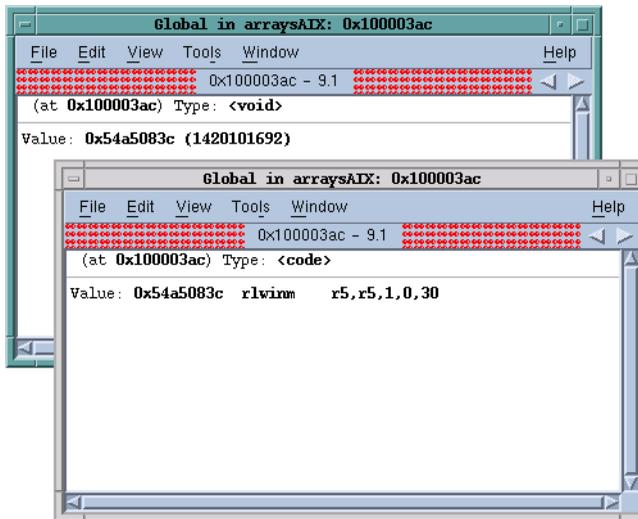


FIGURE 153: Casting Code

Diving in Variable Windows

If the variable being displayed in a Variable Window is a pointer, structure, or array, you can dive into a value shown in the Variable Window. This new dive, which is called a *nested dive*, tells TotalView to replace the information in the Variable Window with information about the selected variable. If this information contains non-simple data types, you can dive on these data types. While a typical data structure doesn't have too many levels, repeatedly diving on data lets you follow pointer chains. That is, diving allows you to see the elements of a linked list.

TotalView remembers your dives. This means you can use the "undive" and "redive" buttons  as a convenient way to see other dive results.

Figure 154 shows a Variable Window after diving into a pointer variable named `sp` with a type of `simple*`. The first Variable Window, which is called the *base window*, displays the value of `sp`. (This is the window in the upper left corner.)

The nested dive window—displayed in the bottom right corner of the figure—shows the structure referenced by the `simple*` pointer.

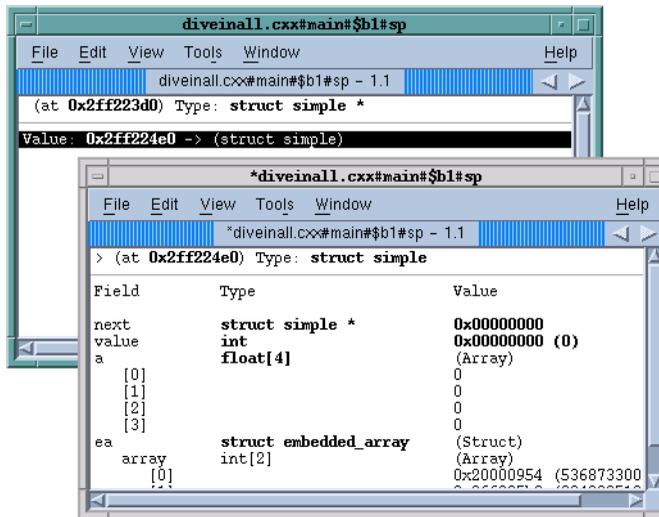


FIGURE 154: Nested Dives

You can manipulate Variable Windows and nested dive windows in the following ways:

- To “undive” from a nested dive, select the left-facing arrow in the upper right-hand corner of the Variable Window. After clicking on the arrow, the previous contents of the Variable Window appears.
- To “redive” after you “undive,” select the right-facing arrow in the upper right-hand corner of the Variable Window. After clicking on the arrow, TotalView performs a previously executed dive operation.
- If you dive into a variable that already has a Variable Window open, the Variable Window pops to the top of the window display.
- If you have performed several nested dives and want to create a new copy of the base window, select the **Window > Duplicate Base** command.
- If you select the **Window > Duplicate** command, a new Variable Window appears that is a duplicate of the current Variable Window. It differs internally as it has an empty dive stack.

Displaying Array of Structure Elements

The **View > Dive In All** command (which is also available when you right-click on a field) allows you to display an element in an array of structures as if it were a simple array. For example, suppose you have the following Fortran definition:

```
type embedded_array
  real r
  integer, pointer :: ia(:)
end type embedded_array
```

```
type(embedded_array) ea (3)
```

After selecting an `r` element, select the **View > Dive In All** command, TotalView displays all three `r` elements of the `ea` array as if it were a single array. (See Figure 155 on page 293.)

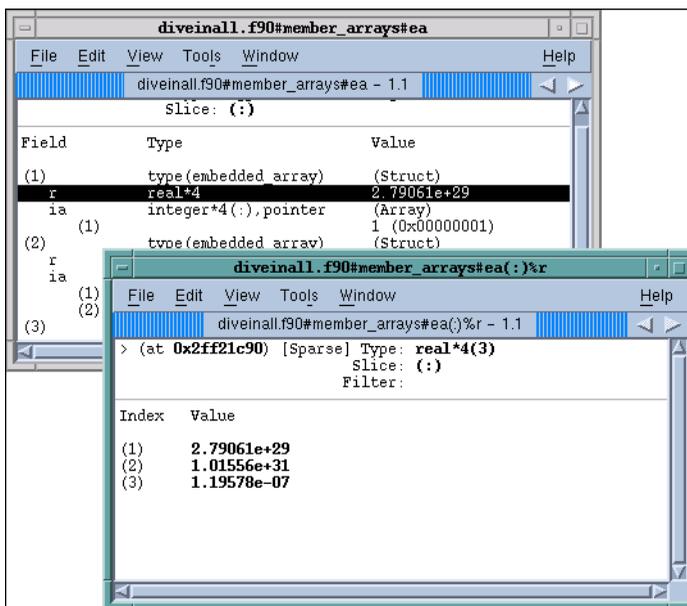


FIGURE 155: Displaying a Fortran Structure

The **View > Dive in All** command can also display the elements of C array of structures as arrays. Figure 156 on page 294 shows TotalView displaying a unified array of structures and a multidimensional array in a structure.

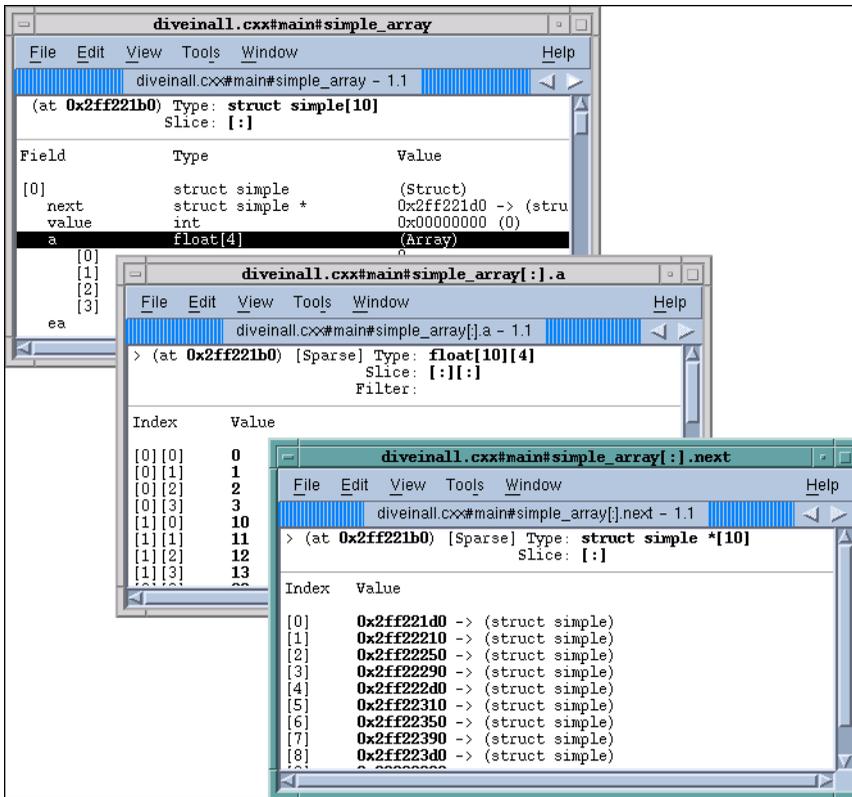


FIGURE 156: Displaying C Structures and Arrays

NOTE As TotalView's array manipulation commands (which are described in Chapter 8) work on what's displayed and not what is stored in memory, you can operate on an array created by this command in the same manner as any other array. For example, you can visualize the array, obtain statistics about it, filter elements in it, and so on.

Scoping and Symbol Names

Many CLI and some GUI commands have arguments whose elements are variables and other things found in your program. TotalView assigns a unique name to all of your program's element based on the scope in which the element exists. A *scope* defines what part of a program knows about a symbol. For example, the scope of a variable that is defined at the beginning of a subroutine is all statements in the subroutine. The variable's scope does not extend outside of this subroutine. A pro-

gram consists of *scopes*. Of course, a block contained in the subroutine could have its own definition of the same variable. This would *hide* the definition in the enclosing scope.

All scopes are defined by your program's structure. Except for the most trivial program, scopes are embedded in other scopes. The exception is, of course, the top-most scope. Every element in a program is associated with a scope.

Whenever you tell the CLI or the GUI to execute a command, TotalView consults the program's symbol table to discover what object you are referring to—this process is known as *symbol lookup*. A symbol lookup is performed with respect to a particular context, and each context uniquely identifies the scope to which a symbol name refers.

Qualifying Symbol Names

The way you describe a scope is similar to the way you specify a file. The scopes in a program form a tree, with the outermost scope, which is your program, as the root. At the next level are executable files and dynamic libraries; further down are compilation units (source files), procedures, modules, and other scoping units (for example, blocks) supported by the programming language. Qualifying a symbol is equivalent to describing the path to a file in UNIX file systems.

A symbol is fully scoped when you name all levels of its tree. The following example shows how this is done. It also indicates parts that are optional.

```
[#executable-or-lib#][file#][procedure-or-line#]symbol
```

The pound sign (#) separates elements of the fully qualified name.

NOTE Because of the number of different kinds of things that can appear in your program, a formal specification of what can appear and the order in which things can appear complicated, and, unreadable. After you see the name, in the Stack Frame Pane, you'll know a variable's scoped name.

TotalView interprets the components as follows:

- Just as file names need not be qualified with a full path, you do not need to use all levels in a symbol's scoping tree.
- If a qualified symbol begins with #, the name that follows indicates the name of the executable or shared library (just as an absolute file path begins with a direc-

tory immediately within the root directory). If you omit the executable or library component, the qualified symbol doesn't begin with #.

- The source file's name may appear after the (possibly omitted) executable or shared library.
- Because programming languages typically do not let you name blocks, that portion of the qualifier is specified using the letter **b** followed by a number indicating which block. For example, the first unnamed block is named **#b1**, the second as **#b2**, and so on.

You can omit any part of the scope specification that TotalView doesn't need to uniquely identify the symbol. Thus, **foo#x** identifies the symbol **x** in the procedure **foo**. In contrast, **#foo#x** identifies either procedure **x** in executable **foo** or variable **x** in a scope from that executable.

Similarly, **#foo#bar#x** could identify variable **x** in procedure **bar** in executable **foo**. If **bar** were not unique within that executable, the name would be ambiguous unless you further qualified it by providing a file name. Ambiguities can also occur if a file-level variable (common in C programs) has the same name as variables declared within functions in that file. For instance, **bar.c#x** refers to a file-level variable, but the name can be ambiguous when there are different definitions of **x** embedded in functions occurring in the same file. In this case, you would need to say **bar.c#b1#x** to identify the scope that corresponds to the "outer level" of the file (that is, the scope containing line 1 of the file).

Changing the Values of Variables

You can change the value of any variable or the contents of any memory location displayed in a Variable Window by selecting the value and typing the new value. In addition to typing a value, you can also type an expression. For example, you can enter $1024 * 1024$ as shown in Figure 157 on page 297. You can include logical operators in all TotalView expressions.

```
CLI EQUIVALENT:  set my_var [expr 1024*1024]
                  dassign int8_array(3) $my_var
```

If a value is displayed in bold in the Stack Frame Pane, you can edit the value.

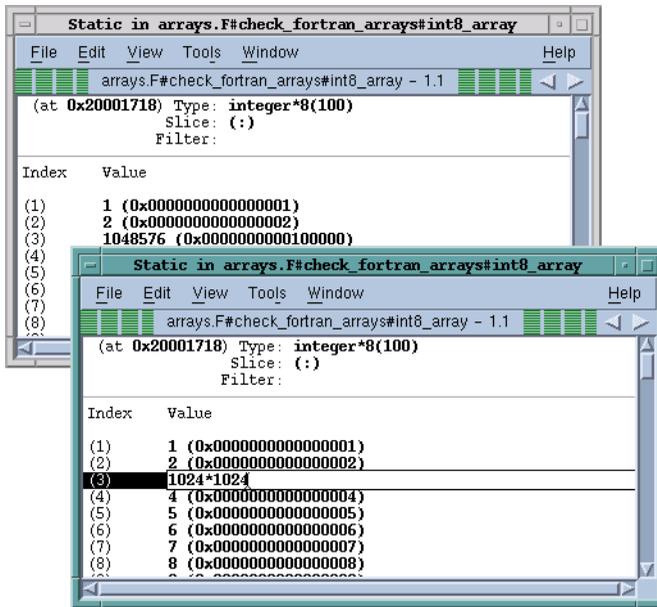


FIGURE 157: Using an Expression to Change a Value

While TotalView does not let you directly change the value of bit fields, the **Tools > Evaluate** Window lets you assign a value to a bit field. See “*Evaluating Expressions*” on page 371. Similarly, you cannot directly change the value of fields in nested structures; you must first dive into the value. When TotalView displays a value in bold, it is ready to be edited.

CLI EQUIVALENT: Tcl lets you use operators such as **&** and **|** to manipulate bit fields on Tcl values.

Changing the Data Type of Variables

The data type declared for the variable determines its format and size (amount of memory). For example, if you declare an **int** variable, TotalView displays the variable as an integer.

Topics in this section are:

- "Displaying C Data Types" on page 298
- "Pointers to Arrays" on page 299
- "Arrays" on page 299
- "Typedefs" on page 300
- "Structures" on page 300
- "Unions" on page 301
- "Built-In Types" on page 302
- "Type Casting Examples" on page 304

You can change the way TotalView displays data in the Variable Window by editing its data type. This is known as *casting*. TotalView assigns types to all data types, and in most cases, they are identical to their programming language counterparts.

- When a C variable is displayed in TotalView, the data types are identical to C type representations, except for pointers to arrays. TotalView uses a simpler syntax for pointers to arrays. (See "Pointers to Arrays" on page 299.)
- When Fortran is displayed in TotalView, the types are identical to Fortran type representations for most data types including **INTEGER**, **REAL**, **DOUBLE PRECISION**, **COMPLEX**, **LOGICAL**, and **CHARACTER**.

If the window contains a structure with a list of fields, you can edit the data types of the fields listed in the window.

NOTE When you edit a data type, TotalView changes how it displays the variable in the current window. Other windows listing the variable do not change.

Displaying C Data Types

TotalView's syntax for displaying data is identical to C Language cast syntax for all data types except pointers to arrays. That is, you should use C Language cast syntax for **int**, **short**, **unsigned**, **float**, **double**, **union**, and all named **struct** types. In addition, TotalView has a built-in type called **<string>**. Unless you tell it otherwise, it maps **char** arrays into this type.

TotalView types are read from right to left. For example, **<string>*[20]*** is a pointer to an array of 20 pointers to **<string>**.

Table 13 shows some common data types.

Table 13: Common Types

Data Type String	Meaning
<code>int</code>	Integer
<code>int*</code>	Pointer to integer
<code>int[10]</code>	Array of 10 integers
<code><string></code>	Null-terminated character string
<code><string>**</code>	Pointer to a pointer to a null-terminated character string
<code><string>*[20]*</code>	Pointer to an array of 20 pointers to null-terminated strings

You can also enter C Language cast syntax verbatim in the type field for any type.

The following sections discuss the more complex types.

Pointers to Arrays

Suppose you declared a variable `vbl` as a pointer to an array of 23 pointers to an array of 12 objects of type `mytype_t`. The C language declaration for this is:

```
mytype_t ((*vbl)[23])[12];
```

Here is how you would cast the `vbl` variable to this type:

```
(mytype_t ((*)[23])[12])vbl
```

The TotalView cast for `vbl` is:

```
mytype_t[12]*[23]*
```

Arrays

Array type names can include a lower and upper bound separated by a colon (:).

By default, the lower bound for a C or C++ array is 0, and the lower bound for Fortran is 1. In the following example, an array of ten integers is declared in C and then in Fortran:

```
int a[10];
integer a(10)
```

The elements of the array range from `a[0]` to `a[9]` in C, while the elements of the equivalent Fortran array range from `a(1)` to `a(10)`.

When the lower bound for an array dimension is the default for the language, TotalView displays only the extent (that is, the number of elements in the dimension). Consider the following Fortran array declaration:

```
integer a(1:7, 1:8)
```

Since both dimensions of the array use the default lower bound for Fortran, which is 1, TotalView displays the data type of the array by using only the extent of each dimension, as follows:

```
integer (7, 8)
```

If an array declaration doesn't use the default lower bound, TotalView displays both the lower bound and upper bound for each dimension of the array. For example, in Fortran, you would declare an array of integers with the first dimension ranging from -1 to 5 and the second dimension ranging from 2 to 10, as follows:

```
integer a(-1:5, 2:10)
```

TotalView displays this in exactly the same way.

When editing an array's dimension, you can enter just the extent (if using the default lower bound) or you can enter the lower and upper bounds separated by a colon.

TotalView also lets you display a subsection of an array, or filter a scalar array for values matching a filter expression. Refer to "Displaying Array Slices" on page 319 and "Array Data Filtering" on page 324 for further information.

Typedefs

TotalView recognizes the names defined with **typedef**, and displays these user-defined types. For example:

```
typedef double *dptr_t;  
dptr_t p_vbl;
```

TotalView will display the type for **p_vbl** as **dptr_t**.

Structures

TotalView lets you use the **struct** keyword as part of a type string. In all cases, it is usually optional. If you have a structure and another data type with the same name,

however, you must include the **struct** keyword so that TotalView can distinguish between the two data types.

If you name a structure using **typedef**, the debugger uses the **typedef** name as the type string. Otherwise, the debugger uses the structure tag for the **struct**.

For example, consider the structure definition:

```
typedef struct mystruc_struct {
    int field_1;
    int field_2;
} mystruc_type;
```

TotalView displays **mystruc_type** as the type for **struct mystruc_struct**.

Unions

TotalView displays a union in the same way that it displays a structure. Even though the fields of a union are overlaid in storage, TotalView displays the fields on separate lines. (See Figure 158.)

CLI EQUIVALENT: `dprint variable`

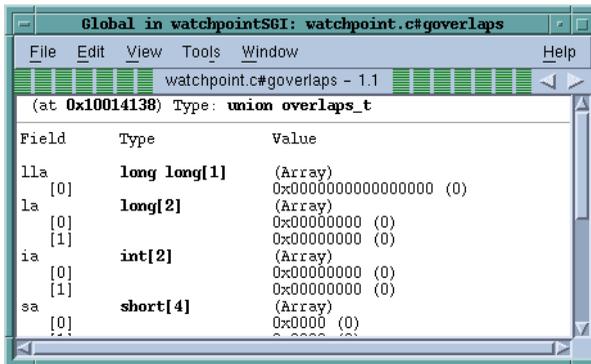


FIGURE 158: Displaying a Union

When TotalView displays some complex arrays and structures, it displays the compound object or array types in the Variable Window.

NOTE Editing the compound object or array types could yield undesirable results.

Built-In Types

TotalView provides a number of predefined types. These types are enclosed in angle brackets (<>) to avoid conflict with types contained in a programming language. You can use these built-in types anywhere you can use ones defined in your programming language. These types are also useful when debugging executables with no debugging symbol table information. The following table lists the built-in types.

Table 14: Built-In Types

Type String	Language	Size	Meaning
<address>	C	void*	Void pointer (address)
<char>	C	char	Character
<character>	Fortran	character	Character
<code>	C	<i>architecture-dependent</i>	Machine instructions The size used here is the number of bytes required to hold the shortest instruction for your computer.
<complex>	Fortran	complex	Single-precision floating-point complex number. complex types contain a real part and an imaginary part, which are both of type real .
<complex*8>	Fortran	complex*8	real*4 -precision floating-point complex number complex*8 types contain a real part and an imaginary part, which are both of type real*4 .
<complex*16>	Fortran	complex*16	real*8 -precision floating-point complex number complex*16 types contain a real part and an imaginary part, which are both of type real*8 .
<double>	C	double	Double-precision floating-point number
<double precision>	Fortran	double precision	Double-precision floating-point number

Table 14: Built-In Types (cont.)

Type String	Language	Size	Meaning
<extended>	C	long double	Extended-precision floating-point number Extended-precision numbers must be supported by the target architecture.
<float>	C	float	Single-precision floating-point number
<int>	C	int	Integer
<integer>	Fortran	integer	Integer
<integer*1>	Fortran	integer*1	One-byte integer
<integer*2>	Fortran	integer*2	Two-byte integer
<integer*4>	Fortran	integer*4	Four-byte integer
<integer*8>	Fortran	integer*8	Eight-byte integer
<logical>	Fortran	logical	Logical
<logical*1>	Fortran	logical*1	One-byte logical
<logical*2>	Fortran	logical*2	Two-byte logical
<logical*4>	Fortran	logical*4	Four-byte logical
<logical*8>	Fortran	logical*8	Eight-byte logical
<long>	C	long	Long integer
<long long>	C	long long	Long long integer
<real>	Fortran	real	Single-precision floating-point number NOTE When using a value such as real, be careful that the actual data type used by your computer is not real*4 or real*8 as different results could occur.
<real*4>	Fortran	real*4	Four-byte floating-point number
<real*8>	Fortran	real*8	Eight-byte floating-point number
<real*16>	Fortran	real*16	Sixteen-byte floating-point number
<short>	C	short	Short integer
<string>	C	char	Array of characters
<void>	C	long	Area of memory

The next sections contain more information about the following built-in types:

- Character Arrays (<string> Data Type)
- Areas of Memory (<void> Data Type)
- Instructions (<code> Data Type)

Character Arrays (<string> Data Type)

If you declare a character array as `char vbl[n]`, TotalView automatically changes the type to `<string>[n]`; that is, a null-terminated, quoted string with a maximum length of n . This means that TotalView will display an array as a quoted string of n characters, terminated by a null character. Similarly, TotalView changes `char*` declarations to `<string>*` (a pointer to a null-terminated string).

Since most C character arrays represent strings, the `<string>` type can be very convenient. If this isn't what you want, you can edit the `<string>` back to a `char` (or `char[n]`) to display the variable as you declared it.

Areas of Memory (<void> Data Type)

TotalView uses the `<void>` type for data of an unknown type, such as the data contained in registers or in an arbitrary block of memory. The `<void>` type is similar to the `int` in the C language.

If you dive into registers or display an area of memory, TotalView lists the contents as a `<void>` data type. Furthermore, if you display an array of `<void>` variables, the index for each object in the array is the address, not an integer. This address can be useful in displaying large areas of memory.

If desired, you can change a `<void>` into another type. Similarly, you can change any type into a `<void>` to see the variable in decimal and hexadecimal formats.

Instructions (<code> Data Type)

TotalView uses the `<code>` data type to display the contents of a location as machine instructions. Thus, to look at disassembled code stored at a location, dive on the location and change the type to `<code>`. To specify a block of locations, use `<code>[n]`, where n is the number of locations being displayed.

Type Casting Examples

This section contains three type casting examples:

- Displaying the argv Array
- Displaying Declared Arrays
- Displaying Allocated Arrays

Displaying the argv Array

Typically, `argv` is the second argument passed to `main()`, and it is either a `char **argv` or `char *argv[]`. Suppose `argv` points to an array of three pointers to character strings. Here is how you can edit its type to display an array of three pointers:

- 1 Select the type string for `argv`.

CLI EQUIVALENT: `dprint argv`

- 2 Edit the type string using the field editor commands. Change it to: `<string>* [3] *`

CLI EQUIVALENT: `dprint (<string>* [3] *)argv`

- 3 To display the array, dive into the value field for `argv`. (See Figure 159 on page 305.)

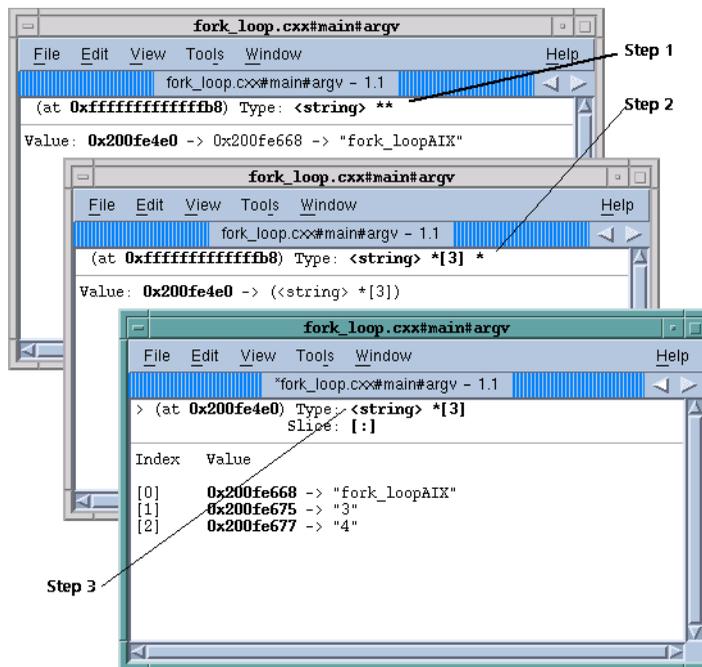


FIGURE 159: Editing argv

Displaying Declared Arrays

TotalView displays arrays in the same way as it displays local and global variables. In the Stack Frame or Source Pane, dive into the declared array. A Variable Window displays the elements of the array.

CLI EQUIVALENT: `dprint array-name`

Displaying Allocated Arrays

The C language uses pointers for dynamically allocated arrays. For example:

```
int *p = malloc(sizeof(int) * 20);
```

Because TotalView doesn't know that `p` actually points to an array of integers, here is how you would display the array:

- 1 Dive on the variable `p` of type `int*`.
- 2 Change its type to `int[20]*`.
- 3 Dive on the value of the pointer to display the array of 20 integers.

Working with Opaque Data

An opaque type is a data type that isn't fully specified, is hidden, or whose declaration is deferred. For example, the following C declaration defines the data type for `p` as a pointer to `struct foo`, which is not yet defined:

```
struct foo;  
struct foo *p;
```

When TotalView encounters this kind of information, it may indicate that `foo`'s data type is `<opaque>`. For example:

```
struct foo <opaque>
```

Changing the Address of Variables

You can edit the address of a variable in a Variable Window. When you edit the address, the Variable Window shows the contents of the new location.

You can also enter an address expression, such as `0x10b8 - 0x80`.

Changing Types to Display Machine Instructions

Here is how you can display machine instructions in a Variable Window:

- 1 Select the type string at the top of the Variable Window.
- 2 Change the type string to be an array of `<code>` data types, where *n* indicates the number of instructions to be displayed. For example:

`<code> [n]`

TotalView displays the contents of the current variable, register, or area of memory as machine-level instructions.

The Variable Window (shown in Figure 152 on page 290) lists the following information about each machine instruction:

Address	The machine address of the instruction.
Value	The hexadecimal value stored in the location.
Disassembly	The instruction and operands stored in the location.
Offset+Label	The symbolic address of the location as a hexadecimal offset from a routine name.

You can also edit the value listed in the **Value** field for each machine instruction.

Displaying C++ Types

Classes

TotalView displays C++ classes and accepts **class** as a keyword. When you debug C++, TotalView also accepts the *unadorned* name of a **class**, **struct**, **union**, or **enum** in the type field. TotalView displays nested classes that use inheritance, showing derivation by indentation.

NOTE Some C++ compilers do not output accessibility information. In these cases, the information is omitted from the display.

For example, Figure 160 displays an object of a **class c**.

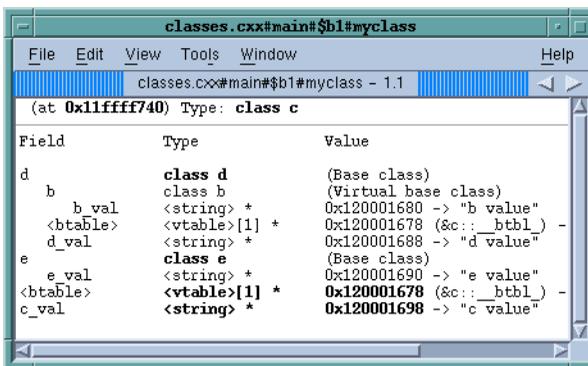


FIGURE 160: Displaying C++ Classes That Use Inheritance

The definition is as follows:

```
class b {
    char * b_val;
public:
    b() {b_val = "b value";}
};

class d : virtual public b {
    char * d_val;
public:
    d() {d_val = "d value";}
};

class e {
    char * e_val;
public:
    e() {e_val = "e value";}
};

class c : public d, public e {
    char * c_val;
public:
    c() {c_val = "c value";}
};
```

Changing Class Types in C++

TotalView tries to display the correct data when you change the type of a Variable Window as it and you move up or down the derivation hierarchy.

If a change in the data type also requires a change in the address of the data being displayed, TotalView asks you about changing the address. For example, if you edit a Variable Window's **Type** field from **class c** to **class e**, TotalView displays the following dialog box.

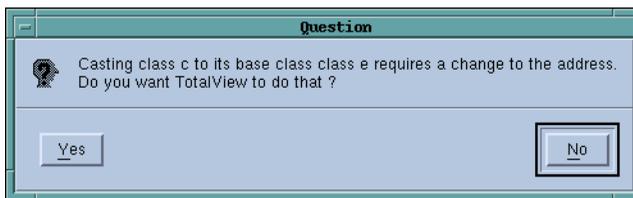


FIGURE 161: C++ Type Cast to Base Class Question Window

Selecting **Yes** tells TotalView to change the address to ensure that it displays the correct base class member. Selecting **No** tells TotalView to display the memory area as if it were an instance of the type to which it is being cast, leaving the address unchanged.

Similarly, if you change a data type in the Variable Window because you want to cast a base class to a derived class, and that change requires an address change, TotalView asks that you confirm the address change. For example, Figure 162 shows the dialog posted if you cast from **class e** to **class c**.

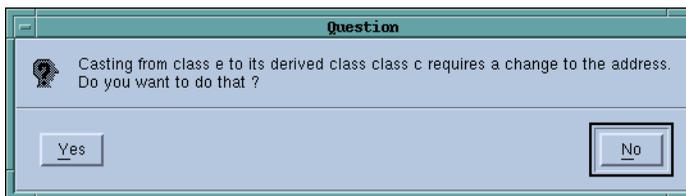


FIGURE 162: C++ Type Cast to Derived Class Question Window

Displaying Fortran Types

TotalView allows you to display FORTRAN 77 and Fortran 90 data types.

Topics in this section are:

- "Displaying Fortran Common Blocks" on page 310
- "Displaying Fortran Module Data" on page 310
- "Debugging Fortran 90 Modules" on page 312
- "Fortran 90 User-Defined Types" on page 314
- "Fortran 90 Deferred Shape Array Types" on page 314
- "Fortran 90 Pointer Types" on page 315
- "Displaying Fortran Parameters" on page 315

Displaying Fortran Common Blocks

For each common block defined within the scope of a subroutine or function, TotalView creates an entry in that function's common block list. The Stack Frame Pane displays the name of each common block for a function. The names of common block members have function scope, not global scope.

CLI EQUIVALENT: `dprint variable-name`

If you dive on a common block name in the Stack Frame Pane, TotalView displays the entire common block in a Variable Window, as shown in Figure 163 on page 311.

The top-left pane shows a common block list in a Stack Frame Pane. The bottom right window shows the results of diving on the common block to see its elements.

If you dive on a common block member name, TotalView searches all common blocks in the function's scope for a matching member name and displays the member in a Variable Window.

Displaying Fortran Module Data

TotalView tries to locate all data associated with a Fortran module and provide a single display that contains all of it. For functions and subroutines defined in a

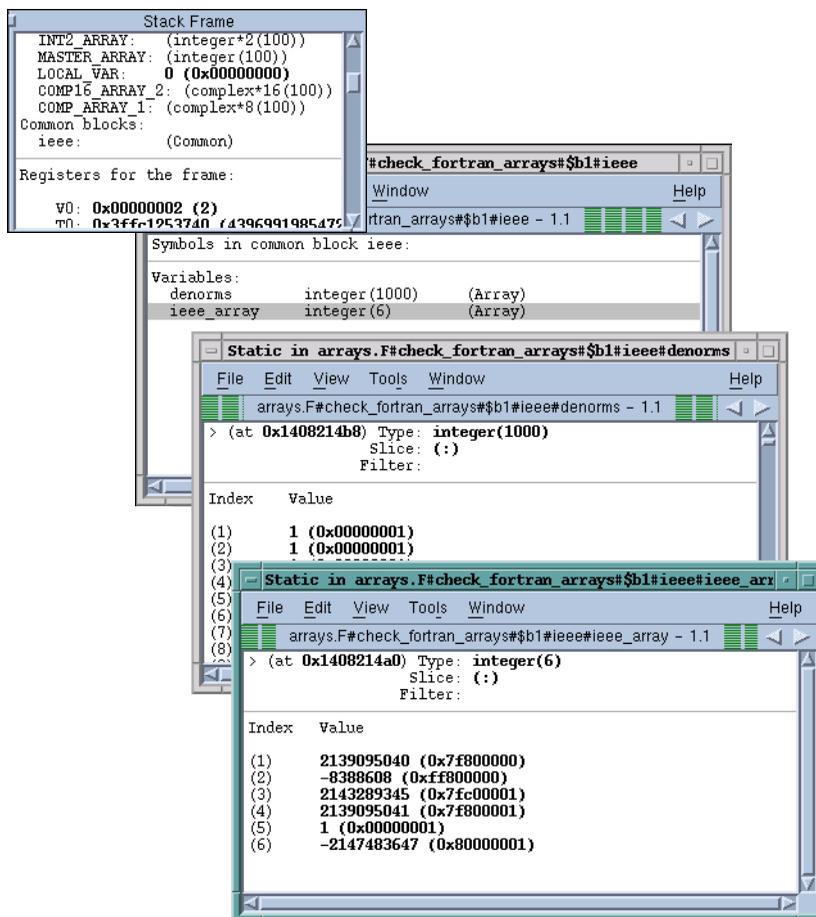


FIGURE 163: Diving into a Common Block List in the Stack Frame Pane

module, TotalView adds the full module data definition to the list of modules displayed in the Stack Frame Pane.

CLI EQUIVALENT: **dprint** *variable-name*

NOTE TotalView only displays a module if it contains data. Also, the amount of information that your compiler gives TotalView may restrict what's displayed.

Although a function may use a module, TotalView doesn't always know if the module was used or what the true names of the variables in the module are. If this happens, either:

- Module variables appear as local variables of the subroutine.
- A module appears on the list of modules in the Stack Frame Pane that contains (with renaming) only the variables used by the subroutine.



Alternatively, you can view a list of all the known modules by using the **Tools > Fortran Modules** command. Like in any Variable Window, you can dive through an entry to display the actual module data, as shown in Figure 164 on page 313.

NOTE If you are using the SUNPro compiler, TotalView can only display module data if you force it to read the debug information for a file that contains the module definition or a module function. For more information, see "*Finding the Source Code for Functions*" on page 213.

Debugging Fortran 90 Modules

Fortran 90 and Fortran 95 let you place functions, subroutines, and variables inside modules. These modules can then be included elsewhere using a **USE** command. When doing this, the names in the module become available in the *using* compilation unit, unless you exclude them with a **USE ONLY** statement, or rename them. This means that you don't need to explicitly qualify the name of a module function or variable from the Fortran source code.

When debugging this kind of information, you will need to know the location of the function being called. Consequently, TotalView uses the following syntax when it displays a function contained in a module:

modulename`functionname

You can also use this syntax in the **File > New Program** and **View > Lookup Variable** commands.

Fortran 90 also introduced the idea of a contained function that is only visible in the scope of its parent and siblings. There can be many contained functions in a program, all using the same name. If the compiler gave TotalView the function name for a nested function, TotalView displays it using the following syntax:

Dive on module name to see Variable Window containing module variables

The screenshot displays several windows from the TotalView IDE:

- Stack Frame:** Shows the function "stuff`testarrays" with parameters n=100 and Modules: stuff: (Module). It also lists registers R0 through R4.
- Fortran Modules:** A window titled "Fortran Modules - 1 'f90stepAIX1'" listing modules from process '1 "f90stepAIX1": drivers, setup_module, parameter_module, stuff, init_module, and several _hpf_random_module instances.
- Symbols in module stuff:** A window showing a list of variables: arr1, arr2, arr3 (all real*8 arrays), typ1, typ2, typ3 (all type(whopper) arrays).
- Global in f90stepSGI: f90step.f#stuff testtypes#stuff:** A window showing the structure of a variable. It indicates the type is type(WHOPPER) and shows a table of fields:

Field	Type	Value
(1)	type(WHOPPER)	(Struct)
FLAGS	enum LOGICAL_4(1000)	(Array)
DPSA	REAL_8(1000)	(Array)
DPPA	REAL_8(:,)pointer	0x1001a2c8 -> (REAL_8)
<padding>	<char>[64]	(Array)
(2)	type(WHOPPER)	(Struct)
FLAGS	enum LOGICAL_4(1000)	(Array)
DPSA	REAL_8(1000)	(Array)

Arrows indicate the flow of information: from the Stack Frame to the Fortran Modules window, and from the Symbols in module stuff window to the Global in f90stepSGI window.

or

Dive on module variable to see a Variable Window with more detail

FIGURE 164: Fortran Modules Window

parentfunction() `containedfunction

CLI EQUIVALENT: **dprint** *module_name*'*variable_name*

Fortran 90 User-Defined Types

A Fortran 90 user-defined type is similar to a C structure. TotalView displays a user-defined type as `type(name)`, which is the same syntax used in Fortran 90 to create a user-defined type. For example, here is a code fragment that defines a variable `typ2` of `type(whopper)`:

```
TYPE WHOPPER
  LOGICAL, DIMENSION(ISIZE) :: FLAGS
  DOUBLE PRECISION, DIMENSION(ISIZE) :: DPSA
  DOUBLE PRECISION, DIMENSION(:), POINTER :: DPPA
END TYPE WHOPPER
```

```
TYPE(WHOPPER), DIMENSION(:), ALLOCATABLE :: TYP2
```

TotalView displays this code as shown in Figure 165.

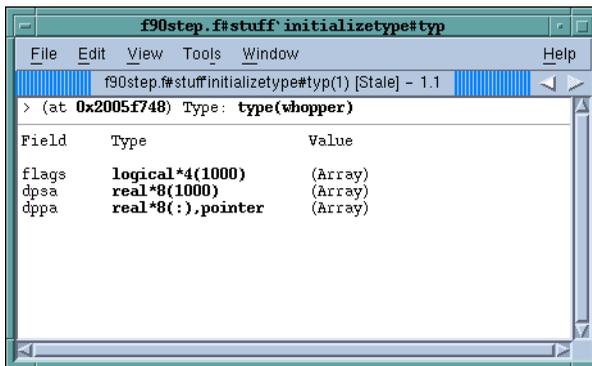


FIGURE 165: Fortran 90 User-Defined Type

Fortran 90 Deferred Shape Array Types

Fortran 90 allows you to define deferred shape arrays and pointers. The actual bounds of the array are not determined until the array is allocated, the pointer is assigned, or, in the case of an assumed shape argument to a subroutine, the subroutine is called. TotalView displays the type of deferred shape arrays as `type(:)`.

When TotalView displays the data for a deferred shape array, it displays the type used in the definition of the variable and the actual type that this instance of the variable has. The actual type is not editable since you can achieve the same effect

by editing the definition's type. The following example shows the type of a deferred shape rank 2 array of **real** data with runtime lower bounds of -1 and 2, and upper bounds of 5 and 10:

```
Type: real(:, :)
Actual Type: real(-1:5, 2:10)
Slice: (:, :)
```

Fortran 90 Pointer Types

A Fortran 90 pointer type allows you to point to scalar or array types.

TotalView implicitly handles slicing operations that set up a pointer or assumed shape subroutine argument so that indices and values it displays in a Variable Window are the same as you would see in the Fortran code. For example:

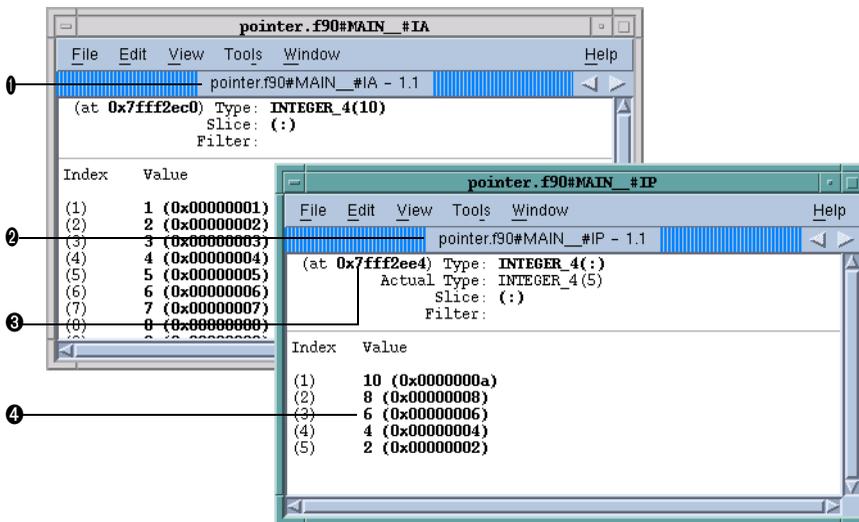
```
integer, dimension(10), target :: ia
integer, dimension(:), pointer :: ip
do i = 1, 10
    ia(i) = i
end do
ip => ia(10:1:-2)
```

After diving through the **ip** pointer, TotalView displays the window shown in Figure 166 on page 316.

Notice that the address displayed is not that of the array's base. Since the array's stride is negative, array elements that follow are at lower absolute addresses. Consequently, the address displayed is that of the array element having the lowest index. This may not be the first displayed element if you used a slice to display the array with reversed indices.

Displaying Fortran Parameters

A Fortran **PARAMETER** defines a named constant. Most compilers do not generate information that TotalView can use to determine what a **PARAMETER**'s value is. This means that you must make a few changes to your program if you want to see this kind of information.



- 1 Target array ia
- 2 Pointer ip into array ia
- 3 Address of ip(1)
- 4 Values reflect slice

FIGURE 166: Fortran 90 Pointer Value

If you're using Fortran 90, you can define variables in a module that you initialize to the value of these **PARAMETER** constants. For example:

```
INCLUDE 'PARAMS.INC'

MODULE CONSTS
  SAVE
  INTEGER PI_C = PI
  ...
END MODULE CONSTS
```

The **PARAMS.INC** file will contain your parameter definitions. You would then use these parameters to initialize variables in a module. After you compile and link this module into your program, the value of these *parameter variables* are visible.

If you're using Fortran 77, you could achieve the same results if you make the assignments in a common block and then include the block in **main()**. You could also use a block data subroutine to access this information.

Displaying Thread Objects

On HP Alpha Tru64 UNIX and IBM AIX systems, TotalView can display information about mutexes and conditional variables. In addition, TotalView can display information on read/write locks and data keys on IBM AIX. You can obtain this information by selecting the **Tools > Thread Objects** command. After selecting this command, TotalView displays a window that will either contain two tabs (HP Alpha) or four tabs (IBM). Figure 167 on page 318 shows some AIX examples.

Diving on any line in these windows displays a Variable Window containing additional information about the item.

Here are some things you should know:

- If you're displaying data keys, many applications initially set keys to 0 (the NULL pointer value). TotalView doesn't display a key's information, however, until a thread sets a non-NULL value to the key.
- If you select a thread ID in a data key window, you can dive on it using the **View > Dive Thread** and **View > Dive Thread New** commands to display a Process Window for that thread ID.

The online Help contains considerable information on the contents of these windows.

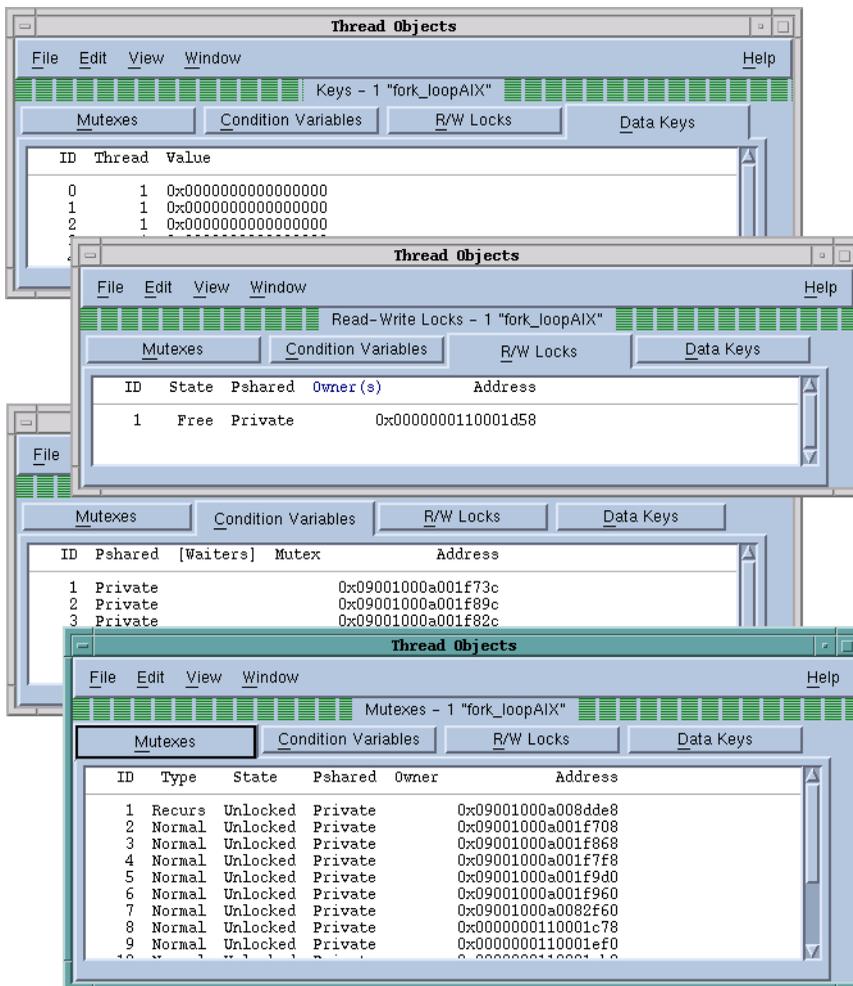


FIGURE 167: Thread Objects Page on an IBM AIX machine

Chapter 13

Examining Arrays

This chapter explains how to examine and change data as you debug your program. You will learn about the following:

- “*Examining and Analyzing Arrays*” on page 319
- “*Displaying a Variable in All Processes or Threads*” on page 333
- “*Visualizing Array Data*” on page 335

Examining and Analyzing Arrays

TotalView can quickly display very large arrays in Variable Windows. An array can be the elements that you’ve defined in your program or it can be an area of memory that you’ve cast into an array.

If an array overlaps nonexistent memory, the initial portion of the array is correctly formatted. If memory isn’t allocated for an array element, TotalView displays **Bad Address** in the element’s subscript.

Topics in this section are:

- “*Displaying Array Slices*” on page 319
- “*Array Data Filtering*” on page 324
- “*Sorting Array Data*” on page 330
- “*Obtaining Array Statistics*” on page 331

Displaying Array Slices

TotalView lets you display array subsections by editing the *slice field* in an array’s Variable Window. (An array subsection is called a *slice*.) The slice field contains place-

holders for all array dimensions. For example, here is a C declaration for a three-dimensional array:

```
integer an_array[10][20][5]
```

Because this is a three-dimensional array, the initial slice definition is `[:][:][:]`. This lets you know that the array has three dimensions and that TotalView is displaying all array elements.

Here is a deferred shape array definition for a two-dimensional array variable:

```
integer, dimension (:,:) :: another_array
```

Its TotalView slice definition is `(:,:)`.

As you can see, TotalView displays as many colons (`:`) as there are array dimensions. For example, the slice definition for a one-dimensional array (a vector) is `[:]` for C arrays and `(:)` for Fortran arrays.

CLI EQUIVALENT: `dprint an_array[n:m,p:q]`
`dprint an_array(n:m,p:q)`

Using Slices and Strides

A slice definition has the following form:

```
lower_bound:upper_bound:stride
```

(The *stride* tells TotalView that it should skip over elements and not display them. Adding a *stride* to a slice tells TotalView to display every *stride* element of the array, starting at the *lower_bound* and continuing through the *upper_bound*, inclusive.

For example, a slice of `[0:9:9]` used on a 10-element C array tells TotalView to display the first element and last element, which is the ninth element beyond the lower bound.

If the stride is negative and the upper bound is greater than the lower bound, TotalView lets you view a dimension with reversed indexing. That is, TotalView treats the slice as if it were:

```
[ub : lb : stride]
```

CLI EQUIVALENT: `dprint an_array(n:m,p,q:r:s)`

For example, the following definition tells TotalView to display an array beginning at its last value and moving to its first:

```
[::-1]
```

In contrast, Fortran 90 requires that you explicitly enter the upper and lower bounds when you're reversing the order in which it displays array elements.

Because the default value for the stride is 1, you can omit the stride (and the colon that precedes it) if your stride value is 1. For example, the following two definitions display array elements 0 through 9:

```
[0:9:1]
[0:9]
```

If the lower and upper bounds are the same, just use a single number. For example, the following two definitions tell TotalView to display array element 9:

```
[9:9:1]
[9]
```

NOTE The `lower_bound`, `upper_bound`, and `stride` can only be constants.

Example 1: A slice declaration of `[::2]` for a C or C++ array (with a default lower bound of 0) tells TotalView to display elements with even indices of the array: 0, 2, 4, and so on. However, if this were defined for a Fortran array (where the default lower bound is 1), TotalView displays elements with odd indices of the array: 1, 3, 5, and so on.

Example 2: Figure 168 on page 322 displays a slice of `(::9,::9)`. This definition displays the four corners of a 10-element by 10-element Fortran array.

Example 3: You can use a stride to invert the order *and* skip elements. For example, here is a slice that begins with the upper bound of the array and display every other element until it reaches the lower bound of the array: `(::-2)`. Thus, using `(::-2)` with a Fortran `integer(10)` array tells TotalView to display the elements 10, 8, 6, 4, and 2.

Example 4: You can simultaneously invert the array's order and limit its extent to display a small section of a large array. Figure 169 shows how to specify a `(2:3,7::-1)` slice with an `integer*4(-1:5,2:10)` Fortran array.

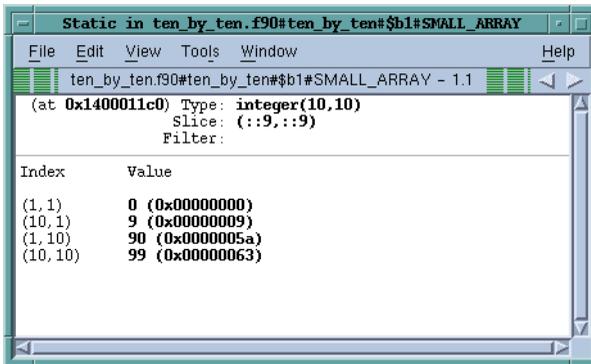


FIGURE 168: Slice Displaying the Four Corners of an Array

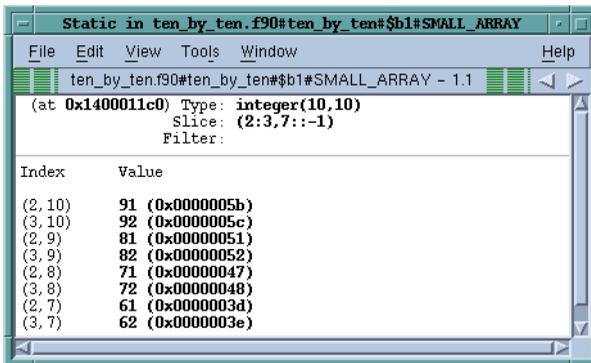


FIGURE 169: Fortran Array with Inverse Order and Limited Extent

After you enter this slice value, TotalView only shows elements in rows 2 and 3 of the array, beginning with column 10 and ending with column 7.

Using Slices in the Lookup Variable Command

When you use the **View > Lookup Variable** command to display a Variable Window, you can include a slice expression as part of the variable name. Specifically, if you type an array name followed by a set of slice descriptions in the **View > Lookup Variable** command's dialog box, TotalView initializes the slice field in the Variable Window to this slice descriptions.

If you add subscripts to an array name in the **View > Lookup Variable** command's dialog box, TotalView interprets these subscripts as a slice description rather than

as a request to display an individual value of the array. As a result, you can display different values of the array by changing the slice expression.

For example, suppose that you have a 10-element by 10-element Fortran array named `small_array`, and you want to display element (5,5). Using the **View > Lookup Variable** command, type `small_array(5,5)`. This sets the initial slice to (5:5,5:5). This is the top-left screen in Figure 170.

CLI EQUIVALENT: `dprint small_array(5,5)`

You can tell TotalView to display one of the array's values by enclosing the array name and subscripts (that is, the information normally included in a slice expression) within parentheses, such as `(small_array(5,5))`.

CLI EQUIVALENT: `dprint (small_array(5,5))`

In this case, the Variable Window just displays the type and value of the element and doesn't show its array index. This is shown in the center screen in Figure 170.

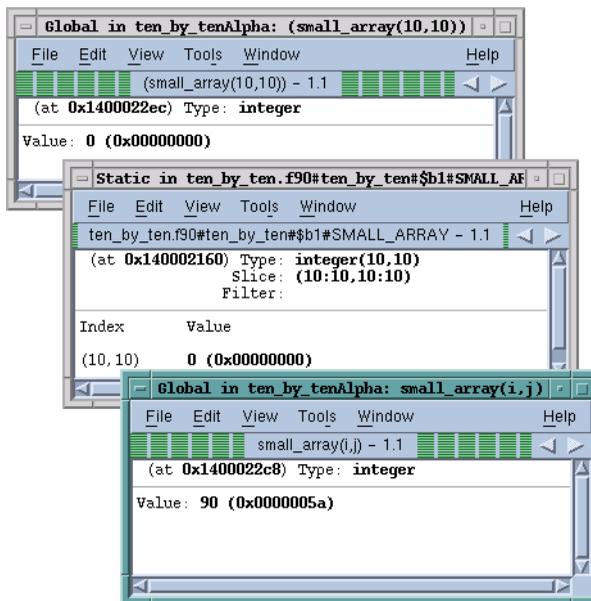


FIGURE 170: Variable Window for `small_array`

Perhaps the most interesting of the screens in Figure 170 on page 323 is the one in the bottom-right corner. This was created by doing a **View > Lookup Variable** with a value of `small_array(i,j)`. Here, TotalView evaluated the values of `i` and `j` before it displayed the window. If you do this, you should know that the values of `i` and `j` are just computed once. This means that if the values of `i` and `j` change, the displayed value will not change.

Array Data Filtering

You can restrict what TotalView displays in a Variable Window by adding a filter to the window. You can filter arrays of type character, integer, or floating point. Your filtering options are:

- Arithmetic comparison to a constant value
- Equal or not equal comparison to IEEE NANs, INFs, and DENORMs
- Within a range of values, inclusive or exclusive
- General expressions

When an element of an array matches the filter expression, TotalView includes the element in the Variable Window display.

Topics in this section are:

- “*Filtering Array Data*” on page 324
- “*Filtering by Comparison*” on page 325
- “*Filtering for IEEE Values*” on page 326
- “*Filtering By a Range of Values*” on page 327
- “*Creating Array Filter Expressions*” on page 329
- “*Using Filter Comparisons*” on page 329

Filtering Array Data

The procedure for filtering an array is quite simple: select the **Filter** field, enter the array filter expression, and then press Return.

TotalView updates the Variable Window to exclude only the elements that do not match the filter expression.

TotalView only displays an element if its value matches the filter expression and the slice operation.

If necessary, TotalView converts the array element before evaluating the filter expression. The following conversion rules apply:

- If the filter operand or array element type is floating point, TotalView converts it to a double-precision floating-point value. TotalView truncates extended-precision values to double precision. Converting integer or unsigned integer values to double-precision values may result in a loss of precision. TotalView converts unsigned integer values to non-negative double-precision values.
- If the filter operand or the array element is an unsigned integer, TotalView converts the values to an unsigned 64-bit integer.
- If both the filter operand and array element are of type integer, TotalView converts the values to type 64-bit integer.

These conversions modify a copy of the array's elements—they never alter the actual array elements.

To stop filtering an array, delete the contents of the **Filter** field in the Variable Window and press Return. TotalView will then update the Variable Window so that it includes all elements.

Filtering by Comparison

The simplest filters are ones whose formats are:

operator value

where *operator* is either a C/C++ or Fortran-style comparison operator, and *value* is a signed or unsigned integer constant, or a floating-point number. For example, here's the filter for displaying all values greater than 100:

> 100

Table 15 lists the comparison operators.

Table 15: Array Data Filtering Comparison Operators

Comparison	C/C++ Operator	Fortran Operator
Equal	==	.eq.
Not equal	!=	.ne.
Less than	<	.lt.

Table 15: Array Data Filtering Comparison Operators (cont.)

Comparison	C/C++ Operator	Fortran Operator
Less than or equal	<code><=</code>	<code>.le.</code>
Greater than	<code>></code>	<code>.gt.</code>
Greater than or equal	<code>>=</code>	<code>.ge.</code>

Figure 171 shows an array whose filter is "`< 0`". This indicates that TotalView should only display array elements whose value is less than 0 (zero).

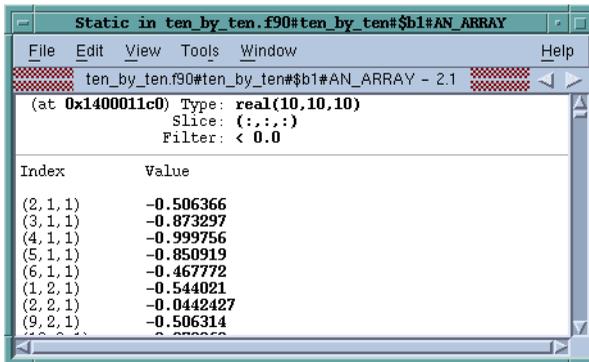


FIGURE 171: Array Data Filtering by Comparison

If the *value* you're using in the comparison is an integer constant, TotalView performs a signed comparison. If you add a **u** or **U** to the constant, TotalView performs an unsigned comparison.

Filtering for IEEE Values

You can filter IEEE NaN, infinity, or denormalized floating-point values by specifying a filter in the following form:

operator ieee-tag

The only comparison operators you can use are *equal* and *not equal*.

The *ieee-tag* represents an encoding of IEEE floating-point values, as explained in the following table:

Table 16: Array Data Filtering IEEE Tag Values

IEEE Tag Value	Meaning
\$nan	NaN (Not a number), either Quiet or Signaling
\$nanq	Quiet NaN
\$nans	Signaling NaN
\$inf	Infinity, either Positive or Negative
\$pinf	Positive Infinity
\$ninf	Negative Infinity
\$denorm	Denormalized number, either positive or negative
\$pdenorm	Positive denormalized number
\$ndenorm	Negative denormalized number

Figure 172 on page 328 shows an example of filtering an array for IEEE values. The bottom left Variable Window shows how TotalView displays the unfiltered array. Notice the NANQ, and NANS, INF, and -INF values. Then other two windows show filtered displays. The top left window only shows infinite values. The center window only shows the values of denormalized numbers.

Filtering By a Range of Values

Specify ranges as follows:

```
[>] low-value : [<] high-value
```

where *low-value* specifies the lowest value to include, and *high-value* specifies the highest value to include, separated by a colon. The high and low values are inclusive unless you use < and > symbols. If you specify a > before *low-value*, the low value is exclusive. Similarly, a < before *high-value* makes it exclusive.

low-value and *high-value* must be constants of type integer, unsigned integer, or floating point. The data type of *low-value* must be the same as the type of *high-value*, and *low-value* must be less than *high-value*. If *low-value* and *high-value* are integer constants, you can append a **u** or **U** to the value to force an unsigned comparison. Figure 173 on page 328 shows a filter that tells TotalView that it should only display values equal to or greater than 64 but less than 512.

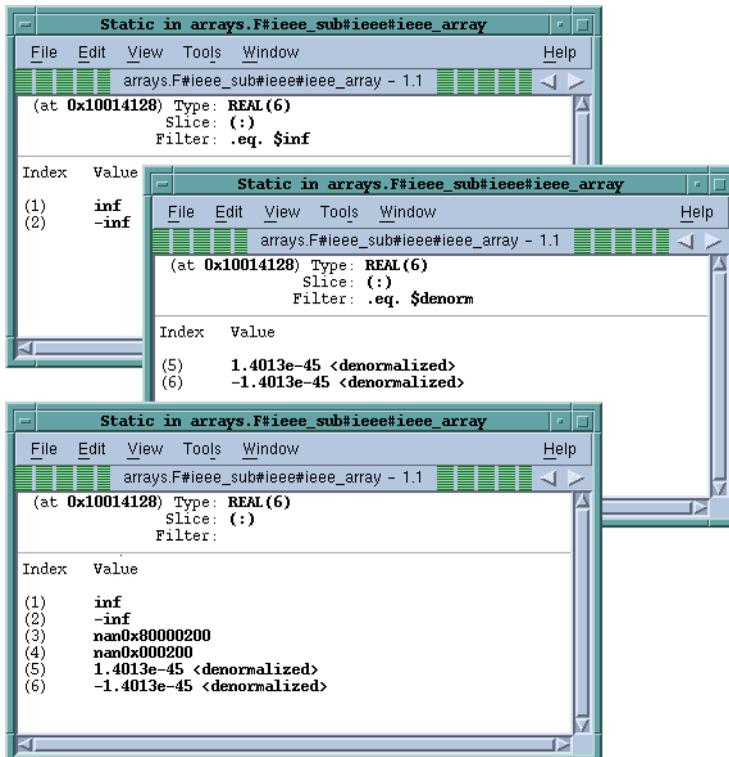


FIGURE 172: Array Data Filtering for IEEE Values

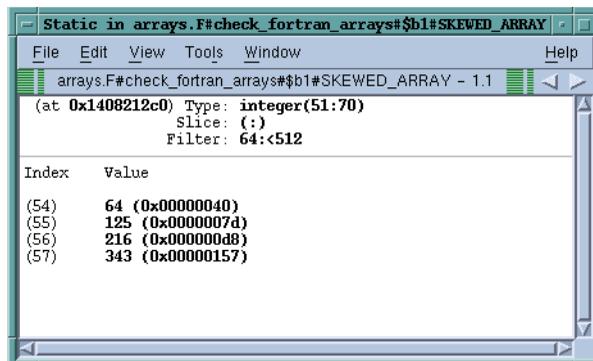


FIGURE 173: Array Data Filtering by Range of Values

Creating Array Filter Expressions

The filtering capabilities described in the previous sections are those that you will most often use. In some circumstances, you may need to create a more general expression. When you create a filter expression, you're creating a Fortran or C Boolean expression that TotalView evaluates for every element in the array or the array slice. For example, here is an expression that displays all array elements whose contents are greater than 0 and less than 50 or greater than 100 and less than 150.

```
($value > 0 && $value < 50) ||
($value > 100 && $value < 150)
```

Here's the Fortran equivalent:

```
($value .gt. 0 && $value .lt. 50) .or.
($value .gt. 100 .and. $value .lt.150)
```

\$value is a special TotalView variable that represents the current array element. You can now use this value when creating expressions.

Notice also the use of the **and** and **or** operators within the expression. The way in which TotalView computes the results of an expression is identical to the way it computes values at an evaluation point. For more information, see "*Defining Evaluation Points and Conditional Breakpoints*" on page 354.

NOTE You cannot use any of the IEEE tag values described in "*Filtering for IEEE Values*" on page 326 in these kinds of expressions.

Using Filter Comparisons

TotalView lets you filter array information in a variety of ways. This means that you can do the same thing in more than one way. For example, the following two filters display the same array items:

```
> 100
$value > 100
```

Similarly, the following expression displays the same array items:

```
>0:<100
$value > 0 && $value < 100
```

The only difference is that the first method is easier to type than the second. In general, you'd only use the second method when you're creating more complicated expressions.

Sorting Array Data

TotalView lets you sort the displayed array data into ascending or descending order. (It does not, of course, sort the actual data.)

If you select the Variable Window's **View > Sort > Ascending** command, TotalView places all of the array's elements in ascending order. (See Figure 174 for an example.)

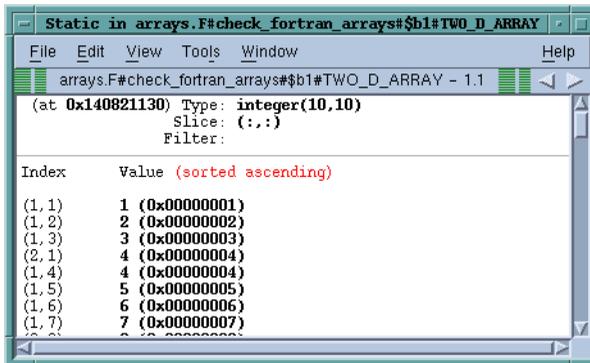


FIGURE 174: **Sorted Variable Window**

As you would expect, **View > Sort > Descending** places array elements into descending order. The **View > Sort > None** command returns the array to its original order.

The sort commands only manipulate the displayed elements. This means that if you limit the number of elements by defining a slice or a filter, TotalView only sorts the result of the filtering and slicing operations.

Obtaining Array Statistics

The **Tools > Statistics** command displays a window containing information about your array. Figure 175 shows an example.

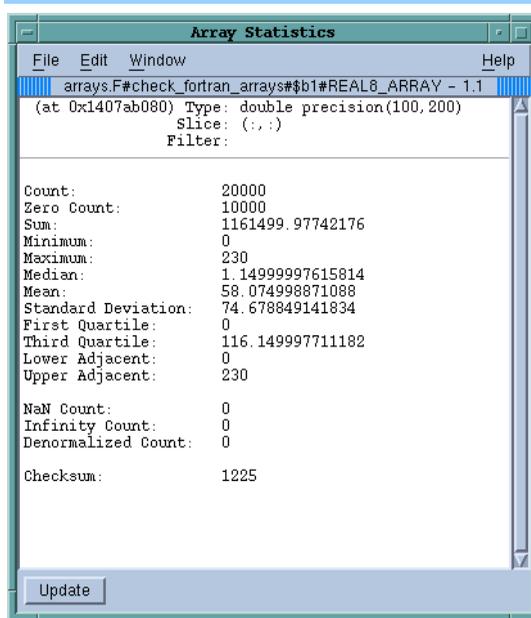


FIGURE 175: **Array Statistics Window**

If you have added a filter or a slice, these statistics only describe the information currently being displayed; they do not describe the entire unfiltered array. For example, if 90% of an array's values are less than 0 and you filter the array to show only values greater than zero, the median value will be positive even though the array's real median value is less than zero.

The statistics TotalView displays are as follows:

■ Checksum

A checksum value for the array elements.

■ Count

The total number of displayed array values. If you're displaying a floating-point array, this number doesn't include NaN or Infinity values.

■ Denormalized Count

A count of the number of denormalized values found in a floating-point array. This includes both negative and positive denormalized values as defined in the IEEE floating-point standard. Unlike other floating-point statistics, these elements participate in the statistical calculations.

■ Infinity Count

A count of the number of infinity values found in a floating-point array. This includes both negative and positive infinity as defined in the IEEE floating-point standard. These elements don't participate in statistical calculations.

■ Lower Adjacent

This value provides an estimate of the lower limit of the distribution. Values below this limit are called *outliers*. The lower adjacent value is the first quartile value minus 1.5 times the difference between the first and third quartiles.

■ Maximum

The largest array value.

■ Mean

The average value of array elements.

■ Median

The middle value. Half of the array's values are less than the median, and half are greater than the median.

■ Minimum

The smallest array value.

■ NaN Count

A count of the number of NaN values found in a floating-point array. This includes both signaling and quiet NaNs as defined in the IEEE floating-point standard. These elements don't participate in statistical calculations.

■ Quartiles, First and Third

Either the 25th or 75th percentile values. The first quartile value means that 25% of the array's values are less than this value and 75% are greater than this value. In contrast, the fourth quartile value means that 75% of the array's values are less than this value and 25% are greater.

■ Standard Deviation

The standard deviation for the array's values.

■ Sum

The sum of all of the displayed array's values.

■ Upper Adjacent

This value provides an estimate of the upper limit of the distribution. Values above this limit are called *outliers*. The upper adjacent value is the third quartile value plus 1.5 times the difference between the first and third quartiles.

■ Zero Count

The number of elements whose value is 0.

Displaying a Variable in All Processes or Threads

When you're debugging a parallel program that is running many instances of the same executable, you usually need to view or update the value of a variable in all of the processes or threads at once.

Before displaying a variable's value in all threads or processes, you must display an instance of the variable in a Variable Window. After TotalView displays this window, use one of the following commands:

- **View > Laminare > Process**, which displays the value of the variable in all of the processes.
- **View > Laminare > Thread**, which displays the value of a variable in all threads within a single process.

NOTE You cannot simultaneously laminate across processes and threads in the same Variable Window.

After using one of these commands, the Variable Window switches to "laminated" mode, and displays the value of the variable in each process or thread. Figure 176 on page 334 shows a simple, scalar variable in each of the processes in an OpenMP program. Notice that the first six have a variable in a matching call frame. The corresponding variable can't be found for the seventh thread.

If you decide that you no longer want the pane to be laminated, select the **View > Laminare > None** command to delaminate it.

When looking for a matching call frame, TotalView matches frames starting from the top frame, and considers calls from different memory or stack locations to be different calls. For example, the following definition of **recurse** contains two additional calls to **recurse**. Each of these generate nonmatching call frames.

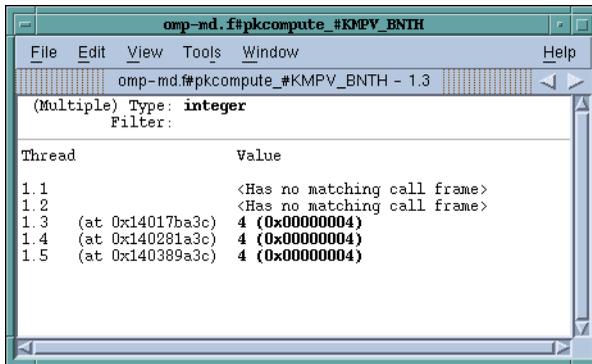


FIGURE 176: Laminated Scalar Variable

```

void recurse(int i) {
    if (i <= 0)
        return;
    if (i & 1)
        recurse(i - 1);
    else
        recurse(i - 1);
}

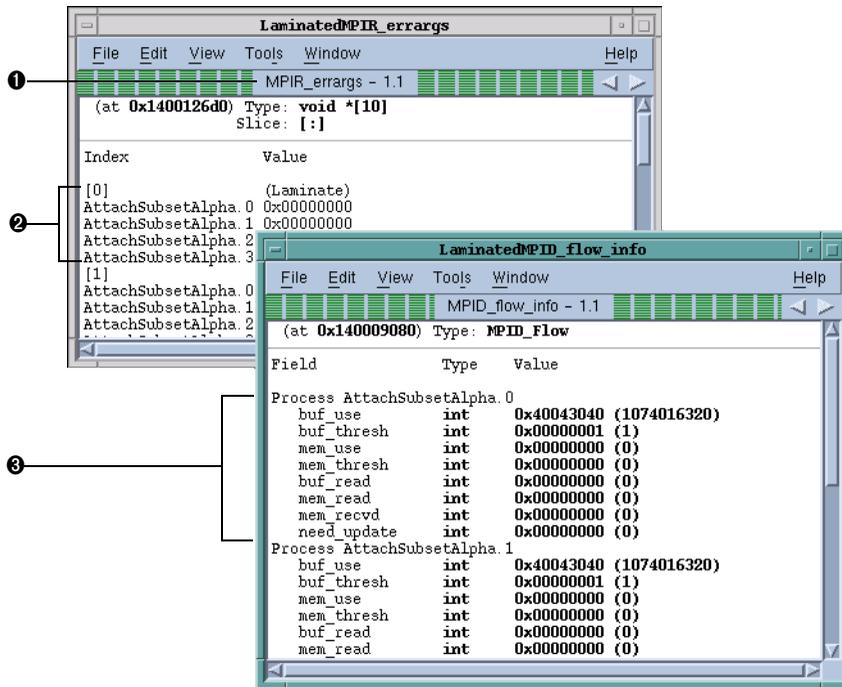
```

If the variables are at different addresses in the different processes or threads, the address field at the top of the pane displays **(Multiple)** and the unique addresses are displayed with each data item, as was shown in Figure 176.

TotalView also allows you to laminate arrays and structures. When you laminate an array, TotalView displays each element in the array across all processors. You can use a slice to select elements to be displayed in laminated windows. Figure 177 on page 335 shows an example of a laminated array and a laminated structure. You can also laminate an array of structures.

Diving in a Laminated Pane

You can dive through pointers in a laminated Variable Window, and the dive will apply to the associated pointer in each process or thread.



- ① Laminated array
- ② Element [0] for each of the processes
- ③ Structure elements for one process

FIGURE 177: Laminated Array and Structure

Editing a Laminated Variable

If you edit a value in a laminated Variable Window, TotalView asks if it should apply this change to all of the processes or threads or only the one in which you made a change. This is an easy way to update a variable in all processes.

Visualizing Array Data

The TotalView Visualizer lets you create graphic images of array data. This visualization lets you see your data in one glance and can help you quickly find problems with your data while you are debugging your programs.

You can execute the Visualizer from within TotalView or you can run it from the command line to visualize data dumped to a file in a previous TotalView session.

For information about running the TotalView Visualizer, see Chapter 7, “*Visualizing Programs and Data*” on page 159.

Visualizing a Laminated Variable Window

You can export data from a laminated Variable Window to the Visualizer by using the **Tools > Visualize** command. When visualizing laminated data, the process (or thread) index is the first axis of the visualization. This means that you must use one less data dimension than you normally would. If you do not want the process/thread axis to be significant, you can use a normal Variable Window since all of the data must be in one process.

Chapter 14

Setting Action Points

This chapter explains how to use action points. TotalView supports four kinds of action points: breakpoints, barrier points, evaluation points, and watchpoints. A *breakpoint* stops execution of processes and threads that reach it. A *barrier* point synchronizes a set of threads or processes at a location. An *evaluation point* causes a code fragment to execute when it is reached. A *watchpoint* lets you monitor a location in memory and stop execution when it changes.

Topics in this chapter are:

- "Action Points Overview" on page 337
- "Setting Breakpoints and Barriers" on page 339
- "Defining Evaluation Points and Conditional Breakpoints" on page 354
- "Using Watchpoints" on page 363
- "Saving Action Points to a File" on page 370
- "Evaluating Expressions" on page 371
- "Writing Code Fragments" on page 373

Action Points Overview

Action points allow you to specify an action that TotalView will perform when a thread or process reaches a source line or machine instruction in your program. Here are the different kinds of action points that you can use:

■ Breakpoints

When a thread encounters a breakpoint, it stops at the breakpoint. Other threads in the process will also stop. You can also indicate that you want other related processes to stop.

Breakpoints are the simplest kind of action point.

■ Barrier points

Barrier points are similar to simple breakpoints, differing in that you use them to synchronize a group of processes or threads. They hold each thread or process that reaches it until all threads or processes reach it. Barrier points work together with the TotalView hold and release feature. TotalView supports thread barrier and process barrier points.

■ Evaluation points

An evaluation point is a breakpoint that has a code fragment associated with it. When a thread or process encounters an evaluation point, it executes this code. You can use evaluation points in a variety of ways, including conditional breakpoints, thread-specific breakpoints, countdown breakpoints, and patching code fragments into and out of your program.

■ Watchpoints

A watchpoint tells TotalView that it should either stop the thread so that you can interact with your program (unconditional watchpoint) or evaluate an expression (conditional watchpoint).

All action points share some common properties.

- You can independently enable or disable them. A disabled action isn't deleted; however, when your program reaches a disabled action point, TotalView ignores it.
- You can share action points across multiple processes, or set them in individual processes.
- Action points apply to the process, so in a multithreaded process, the action point applies to all of the threads contained in the process.
- TotalView assigns unique ID numbers to each action point. These IDs appear in several places, including the Root Window, the Action Points Pane of the Process Window, and the **Action Point > Properties** Dialog Box.

Each type of action point has a unique symbol. Figure 178 on page 339 shows some of them.

CLI EQUIVALENT: **dactions** shows information about action points.

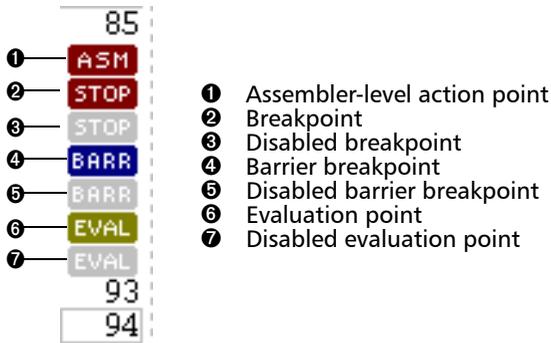


FIGURE 178: Action Point Symbols

The **ASM** icon lets you know that there are one or more assembler-level action points associated with the source line.

CLI EQUIVALENT: All action points display as “@” when you use the `dlist` command to display your source code. Use the `dactions` command to see what kind of action point is set.

Setting Breakpoints and Barriers

TotalView has several options for setting breakpoints. You can set:

- Source-level breakpoints
- Breakpoints that are shared among all processes in multiprocess programs
- Assembler-level breakpoints

You can also control whether or not TotalView stops all processes in the control group when a single member reaches a breakpoint.

Topics in this section are:

- “*Setting Source-Level Breakpoints*” on page 340
- “*Setting and Deleting Breakpoints at Locations*” on page 340
- “*Displaying and Controlling Action Points*” on page 343

Setting Source-Level Breakpoints



Typically, you set and clear breakpoints before you start a process. To set a source-level breakpoint, select a boxed line number in the Process Window. (A boxed line number indicates that the line is associated with executable code.) A **STOP** icon lets you know that a breakpoint is set immediately before the source statement.

CLI EQUIVALENT: **@ next to the line number**

You can also set a breakpoint while a process is running by selecting a boxed line number in the Process Window.

CLI EQUIVALENT: **Use dbreak whenever the CLI is displaying a prompt.**

Choosing Source Lines

If you're using C++ templates, TotalView will set a breakpoint in all instantiations of that template if **Plant in share group** is selected. If this isn't what you want, clear the button and then select the **Addresses** button in the Action Point Properties Dialog Box. You can now clear locations where the action point shouldn't be set. (See the top portion Figure 179 on page 341.)

Similarly, in a multiprocess program, you may not want to set the breakpoint in all processes. If this is the case, select the **Process** button. (See the bottom portion on Figure 179 on page 341.)

Setting and Deleting Breakpoints at Locations

You can set or delete a breakpoint at a specific function or source-line number without having to first find the function or source line in the Source Pane. All you need do is enter a line number or function name in the **Action Point > At Location** Dialog Box. (See Figure 180 on page 342.)

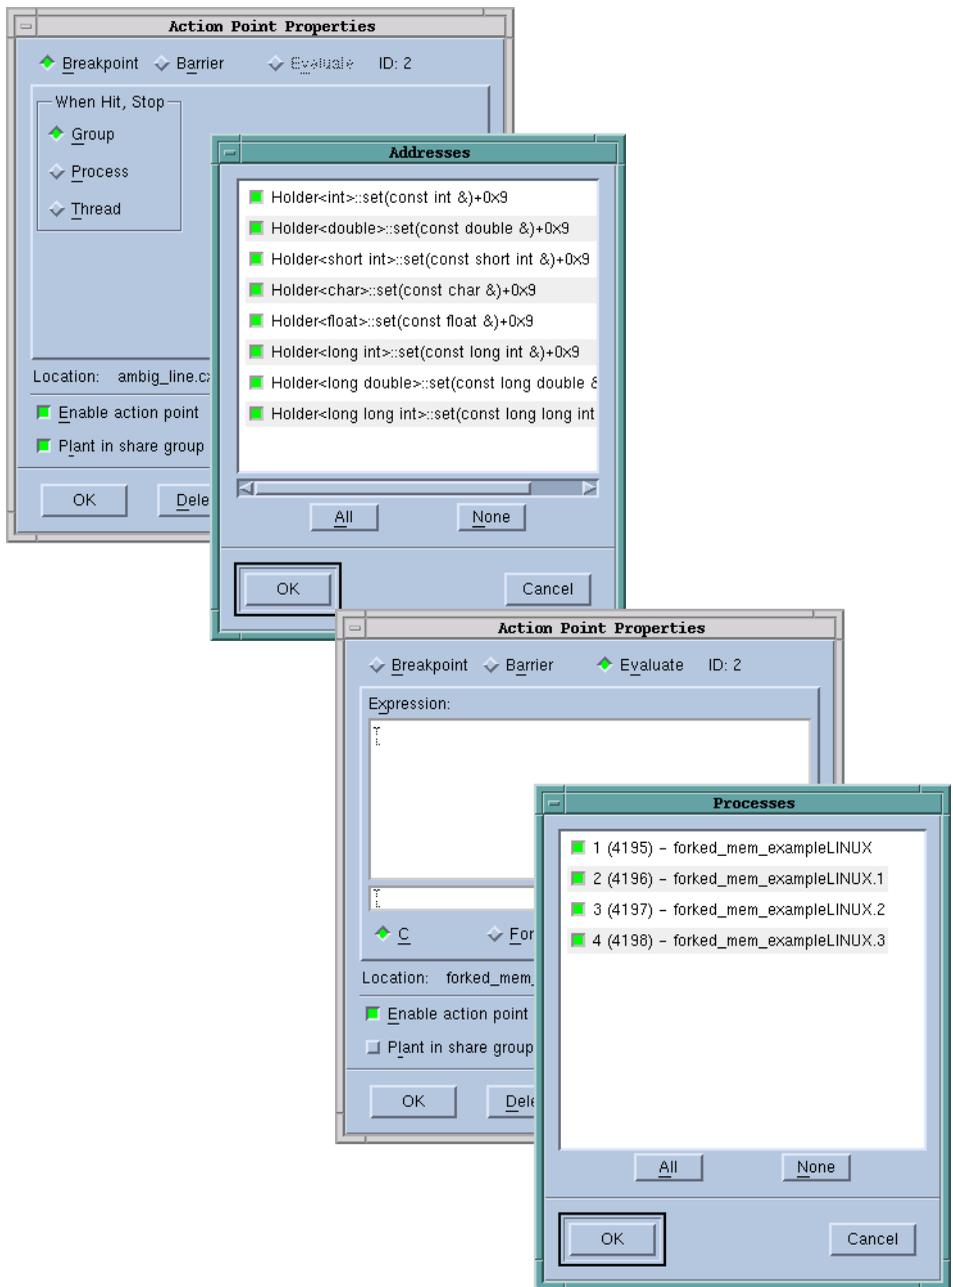
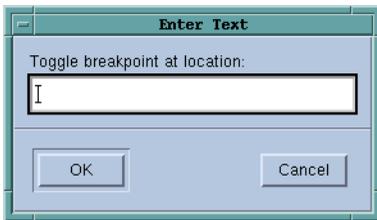


FIGURE 179: Setting Breakpoints on Multiple Similar Addresses and on Processes

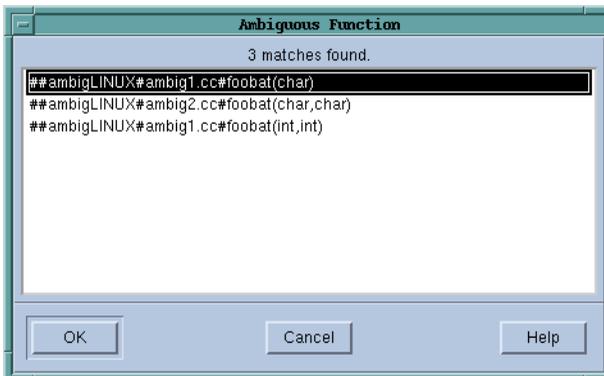
FIGURE 180: **Action Point > At Location Dialog Box**

When you're done, TotalView sets a breakpoint at the location. If you enter a function name, TotalView sets the breakpoint at the function's first executable line. In either case, if a breakpoint already exists at a location, TotalView deletes it.

CLI EQUIVALENT: **dbreak sets a breakpoint.**
ddelete deletes a breakpoint.



If you enter an ambiguous function name using the **Action Point > At Location** command, TotalView displays its **Ambiguous Function Dialog Box**. See Figure 181.

FIGURE 181: **Ambiguous Function Dialog Box**

The procedure for resolving ambiguous function names is similar to the procedure described in "Choosing Source Lines" on page 340.

Displaying and Controlling Action Points

The **Action Point > Properties** Dialog Box lets you set and control an action point. (See Figure 182.) Controls in this dialog box also allows you to set an action point's type to breakpoint, barrier point, or evaluation point. You can also define what will happen to other threads and processes when execution reaches this action point.

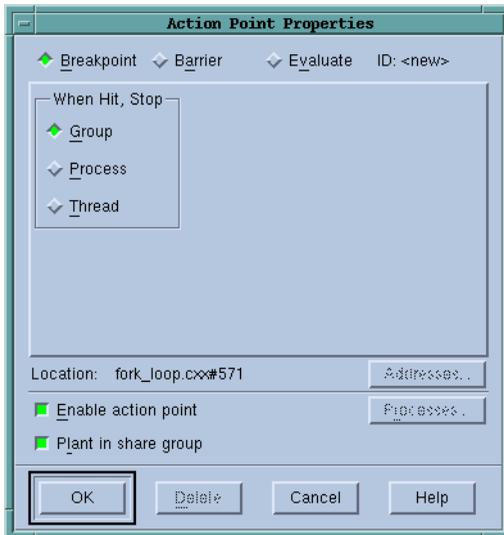


FIGURE 182: **Action Point > Properties Dialog Box**

The following sections explain how you can control action points by using the Process Window and the **Action Point > Properties** Dialog Box.

CLI EQUIVALENT: **dset** SHARE_ACTION_POINT
dset STOP_ALL
ddisable *action-point*

Disabling

TotalView can retain an action point's definition and ignore it while your program is executing. That is, disabling an action point deactivates it without removing it.

CLI EQUIVALENT: **ddisable** *action-point*

You can disable an action point by:

- Clearing **Enable action point** in the **Properties** Dialog Box.
- Selecting the **STOP** or **BARR** symbol in the Action Points Pane.
- Using the context (right-click) menu.
- Clicking on a disable command in the menubar.

Deleting

You can permanently remove an action point by selecting the **STOP** or **BARR** symbol or selecting the **Delete** button in the **Action Point > Properties** Dialog Box.

To delete all breakpoints and barrier points, use the **Action Point > Delete All** command.

CLI EQUIVALENT: **ddelete**

Enabling

You can activate an action point that was previously disabled by selecting a dimmed **STOP**, **BARR**, or **EVAL** symbol in the Source or Action Points Pane, or by selecting **Enable action point** in the **Properties** Dialog Box.

CLI EQUIVALENT: **denable**

Suppressing

You can tell TotalView to ignore action points by using the **Action Point > Suppress All** command.

CLI EQUIVALENT: **ddisable -a**

When you suppress action points, you disable them. If you have suppressed action points, you cannot update existing action points or create new ones.

You can make previously suppressed action points active and allow the creation of new ones by reusing the **Action Point > Suppress All** command.

CLI EQUIVALENT: **denable -a**

Setting Machine-Level Breakpoints

To set a machine-level breakpoint, you must first display assembler code. (Refer to “Viewing the Assembler Version of Your Code” on page 216 for information.) You can now select an instruction’s line number. The line number must be replaced with a dotted box (:::)—this indicates the line is the beginning of a machine instruction. Since instruction sets on some platforms support variable-length instructions, you might see more than one line associated with a single line contained in the dotted box. The **STOP** icon appears, indicating that the breakpoint occurs before the instruction executes.

If you set a breakpoint on the first instruction after a source statement, however, TotalView assumes that you are creating a source-level breakpoint, not an assembler-level one.

116	0x0804b213:	0xeb	jmp	0x804b215
	0x0804b214:	0x00		
117	0x0804b215:	0x89	movl	%ebp, %esp
	0x0804b216:	0xec		
:::	0x0804b217:	0x5d	popl	%ebp
:::	0x0804b218:	0xc3	ret	
:::	0x0804b219:	0x8d	leal	0(%esi), %esi
	0x0804b21a:	0x76		
	0x0804b21b:	0x00		
	MB_fork_notify_breakpoint_here:	0x55	pushl	%ebp
	0x0804b21d:	0x89	movl	%esp, %ebp
	0x0804b21e:	0xe5		
124	0x0804b21f:	0xeb	jmp	0x804b221
	0x0804b220:	0x00		

FIGURE 183: Breakpoint at Assembler Instruction

If you set machine-level breakpoints on one or more instructions generated from a single source line and then display source code in the Source Pane, TotalView displays an **ASM** icon (see Figure 178 on page 339) on the line number. To see the actual breakpoint, you must redisplay assembler instructions.

When a process reaches a breakpoint, TotalView:

- Suspends the process.
- Displays the PC arrow icon () over the stop sign to indicate that the PC is at the breakpoint. (See Figure 184 on page 346.)
- Displays **At Breakpoint** in the Process Window title bar and other windows.
- Updates the Stack Trace and Stack Frame Panes and all Variable Windows.

```

80      do 40 i = 1, 500
81          denorms(i) = x'00000001'
82      40  continue
83  STOP  do 42 i = 500, 1000
84          denorms(i) = x'80000001'
85  → 42  continue
86      local_var=100
87      ieee_array(1) = x'7f800000'          ! infinity
88      ieee_array(2) = x'ff800000'          ! -infinity
89      ieee_array(3) = x'7fc00001'          ! NANQ
90      ieee_array(4) = x'7f800001'          ! NANQ
91      ieee_array(5) = x'00000001'          ! positive denormalized number
92      ieee_array(6) = x'80000001'          ! negative denormalized number

```

FIGURE 184: PC Arrow Over a Stop Icon

Setting Breakpoints for Multiple Processes

In all programs including multiprocess programs, you can set breakpoints in parent and child processes before you start the program and while the program is executing. Do this using the **Action Point > Properties** Dialog Box. (See Figure 185.)

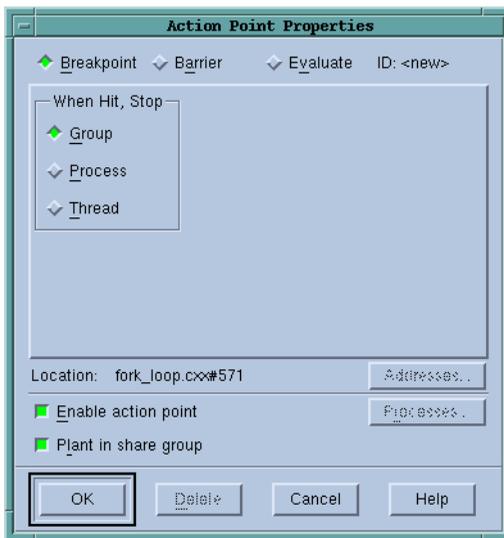


FIGURE 185: Action Point > Properties Dialog Box

This dialog box provides the following controls for setting breakpoints:

■ **When Hit, Stop**

When your thread hits a breakpoint, TotalView can also stop the thread's control group or the process in which it is running.

```
CLI EQUIVALENT:  dset STOP_ALL
                  dbreak -p | -g | -t
```

■ **Plant in share group**

If this is selected, TotalView enables the breakpoint in all members of this thread's share group at the same time. If this isn't selected, you must individually enable and disable breakpoints in each member of the share group.

```
CLI EQUIVALENT:  dset SHARE_ACTION_POINT
```

The **Process** button lets you indicate which process in a multiprocess program will have enabled breakpoints. Note that if **Plant in share group** is selected, this button won't be enabled because you've told TotalView to set the breakpoint in all of the processes.

You can preset many of the properties in this dialog box by using TotalView preferences, as shown in Figure 186 on page 348.

You can find additional information about this dialog box within the online Help.

If you select the **Evaluate** button in the **Action Point > Properties** Dialog Box, you can add an expression to the action point. This expression will be attached to control and share group members. Refer to "Writing Code Fragments" on page 373 for more information.

If you're trying to synchronize your program's threads, you will want to set a barrier point. For more information, see "Barrier Points" on page 350.

Setting Breakpoints When Using fork()/execve()

You must link with the **dbfork** library before debugging programs that call **fork()** and **execve()**. See "Compiling Programs" on page 40.

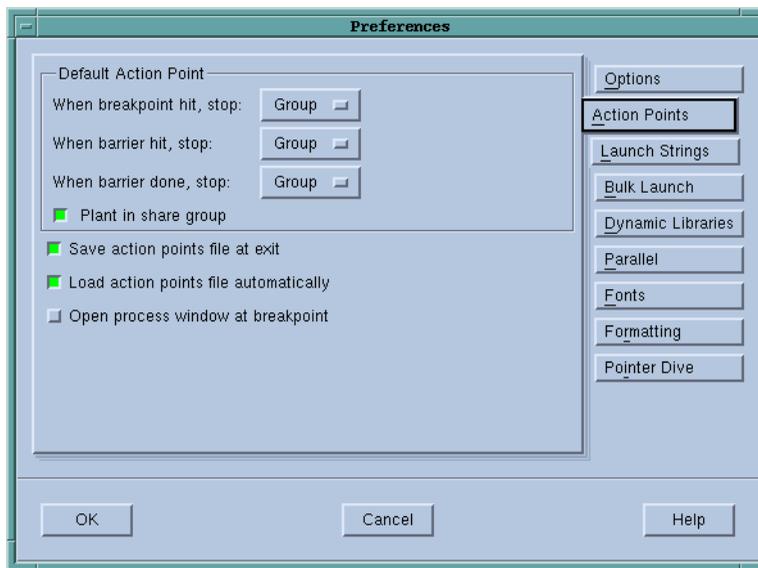


FIGURE 186: **File > Preferences: Action Points Page**

Processes That Call `fork()`

By default, TotalView places breakpoints in all processes in a share group. (For information on share groups, see “*Organizing Chaos*” on page 25.) When any process in the share group reaches a breakpoint, TotalView stops all processes in the control group. This means that TotalView stops the control group containing the share group. This control can, of course, contain more than one share group. To override these defaults:

- 1 Dive into the line number to display the **Action Point > Properties** Dialog Box.
- 2 Clear the **Plant in share group** check box and make sure that the **Group** radio button is selected.

CLI EQUIVALENT: `dset SHARE_ACTION_POINT false`

Processes That Call `execve()`

Shared breakpoints are not set in children having different executables. To set the breakpoints for children that call `execve()`:

- 1 Set the breakpoints and breakpoint options in the parent and the children that do not call `execve()`.
- 2 Start the multiprocess program by displaying the **Group > Go** command. When the first child calls `execve()`, TotalView displays the following message:

```
Process name has exec'd name.
Do you want to stop it now?
```

CLI EQUIVALENT: **G**

- 3 Answer **Yes**. TotalView opens a Process Window for the process. (If you answer **No**, you won't have an opportunity to set breakpoints.)
- 4 Set breakpoints for the process. After you set breakpoints for the first child using this executable, TotalView won't prompt when other children call `execve()`. This means that if you do not want to share breakpoints in children using the same executable, dive into the breakpoints and set the breakpoint options.
- 5 Select the **Group > Go** command.

Example: Multiprocess Breakpoint

The following program excerpt illustrates the places where you can set breakpoints in a multiprocess program:

```

1  pid = fork();
2  if (pid == -1)
3      error ("fork failed");
4  else if (pid == 0)
5      children_play();
6  else
7      parents_work();

```

Here's what happens when you set a breakpoint at different places:

Line Number	Result
1	Stops the parent process before it forks.
2	Stops both the parent and child processes.
3	Stops the parent process if fork() failed.
5	Stops the child process.
7	Stops the parent process.

Barrier Points

A barrier breakpoint is similar to a simple breakpoint, differing in that it holds processes and threads that reach the barrier point. Other processes and threads continue to run. TotalView holds these processes or threads until all processes or threads defined in the barrier point reach this same place. When the last one reaches a barrier point, TotalView releases all the held processes or threads.

CLI EQUIVALENT: **dbarrier**

Topics in this section are:

- "Barrier Breakpoint States" on page 350
- "Setting a Barrier Breakpoint" on page 351
- "Creating a Satisfaction Set" on page 353
- "Hitting a Barrier Point" on page 353
- "Releasing Processes from Barrier Points" on page 353
- "Deleting a Barrier Point" on page 353
- "Changes When Setting and Disabling a Barrier Point" on page 354

Barrier Breakpoint States

Processes and threads at a barrier point are held or stopped, as follows:

Held	A held process or thread cannot execute until all the processes or threads in its group are at the barrier, or until you manually release it. The various <i>go</i> and <i>step</i> commands from the Group , Process , and Thread menus will not start held processes.
Stopped	When all processes in the group reach a barrier point, TotalView automatically releases them. They remain stopped at the barrier point until you tell them to resume executing.

You can manually release held processes and threads with the **Hold** and **Release** commands contained in the **Group**, **Process**, and **Thread** menus. When you manually release a process, the *go* and *step* commands become available again.

CLI EQUIVALENT: **dfocus ... dhold**
dfocus ... dunhold

You can reuse the **Hold** command to again toggle the hold state of the process or thread. See “*Holding and Releasing Processes and Threads*” on page 221 for more information.

When a process or a thread is held, TotalView displays an **H** (for a held process) or an **h** (for a held thread) in the process’s or thread’s entry in the Root Window.

Setting a Barrier Breakpoint

You can set a barrier breakpoint by using the **Action Point > Set Barrier** command or from the **Action Point > Properties** Dialog Box. (See Figure 187.) As an alternative, you can right-click on the line. From the displayed context menu, you can select the **Set Barrier** command.

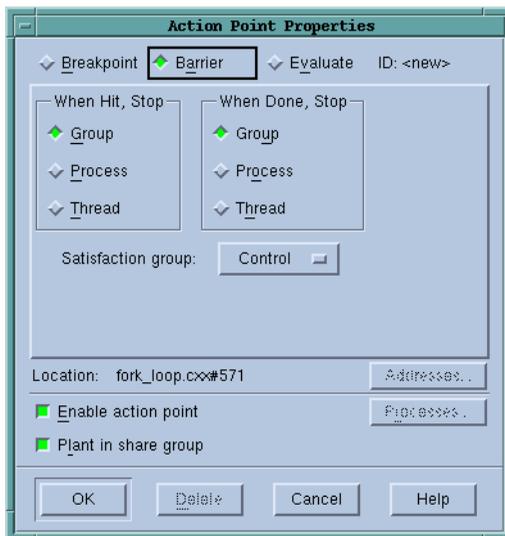


FIGURE 187: **Action Point > Properties** Dialog Box

Barrier points are most often used to synchronize a set of threads. When a thread reaches a barrier, it stops, just as it does for a breakpoint. The difference is that TotalView prevents—that is, holds—each thread reaching the barrier from responding to resume commands (for example, *step*, *next*, or *go*) until all threads in the affected set arrive at the barrier. When all threads reach the barrier, TotalView considers the barrier to be *satisfied* and releases all of the threads being held there. *They are just released; they are not continued.* That is, they are left stopped at the barrier. If you now continue the process, those threads will also run.

If a process is stopped and then continued, the held threads, including the ones waiting at an unsatisfied barrier, do not run. Only unheld threads run.

The **When Hit, Stop** radio buttons indicate what other threads TotalView will stop when execution reaches the breakpoint, as follows:

Scope	TotalView will:
Group	Stop all threads in the current thread's control group.
Process	Stop all threads in the current thread's process.
Thread	Only stop this thread.

CLI EQUIVALENT: **`dbarrier -stop_when_hit`**

After all processes or threads reach the barrier, TotalView releases all held threads. *Released* means that these threads and processes can now run.

The **When Done, Stop** radio buttons tell TotalView what else it should stop, as follows:

Scope	TotalView will:
Group	Stop all threads in the current thread's control group.
Process	Stop all threads in the current thread's process.
Thread	Only stop this thread.

CLI EQUIVALENT: **`dbarrier -stop_when_done`**

Creating a Satisfaction Set

For even more control over what TotalView will stop, you can select a *satisfaction set*. This set tells TotalView which threads must be held before it can release the group of threads. That is, the barrier is *satisfied* when TotalView has held all of the indicated threads. Use the **Satisfaction group** items to tell TotalView that the satisfaction set consists of all threads in the current thread's **Control**, **Workers**, or **Lockstep** group.

When you set a barrier point, TotalView places it in every process in the share group.

Hitting a Barrier Point

If you run one of the processes or threads in a group and it hits a barrier point, you will see an **H** next to the process or thread name in the Root Window and the word **[Held]** in the title bar in the main Process Window. Barrier points are always shared.

CLI EQUIVALENT: `dstatus`

If you create a barrier and all the process's threads are already at that location, TotalView won't hold any of them. However, if you create a barrier and all of the processes and threads are not at that location, TotalView will hold any thread that is already there.

Releasing Processes from Barrier Points

TotalView automatically releases processes and threads from a barrier point when they hit that barrier point and all other processes or threads in the group are already held at it.

Deleting a Barrier Point

You can delete a barrier point in two ways:

- Using the **Action Point > Properties** Dialog Box.
- Clicking on the **BARR** icon in the line number area.

CLI EQUIVALENT: `ddelete`

Changes When Setting and Disabling a Barrier Point

Setting a barrier point at the current PC for a *stopped* process or thread holds the process there. If, however, all other processes or threads affected by the barrier point are at the same PC, TotalView doesn't hold them. Instead, TotalView treats the barrier point as if it was an ordinary breakpoint.

TotalView releases all processes and threads that are held and which have threads at the barrier point when you disable the barrier point. You can disable the barrier point in the **Action Point > Properties** Dialog Box by clicking on **Enable action point** at the bottom of the dialog box.

CLI EQUIVALENT: **ddisable**

Defining Evaluation Points and Conditional Breakpoints

TotalView lets you define *evaluation points*. These are action points at which you have added a code fragment that TotalView will execute. You can write the code fragment in C, Fortran, or assembler.

NOTE Assembler support is currently available on the HP Alpha Tru64 UNIX, IBM AIX, and SGI IRIX operating systems. While any user can enable or disable TotalView's ability to compile evaluation points, they must be enabled if you are entering assembler code.

Topics in this section are:

- "Setting Evaluation Points" on page 356
- "Creating Conditional Breakpoint Examples" on page 356
- "Patching Programs" on page 357
- "Interpreted vs. Compiled Expressions" on page 358
- "Allocating Patch Space for Compiled Expressions" on page 360

By using evaluation points, you can:

- Include instructions that stop a process and its relatives. If the code fragment can make a decision whether it should stop execution, it is called a *conditional breakpoint*.
- Test potential fixes for your program.
- Set the values of your program's variables.

- Automatically send data to the Visualizer. This can produce animated displays of the changes in your program's data.

You can set an evaluation point at any source line that generates executable code (marked with a boxed line number surrounding a line number) or a line containing assembler-level instructions. This means that if you can set a breakpoint, you can set an evaluation point.

At each evaluation point, TotalView or your program executes the code contained in the evaluation point before your program executes the code on that line. While your program can then go on to execute this source line or instruction, it can:

- Include a branching instruction (such as **goto** in C or Fortran). The instruction can transfer control to a different point in the target program, enabling you to test program patches.
- Execute a TotalView function. TotalView's functions let you stop execution, create barriers, and count down breakpoints. For more information on these statements, refer to Table 20 "Built-In Statements Used in Expressions" on page 375.

TotalView evaluates code fragments in the context of the target program. This means that you can refer to program variables and branch to places in your program.

For complete information on what you can include in code fragments, refer to "Writing Code Fragments" on page 373.

Evaluation points only modify the processes being debugged—they do not modify your source program or create a permanent patch in the executable. If you save a program's evaluation points, however, TotalView reapplies them whenever you start a debugging session for that program. To save your evaluation points, refer to "Saving Action Points to a File" on page 370.

NOTE You should stop a process before setting an evaluation point in it. This ensures that the evaluation point is set in a stable context.

Setting Evaluation Points

To set an evaluation point:

- 1 Display the **Action Point > Properties** Dialog Box. You can do this, for example, by right-clicking on a **STOP** icon and selecting **Properties** or by selecting a line and then invoking the command from the menu bar.
- 2 Select the **Evaluate** button.
- 3 Select the button (if it isn't already selected) for the language in which you will code the fragment.
- 4 Type the code fragment. For information on supported C, Fortran, and assembler language constructs, refer to "Writing Code Fragments" on page 373.
- 5 For multiprocess programs, decide whether to share the evaluation point among all processes in the program's share group. By default, TotalView selects the **Plant in share group** check box for multiprocess programs, but you can override this by clearing it.
- 6 Select the **OK** button to confirm your changes. If the code fragment has an error, TotalView displays an error message. Otherwise, it processes the code, closes the dialog box, and places an **EVAL** icon.

CLI EQUIVALENT: **dbreak -e**
 dbarrier -e

Creating Conditional Breakpoint Examples

Here are some examples:

- To define a breakpoint that is reached whenever the **counter** variable is greater than 20 but less than 25:

```
if (counter > 20 && counter < 25)  
  $stop;
```
- To define a breakpoint that will stop execution every tenth time that TotalView executes the **\$count** function

```
$count 10
```
- To define a breakpoint with a more complex expression, consider:

```
$count my_var * 2
```

When the `my_var` variable equals 4, the process stops the eighth time it executes the `$count` function. After the process stops, TotalView reevaluates the expression. If `my_var` now equals 5, the process will stop again after the process executes the `$count` function ten more times.

For complete descriptions of the `$stop` and `$count` statements, refer to “Built-In Statements” on page 375.

Patching Programs

You can use expressions in evaluation points to patch your code if you use the `goto` (C) and `GOTO` (Fortran) statements to jump to a different program location. This lets you:

- Branch around code that you don’t want your program to execute.
- Add new pieces of code.

In many cases, correcting an error means that you will do both operations: you patch out incorrect lines and patch in corrections.

Conditionally Patching Out Code

The following example contains a logic error where the program dereferences a null pointer:

```
1  int check_for_error (int *error_ptr)
2  {
3      *error_ptr = global_error;
4      global_error = 0;
5      return (global_error != 0);
6  }
```

The error occurs because the routine calling this function assumes that the value of `error_ptr` can be 0. The `check_for_error()` function, however, assumes that `error_ptr` isn’t null, which means that line 3 can dereference a null pointer.

You can correct this error by setting an evaluation point on line 3 and entering:

```
if (error_ptr == 0) goto 4;
```

If the value of `error_ptr` is null, line 3 isn’t executed. Notice that you are not naming a label used in your program. Instead, you are naming one of the TotalView-generated line numbers.

Patching in a Function Call

Instead of routing around the problem, you could patch in a `printf()` statement that displays the value of the `global_error` variable created in the preceding program.

You would set an evaluation point on line 4 and enter:

```
printf ("global_error is %d\n", global_error);
```

TotalView executes this code fragment before the code on line 4; that is, it is executed before `global_error` is set to 0.

Correcting Code

The next example contains a coding error: the function returns the maximum value instead of the minimum value:

```
1 int minimum (int a, int b)
2 {
3     int result; /* Return the minimum */
4     if (a < b)
5         result = b;
6     else
7         result = a;
8     return (result);
9 }
```

In his example, you would correct this error by adding the following code to an evaluation point at line 4:

```
if (a < b) goto 7; else goto 5;
```

This effectively replaces the `if` statement on line 4 with the code in the evaluation point.

Interpreted vs. Compiled Expressions

On most platforms, TotalView executes interpreted expressions. TotalView can also execute compiled expressions on the HP Alpha Tru64 UNIX, IBM AIX, and SGI IRIX platforms. On HP Alpha Tru64 UNIX and IBM AIX platforms, compiled expressions are enabled by default.

You can use the `TV::compile_expressions` CLI variable to enable or disable compiled expressions. See "Operating Systems" in the TOTALVIEW REFERENCE GUIDE to find out how TotalView handles expressions on specific platforms.

NOTE Using any of the following functions forces TotalView to interpret the evaluation point instead of compiling it: `$clid`, `$duid`, `$nid`, `$processduid`, `$systid`, `$tid`, and `$visualize`. In addition, `$pid` forces interpretation on AIX.

Interpreted Expressions

Interpreted expressions are interpreted by TotalView. As is always the case, interpreted expressions run slower (and possibly much slower) than compiled expressions. With multiprocess programs, interpreted expressions run even more slowly because processes may need to wait for TotalView to execute the expression.

When you're debugging remote programs, interpreted expressions always run slower because the TotalView process on the host, not the TotalView debugger server (`tvdsvr`) on the client, interprets the expression. For example, an interpreted expression could require that 100 remote processes wait for the TotalView debugger process on the host machine to evaluate one interpreted expression. In contrast, if TotalView compiles the expression, it evaluates them on each remote process.

NOTE Whenever a thread hits an interpreted patch point, TotalView stops execution. This means that TotalView will create a new set of lockstep groups. Consequently, if goal threads contain interpreted patches, the results are unpredictable.

Compiled Expressions

TotalView compiles, links, and patches expressions into the target process. Because the target thread executes this code, evaluation points and conditional breakpoints execute very quickly. (Note that conditional watchpoints are always interpreted.) And, more importantly, this code doesn't need to communicate with the TotalView host process until it needs to.

If the expression executes a `$stop` function, TotalView stops executing the compiled expression. At this time, you can single-step through it and continue executing the expression as you would the rest of your code. See Figure 188 on page 360.

If you will be using many compiled expressions or your expressions are long, you may need to think about allocating patch space. For more information, see "Allocating Patch Space for Compiled Expressions" on page 360.

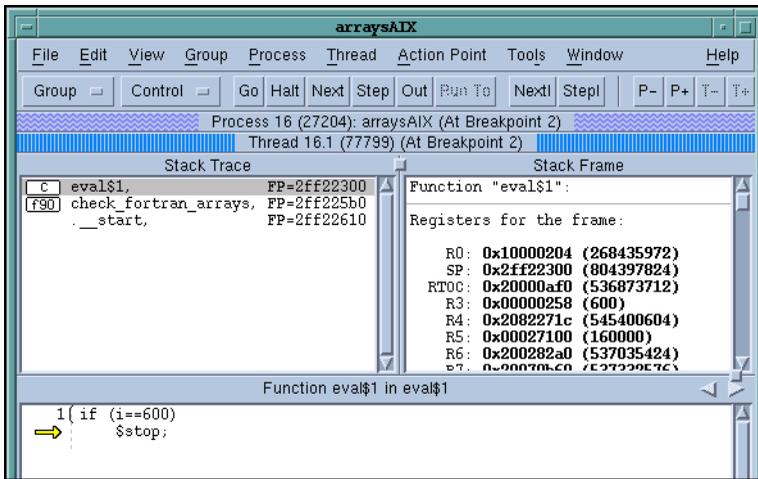


FIGURE 188: Stopped Execution of Compiled Expressions

Allocating Patch Space for Compiled Expressions

TotalView must either allocate or find space in your program to hold the code it generates for compiled expressions. Since this patch space is part of your program's address space, the location, size, and allocation scheme that TotalView uses may conflict with your program. As a result, you may need to change how TotalView allocates this space.

You can choose one of the following patch space allocation schemes:

- **Dynamic patch space allocation:** Tells TotalView to dynamically find the space for your expression's code.
- **Static patch space allocation:** Tells TotalView to use a statically allocated area of memory.

Dynamic Patch Space Allocation

Dynamic patch space allocation means that TotalView dynamically allocates patch space for code fragments. If you do not specify the size and location for this space, TotalView allocates 1 MB. TotalView creates this space using system calls.

TotalView allocates memory for read, write, and execute access in the following addresses:

Table 17: Dynamic Patch Space Allocation Default Addresses

Platform	Address range
HP Alpha Tru64 UNIX	0xFFFFF00000 - 0xFFFFFFFFFFFF
IBM AIX	0xCFF00000 - 0xCFFFFFFF
SGI IRIX (-n32)	0x4FF00000 - 0x4FFFFFFF
SGI IRIX (-64)	0x8FF00000 - 0x8FFFFFFF

NOTE You can only allocate dynamic patch space for these machines.

If the default address range conflicts with your program, or you would like to change the size of the dynamically allocated patch space, you can change:

- Patch space base address by using the `-patch_area_base` command-line option.
- Patch space length by using the `-patch_area_length` command-line option.

Static Patch Space Allocation

TotalView can statically allocate patch space if you add a specially named array to your program. When TotalView needs to use patch space, it uses this space created for this array.

You can include, for example, a 1 MB statically allocated patch space in your program by adding the `TVDB_patch_base_address` data object in a C module. Because this object must be 8-byte aligned, declare it as an array of doubles. For example:

```

/* 1 megabyte == size TV expects */
#define PATCH_LEN 0x100000
double TVDB_patch_base_address [PATCH_LEN /
sizeof(double)]

```

If you need to use a static patch space size that differs from the 1 MB default, you must use assembler language. Table 18 shows sample assembler code for three platforms that support compiled patch points.

Table 18: Static Patch Space Assembler Code

Platform	Assembler Code
HP Alpha Tru64 UNIX	<pre>.data .align 3 .globl TVDB_patch_base_address .globl TVDB_patch_end_address TVDB_patch_base_address: .byte 0x00 : PATCH_SIZE TVDB_patch_end_address:</pre>
IBM AIX	<pre>.csect .data{RW}, 3 .globl TVDB_patch_base_address .globl TVDB_patch_end_address TVDB_patch_base_address: .space PATCH_SIZE TVDB_patch_end_address:</pre>
SGI IRIX	<pre>.data .align 3 .globl TVDB_patch_base_address .globl TVDB_patch_end_address TVDB_patch_base_address: .space PATCH_SIZE TVDB_patch_end_address:</pre>

Here is how you would use the static patch space assembler code:

- 1 Use an ASCII editor and place the assembler code into a file named **tvdb_patch_space.s**.
- 2 Replace the **PATCH_SIZE** tag with the decimal number of bytes you want. This value must be a multiple of 8.
- 3 Assemble the file into an object file by using a command such as:


```
cc -c tvdb_patch_space.s
```

 On SGI IRIX, use **-n32** or **-64** to create the correct object file type.
- 4 Link the resulting **tvdb_patch_space.o** into your program.

Using Watchpoints

TotalView lets you monitor the changes that occur to memory locations by creating a special kind of action point called a *watchpoint*. Watchpoints are most often used to find a statement in your program that is writing to places where it shouldn't be writing. This can occur, for example, when processes share memory and more than one process writes to the same location. It can also occur when your program writes off the end of an array or when your program has a dangling pointer.

Topics in this section are:

- "Architectures" on page 363
- "Creating Watchpoints" on page 365
- "Watching Memory" on page 366
- "Triggering Watchpoints" on page 367
- "Using Conditional Watchpoints" on page 368

TotalView watchpoints are called *modify watchpoints* because TotalView only *triggers* a watchpoint when your program modifies a memory location. If a program writes a value into a location that is the same as what is already stored, TotalView doesn't trigger the watchpoint because the location's value did not change.

For example, if location 0x10000 has a value of 0 and your program writes a 0 into this location, TotalView doesn't trigger the watchpoint even though your program wrote data into the memory location. See "Triggering Watchpoints" on page 367 for more details on when watchpoints trigger.

You can also create *conditional watchpoints*. A conditional watchpoint is similar to a conditional breakpoint in that TotalView will evaluate the expression when the watchpoint triggers. You can use conditional watchpoints for a number of purposes. For example, you can use one to test if a value changes its sign—that is, it becomes positive or negative—or if a value moves above or below some threshold value.

Architectures

The number of watchpoints, their size, and alignment restrictions differ from platform to platform. This is because TotalView relies on the operating system and its hardware to implement watchpoints.

NOTE Watchpoints are not available on Alpha Linux and HP.

The following list describes constraints that exist on each platform:

HP Alpha Tru64 Tru64 places no limitations on the number of watchpoints that you can create, and there are no alignment or size constraints. However, watchpoints can't overlap, and you can't create a watchpoint on an already write-protected page.

Watchpoints use a page protection scheme. Because the page size is 8,192 bytes, watchpoints can degrade performance if your program frequently writes to pages containing watchpoints.

IBM AIX You can create one watchpoint on AIX 4.3.3.0-2 (AIX 4.3R) or later systems running 64-bit chips. These are Power3 and Power4 systems. (AIX 4.3R is available as APAR IY06844.) A watchpoint cannot be longer than 8 bytes, and you must align it within an 8-byte boundary.

IRIX6 MIPS Watchpoints are implemented on IRIX 6.2 and later operating systems. These systems allow you to create about 100 watchpoints. There are no alignment or size constraints. However, watchpoints can't overlap.

Linux x86 You can create up to four watchpoints and each must be 1, 2, or 4 bytes in length, and a memory address must be aligned for the byte length. That is, you must align a 4-byte watchpoint on a 4-byte address boundary, and a 2-byte watchpoint must be aligned on a 2-byte boundary, and so on.

Solaris SPARC TotalView supports watchpoints on Solaris 2.6 or later operating systems. These operating system allow you to create hundreds of watchpoints, and there are no alignment or size constraints. However, watchpoints can't overlap.

Typically, a debugging session doesn't use many watchpoints. In most cases, you are only monitoring one memory location at a time. So, restrictions on the number of values you can watch are seldom an issue.

Creating Watchpoints

Watchpoints are created by using the **Tools > Watchpoint** Dialog Box, which is only invocable from a Variable Window (If your platform doesn't support watchpoints, this menu item is dimmed.) See Figure 189.

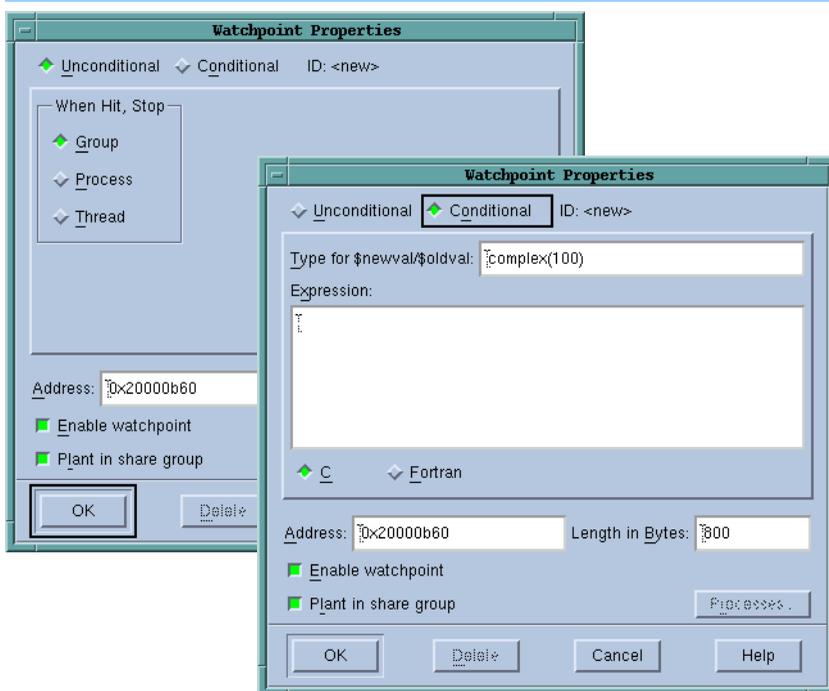


FIGURE 189: **Tools > Watchpoint** Dialog Boxes

Controls within this dialog box let you create unconditional and conditional watchpoints. You will be setting a watchpoint for all of the information displayed in the Variable Window, not just for an element of it. If, for example, the Variable Window is displaying an array, you will be setting a watchpoint for the entire array (or as much of it as TotalView can watch.) If you only want to watch one element, dive on the element and then set the watchpoint. Similarly, if the Variable Window is displaying a structure and you only want to watch one element, dive on the element before setting the watchpoint.

The online Help contains information on the fields in this dialog box.

Displaying Watchpoints

The watchpoint entry, indicated by UDWP (Unconditional Data Watchpoint) and CDWP (Conditional Data Watchpoint), displays the action point ID, the amount of memory being watched, and the location being watched.

If you dive into a watchpoint, TotalView displays the **Watchpoint Properties** Dialog Box.

If you select a watchpoint, TotalView will toggle the enabled/disabled state of the watchpoint.

Watching Memory

A watchpoint tracks a memory location—it does not track a variable. This means that a watchpoint might not perform as you would expect it to when watching stack or automatic variables. For example, assume that you want to watch a variable in a subroutine. When control exits from the subroutine, the memory allocated on the stack for this subroutine is deallocated. At this time, TotalView is watching unallocated stack memory. When the stack memory is reallocated to a new stack frame, TotalView is still watching this same position. This means that TotalView triggers the watchpoint when something changes this newly allocated memory.

Also, if your program reinvokes a subroutine, it usually executes in a different stack location. So, TotalView will not be able to monitor changes to the variable because it is at a different memory location.

All of this means that in most circumstances, you can't place a watchpoint on a stack variable. If you need to watch a stack variable, you will need to create and delete the watchpoint each time your program invokes the subroutine.

NOTE In some circumstances, a subroutine will always be called from the same location. This means that its local variables will probably be in the same location, so trying can't hurt.

This doesn't mean you can't place a watchpoint on a stack or heap variable. It just means that what happens is undefined after this memory is released. For example, after you enter a routine, you can be assured that memory locations are always tracked accurately until the memory is released.

If you place a watchpoint on a global or static variable that is always accessed by reference (that is, the value of a variable is always accessed using a pointer to the variable), you can set a watchpoint on it because the memory locations used by the variable are not changing.

Triggering Watchpoints

When a watchpoint triggers, the thread's program counter (PC) points to the instruction *following* the instruction that caused the watchpoint to trigger. If the memory store instruction is the last instruction in a source statement, the PC will be pointing to the source line *following* the statement that triggered the watchpoint. (Breakpoints and watchpoints work differently. A breakpoint stops *before* an instruction executes. In contrast, a watchpoint stops *after* an instruction executes.)

Using Multiple Watchpoints

If a program modifies more than one byte with one program instruction or statement (which is normally the case when storing a word), TotalView triggers the watchpoint with the lowest memory location in the modified region. Although the program may be modifying locations monitored by other watchpoints, only the watchpoint for the lowest memory location is triggered. This can occur when your watchpoints are monitoring adjacent memory locations and a single store instruction modifies these locations.

For example, assume that you have two 1-byte watchpoints, one on location 0x10000 and the other on location 0x10001. Also assume that your program uses a single instruction to store a 2-byte value at locations 0x10000 and 0x10001. If the 2-byte storage operation modifies both bytes, the watchpoint for location 0x10000 triggers. The watchpoint for location 0x10001 does not and will not trigger at this time.

Here's a second example. Assume that you have a 4-byte integer that uses storage locations 0x10000 through 0x10003 and you set a watchpoint on this integer. If a process modifies location 0x10002, TotalView triggers the watchpoint. Now assume that you're watching two adjacent 4-byte integers that are stored in locations 0x10000 through 0x10007. If a process writes to locations 0x10003 and 0x10004

(that is, one byte in each), TotalView triggers the watchpoint associated with location 0x10003. The watchpoint associated with location 0x10004 does not trigger.

Data Copies

TotalView keeps an internal copy of data in the watched memory locations for each process sharing the watchpoint. If you create watchpoints that cover a large area of memory or if your program has a large number of processes, you will increase TotalView's virtual memory requirements. Furthermore, TotalView refetches data for each memory location whenever it continues the process or thread. This can affect TotalView's performance.

Using Conditional Watchpoints

If you associate an expression with a watchpoint (by selecting the **CDWP** icon in the **Tools > Watchpoint** Dialog Box and entering an expression), TotalView will evaluate the expression after the watchpoint triggers. The programming statements that you can use are identical to those used when creating an evaluation point, except that you can't call functions from a watchpoint expression.

The variables used in watchpoint expressions must be global. This is because the watchpoint can be triggered from any procedure or scope in your program.

Because memory locations are not scoped, the variable used in your expression must be globally accessible.

NOTE Fortran does not have global variables. Consequently, you can't directly refer to your program's variables.

TotalView has two function variables that are specifically designed to be used with conditional watchpoint expressions:

\$oldval The value of the memory locations before a change is made.

\$newval The value of the memory locations after a change is made.

Here is an expression that uses these values:

```
if (iValue != 42 && iValue != 44) {  
    iNewValue = $newval; iOldValue = $oldval; $stop;}
```

When the value of the **iValue** global variable is neither 42 nor 44, TotalView will store the new and old memory values in the **iNewValue** and **iOldValue** variables. These variables are defined in the program. (Storing the old and new values is a convenient way of letting you monitor the changes made by your program.)

Here is a condition that triggers a watchpoint when a memory location's value becomes negative:

```
if ($oldval >= 0 && $newval < 0) $stop
```

And here's a condition that triggers a watchpoint when the sign of the value in the memory location changes:

```
if ($newval * $oldval <= 0) $stop
```

Both of these examples require that you set the **Type for \$oldval/\$newval** field in the **Watchpoint Properties** Dialog Box.

For more information on writing expressions, see "*Writing Code Fragments*" on page 373.

If a watchpoint has the same length as the **\$oldval** or **\$newval** data type, the value of these variables is apparent. However, if the data type is shorter than the length of the watch region, TotalView searches for the first changed location in the watched region and uses that location for the **\$oldval** and **\$newval** variables. (It aligns data within the watched region based on the size of the data's type. For example, if the data type is a 4-byte integer and byte 7 in the watched region changes, TotalView uses bytes 4 through 7 of the watchpoint when it assigns values to these variables.)

For example, suppose you're watching an array of 1000 integers called **must_be_positive** and you want to trigger a watchpoint as soon as one element becomes negative. You would declare the type for **\$oldval** and **\$newval** to be **int** and use the following condition:

```
if ($newval < 0) $stop;
```

When your program writes a new value to the array, TotalView triggers the watchpoint, sets the values of **\$oldval** and **\$newval**, and evaluates the expression. When **\$newval** is negative, the **\$stop** statement halts the process.

This can be a very powerful technique for range checking all the values your program writes into an array. (Because of byte length restrictions, you can only use this technique on IRIX and Solaris.)

NOTE TotalView always interprets conditional watchpoints; it never compiles them. And, because interpreted watchpoints are single threaded in TotalView, every process or thread that writes to the watched location must wait for other instances of the watchpoint to finish executing. This can adversely affect performance.

Saving Action Points to a File

You can save a program's action points into a file. TotalView will then use this information to reset these points when you restart the program. When you save action points, TotalView creates a file named *program.TVD.breakpoints*, where *program* is the name of your program.

NOTE Watchpoints are not saved.

Use the **Action Point > Save All** command to save your action points to a file. TotalView places the action points file in the same directory as your program.

CLI EQUIVALENT: **dactions** *–save filename*

If you're using a preference to automatically save breakpoints, TotalView will automatically save action points to a file. Alternatively, starting TotalView with the **–sb** option (see "TotalView Command Syntax" in the TOTALVIEW REFERENCE GUIDE) also tells TotalView to save your breakpoints.

At any time, you can restore saved action points if you use the **Action Points > Load All** command.

CLI EQUIVALENT: **dactions** *–load filename*

Automatic saving and loading is controlled by preferences (see **File > Preferences** in the online Help for more information).

CLI EQUIVALENT: **dset TV::auto_save_breakpoints**

Evaluating Expressions

TotalView lets you open a window for evaluating expressions in the context of a particular process and evaluate expressions in C, Fortran, or assembler.

NOTE Not all platforms let you use assembler constructs; see “*Architectures*” in the TotalView Reference Guide for details.

You can use the **Tools > Evaluate** Dialog Box in many different ways, but here are two examples:

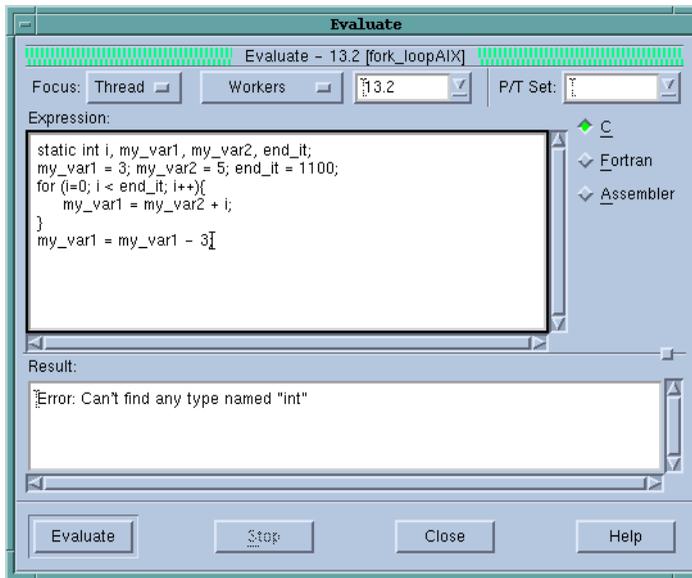
- Expressions can contain loops, so you could use a **for** loop to search an array of structures for an element set to a certain value. In this case, you use the loop index at which the value is found as the last expression in the expression field.
- Because you can call subroutines, you can test and debug a single routine in your program without building a test program to call it.

To evaluate an expression:

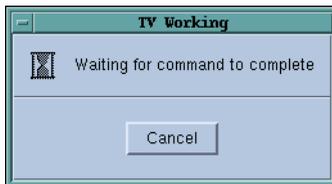
- 1** Tell TotalView to display the **Evaluate** Dialog Box by selecting the **Tools > Evaluate** command. An **Evaluate** Dialog Box appears. If your program hasn't yet been created, you won't be able to use any of the program's variables or call any of its functions.
- 2** Select a button for the programming language you're writing the expression in (if it isn't already selected).
- 3** Move to the **Expression** field and enter a code fragment. For a description of the supported language constructs, see “*Writing Code Fragments*” on page 373.

TotalView returns the value of the last statement in the code fragment. This means that you don't have to assign the expression's return value to a variable. Figure 190 on page 372 shows a sample expression. The last statement in this example assigns the value of **my_var1-3** back to **my_var1**. Because this is the last statement in the code fragment, the value placed in the **Result** field would be the same if you had just typed **my_var1-3**.

- 4** Select the **Evaluate** button. If TotalView finds an error, it places the cursor on the incorrect line and displays an error message. Otherwise, it interprets (or on some platforms, compiles and executes) the code, and displays the value of the last expression in the **Result** field.

FIGURE 190: **Tools > Evaluate Dialog Box**

While the code is being executed, you can't modify anything in the dialog box. TotalView may also display a message box that tells you that it is waiting for the command to complete. (See Figure 191.)

FIGURE 191: **Waiting to Complete Message Box**

If you select **Cancel**, TotalView stops execution.

Since TotalView evaluates code fragments in the context of the target process, it evaluates stack variables according to the currently selected stack frame. If the fragment reaches a breakpoint (or stops for any other reason), TotalView stops evaluating your expression. Assignment statements in an expression can affect the target process because they can change a variable's value.

The controls at the top of the dialog box let you refine the scope at which TotalView evaluates the information you enter. For example, you could evaluate a function in more than one process. Figure 192 shows TotalView displaying the value of a variable in multiple processes and then sending the value as it exists in each process to a function that runs on each of these processes.

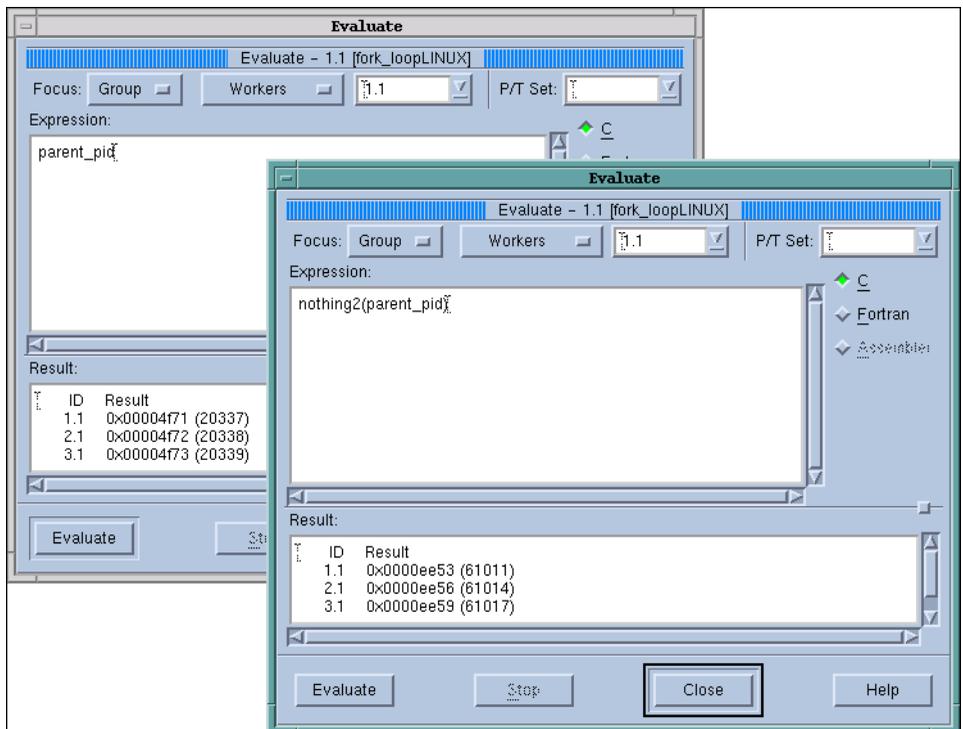


FIGURE 192: Evaluating Information in Multiple Processes

See Chapter 11, "Using Groups, Processes, and Threads" on page 239 for information on using the P/T set controls at the top of this window.

Writing Code Fragments

You can use code fragments in evaluation points and in the **Tools > Evaluate** Dialog Box. This section describes the function variables, built-in statements, and language constructs supported by TotalView.

NOTE While the CLI does not have an “evaluate” command, the information in the following sections does apply to the expression argument of the `dbreak`, `dbarrier`, and `dwatch` commands.

Topics in this section are:

- “*TotalView Variables*” on page 374
- “*Built-In Statements*” on page 375
- “*C Constructs Supported*” on page 377
- “*Fortran Constructs Supported*” on page 378
- “*Writing Assembler Code*” on page 380

TotalView Variables

The TotalView expression system supports built-in variables that allow you to access special thread and process values. All variables are 32-bit integers, which is an `int` or a `long` on most platforms. Table 19 lists TotalView’s built-in variables and their meanings.

Table 19: TotalView Built-in Variables

Name	Returns
<code>\$clid</code>	The cluster ID. (Interpreted expressions only.)
<code>\$duid</code>	The TotalView-assigned Debugger Unique ID (DUID). (Interpreted expressions only.)
<code>\$newval</code>	The value just assigned to a watched memory location. (Watchpoints only.)
<code>\$nid</code>	The node ID. (Interpreted expressions only.)
<code>\$oldval</code>	The value that existed in a watched memory location before a new value modified it. (Watchpoints only.)
<code>\$pid</code>	The process ID.
<code>\$processduid</code>	The DUID of the process. (Interpreted expressions only.)
<code>\$systemid</code>	The system-assigned thread ID. When this is referenced from a process, TotalView throws an error.
<code>\$tid</code>	The TotalView-assigned thread ID. When this is referenced from a process, TotalView throws an error.

TotalView’s built-in variables allow you to create thread-specific breakpoints from the expression system. For example, the `$tid` variable and the `$stop` built-in function let you create a thread-specific breakpoint as follows:

```
if ($tid == 3)
    $stop;
```

This tells TotalView to stop the process only when the third thread evaluates the expression.

You can also create complex expressions by using these variables. For example:

```
if ($pid != 34 && $tid > 7)
    printf ("Hello from %d.%d\n", $pid, $tid);
```

NOTE Using any of the following variables means that the evaluation point is interpreted instead of compiled: \$clid, \$duid, \$nid, \$processduid, \$systid, \$tid, and \$visualize. In addition, \$pid forces interpretation on AIX.

You can't assign a value to a built-in variable or obtain its address.

Built-In Statements

TotalView provides a set of built-in statements that you can use when writing code fragments. The statements are available in all languages, and are shown in the following table.

Table 20: Built-In Statements Used in Expressions

Statement	Use
\$count <i>expression</i>	Sets a process-level countdown breakpoint. When any thread in a process executes this statement for the number of times specified by <i>expression</i> , the process stops. The other processes in the control group continue to execute.
\$countprocess <i>expression</i>	
\$countall <i>expression</i>	Sets a program-group-level countdown breakpoint. All processes in the control group stop when any process in the group executes this statement for the number of times specified by <i>expression</i> .
\$countthread <i>expression</i>	Sets a thread-level countdown breakpoint. When any thread in a process executes this statement for the number of times specified by <i>expression</i> , the thread stops. Other threads in the process continue to execute. If the target system cannot stop an individual thread, this statement performs identically to \$countprocess .

Table 20: Built-In Statements Used in Expressions (cont.)

Statement	Use
	<p>A thread evaluates <i>expression</i> when it executes \$count for the first time. This expression must evaluate to a positive integer. When TotalView first encounters this variable, it determines a value for expression. TotalView will not reevaluate until the expression actually stops the thread. This means that TotalView ignores changes in the value of <i>expression</i> until it hits the breakpoint. After the breakpoint occurs, TotalView reevaluates the expression and sets a new value for this statement.</p> <p>The internal counter is stored in the process and shared by all threads in that process.</p>
\$hold \$holdprocess	<p>Holds the current process. If all other processes in the group are already held at this Eval point, then TotalView will release all of them. If other processes in the group are running, they continue to run.</p>
\$holdstopall \$holdprocessstopall	<p>Exactly like \$hold, except any processes in the group which are running are <i>stopped</i>. Note that the other processes in the group are not automatically held by this call—they are just stopped.</p>
\$holdthread \$holdthreadstop \$holdthreadstopprocess \$holdthreadstopall	<p>Freezes the current thread, leaving other threads running.</p> <p>Exactly like \$holdthread except it <i>stops</i> the <i>process</i>. The other processes in the group are left running.</p> <p>Exactly like \$holdthreadstop except it stops the entire group.</p>
\$stop \$stopprocess	<p>Sets a process-level breakpoint. The process that executes this statement stops; other processes in the control group continue to execute.</p>
\$stopall	<p>Sets a program-group-level breakpoint. All processes in the control group stop when any thread or process in the group executes this statement.</p>
\$stopthread	<p>Sets a thread-level breakpoint. Although the thread that executes this statement stops, all other threads in the process continue to execute. If the target system cannot stop an individual thread, this statement performs identically to \$stopprocess.</p>

Table 20: Built-In Statements Used in Expressions (cont.)

Statement	Use
<code>\$visualize(expression[,slice])</code>	Visualizes the data specified by <i>expression</i> and modified by the optional <i>slice</i> value. <i>Expression</i> and <i>slice</i> must be expressed using the code fragment's language. The expression must return a dataset (after modification by <i>slice</i>) that can be visualized. <i>slice</i> is a quoted string containing a slice expression. For more information on using \$visualize in an expression, see "Visualizing Data Programmatically" on page 169.

C Constructs Supported

When writing code fragments in C, keep these guidelines in mind:

- You can use C-style (*/* comment */*) and C++-style (*// comment*) comments. For example:

```
// This code fragment creates a temporary patch
i = i + 2;           /* Add two to i */
```
- You can omit semicolons if the result isn't ambiguous.
- You can use dollar signs (\$) in identifiers.

Data Types and Declarations

The following list describes the C data types and declarations that you can use:

- The data types that you can use are **char**, **short**, **int**, **float**, **double**, and pointers to any primitive type or any named type in the target program.
- Only simple declarations are permitted. Do not use **struct**, **union**, and array declarations.
- You can refer to variables of any type in the target program.
- Unmodified variable declarations are considered local. References to these declarations override references to similarly named global variables and other variables in the target program.
- (Compiled evaluation points only.) The **global** declaration makes a variable available to other evaluation points and expression windows in the target process.
- (Compiled evaluation points only.) The **extern** declaration references a global variable that was or will be defined elsewhere. If the global variable is not yet defined, TotalView displays a warning.

- Static variables are local and persist even after TotalView evaluates an evaluation point.
- TotalView only evaluates expressions that initialize static and global variables the first time it evaluates a code fragment. In contrast, it initializes local variables each time it evaluates a code fragment.

Statements

The following list describes the C language statements that you can use.

- The statements that you can use are assignment, **break**, **continue**, **if/else** structures, **for**, **goto**, and **while**.
- You can use the **goto** statement to define and branch to symbolic labels. These labels are local to the window. You can also refer to a line number in the program. This line number is the number displayed in the Source Pane. For example, here is a **goto** statement that branches to source line number 432 of the target program:

```
goto 432;
```

- Although function calls are permitted, you can't pass structures.
- Type casting is permitted.

All operators are permitted, with these limitations:

- TotalView doesn't support the **?:** conditional operator.
- While you can use the **sizeof** operator, you can't use it for data types.
- The *(type)* operator can't cast data to fixed-dimension arrays by using C cast syntax.

Fortran Constructs Supported

When writing code fragments in Fortran, keep these guidelines in mind:

- Enter only one statement on a line. You can't continue a statement onto more than one line.
- You can use **GOTO**, **GO TO**, **ENDIF**, and **END IF** statements; while **ELSEIF** isn't allowed, you can use **ELSE IF**.
- Syntax is free-form. No column rules apply.
- You can enter comments in three ways: with a **C** in column 1 or as free-form comments using the **/* ... */** delimiters or the **//** characters. In addition, anything typed after **//** characters is also ignored. The following example shows all three:

```

C I=I+1
/*
I=I+1
J=J+1
ARRAY1(I,J)= I * J
*/
k = 4 // This is also a comment

```

- The space character is significant and is sometimes required. (Some Fortran 77 compilers ignore all space characters.) For example:

Valid	Invalid
DO 100 I=1,10	DO100I=1,10
CALL RINGBELL	CALL RING BELL
X.EQ.1	X.EQ.1

Data Types and Declarations

The following is a list of data types and declarations that you can use in a Fortran expression.

- You can use the following data types: **INTEGER** (assumed to be **long**), **REAL**, **DOUBLE PRECISION**, and **COMPLEX**.
- You can't use implied data types.
- You can only use simple declarations. You can't use a **COMMON**, **BLOCK DATA**, **EQUIVALENCE**, **STRUCTURE**, **RECORD**, **UNION**, or an array declaration.
- You can refer to variables of any type in the target program.

Statements

The following list describes the Fortran language statements that you can use.

- You can use the following statements: assignment, **CALL** (to subroutines, functions, and all intrinsic functions except **CHARACTER** functions in the target program), **CONTINUE**, **DO**, **GOTO**, **IF** (including block **IF**, **ENDIF**, **ELSE**, and **ELSE IF**), and **RETURN** (but not an alternate return).
- A **GOTO** statement can refer to a line number in your program. This line number is the number displayed in the Source Pane. For example, the following **GOTO** statement branches to source line number 432:

```
GOTO $432;
```

You must use a dollar sign (\$) before the line number so that TotalView knows that you're referring to TotalView's source line number rather than a statement label.

- The only expression operators that are not supported are the **CHARACTER** operators and the **.EQV.**, **.NEQV.**, and **.XOR.** logical operators.
- You can't use subroutine function and entry definitions.
- You can't use Fortran 90 array syntax.
- You can't use Fortran 90 pointer assignment (the => operator).
- You can't call Fortran 90 functions that require assumed shape array arguments.

Writing Assembler Code

On HP Alpha Tru64 UNIX, RS/6000 IBM AIX, and SGI IRIX operating systems, TotalView lets you use assembler code in evaluation points, conditional breakpoints, and in the **Tools > Evaluate** Dialog Box. However, if you want to use assembler constructs, you must enable compiled expressions. See "*Interpreted vs. Compiled Expressions*" on page 358 for instructions.

To indicate that an expression in the breakpoint or **Evaluate** Dialog Box is an assembler expression, click on the **Assembler** button in the **Action Point > Properties** Dialog Box, as shown in Figure 193.

Assembler expressions are written in the TotalView Assembler Language. In this language, instructions are written in the target machine's native assembler language. However, the operators available to construct expressions in instruction operands and the set of available pseudo-operators are the same on all machines.

The TotalView assembler accepts instructions using the same mnemonics recognized by the native assembler, and it recognizes the same names for registers that native assemblers recognize.

Some architectures provide extended mnemonics that do not correspond exactly with machine instructions and which represent important, special cases of instructions, or provide for assembling short, commonly used sequences of instructions. The TotalView assembler recognizes these mnemonics if:

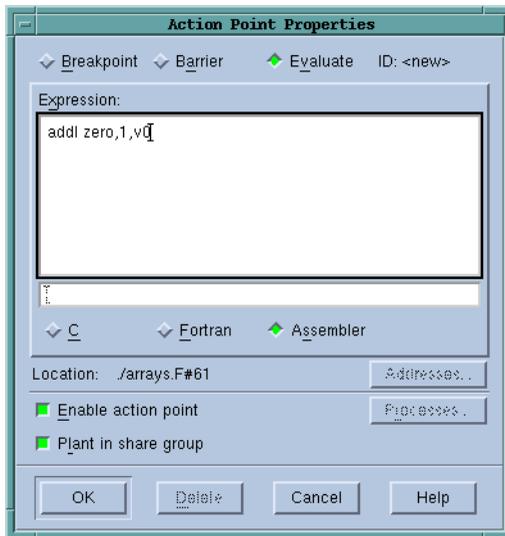


FIGURE 193: Using Assembler

- They assemble to exactly one instruction.
- The relationship between the operands of the extended mnemonics and the fields in the assembled instruction code is a simple one-to-one correspondence.

Assembler language labels are indicated as *name:* and appear at the beginning of a line. Labels can appear alone on a line. The symbols you can use include labels defined in the assembler expression and all program symbols.

The TotalView assembler operators are described in the following table:

TABLE 21: TotalView Assembler Operators

Operators	Definition
+	Plus
-	Minus (also unary)
*	Times
#	Remainder
/	Quotient
&	Bitwise AND
^	Bitwise XOR
!	Bitwise OR NOT (also unary -, bitwise NOT)
	Bitwise OR

TABLE 21: TotalView Assembler Operators (cont.)

Operators	Definition
<i>(expr)</i>	Grouping
<<	Left shift
>>	Right shift
" <i>text</i> "	Text string, 1-4 characters long, is right justified in a 32-bit word
hi16 (<i>expr</i>)	Low 16 bits of operand <i>expr</i>
hi32 (<i>expr</i>)	High 32 bits of operand <i>expr</i>
lo16 (<i>expr</i>)	High 16 bits of operand <i>expr</i>
lo32 (<i>expr</i>)	Low 32 bits of operand <i>expr</i>

The TotalView Assembler pseudo-operations are as follows:

Table 22: TotalView Assembler Pseudo-Ops

Pseudo Ops	Definition
\$debug [0 1]	<i>Internal debugging option.</i> With no operand, toggle debugging; 0 => turn debugging off 1 => turn debugging on
\$hold	Hold the process
\$holdprocess	
\$holdstopall	Hold the process and stop the control group
\$holdprocessstopall	
\$holdthread	Hold the thread
\$holdthreadstop	Hold the thread and stop process
\$holdthreadstopprocess	
\$holdthreadstopall	Hold the thread and stop the control group
\$long_branch <i>expr</i>	Branch to location <i>expr</i> using a single instruction in an architecture-independent way; using registers is not required
\$stop	Stop the process
\$stopprocess	
\$stopall	Stop the control group
\$stopthread	Stop the thread
<i>name=expr</i>	Same as def <i>name,expr</i>
align <i>expr</i> [, <i>expr</i>]	Align location counter to an operand 1 alignment; use operand 2 (or 0) as the fill value for skipped bytes
ascii <i>string</i>	Same as <i>string</i>
asciz <i>string</i>	Zero-terminated string

Table 22: TotalView Assembler Pseudo-Ops (cont.)

Pseudo Ops	Definition
bss <i>name,size-expr[,expr]</i>	Define <i>name</i> to represent <i>size-expr</i> bytes of storage in the bss section with alignment optional <i>expr</i> ; the default alignment depends on the size: if <i>size-expr</i> >= 8 then 8 else if <i>size-expr</i> >= 4 then 4 else if <i>size-expr</i> >= 2 then 2 else 1
byte <i>expr [, expr] ...</i>	Place <i>expr</i> values into a series of bytes
comm <i>name,expr</i>	Define <i>name</i> to represent <i>expr</i> bytes of storage in the bss section; <i>name</i> is declared global; alignment is as in bss without an alignment argument
data	Assemble code into data section (data)
def <i>name,expr</i>	Define a symbol with <i>expr</i> as its value
double <i>expr [, expr] ...</i>	Place <i>expr</i> values into a series of doubles
equiv <i>name,name</i>	Make operand 1 be an abbreviation for operand 2
fill <i>expr, expr, expr</i>	Fill storage with operand 1 objects of size operand 2, filled with value operand 3
float <i>expr [, expr] ...</i>	Place <i>expr</i> values into a series of floating point numbers
global <i>name</i>	Declare <i>name</i> as global
half <i>expr [, expr] ...</i>	Place <i>expr</i> values into a series of 16-bit words
lcomm <i>name,expr[,expr]</i>	Identical to bss
lsym <i>name,expr</i>	Same as def name,expr but allows redefinition of a previously defined name
org <i>expr [, expr]</i>	Set location counter to operand 1 and set operand 2 (or 0) to fill skipped bytes
quad <i>expr [, expr] ...</i>	Place <i>expr</i> values into a series of 64-bit words
string <i>string</i>	Place <i>string</i> into storage
text	Assemble code into text section (code)
word <i>expr [, expr] ...</i>	Place <i>expr</i> values into a series of 32-bit words
zero <i>expr</i>	Fill <i>expr</i> bytes with zeros





Glossary

ACTION POINT: A debugger feature that allows a user to request that program execution stop under certain conditions. Action points include breakpoints, watchpoints, evaluation points, and barriers.

ACTION POINT IDENTIFIER: A unique integer ID associated with an action point.

ADDRESS SPACE: A region of memory that contains code and data from a program. One or more threads can run in an address space. A process normally contains an address space.

ADDRESSING EXPRESSION: A set of instructions that tell TotalView where it can find information. These expressions are only used within the type transformation facility.

AFFECTED P/T SET: The set of process and threads that will be affected by the command. For most commands, this is identical to the target P/T set, but in some cases it may include additional threads. (See "P/T (process/thread) set" on page 394 for more information.)

AGGREGATE DATA: A collection of data elements. For example, a structure or an array is an aggregate.

AGGREGATED OUTPUT: The CLI compresses output from multiple threads when they would be identical except for the P/T identifier.

ARENA: A specifier that indicates the processes, threads, and groups upon which a command executes. Arena specifiers are **p** (process), **t** (thread), **g** (group), **d** (default), and **a** (all).

ARRAY SLICE: A subsection of an array, which is expressed in terms of a lower bound, upper bound, and stride. Displaying a slice of an array can be useful when you are working with very large arrays.

ASYNCHRONOUS: When processes communicate with one another, they send messages. If a process decides that it doesn't want to wait for an answer, it is said to run "asynchronously." For example, in most client/server programs, one program sends an RPC request to a second program and then waits to receive a response from the second program. This is the normal *synchronous* mode of operation. If, however, the first program sends a message and then continues executing, not waiting for a reply, the first mode of operation is said to be *asynchronous*.

AUTOLAUNCHING: When a process begins executing on a remote computer, TotalView can also launch a **tvdsvr** (TotalView Debugger Sever) process on this computer that will send debugging information back to the TotalView process that you are interacting with.

AUTOMATIC PROCESS ACQUISITION: TotalView automatically detects the many processes that parallel and distributed programs run in, and attaches to them automatically so you don't have to attach to them manually. This process is called *automatic process acquisition*. If the process is on a remote machine, automatic process acquisition automatically starts the TotalView Debugger Server (the **tvdsvr**).

BARRIER: An action point specifying that processes reaching a particular location in the source code should stop and wait for other processes to catch up.

BASE WINDOW: The original Process Window or Variable Window before you dive into routines or variables. After diving, you can use a **Reset** or **Undive** command to restore this original window.

BLOCKED: A thread state where the thread is no longer executing because it is waiting for an event to occur. In most cases, the thread is blocked because it is waiting for a mutex or condition state.

BREAKPOINT: A point in a program where execution can be suspended to permit examination and manipulation of data.

CALL FRAME: The memory area containing the variables belonging to a function, subroutine, or other scope division such as a block.

- CALL STACK:** A higher-level view of stack memory, interpreted in terms of source program variables and locations. This is where your program places stack frames.
- CHILD PROCESS:** A process created by another process (*see* "parent process" on page 393) when that other process calls **fork()**.
- CLOSED LOOP:** *see* **closed loop**.
- CLUSTER DEBUGGING:** The action of debugging a program that is running on a cluster of hosts in a network. Typically, the hosts are homogeneous.
- COMMAND HISTORY LIST:** A debugger-maintained list storing copies of the most recent commands issued by the user.
- CONDITION SYNCHRONIZATION:** A process that delays thread execution until a condition is satisfied.
- CONTEXTUALLY QUALIFIED (SYMBOL):** A symbol that is described in terms of its dynamic context, rather than its static scope. This includes process identifier, thread identifier, frame number, and variable or subprocedure name.
- CONTROL GROUP:** All the processes that a program creates. These processes can be local or remote. If your program uses processes that it did not create, TotalView places them in separate control groups. For example, a client/server program has two distinct executables that run independently of one another. Each would be in a separate control group. In contrast, processes created by **fork()** are in the same control group.
- CORE FILE:** A file containing the contents of memory and a list of thread registers. The operating system dumps (creates) a core file whenever a program exits because of a severe error (such as an attempt to store into an invalid address).
- CORE-FILE DEBUGGING:** A debugging session that examines a core file image. Commands that modify program state are not permitted in this mode.
- CROSS-DEBUGGING:** A special case of remote debugging where the host platform and the target platform are different types of machines.
- CURRENT FRAME:** The current portion of stack memory, in the sense that it contains information about the subprocedure invocation that is currently executing.

CURRENT LANGUAGE: The source code language used by the file containing the current source location.

CURRENT LIST LOCATION: The location governing what source code will be displayed in response to a list command.

DATA SET: A set of array elements generated by TotalView and sent to the Visualizer. (See **visualizer process**.)

DBELOG LIBRARY: A library of routines for creating event points and generating event logs from within TotalView. To use event points, you must link your program with both the **dbelog** and **elog** libraries.

DBFORK LIBRARY: A library of special versions of the **fork()** and **execve()** calls used by TotalView to debug multiprocess programs. If you link your program with TotalView's **dbfork** library, TotalView will be able to automatically attach to newly spawned processes.

DEBUGGING INFORMATION: Information relating an executable to the source code from which it was generated.

DEBUGGER INITIALIZATION FILE: An optional file establishing initial settings for debugger state variables, user-defined commands, and any commands that should be executed whenever TotalView or the CLI is invoked. Must be called **.tvdrc**.

DEBUGGER PROMPT: A string printed by the CLI that indicates that it is ready to receive another user command.

DEBUGGER SERVER: See **tvdsvr process**.

DEBUGGER STATE: Information that TotalView or the CLI maintains in order to interpret and respond to user commands. Includes debugger modes, user-defined commands, and debugger variables.

DEPRECATED: A feature that is still available but may be eliminated in a future release.

DISTRIBUTED DEBUGGING: The action of debugging a program that is running on more than one host in a network. The hosts can be homogeneous or heterogeneous. For example, programs written with message-passing libraries such as Parallel Virtual Machine (PVM) or Parallel Macros (PARMACS) run on more than one host.

DIVE STACK: A series of nested dives that were performed in the same Variable Window. The number of greater-than symbols (>) in the upper left-hand corner of a Variable Window indicates the number of nested dives on the dive stack. Each time that you undive, TotalView pops a dive from the dive stack and decrements the number of greater-than symbols shown in the Variable Window.

DIVING: The action of displaying more information about an item. For example, if you dive into a variable in TotalView, a window appears with more information about the variable.

DOPE VECTOR: This is a runtime descriptor that contains all information about an object that requires more information than is available as a single pointer or value. For example, you might declare a Fortran 90 pointer variable that is a pointer to some other object but which has its own upper bound as follows:

```
integer, pointer, dimension (:) :: iptr
```

Assume that you initialize it as follows:

```
iptr => iarray (20:1:-2)
```

iptr is now a synonym for every other element in the first twenty elements of **iarray**, and this pointer array is in reverse order. For example, **iptr(1)** maps to **iarray(20)**, **iptr(2)** maps to **iarray(18)**, and so on.

A compiler represents an **iptr** object using a run time descriptor that contains (at least) elements such as a pointer to the first element of the actual data, a stride value, and a count of the number of elements (or equivalently an upper bound).

DPID: Debugger ID. This is the ID TotalView uses for processes.

EDITING CURSOR: A black rectangle that appears when a TotalView GUI field is selected for editing. You use field editor commands to move the editing cursor.

EVALUATION POINT: A point in the program where TotalView evaluates a code fragment without stopping the execution of the program.

EVENT LOG: A file containing a record of events for each process in a program.

EVENT POINT: A point in the program where TotalView writes an event to the event log for later analysis with TimeScan.

EXECUTABLE: A compiled and linked version of source files, containing a “main” entry point.

EXPRESSION: An expression consists of symbols (possibly qualified), constants, and operators, arranged in the syntax of the current source language. Not all Fortran 90, C, and C++ operators are supported.

EXTENT: The number of elements in the dimension of an array. For example, a Fortran array of integer(7,8) has an extent of 7 in one dimension (7 rows) and an extent of 8 in the other dimension (8 columns).

FIELD EDITOR: A basic text editor that is part of TotalView’s interface. The field editor supports a subset of GNU Emacs commands.

FOCUS: The set of groups, processes, and threads upon which a CLI command acts. The current focus is indicated in the CLI prompt (if you’re using the default prompt).

FRAME: An area in stack memory containing the information corresponding to a single invocation of a subprocedure. See **stack frame**.

FRAME POINTER: See **stack pointer**.

FULLY QUALIFIED (SYMBOL): A symbol is fully qualified when each level of source code organization is included. For variables, those levels are executable or library, file, procedure or line number, and variable name.

GID: The TotalView group ID.

IMAGE: All of the programs, libraries, and other components that make up your executable is called an image.

GOI: The group of interest. This is the group that TotalView uses when it is trying to determine what to step, stop, and the like.

GROUP: When TotalView starts processes, it places related processes in families. These families are called “groups.”

GROUP OF INTEREST: The primary group that is affected by a command. This is the group that TotalView uses when it is trying to determine what to step, stop, and the like.

HEAP: An area of memory that your program uses when it dynamically allocates blocks of memory. It is also how people describe my car.

HOST MACHINE: The machine on which the TotalView debugger is running.

INITIAL PROCESS: The process created as part of a load operation, or that already existed in the runtime environment and was attached by TotalView or the CLI.

INFINITE LOOP: See **loop**, **infinite**.

LVALUE: A symbol name or expression suitable for use on the left-hand side of an assignment statement in the corresponding source language. That is, the expression must be appropriate as the target of an assignment.

LHS EXPRESSION: This is a synonym for **lvalue**.

LOCKSTEP GROUP: All threads that are at the same PC (program counter). This group is a subset of a workers group. A lockstep group only exists for stopped threads. All threads in the lockstep group are also in a workers group. By definition, all members of a lockstep group are in the same workers group. That is, a lockstep group cannot have members in more than one workers group or more than one control group.

LOOP, INFINITE: see **infinite loop**.

LOWER BOUND: The first element in the dimension of an array or the slice of an array. By default, the lower bound of an array is 0 in C and 1 in Fortran, but the lower bound can be any number, including negative numbers.

MACHINE STATE: Convention for describing the changes in memory, registers, and other machine elements as execution proceeds.

MANAGER THREAD: A thread created by the operating system. In most cases, you do not want to manage or examine manager threads.

MESSAGE QUEUE: A list of messages sent and received by message-passing programs.

MIMD: An acronym for "Multiple Instruction, Multiple Data," describing a type of parallel computing.

MISD: An acronym for "Multiple Instruction, Single Data," describing a type of parallel computing.

MPI: This is an acronym for "Message Passing Interface."

MPICH: MPI/Chameleon (Message Passing Interface/Chameleon) is a freely available and portable MPI implementation. MPICH was written as a collaboration between Argonne National Lab and Mississippi State University. For more information, see www.mcs.anl.gov/mpi.

MPMD (MULTIPLE PROGRAM MULTIPLE DATA) PROGRAMS: A program involving multiple executables, executed by multiple threads and processes.

MUTEX (MUTUAL EXCLUSION): Techniques for sharing resources so that different users do not conflict and cause unwanted interactions.

NATIVE DEBUGGING: The action of debugging a program that is running on the same machine as TotalView.

NESTED DIVE: TotalView lets you dive into pointers, structures, or arrays in a variable. When you dive into one of these elements, TotalView updates the display so that the new element is displayed. So, a nested dive is a *dive* within a dive. You can return to the previous display by selecting the left-facing arrow in the top-right corner of the window.

NODE: A machine on a network. Each machine has a unique network name and address.

OUT OF SCOPE: When symbol lookup is performed for a particular symbol name and it isn't found in the current scope or any containing scopes, the symbol is said to be out of scope.

PARALLEL PROGRAM: A program whose execution involves multiple threads and processes.

PARALLEL TASKS: Tasks whose computations are independent of each other, so that all such tasks can be performed simultaneously with correct results.

PARALLELIZABLE PROBLEM: A problem that can be divided into parallel tasks. This may require changes in the code and/or the underlying algorithm.

PARCEL: The number of bytes required to hold the shortest instruction for the target architecture.

PARENT PROCESS: A process that calls `fork()` to spawn other processes (usually called “child processes”).

PARMACS LIBRARY: A message-passing library for creating distributed programs that was developed by the German National Research Centre for Computer Science.

PARTIALLY QUALIFIED (SYMBOL): A symbol name that includes only some of the levels of source code organization (for example, filename and procedure, but not executable). This is permitted as long as the resulting name can be associated unambiguously with a single entity.

PC: This is an abbreviation for *Program Counter*.

PID: Depending on context, this is either the “process ID” or the “program ID.” In most cases, this will be a process ID.

POI: The process of interest. This is the process that TotalView uses when it is trying to determine what to step, stop, and the like.

PROCESS: An executable that is loaded into memory and is running (or capable of running).

PROCESS GROUP: A group of processes associated with a multiprocess program. A process group includes program control groups and share groups.

PROCESS/THREAD IDENTIFIER: A unique integer ID associated with a particular process and thread.

PROCESS OF INTEREST: The primary process that TotalView uses when it is trying to determine what to step, stop, and the like. This is abbreviated as POI.

PROGRAM EVENT: A program occurrence that is being monitored by TotalView or the CLI, such as a breakpoint.

PROGRAM CONTROL GROUP: A group of processes that includes the parent process and all related processes. A program control group includes children that were forked (processes that share the same source code as the parent) and children that were forked with a subsequent call to `execve()` (processes that don't share the same source code as the parent). Contrast with **share group**.

PROGRAM STATE: A higher-level view of the machine state, where addresses, instructions, registers, and such, are interpreted in terms of source program variables and statements.

P/T (PROCESS/THREAD) SET: The set of threads drawn from all threads in all processes of the target program.

PVM LIBRARY: Parallel Virtual Machine library. A message-passing library for creating distributed programs that was developed by the Oak Ridge National Laboratory and the University of Tennessee.

RACE CONDITION: A problem that occurs when threads try to simultaneously access a resource. The result can be a deadlock, data corruption, or a program fault.

REMOTE DEBUGGING: The action of debugging a program that is running on a different machine than TotalView. The machine on which the program is running can be located many miles away from the machine on which TotalView is running.

RESUME COMMANDS: Commands that cause execution to restart from a stopped state: **dstep**, **dgo**, **dcont**, **dwait**.

RHS EXPRESSION: This is a synonym for **rvalue**.

RVALUE: An expression suitable for inclusion on the right-hand side of an assignment statement in the corresponding source language. In other words, an expression that evaluates to a value or collection of values.

SATISFACTION SET: The set of processes and threads that must be held before a barrier can be satisfied.

SATISFIED: A condition indicating that all processes or threads in a group have reached a barrier. Prior to this event, all executing processes and threads are either running because they have not yet hit the barrier or are being held at the barrier because not all of the processes or threads have reached it. After the barrier is *satisfied*, the held processes or threads are released, which means they can now be run. Prior to this event, they could not be run.

SERIAL EXECUTION: Execution of a program sequentially, one statement at a time.

SERIAL LINE DEBUGGING: A form of remote debugging where TotalView and the TotalView Debugger Server communicate over a serial line.

SERVICE THREAD: A thread whose purpose is to “service” or manage other threads.

For example, queue managers and print spoolers are service threads. There are two kinds of service threads: those created by the operating system or runtime system and those created by your program. If a service thread is not created by your program, you won’t be interested in debugging it. If your program is creating a service thread, however, you will probably debug it separately from the rest of your program.

SHARE GROUP: All the processes in a control group that share the same code. In most cases, your program will have more than one share group. Share groups, like control groups, can be local or remote.

SHARED LIBRARY: A compiled and linked set of source files that are dynamically loaded by other executables—and have no “main” entry point.

SIGNALS: Messages informing processes of asynchronous events, such as serious errors. The action the process takes in response to the signal depends on the type of signal and whether or not the program includes a signal handler routine, a routine that traps certain signals and determines appropriate actions to be taken by the program.

SIMD: An acronym for “Single Instruction, Multiple Data,” describing a type of parallel computing.

SISD: An acronym for “Single Instruction, Single Data,” describing a type of parallel computing.

SINGLE STEP: The action of executing a single statement and stopping (as if at a breakpoint).

SLICE: A subsection of an array, which is expressed in terms of a lower bound, upper bound, and stride. Displaying a slice of an array can be useful when you are working with very large arrays.

SOID: An acronym for “symbol object ID”. A soid uniquely identifies all information within TotalView. It also represents a handle by which this information can be accessed.

SOURCE FILE: Program file containing source language statements. TotalView allows you to debug FORTRAN 77, Fortran 90, Fortran 95, C, C++, and assembler.

SOURCE LOCATION: For each thread, the source code line it will execute next. This is a static location, indicating the file and line number; it does not, however, indicate which invocation of the subprocedure is involved.

SPAWNED PROCESS: The process created by a user process executing under debugger control.

SPMD (SINGLE PROGRAM MULTIPLE DATA) PROGRAMS: A program involving just one executable, executed by multiple threads and processes.

STACK: A portion of computer memory and registers used to hold information temporarily. The stack consists of a linked list of stack frames that holds return locations for called routines, routine arguments, local variables, and saved registers.

STACK FRAME: Whenever your program calls a function, it creates a set of information that includes the local variables, arguments, contents of the registers used by an individual routine, a frame pointer pointing to the previous stack frame, and the value of the program counter (PC) at the time the routine was called. The information for one function is called a "stack frame" as it is placed on your program's stack.

When your program begins executing, it has only one frame: the one allocated for function `main()`. As your program calls functions, new frames are allocated. When a function returns to the function from which it is called, the frame is deallocated.

STACK POINTER: A pointer to the area of memory where subprocedure arguments, return addresses, and similar information is stored. This is also called a "frame pointer."

STACK TRACE: A sequential list of each currently active routine called by a program and the frame pointer pointing to its stack frame.

STATIC (SYMBOL) SCOPE: A region of a program's source code that has a set of symbols associated with it. A scope can be nested inside another scope.

STEPPING: Advancing program execution by fixed increments, such as by source code statements.

STL: A TLA for Standard Template Library.

- STOP SET:** A set of threads that should be stopped once an action point has been triggered.
- STOPPED/HELD STATE:** The state of a process whose execution has paused in such a way that another program event (for example, arrival of other threads at the same barrier) will be required before it is capable of continuing execution.
- STOPPED/RUNNABLE STATE:** The state of a process whose execution has been paused (for example, when a breakpoint triggered or due to some user command) but can continue executing as soon as a resume command is issued.
- STOPPED STATE:** The state of a process that is no longer executing, but will eventually execute again. This is subdivided into stopped/runnable and stopped/held.
- STRIDE:** The interval between array elements in a slice and the order in which the elements are displayed. If the stride is 1, every element between the lower bound and upper bound of the slice is displayed. If the stride is 2, every other element is displayed. If the stride is -1, every element between the upper bound and lower bound (reverse order) is displayed.
- SYMBOL:** Entities within program state, machine state, or debugger state.
- SYMBOL LOOKUP:** Process whereby TotalView consults its debugging information to discover what entity a symbol name refers to. Search starts with a particular static scope and occurs recursively so that containing scopes are searched in an outward progression.
- SYMBOL NAME:** The name associated with a symbol known to TotalView (for example, function, variable, data type, and such).
- SYMBOL TABLE:** A table of symbolic names (such as variables or functions) used in a program and their memory locations. The symbol table is part of the executable object generated by the compiler (with the `-g` option) and is used by debuggers to analyze the program.
- SYNCHRONIZATION:** A mechanism that prevents problems caused by concurrent threads manipulating shared resources. The two most common mechanisms for synchronizing threads are mutual exclusion and condition synchronization.
- TARGET MACHINE:** The machine on which the process to be debugged is running.

TARGET PROCESS SET: The target set for those occasions when operations can only be applied to entire processes, not to individual threads in a process.

TARGET PROGRAM: The executing program that is the target of debugger operations.

TARGET P/T SET: The set of processes and threads that a CLI command will act on.

TASK: A logically discrete section of computational work. (This is an informal definition.)

THREAD: An execution context that normally contains a set of private registers and a region of memory reserved for an execution stack. A thread runs in an address space.

THREAD EXECUTION STATE: The convention of describing the operations available for a thread, and the effects of the operation, in terms of a set of predefined states.

THREAD OF INTEREST: The primary thread that will be affected by a command. This is abbreviated as TOI.

TID: The thread ID.

TLA: An acronym for "Three-Letter Acronym." So many things from computer hardware and software vendors are referred to by a three-letter acronym that yet another acronym was created to describe these terms.

TOI: The thread of interest. This is the primary thread that will be affected by a command.

TRIGGER SET: The set of threads that can trigger an action point (that is, the threads upon which the action point was defined).

TRIGGERS: The effect during execution when program operations cause an event to occur (such as, arriving at a breakpoint).

TTF: See type transformation facility.

TVDSVR PROCESS: The TotalView Debugger Server process, which facilitates remote debugging by running on the same machine as the executable and communicating with TotalView over a TCP/IP port or serial line.

TYPE TRANSFORMATION FACILITY: A protocol that allows you to change the way information is displayed. For example, an STL vector can be displayed as an array.

UNDISCOVERED SYMBOL: A symbol that is referred to by another symbol. For example, a **typedef** is a reference to the aliased type.

UNDIVING: The action of displaying the previous contents of a window, instead of the contents displayed for the current dive. To undive, you click on the **undive** icon in the upper right-hand corner of the window.

UPPER BOUND: The last element in the dimension of an array or the slice of an array.

USER THREAD: A thread created by your program.

USER INTERRUPT KEY: A keystroke used to interrupt commands, most commonly defined as **^C** (Ctrl+C).

VARIABLE WINDOW: A TotalView window displaying the name, address, data type, and value of a particular variable.

VISUALIZER PROCESS: A process that works with TotalView in a separate window, allowing you to see a graphic representation of program array data.

WATCHPOINT: An action point specifying that execution should stop whenever the value of a particular variable is updated.

WORKER THREAD: A thread in a workers group. These are threads created by your program that performs the actual "work." However, you might want to distinguish between threads that do the work and threads that assist the work. For example, you might not consider a thread that acts as a queue manager as being a worker thread even though TotalView considers it to be a worker thread. (This kind of thread might be called a "worker-manager" thread.) All worker thread is always part of a workers group.

WORKERS GROUP: All the worker threads in a control group. These threads can reside in more than one share group.



Glossary

workers group – workers group



Index

Symbols

- # scope separator character 295
- \$clid built-in variable 374
- \$count built-in function 6, 356, 359, 375
- \$countall built-in function 375
- \$countthread built-in function 375
- \$debug assembler pseudo op 382
- \$denorm filter 327
- \$duid built-in variable 374
- \$hold assembler pseudo op 382
- \$hold built-in function 376
- \$holdprocess assembler pseudo op 382
- \$holdprocess built-in function 376
- \$holdprocessall built-in function 376
- \$holdprocessstopall assembler pseudo op 382
- \$holdstopall assembler pseudo op 382
- \$holdstopall built-in function 376
- \$holdthread assembler pseudo op 382
- \$holdthread built-in function 376
- \$holdthreadstop assembler pseudo op 382
- \$holdthreadstop built-in function 376
- \$holdthreadstopall assembler pseudo op 382
- \$holdthreadstopall built-in function 376
- \$holdthreadstopall built-in function 376
- \$holdthreadstopprocess assembler pseudo op 382
- \$holdthreadstopprocess built-in function 376
- \$inf filter 327
- \$long_branch assembler pseudo op 382
- \$nan filter 327
- \$nanq filter 327
- \$nans filter 327
- \$ndenorm filter 327
- \$newval built-in function 368
- \$newval built-in variable 374
- \$nid built-in variable 374
- \$ninf filter 327
- \$oldval built-in function 368
- \$oldval built-in variable 374
- \$pdenorm filter 327
- \$pid built-in variable 374
- \$pinf filter 327
- \$processduid built-in variable 374
- \$stop assembler pseudo op 382
- \$stop built-in function 5, 359, 369, 376
- \$stopall assembler pseudo op 382
- \$stopall built-in function 376
- \$stopprocess assembler pseudo op 382
- \$stopprocess built-in function 376
- \$stopthread assembler pseudo op 382
- \$stopthread built-in function 376
- \$stid built-in variable 374
- \$visualize built-in function 169, 170, 377
 - in animations 169
 - using casts 170
- %C server launch replacement characters 80
- %D bulk server launch command 82
- %D single process server launch command 81
- %H bulk server launch command 82
- %L bulk server launch command 82
- %L single process server launch command 81
- %N bulk server launch command 83
- %P bulk server launch command 82
- %P single process server launch command 81
- %R single process server launch command 81
- %t1 bulk server launch command 83
- %t2 bulk server launch command 83

- %V bulk server launch command 82
 - & intersection operator 270
 - . (dot) current set indicator 250, 271
 - . (period), in suffix of process names 224
 - .rhosts file 84, 100
 - .totalview subdirectory 43
 - .tvdrcl initialization files 43
 - .Xdefaults file 45, 69
 - autoLoadBreakpoints 69
 - deprecated resources 69
 - see also*
 - www.etnus.com/Support/docs/xresources/Xresources.html
 - / slash in group specifier 256
 - /usr/lib/array/arrayd.conf file 82
 - : (colon), in array type strings 300
 - : as array separator 320
 - < first thread indicator 250
 - <address> data type 302
 - <char> data type 302
 - <character> data type 302
 - <code> 290
 - <code> data type 302, 304, 307
 - <complex*16> data type 302
 - <complex*8> data type 302
 - <complex> data type 302
 - <double precision> data type 302
 - <double> data type 302
 - <extended> data type 303
 - <float> data type 303
 - <int> data type 303
 - <integer*1> data type 303
 - <integer*2> data type 303
 - <integer*4> data type 303
 - <integer*8> data type 303
 - <integer> data type 303
 - <logical*1> data type 303
 - <logical*2> data type 303
 - <logical*4> data type 303
 - <logical*8> data type 303
 - <logical> data type 303
 - <long long> data type 303
 - <long> data type 303
 - <opaque> data type 306
 - <real*16> data type 303
 - <real*4> data type 303
 - <real*8> data type 303
 - <real> data type 303
 - <short> data type 303
 - <string> data type 298, 299, 303, 304
 - <void> data type 303, 304
 - > (right angle bracket), indicating nested dives 292
 - @ action point marker, in CLI 340
 - difference operator 270
 - | union operator 270
 - ' module separator 312
- A**
- a command-line option 42, 201
 - a passing arguments to program option 42
 - a width specifier 257
 - examples 260
 - general discussion 259
 - abbreviating commands 203
 - absolute addresses, display
 - assembler as 217
 - acquiring processes 102
 - Action Point > At Location
 - command 5, 340, 342
 - Action Point > Delete All
 - command 344
 - Action Point > Properties
 - command 5, 140, 228, 338, 343, 344, 346, 348, 351, 354, 356
 - deleting barrier points 353
 - Action Point > Properties Dialog
 - Box figure 343, 346, 351
 - Action Point > Save All
 - command 370
 - Action Point > Set Barrier
 - command 351
 - Action Point > Suppress All
 - command 344
 - action point files 45
 - action point identifiers 207
 - never reused in a session 207
 - Action Point Properties and Address Dialog Boxes figure 228
 - Action Point Properties Dialog
 - Box figure 6
 - Action Point Symbol figure 339
 - action points 207
 - see also* barrier points
 - see also* eval points
 - common properties 338
 - defined 5
 - definition 337
 - deleting 344
 - disabling 343
 - enabling 344
 - evaluation points 5
 - ignoring 344
 - list of 152
 - multiple addresses 340
 - saving 370
 - suppressing 344
 - unsuppressing 344
 - watchpoint 10
 - Action Points Page 137, 152
 - Action Points Pane 344
 - adapter_use option 99
 - adding command-line arguments 62
 - adding environment variables 70
 - adding members to a group 254
 - adding program arguments 43
 - Address Only (Absolute Addresses) figure 217
 - address range conflicts 361
 - addresses
 - changing 306
 - editing 306
 - of machine instructions 307
 - specifying in variable window 289
 - tracking in variable window 286
 - advancing and holding processes 206
 - advancing program execution 206
 - aliases

- built-in 203
- group 203
- group, limitations 204
- align assembler pseudo op 382
- all width specifier 251
- allocated arrays, displaying 306
- altering groups 275
- Ambiguous Addresses Dialog Box
 - figure 229
- Ambiguous Function Dialog Box 342
- Ambiguous Function Dialog Box
 - figure 213, 342
- ambiguous function names 212, 342
- Ambiguous Line Dialog Box figure 341
- ambiguous locations 342
- ambiguous names 214
- ambiguous scoping 296
- ambiguous source lines 227
- analyzing memory 161
- angle brackets, in windows 292
- animation using \$visualize 169
- areas of memory, data type 304
- arena specifiers 250
 - defined 250
 - incomplete 267
 - inconsistent widths 267
- arenas
 - and scope 240
 - defined 240, 249
 - iterating over 250
- ARGS variable 201
 - modifying 201
- ARGS_DEFAULT variable 43, 62, 201
 - clearing 202
- arguments
 - in server launch command 80, 85
 - passing to program 42
 - replacing 202
 - setting 61
- Arguments page 62
- argv, displaying 305
- Array Data Filter by Range of Values figure 328
- array data filtering
 - by comparison 324
 - by range of values 327
 - for IEEE values 326
- Array Data Filtering by Comparison figure 326
- Array Data Filtering for IEEE Values figure 328
- array data filtering, *see* arrays filtering
- array of structures, displaying 293
- array pointers 287
- array rank 165
- array services handle (ash) 105
- Array Statistics Window figure 331
- Array Visualization figure 10
- arrays
 - array data filtering 324
 - bounds 299
 - character 304
 - checksum statistic 331
 - colon separators 320
 - count statistic 331
 - deferred shape 314, 320
 - denormalized count statistic 332
 - display subsection 300
 - displaying 170, 319
 - displaying allocated 306
 - displaying argv 305
 - displaying contents 153
 - displaying declared 306
 - displaying multiple 170
 - displaying one element 323
 - displaying slices 319
 - diving into 291
 - editing dimension of 300
 - extent 300
 - filter conversion rules 325
 - filter expressions 329
 - filtering 300, 324, 326
 - filtering options 324
 - in C 299
 - in Fortran 300
 - infinity count statistic 332
 - laminating 334
 - limiting display 321
 - lower adjacent statistic 332
 - lower bound of slices 320
 - lower bounds 299, 300
 - maximum statistic 332
 - mean statistic 332
 - median statistic 332
 - minimum statistic 332
 - NaN statistic 332
 - non-default lower bounds 300
 - overlapping nonexistent memory 319
 - pointers to 299
 - quartiles statistic 332
 - skipping elements 321
 - skipping over elements 320
 - slice 323
 - slice example 320, 321
 - slice initializing 187
 - slice refining 170
 - slice, printing 187
 - slices with the variable command 322
 - slicing 8
 - sorting 330
 - standard deviation statistic 332
 - statistics 331
 - stride 320
 - stride elements 320
 - subsections 319
 - sum statistic 332
 - type strings for 299
 - upper adjacent statistic 333
 - upper bound 299
 - upper bound of slices 320
 - visualization 170
 - visualizing 167
 - zero count statistic 333
- arrow over line number 150
- Ascending command 330
- asciz assembler pseudo op 382
- asciz assembler pseudo op 382
- ash (array services handle) 105
- ASM icon 339, 345
- assembler

- absolute addresses 217
 - and `-g` compiler option 153
 - constructs 380
 - displaying 217
 - examining 216
 - expressions 380
 - in code fragment 354
 - symbolic addresses 217
 - Assembler > By Address
 - command 217
 - Assembler > Symbolically
 - command 217
 - Assembler command 216
 - Assembler Only (Symbolic Addresses) figure 218
 - assembler-level action points 339
 - assigning output 200
 - assigning output to variable 200
 - assigning p/t set to variable 252
 - asynchronous processing 18
 - at breakpoint state 55
 - At Location command 5, 340, 342
 - Attach Subsets command 134
 - Attached Page 103, 146, 232, 233
 - Attached Page Showing Process and Thread Status figure 54, 58
 - attached process states 55
 - attached thread states 55
 - attaching
 - restricting 134
 - restricting by communicator 135
 - selective 134
 - to a task 126
 - to all 135
 - to HP MPI job 98
 - to job 102
 - to MPICH application 94
 - to MPICH job 94
 - to none 135
 - to PE 102
 - to poe 102
 - to processes 49, 50, 102, 126, 134, 147
 - to PVM task 126
 - to relatives 52
 - to RMS processes 104
 - to SGI MPI job 105
 - attaching using File > New Program 51
 - Auto Visualize command 172
 - Auto Visualize, in Directory Window 172
 - `auto_array_cast_bounds` variable 288
 - `auto_deref_in_all_c` variable 288
 - `auto_deref_in_all_fortran` variable 288
 - `auto_deref_initial_c` variable 288
 - `auto_deref_initial_fortran` variable 288
 - `auto_deref_nested_c` variable 288
 - `auto_deref_nested_fortran` variable 288
 - `auto_save_breakpoints` variable 370
 - autolaunch 73, 74
 - changing 84
 - defined 48
 - disabling 48, 74, 76, 84
 - launch problems 79
 - sequence 86
 - `autoLoadBreakpoints.Xdefault` 69
 - automatic dereferencing 287
 - automatic process acquisition 94, 99, 124
 - averaging data points 179
 - averaging surface display 179
 - axis, transposing 175
- ## B
- B state 55
 - backtick separator 312
 - backward icon 154
 - barrier points
 - see also* process barrier
 - breakpoint 13, 35, 221, 233, 350, 352
 - clearing 344
 - defined 207
 - defined (again) 350
 - deleting 353
 - satisfying 353
 - states 350
 - stopped process 354
 - baud rate, for serial line 88
 - bit fields 297
 - block scoping 294
 - blocking send operations 112
 - blocks, naming 295
 - bold data 7
 - Both command 216, 237
 - Both Source and Assembler (Symbolic Addresses) figure 218
 - bounds for arrays 299
 - boxed line number 150, 241, 340
 - Breakpoint at Assembler Instruction figure 345
 - breakpoint files 45
 - breakpoint operator 270
 - breakpoints
 - and `MPI_Init()` 101
 - apply to all threads 338
 - automatically copied from master process 94
 - behavior when reached 345
 - changing for parallelization 137
 - clearing 146, 241, 344
 - conditional 354, 356, 375
 - copy, master to slave 94
 - countdown 356, 375
 - counting down 375
 - default stopping action 137
 - defined 207, 337
 - deleting 344
 - disabling 343
 - enabling 344
 - entering 105
 - example setting in multiprocess program 349
 - `fork()` 348
 - ignoring 344
 - in child process 346
 - in parent process 346

- in spawned process 125
 - listing 152
 - machine-level 345
 - multiple processes 346
 - not shared in separated
 - children 348
 - placing 150
 - reloading 101
 - removed when detaching 53
 - removing 146
 - saving 370
 - set while a process is running 340
 - set while running parallel tasks 101
 - setting 101, 146, 189, 241, 339, 340, 346
 - shared by default in processes 348
 - sharing 347, 348
 - stop all related processes 347
 - suppressing 344
 - thread-specific 374
 - toggling 340
 - while stepping over 229
 - bss assembler pseudo op 383
 - built-in aliases 203
 - built-in functions
 - \$count 6, 356, 359, 375
 - \$countall 375
 - \$countthread 375
 - \$hold 376
 - \$holdprocess 376
 - \$holdprocessall 376
 - \$holdstopall 376
 - \$holdthread 376
 - \$holdthreadstop 376
 - \$holdthreadstopall 376
 - \$holdthreadstopprocess 376
 - \$stop 5, 359, 369, 376
 - \$stopall 376
 - \$stopprocess 376
 - \$stopthread 376
 - \$visualize 169, 170, 377
 - forcing interpretation 359
 - built-in type strings 302
 - built-in variables 374
 - \$clid 374
 - \$duid 374
 - \$newval 374
 - \$nid 374
 - \$oldval 374
 - \$pid 374
 - \$processduid 374
 - \$systid 374
 - \$tid 374
 - forcing interpretation 375
 - Bulk Launch Page 79
 - bulk server launch 73, 76
 - command 77
 - connection timeout 78
 - enabling 76
 - on HP Alpha 83
 - on IBM RS/6000 83
 - on SGI MIPS 81
 - bulk server launch command
 - %D 82
 - %H 82
 - %L 82
 - %N 83
 - %P 82
 - %t1 83
 - %t2 83
 - %V 82
 - callback_host 82
 - callback_ports 82
 - set_pws 82
 - verbosity 82
 - working_directory 82
 - bulk_incr_timeout variable 78
 - bulk_launch_base_timeout variable 78
 - bulk_launch_enabled variable 77, 79
 - bulk_launch_string variable 77
 - bulk_launch_tmpefile1_trailer_line variable 78
 - bulk_launch_tmpefile2_trailer_line variable 78
 - bulk_launch_tmpfile1_header_line variable 78
 - bulk_launch_tmpfile1_host_lines variable 78
 - bulk_launch_tmpfile2_header_line variable 78
 - bulk_launch_tmpfile2_host_lines variable 78
 - By Address command 217
 - byte assembler pseudo op 383
- ## C
- C control group specifier 255, 257
 - C language
 - array bounds 299
 - arrays 299
 - filter expression 329
 - how data types are displayed 298
 - in code fragment 354
 - in evaluation points 377
 - type strings supported 298
 - C++
 - changing class types 309
 - display classes 307
 - C++ Type Cast to Base Class
 - Question Window figure 309
 - C++ Type Cast to Derived Class
 - Question Window figure 309
 - call stack 150
 - call tree
 - updating display 159
 - Call Tree command 159
 - callback command-line option 84
 - callback_host bulk server launch command 82
 - callback_option single process server launch command 81
 - callback_ports bulk server launch command 82
 - capture command 200, 201
 - casting 297, 299
 - examples 304
 - to type 290
 - types of variable 297
 - Casting Code figure 291
 - CGROUP variable 254, 262
 - ch_lfshmem device 92
 - ch_mpl device 92
 - ch_p4 device 92, 95, 140

- ch_shmem device 92, 95
- changing autolaunch options 74
- changing command-line arguments 62
- changing groups 275
- changing precision 280
- changing process thread set 248
- changing program state 195
- changing remote shell 84
- changing size 280
- changing values 156
- changing variables 296
- character arrays 304
- chasing pointers 287, 291
- Checkpoint and Restart Dialog
 - Boxes figure 236
- checksum array statistic 331
- child process names 224
- children calling `execve()`, *see* `execve()`
- classes, displaying 307
- Clear All STOP and EVAL command 344
- clearing
 - breakpoints 146, 241, 344, 346
 - continuation signal 234
 - evaluation points 146
- CLI
 - and Tcl 193
 - components 193
 - in startup file 197
 - initialization 197
 - interface 195
 - introduced 14
 - invoking program from shell
 - example 197
 - not a library 193
 - output 200
 - relationship to TotalView 194
 - starting 41, 42, 196
 - starting from command prompt 196
 - starting from TotalView GUI 196
 - starting program using 198
 - CLI and Tcl relationship 195
 - CLI and TotalView figure 194
- CLI commands
 - abbreviating 203
 - assigning output to variable 200
 - capture 200, 201
 - dactions 338
 - dactions `-load` 101, 370
 - dactions `-save` 101, 370
 - dassign 296
 - dattach 42, 47, 49, 51, 53, 94, 102, 103, 206
 - dattach mprun 106
 - dbarrier 350, 352
 - dbarrier `-e` 356
 - dbarrier `-stop_when_hit` 140
 - dbreak 340, 342, 347
 - dbreak `-e` 356
 - dcheckpoint 235
 - ddelete 113, 342, 344, 353
 - ddetach 53
 - ddisable 343, 344, 354
 - ddown 231
 - default focus 249
 - denable 344
 - dfocus 229, 248, 249
 - dgo 97, 101, 105, 138, 226, 268
 - dgroups `-add` 254, 262
 - dhalt 138, 220, 229
 - dhold 222, 351
 - dhold `-thread` 223
 - dkill 139, 199, 206, 234
 - dlist 125
 - dload 42, 47, 48, 49, 80, 198, 206
 - dmstat 162
 - dnnext 138, 227, 231
 - dnnexti 227, 231
 - dout 232, 242
 - dprint 117, 119, 213, 238, 282, 286, 289, 301, 306, 310, 311, 313, 320, 323
 - dptsets 55, 224
 - drerun 198, 235
 - redirecting I/O 63
 - drestart 235
 - drun 198, 201, 202
 - redirecting I/O 63
 - dset 201, 203
 - dstatus 55, 232, 353
 - dstep 138, 227, 230, 241, 250, 252, 268
 - dstepi 227, 230
 - dunhold 222, 351
 - dunhold `-thread` 223
 - dunset 202
 - duntil 227, 231, 241, 244
 - dup 231, 286
 - dwhere 251, 268, 286
 - exit 46
 - run when starting TotalView 44
- CLI prompt 198
- CLI variables
 - ARGS 201
 - ARGS, modifying 201
 - ARGS_DEFAULT 43, 62, 201
 - clearing 202
 - auto_array_cast_bounds 288
 - auto_deref_in_all_c 288
 - auto_deref_in_all_fortran 288
 - auto_deref_initial_c 288
 - auto_deref_initial_fortran 288
 - auto_deref_nested_c 288
 - auto_deref_nested_fortran 288
 - auto_save_breakpoints 370
 - bulk_incr_timeout 78
 - bulk_launch_base_timeout 78
 - bulk_launch_enabled 77, 79
 - bulk_launch_string 77
 - bulk_launch_tmpefile1_trailer_line 78
 - bulk_launch_tmpefile2_trailer_line 78
 - bulk_launch_tmpfile1_header_line 78
 - bulk_launch_tmpfile1_host_lines 78
 - bulk_launch_tmpfile2_header_line 78
 - bulk_launch_tmpfile2_host_lines 78
 - EXECUTABLE_PATH 48, 51
 - LINES_PER_SCREEN 201

- parallel_attach 137
- parallel_stop 136
- pop_at_breakpoint 59
- pop_on_error 58
- process_load_callbacks 45
- PROMPT 203
- search_path 61
- server_launch_enabled 75, 79, 84
- server_launch_string 75
- server_launch_timeout 76
- SHARE_ACTION_POINT 343, 347, 348
- STOP_ALL 343, 347
- warn_step_throw 58
- CLI xterm Window figure 196
- \$clid built-in variable 374
- Close command 154, 290
- Close command (Visualizer) 172
- Close Relatives command 154
- Close Similar command 154, 290
- Close, in Data Window 172
- closed loop, *see* closed loop
- closing similar windows 154
- closing variable windows 290
- closing windows 154
- cluster ID 374
- code constructs supported
 - Assembler 380
 - C 377
 - Fortran 378
- <code> data type 304, 307
- code fragments 354, 372, 373
 - modifying instruction path 355
 - when executed 355
 - which programming languages 354
- colons as array separators 320
- comm assembler pseudo op 383
- command arguments 201
 - clearing example 201
 - passing defaults 202
 - setting 201
- command line arguments 61, 198
 - passing to TotalView 42
- Command Line command 42, 196
- command line-options
 - launch Visualizer 180
- command prompts 203
 - default 203
 - format 203
 - setting 203
 - starting the CLI from 196
- command-line options
 - a 42, 201
 - remote 43
 - s startup 197
- commands 41
 - Action Point > At Location 5, 340
 - Action Point > Delete All 344
 - Action Point > Properties 140, 344, 346, 348, 351, 354
 - Action Point > Save All 370
 - Action Point > Set Barrier 351
 - Action Point > Suppress All 344
- arguments 61
- Auto Visualize (Visualizer) 172
- change Visualizer launch 167
- Clear All STOP and EVAL 344
- CLI, *see* CLI commands
- dmpirun 96, 97
- dpvm 124
- Edit > Copy 156
- Edit > Cut 156
- Edit > Delete 156
- Edit > Find 5, 212
- Edit > Find Again 212
- Edit > Paste 156
- File > Close 154, 290
- File > Close (Visualizer) 172
- File > Close Similar 154, 290
- File > Delete (Visualizer) 172
- File > Directory (Visualizer) 172
- File > Edit Source 214, 218
- File > Exit (Visualizer) 172
- File > New Base Window (Visualizer) 172
- File > New Program 46, 49, 51, 53, 80, 84, 88
- File > Options (Visualizer) 172, 174
- File > Preferences 64
- File > Save Pane 157
- File > Search Path 48, 51, 59, 61, 102
- File > Signals 57
- Group > Attach Subsets 134
- Group > Control > Go 221
- Group > Delete 113, 234
- Group > Edit 254, 275
- Group > Go 101, 138, 226, 349
- Group > Halt 138, 229
- Group > Hold 222
- Group > Next 138
- Group > Release 222
- Group > Restart 235
- Group > Run To 138
- Group > Share > Halt 220
- Group > Step 138
- Group > Workers > Go 225
- group or process 138
- input and output files 62
- interrupting 195
- mpirun 98, 104
- poe 93, 99
- Process > Create 226
- Process > Detach 53
- Process > Go 97, 98, 101, 103, 105, 138, 226
- Process > Halt 138, 220, 229
- Process > Hold 222
- Process > Next 227
- Process > Next Instruction 227
- Process > Out 242
- Process > Run To 227, 241
- Process > Startup 43
- Process > Startup Parameters 62, 63
- Process > Step 227
- Process > Step Instruction 227
- Process Startup Parameters, Environment Page 70
- prun 103
- pvm 122, 124
- Quit Debugger 46

- remsh 84
- rsh 84, 100
- server launch, arguments 80
- Set Signal Handling Mode 123
- single-stepping 229
- Startup 43
- step 4
- Thread > Continuation Signal 52, 233
- Thread > Go 226
- Thread > Hold 222
- Thread > Set PC 237
- Tools > Call Tree 159
- Tools > Command Line 196
- Tools > Create Checkpoint 235
- Tools > Evaluate 165, 170, 371
- Tools > Memory Statistics 161
- Tools > Message 107
- Tools > Message Queue 108, 109
- Tools > P/T Set Browser 271
- Tools > PVM Tasks 126
- Tools > Restart 235
- Tools > Statistics 331
- Tools > Thread Objects 317
- Tools > Variable Browser 283
- Tools > Visualize 10, 169
- Tools > Watchpoint 10, 368
- totalview 41, 96, 100, 104
 - core files 41, 53
 - totalviewcli 41, 42, 104
 - tvdsvr 73
 - launching 80
- View > Assembler > By Address 217
- View > Assembler > Symbolically 217
- View > Dive Anew 283
- View > Dive In All 293
- View > Dive Thread 317
- View > Dive Thread New 317
- View > Graph (Visualizer) 172
- View > Laminate > None 333
- View > Laminate > Process 333
 - View > Laminate > Thread 333
 - View > Lookup Function 212, 213, 215, 216
 - View > Lookup Variable 282, 286, 289, 312, 322
 - View > Reset 214, 216
 - View > Reset (Visualizer) 176, 179
 - View > Sort > Ascending 330
 - View > Sort > Descending 330
 - View > Sort > None 330
 - View > Source As > Assembler 216
 - View > Source As > Both 216, 237
 - View > Source As > Source 216
 - View > Surface (Visualizer) 172
 - View > Variable 117
 - View > Lookup 125
 - Visualize 10
 - visualize 167, 180
 - Window > Duplicate 154, 292
 - Window > Duplicate Base 154, 292
 - Window > Memorize 155
 - Window > Memorize All 155
 - Window > Update 103, 221, 233
 - Windows > Update (PVM) 126
- common block
 - displaying 310
 - diving on 310
 - members have function scope 310
- compiled expressions 358, 359
 - allocating patch space for 360
 - performance 359
- compilers
 - mpcc_r 109
 - mpxlf_r 109
 - mpxlf90_r 109
- compiling
 - considerations 40
 - g compiler option 40
 - multiprocess programs 40
 - O option 40
 - optimization 40
 - programs 40
- compiling programs 3
- completion rules for arena
 - specifiers 267
- compound objects 301
- conditional breakpoints 354, 356, 375
- conditional watchpoints, *see* watchpoints
- conf file 82
- configure command 92
- configuring the Visualizer 165
- connection for serial line 87
- connection timeout 76, 78
 - altering 74
- connection timeout, bulk server
 - launch 78
- contained functions 312
 - displaying 313
- context menus 146
- continuation signal 234
 - clearing 234
- Continuation Signal command 52, 233
- continuing with a signal 233
- continuous execution 195
- contour lines 179
- contour settings 177
- control groups 26, 223
 - defined 25
 - discussion 223
 - overview 254
 - specifier for 255
- control in parallel environments 206
- control in serial environments 206
- control registers 237
 - interpreting 237
- controlling program execution 206
- conversion rules for filters 325
- Copy command 156
- copying 156

- copying between windows 156
 - core dump, naming the signal
 - that caused 53
 - core files
 - can only debug local 53
 - debugging 42
 - examining 53
 - in totalview command 41, 53
 - loading 47
 - correcting programs 358
 - count array statistic 331
 - \$count built-in function 375
 - \$countall built-in function 375
 - countdown breakpoints 356, 375
 - \$countthread built-in function 375
 - CPU registers 237
 - cpu_use option 99
 - Create Checkpoint command 235
 - creating groups 29, 226
 - creating new processes 199
 - creating process (without starting it) 226
 - creating processes 61, 225
 - and starting them 225
 - using Step 227
 - without starting them 226
 - creating threads 20
 - cr0.o module 125
 - Ctrl+C 195
 - current focus 272
 - current location of program
 - counter 150
 - current queue state 107
 - current set indicator 250, 271
 - current stack frame 216
 - current working directory 60, 61
 - Cut command 156
- D**
- D control group specifier 255
 - dactions command 338
 - load 101, 370
 - save 101, 370
 - daemons 18, 20
 - dassign command 296
 - data
 - editing 7
 - examining 6
 - filtering 8
 - slicing 8
 - viewing, from Visualizer 167
 - data assembler pseudo op 383
 - data filtering, *see* arrays
 - filtering
 - data pane, laminated 335
 - data precision, changing display 65
 - data segment 162
 - data segment memory 162
 - data types
 - see also* TotalView data types
 - C++ 307
 - changing 297
 - changing class types in C++ 309
 - for visualization 167
 - int 299
 - int* 299
 - int|| 299
 - opaque data 306
 - pointers to arrays 299
 - predefined 302
 - to visualize 167
 - data watchpoints, *see* watchpoints
 - data window (Visualizer) 172
 - display commands 173
 - scaling 176
 - translating 176
 - zooming 176
 - data window, *see* Variable Window
 - data_format variables 281
 - dataset
 - defined for Visualizer 167
 - deleting 172
 - selecting 171
 - showing parameters 179
 - dattach
 - mprun command 106
 - dattach command 42, 47, 49, 51, 53, 94, 102, 103, 206
 - dbarrier command 350, 352
 - e 356
 - stop_when_hit 140
 - dbfork library 41, 347
 - linking with 41
 - dbreak command 340, 342, 347
 - e 356
 - dcheckpoint command 235
 - ddelete command 113, 342, 344, 353
 - ddetach command 53
 - ddisable command 343, 344, 354
 - ddown command 231
 - deadlocks 244
 - message passing 109
 - \$debug assembler pseudo op 382
 - debug, using with MPICH 113
 - debugger initialization 197
 - debugger PID 205
 - debugger server 73
 - see also*, tvdsvr
 - starting manually 79
 - Debugger Unique ID (DUID) 374
 - debugging
 - executable file 41
 - multiprocess programs 41
 - not compiled with –g 40
 - OpenMP applications 113
 - programs that call execve 41
 - programs that call fork 41
 - PVM applications 121
 - QSW RMS 103
 - SHMEM library code 128
 - debugging a core file 42
 - debugging Fortran modules 312
 - debugging on a remote host 48
 - debugging over a serial line 86
 - debugging PE applications 99
 - debugging PVM applications 122
 - debugging session 206
 - debugging techniques 34
 - declared arrays, displaying 306
 - def assembler pseudo op 383
 - default address range conflicts 361
 - default control group specifier 255
 - default focus 263
 - default process/thread set 248

- default text editor 214
- default width specifier 251
- deferred shape array
 - definition 320
 - types 314
- Delete All command 344
- Delete command 139, 156, 234
- Delete command (Visualizer) 172
- Delete, in Data Window 172
- deleting
 - action points 344
 - datasets 172
 - groups 275
 - processes 355
 - programs 234
- denable command 344
- denorm filter 327
- denormalized count array
 - statistic 332
- DENORMs 324
- deprecated X defaults 69
- dereferencing
 - automatic 287
 - controlling 65
 - pointers 287
- Descending command 330
- Detach command 53
- detaching 135
- detaching from processes 52
- detaching removes all
 - breakpoints 53
- determining scope 240
- dfocus command 229, 248
 - example 249
- dgo command 97, 101, 105, 138, 226, 268
- dgroups command
 - add 262
 - add command 254
 - remove 35
- dhalt command 138, 220, 229
- dhold command 222, 351
 - process 223
 - thread 223
- difference operator 270
- dimmed information, in the Root Window 232
- Dimmed Process Information in the Root Window figure 233
- directories, setting order of search 59
- Directory command (Visualizer) 172
- directory search path 123
- Directory Window, menu commands 172
- Directory, in Data Window 172
- disabling
 - action points 343
 - autolaunch 74, 84
 - autolaunch feature 76
 - visualization 165
- disassembled machine code 213
 - in variable window 307
- discard dive stack 214
- discard mode for signals 59
- discarding signal problem 59
- disconnected processing 18
- Display of Random Data figure 176
- displaying 153
 - areas of memory 289
 - argv array 305
 - array data 153
 - arrays 319
 - common blocks 310
 - declared and allocated arrays 306
 - Fortran data types 310
 - Fortran module data 310
 - global variables 282, 283
 - machine instructions 290, 307
 - memory 289
 - pointer 153
 - pointer data 153
 - registers 284
 - remote hostnames 147
 - stack trace pane 153
 - structs 300
 - subroutines 153
 - thread objects 317
 - typedefs 300
 - unions 301
 - variable 153
- Variable Windows 281
- Displaying a Fortran Structure figure 293
- displaying a process window 152
- Displaying a Union figure 301
- Displaying C Structures and Arrays figure 294
- Displaying C++ Classes that Use Inheritance figure 308
- Displaying Long STL Names figure 288
- displaying long variable names 286
- Displaying Scoped Variables figure 283
- distributed debugging
 - see also* PVM applications
 - remote server 73
- dive icon 154, 291
- Dive In All command 293
- dive stack 292
 - retaining 154
- Dive Thread command 317
- Dive Thread New command 317
- dividing work up 18
- diving 102, 108, 146, 152
 - defined 7
 - from groups page 224
 - in a laminated pane 334
 - in a variable window 291
 - in source code 214
 - into a pointer 153, 291
 - into a process 153
 - into a stack frame 153
 - into a structure 291
 - into a thread 153
 - into a variable 7, 153
 - into an array 291
 - into formal parameters 284
 - into Fortran common blocks 310
 - into function name 214
 - into global variables 282, 283
 - into local variables 284
 - into MPI buffer 111
 - into MPI processes 111
 - into parameters 284

- into pointer 153
 - into processes 50, 153
 - into PVM tasks 126
 - into registers 284
 - into routines 153
 - into the PC 290
 - into threads 150, 153
 - into variables 153
 - nested 153
 - nested dive defined 291
 - using middle mouse button 156
 - Diving into Common Block List in Stack Frame Pane figure 311
 - Diving into Local Variables and Registers figure 285
 - dkill command 139, 199, 206, 234
 - dlist command 125
 - dload command 42, 47, 48, 49, 80, 198, 206
 - returning process ID 200
 - DMPI 109
 - dmpirun command 96, 97
 - dmstat command 162
 - dnex command 138, 227, 231
 - dnex command 227, 231
 - double assembler pseudo op 383
 - dout command 232, 242
 - dpid 205
 - dprint command 117, 119, 213, 238, 282, 286, 289, 301, 306, 310, 311, 313, 320, 323
 - dptsets command 55, 224
 - DPVM
 - see also* PVM
 - enabling support for 124
 - must be running before TotalView 124
 - starting session 123
 - dpvm command 124
 - drawing options 175
 - drerun command 198, 235
 - redirecting I/O 63
 - drestart command 235
 - drun command 198, 201, 202
 - redirecting I/O 63
 - dset command 201, 203
 - dstatus command 55, 232, 353
 - dstep command 227, 230, 241, 250, 252, 268
 - dstep commands 138
 - dstepi command 227, 230
 - DUID 374
 - of process 374
 - \$duid built-in variable 374
 - dunhold command 222, 351
 - thread 223
 - dunset command 202
 - duntil command 227, 231, 241, 244
 - dup command 231
 - dup commands 286
 - Duplicate Base command 154, 292
 - Duplicate command 154, 292
 - dwhere command 251, 268, 286
 - dynamic call tree 159
 - dynamic patch space allocation 360
 - dynamically linked, stopping after start() 125
- ## E
- E state 55
 - Edit > Copy command 156
 - Edit > Cut command 156
 - Edit > Delete command 156
 - Edit > Find Again command 212
 - Edit > Find command 5, 212
 - Edit > Find Dialog Box figure 154, 212
 - Edit > Paste command 156
 - edit mode 146
 - Edit Source command 214, 218
 - editing
 - addresses 306
 - compound objects or arrays 301
 - laminated pane 335
 - source text 218
 - text 156
 - type strings 297
 - Editing argv figure 305
 - Editing Cursor figure 156
 - EDITOR environment variable 214
 - editor launch string 218
 - effects of parallelism on debugger behavior 204
 - Enable action point 344
 - Enable Single Debug Server Launch check box 84
 - Enable Visualizer Launch check box 165
 - enabling
 - action points 344
 - enabling action points 344
 - Environment Page 70
 - environment variables 70
 - adding 70
 - before starting poe 99
 - EDITOR 214
 - how to enter 70
 - LC_LIBRARY_PATH 45
 - LM_LICENSE_FILE 45
 - MP_ADAPTER_USE 99
 - MP_CPU_USE 99
 - MP_EUIDEVELOP 112
 - PATH 51, 59, 60
 - SHLIB_PATH 45
 - TOTALVIEW 93, 139
 - TVDSVRLAUNCHCMD 80
 - equiv assembler pseudo op 383
 - error operators 270
 - error state 55
 - errors, in multiprocess program 58
 - EVAL icon 146
 - for evaluation points 146
 - eval points
 - see* evaluation points
 - Evaluate command 165, 170, 371, 373
 - evaluating an expression in a watchpoint 363
 - evaluating expressions 371
 - evaluating state 207
 - evaluation points 5, 354
 - assembler constructs 380
 - C constructs 377
 - clearing 146
 - commands 375

- defined 207, 338
 - defining 354
 - examples 356
 - Fortran constructs 378
 - listing 152
 - lists of 152
 - machine level 355
 - patching programs 6
 - printing from 5
 - saving 355
 - setting 146, 189, 356
 - using \$stop 5
 - where generated 355
 - event log window 71
 - event points listing 152
 - examining
 - core files 53
 - data 6
 - process groups 224
 - processes 223
 - source and assembler code 216
 - stack trace and stack frame 284
 - status and control registers 237
 - Example of Control Groups and Share Groups figure 224
 - exception enable modes 238
 - executable, specifying name in scope 295
 - EXECUTABLE_PATH variable 48, 51
 - setting 185
 - executables
 - debugging 41
 - loading 47
 - executing
 - out of function 232
 - to the completion of a function 232
 - executing a startup file 44
 - execution
 - controlling 206
 - resuming 221
 - execution models 12
 - execve() 41, 223, 347, 348
 - attaching to processes 49
 - debugging programs that call 41
 - setting breakpoints with 348
 - existent operator 270
 - exit CLI command 46
 - Exit command 46
 - Exit command (Visualizer) 172
 - exiting TotalView 46
 - expression evaluation window
 - compiled and interpreted expressions 358
 - discussion 371
 - expressions 270, 347
 - can contain loops 371
 - compiled 359
 - evaluating 371
 - performance of 359
 - extent of arrays 300
- F**
- figures
 - Action Point > Properties Dialog Box 343, 346, 351
 - Action Point Properties and Address Dialog Boxes 228
 - Action Point Properties Dialog Box 6
 - Action Point Symbol 339
 - Address Only (Absolute Addresses) 217
 - Ambiguous Addresses Dialog Box 229
 - Ambiguous Function Dialog Box 213, 342
 - Ambiguous Line Dialog Box 341
 - Array Data Filter by Range of Values 328
 - Array Data Filtering by Comparison 326
 - Array Data Filtering for IEEE Values 328
 - Array Statistics Window 331
 - Array Visualization 10
 - Assembler Only (Symbolic Addresses) 218
 - Attached Page Showing
 - Process and Thread Status 54, 58
 - Both Source and Assembler (Symbolic Addresses) 218
 - Breakpoint at Assembler Instruction Dialog Box 345
 - C++ Type Cast to Base Class Question Window 309
 - C++ Type Cast to Derived Class Question Window 309
 - Casting Code 291
 - Checkpoint and Restart Dialog Boxes 236
 - CLI xterm Window 196
 - Dimmed Process Information in the Root Window 233
 - Display of Random Data 176
 - Displaying a Fortran Structure 293
 - Displaying a Union 301
 - Displaying C Structures and Arrays 294
 - Displaying C++ Classes that Use Inheritance 308
 - Displaying Long STL Names 288
 - Displaying Scoped Variables 283
 - Diving into Common Block List in Stack Frame Pane 311
 - Diving into Local Variables and Registers 285
 - Edit > Find Dialog Box 154, 212
 - Editing argv 305
 - Editing Cursor 156
 - Example of Control Groups and Share Groups 224
 - File > Exit Dialog Box 46
 - File > New Program 51
 - File > New Program Dialog Box Page 47
 - File > Preferences
 - Parallel Page 136

- File > Preferences Dialog Box:
 - Action Points Page 65
- File > Preferences Dialog Box:
 - Bulk Launch Page 66
- File > Preferences Dialog Box:
 - Dynamic Libraries Page 67
- File > Preferences Dialog Box:
 - Fonts Page 68
- File > Preferences Dialog Box:
 - Launch Strings Page 66
- File > Preferences Dialog Box:
 - Options Page 64
- File > Preferences Dialog Box:
 - Parallel Page 67
- File > Preferences Dialog Box:
 - Pointer Dive Page 69
- File > Preferences Launch
 - Strings Page 166
- File > Preferences: Action
 - Points Page 348
- File > Preferences: Bulk
 - Launch Page 77
- File > Preferences: Formatting
 - Page 280
- File > Preferences: Server
 - Launch Strings Page 75
- File > Save Pane Dialog Box
 - 157
- File > Search Path Dialog Box
 - 60
- Five Processes and Their
 - Groups on Two
 - Computers 29
- Five Processors and Processor
 - Groups (Part 1) 27
- Five Processors and Processor
 - Groups (Part 2) 28
- Fortran 90 Pointer Value 316
- Fortran 90 User-Defined Type
 - 314
- Fortran Array with Inverse
 - Order and Limited Extent
 - 322
- Fortran Modules Window 313
- Four Processor Computer 21
- Four-Processor Computer
 - Networks 22
- Global Variables Window 287
- Graph Options Dialog Box 175
- Group > Attach Subset Dialog
 - Box 134
- Group > Edit Group 275
- Laminated Array and Structure
 - 335
- Laminated Scalar Variable 334
- Laminated UPC Variable
 - Window 132
- Laminated Variable Window 12
- List and Vector
 - Transformations 279
- Mail with Daemon 19
- Manual Launching of
 - Debugger Server 80
- Message Queue Graph 13
- Message Queue Graph window
 - 107
- Message Queue Window 110
- Message Queue Window
 - Showing Pending Receive
 - Operation 111
- More Conditions 7
- Nested Dive 153
- Nested Dives 292
- OpenMP Shared Variable 118
- OpenMP Stack Parent Token
 - Line 120
- OpenMP THREADPRIVATE
 - Common Block Variable
 - 120
- P/T Set Browser Window 272
- P/T Set Browser Windows (Part
 - 1) 273
- P/T Set Browser Windows (Part
 - 2) 274
- P/T Set Control in the Process
 - Window 246
- P/T Set Control in the Tools >
 - Evaluate Window 246
- PC Arrow Over a Stop Icon 346
- Pointer to a Shared Variable
 - 133
- Process > Startup Parameters
 - Dialog Box
 - Arguments Page 62
- Process > Startup Parameters
 - Dialog Box: Environment
 - Page 71
- Process > Startup Parameters
 - Dialog Box: Standard I/O
 - page 63
- Process and Thread Labels in
 - the Process Window 55
- Process and Thread Switching
 - Icons 12
- Process Window 4, 151
- Process Window Tag Field 152
- Program and Daemons 18
- Program Browser and Variables
 - Window (Part 2) 285
- PVM Tasks and Configuration
 - Window 127
- Resizing (and Sometimes Its
 - Consequences) 155
- Resolving Ambiguous Function
 - Names Dialog Box 215
- Root Widow: Group Page 225
- Root Window 11
- Root Window Attached Page
 - 147
- Root Window Groups Page 149
- Root Window Log Page 72, 149
- Root Window Showing Process
 - and Thread Status 86
- Root Window Showing Remote
 - 148
- Root Window Unattached Page
 - 148
- Root Window: Unattached
 - Page 95
- Root Window's Group Page 14
- Rotating and Querying 174
- Sample OpenMP Debugging
 - Session 116
- Sample Visualizer Data
 - Windows 173
- Sample Visualizer Window 171
- Select Directory Dialog Box 61
- SHMEM Sample Session 129

- Sliced UPC Array 131
- Sorted Variable Window 330
- Startup and Initialization Sequence 44
- Step 1: A Program Starts 29
- Step 2: Forking a Process 30
- Step 3: Exec'ing a Process 31
- Step 5: Creating a Second Version 31
- Step 6: Creating a Remote Process 32
- Step 7: A Thread is Created 33
- Stop Before Going Parallel Question Dialog Box 136
- Stopped Execution of Compiled Expressions 360
- Surface Options Dialog Box 178
- The CLI and TotalView 194
- Thread > Continuation Signal Dialog Box 52, 234
- Thread Objects Page on an IBM AIX machine 318
- Threads 20, 23
- Three Dimensional Array Sliced to Two Dimensions 167
- Three Dimensional Surface Visualizer Data Display 178
- Toolbar 219
- Toolbar with Pulldown 14
- Tools > Call Tree Dialog Box 160
- Tools > Evaluate Dialog Box 372, 373
- Tools > Memory Usage Window 161, 163
- Tools > Watchpoint Dialog Box 365
- TotalView Debugging Session Over a Serial Line 87
- TotalView Visualizer Connection 165
- TotalView Visualizer Relationships 164
- Two Computers Working on One Problem 19
- Two Dimensional Surface Visualizer Data Display 177
- Two More Variable Windows 9
- Two Variable Windows 9
- Unattached Page 50
- Undive/Dive Controls 214
- Uniprocessor 18
- UPC Laminated Variable 133
- UPC Variable Window Showing Nodes 131
- User Threads and Service Threads 24
- User, Service, and Manager Threads 24
- Using an Expression to Change a Value 297
- Using Assembler 381
- Variable Window 168
- Variable Window for a Global Variable 282
- Variable Window for Area of Memory 289
- Variable Window for small_array 323
- Variable Window with Machine Instructions 290
- View > Lookup Function Dialog Box 214, 215
- View > Lookup Variable Dialog Box 213
- Visualizer Graph Data Window 175
- Waiting to Complete Message Box 372
- Width Specifiers 252
- Zooming, Rotating, About an Axis 181
- file
 - for start up 44
- File > Close command 154, 290
- File > Close command (Visualizer) 172
- File > Close Relatives command 154
- File > Close Similar command 154, 290
- File > Delete command (Visualizer) 172
- File > Directory command (Visualizer) 172
- File > Edit Source command 214, 218
- File > Exit command 46
- File > Exit command (Visualizer) 172
- File > Exit Dialog Box figure 46
- File > New Base Window (Visualizer) 172
- File > New Program command 42, 46, 49, 51, 53, 80, 84, 88
- File > New Program Dialog Box figure 47, 51
- File > Options command (Visualizer) 172, 174
- File > Preferences 64
 - Action Points page 59, 64, 137
 - Bulk Launch page 65, 76, 79
 - Dynamic Libraries page 65
 - Fonts page 65
 - Formatting page 65
 - Launch Strings Page 74
 - Launch Strings page 65, 219
 - Options page 58
 - Parallel Page 135
 - Parallel page 65
 - Pointer Dive Page 65
- File > Preferences Dialog Box: Action Points Page figure 65
- File > Preferences Dialog Box: Bulk Launch Page figure 66
- File > Preferences Dialog Box: Dynamic Libraries Page figure 67
- File > Preferences Dialog Box: Fonts Page figure 68
- File > Preferences Dialog Box: Launch Strings Page figure 66
- File > Preferences Dialog Box: Options Page figure 64

- File > Preferences Dialog Box:
 - Parallel Page figure 67
 - File > Preferences Dialog Box:
 - Pointer Dive Page figure 69
 - File > Preferences Launch Strings
 - Page figure 166
 - File > Preferences: Action Points
 - Page figure 348
 - File > Preferences: Bulk Launch
 - Page figure 77
 - File > Preferences: Formatting
 - Page figure 280
 - File > Preferences: Parallel Page
 - figure 136
 - File > Preferences: Server Launch
 - Strings Page figure 75
 - File > Save Pane command 157
 - File > Save Pane Dialog Box
 - figure 157
 - File > Search Path command 48,
 - 51, 59, 61, 102
 - search order 59
 - File > Search Path Dialog Box
 - figure 60
 - File > Signals command 57
 - file command-line option to
 - Visualizer 167, 180
 - files
 - .rhosts 100
 - hosts.equiv 100
 - fill assembler pseudo op 383
 - filter expression, matching 324
 - filtering 8
 - array data 324
 - array expressions 329
 - by comparison 325
 - conversion rules 325
 - example 326
 - IEEE values 326
 - in sorts 330
 - options 324
 - ranges of values 327
 - filters
 - \$denorm 327
 - \$inf 327
 - \$nan 327
 - \$nanq 327
 - \$nans 327
 - \$ninf 327
 - \$pdenorm 327
 - \$pinf 327
 - Find Again command 212
 - Find command 5, 212
 - finding
 - functions 213
 - source code 213, 215
 - source code for functions 213
 - first thread indicator of < 250
 - Five Processes and Their Groups
 - on Two Computers figure 29
 - Five Processors and Processor
 - Groups (Part 1) figure 27
 - Five Processors and Processor
 - Groups (Part 2) figure 28
 - float assembler pseudo op 383
 - focus
 - changing 248
 - pushing 249
 - restoring 249
 - for loop 371
 - Force window positions (disables
 - window manager placement
 - modes) check box 155
 - fork() 41, 223, 347
 - debugging programs that call
 - 41
 - setting breakpoints with 348
 - fork_loop.tvd example program
 - 197
 - Fortran
 - array bounds 299
 - arrays 300
 - common blocks 310
 - contained functions 312
 - data types, displaying 310
 - debugging modules 312
 - deferred shape array types 314
 - filter expression 329
 - in code fragment 354
 - in evaluation points 378
 - module data, displaying 310
 - modules 310, 312
 - pointer types 315
 - type strings supported by
 - TotalView 298
 - user defined types 314
 - Fortran 90 Pointer Value figure
 - 316
 - Fortran 90 User-Defined Type
 - figure 314
 - Fortran Array with Inverse Order
 - and Limited Extent figure
 - 322
 - Fortran Modules command 312
 - Fortran Modules Window figure
 - 313
 - forward icon 154
 - four linked processors 21
 - 4142 default port 79
 - Four Processor Computer figure
 - 21
 - Four-Processor Computer
 - Networks figure 22
 - frame pointer 231
 - function visualization 159
 - functions
 - finding 213
 - locating 212
 - returning from 232
- ## G
- g compiler option 40, 153
 - g width specifier 257, 263
 - generating a symbol table 40
 - global assembler pseudo op 383
 - global variables
 - changing 227
 - displaying 227
 - diving into 282, 283
 - Global Variables Window figure
 - 287
 - Go command 5, 97, 101, 103,
 - 105, 138, 225
 - GOI defined 239
 - goto statements 355
 - Graph command (Visualizer) 172
 - Graph Data Window 173
 - graph markers 174
 - Graph Options Dialog Box figure
 - 175

- Graph visualization menu 171
 - graph window, creating 172
 - Graph, in Directory Window 172
 - graphs
 - manipulating, in Visualizer 176
 - two dimensional 173
 - group
 - process 245
 - thread 245
 - Group > Attach Subset Dialog
 - Box figure 134
 - Group > Attach Subsets
 - command 134
 - Group > Control > Go command 221
 - Group > Delete command 113, 139, 234
 - Group > Edit command 254
 - Group > Edit Group command 275
 - Group > Edit Group figure 275
 - Group > Go command 101, 138, 226, 349
 - Group > Halt command 138, 229
 - Group > Hold command 222
 - Group > Next command 138
 - Group > Release command 222
 - Group > Restart command 235
 - Group > Run To command 138
 - Group > Share > Halt command 220
 - Group > Step command 138
 - Group > Workers > Go
 - commands 225
 - group aliases 203
 - limitations 204
 - group commands 138
 - group name 256
 - group number 256
 - group stepping 242
 - group syntax 255
 - group number 256
 - naming names 256
 - predefined groups 255
 - GROUP variable 262
 - group width specifier 251
 - group_indicator
 - defined 255
 - groups 122
 - see also* processes and barriers 13
 - behavior 242
 - changing 275
 - creating 29, 226
 - defined 25, 26
 - deleting 275
 - examining 223
 - holding processes 222
 - listing 149
 - named 275
 - overview 25
 - process 244
 - relationships 252
 - releasing processes 222
 - running 135
 - setting 261
 - starting 226
 - stopping 135
 - thread 244
 - updating 275
 - Groups page 14, 149, 224
 - GUI namespace 202
- ## H
- h held indicator 222
 - h localhost option for HP MPI 98
 - half assembler pseudo op 383
 - Halt command 138, 220, 229
 - halt commands 220
 - halting
 - groups 220
 - processes 220
 - threads 220
 - handler routine 56
 - handling signals 56, 57, 123, 124
 - heap memory 162
 - held indicator 222
 - held operator 270
 - held processes
 - defined 350
 - hexadecimal address, specifying
 - in variable window 289
 - hi16 assembler operator 382
 - hi32 assembler operator 382
 - hold and release 221
 - \$hold assembler pseudo op 382
 - \$hold built-in function 376
 - Hold command 222
 - hold state 222
 - hold state, toggling 351
 - Hold Threads command 222
 - holding and advancing processes 206
 - holding threads 245
 - \$holdprocess assembler pseudo op 382
 - \$holdprocess built-in function 376
 - \$holdprocessall built-in function 376
 - \$holdprocessstopall assembler pseudo op 382
 - \$holdstopall assembler pseudo op 382
 - \$holdstopall built-in function 376
 - \$holdthread assembler pseudo op 382
 - \$holdthread built-in function 376
 - \$holdthreadstop assembler pseudo op 382
 - \$holdthreadstop built-in function 376
 - \$holdthreadstopall assembler pseudo op 382
 - \$holdthreadstopall built-in function 376
 - \$holdthreadstopprocess assembler pseudo op 382
 - \$holdthreadstopprocess built-in function 376
 - hostname
 - abbreviated in Root Window 147
 - for tvdsvr 43
 - in square brackets 147
 - hosts.equiv file 100
 - how TotalView determines share group 225
 - hung processes 49

I

- I state 56
- IBM MPI 99
- IBM SP machine 92, 93
- idle state 56
- Ignore mode warning 59
- ignoring action points 344
- implicitly defined process/thread set 248
- incomplete arena specifier 267
- inconsistent widths 267
- indicator 50
- inf filter 327
- infinite loop, *see* loop, infinite
- infinity count array statistic 332
- INFs 324
- initial process 204
- initialization search paths 44
- initialization subdirectory 43
- initializing an array slice 187
- initializing debugging state 44
- initializing the CLI 197
- initializing TotalView 43
- input files, setting 62
- instructions
 - data type for 304
 - displaying 290, 307
- int data type 299
- int* data type 299
- int[] data type 299
- interactive CLI 193
- interface to CLI 195
- interpreted expressions 358
 - performance 359
- interrupting commands 195
- intersection operator 270
- intrinsic, *see* built-in functions
- inverting array order 321
- inverting axis 175
- invoking CLI program from shell
 - example 197
- IP over the switch 99
- iterating
 - over a list 268
 - over arenas 250

K

- K state, unviewable 55
- KeepSendQueue command-line option 112
- kernel 55
- killing processes when exiting 49
- killing programs 234
- ksq command-line option 112

L

- L lockstep group specifier 256, 257
- labels, for machine instructions 307
- Laminate > None command 333
- Laminate > Process command 333
- Laminate > Thread command 333
- Laminate None command 333
- Laminate Thread command. 119
- Laminated Array and Structure figure 335
- laminated data view 12
- Laminated Scalar Variable figure 334
- Laminated UPC Variable Window figure 132
- Laminated Variable Window figure 12
- laminating data pane 335
- lamination
 - arrays and structures 334
 - data panes and Visualizer 169
 - diving in pane 334
 - editing a pane 335
 - variables 333
- launch
 - configuring Visualizer 165
 - options for Visualizer 165
 - TotalView Visualizer from command line 180
 - tvdsrv 73
- Launch Strings Page 74, 84, 165, 219
- launching processes 135
- lcomm assembler pseudo op 383

- LD_LIBRARY_PATH environment variable 45
- left margin area 150
- left mouse button 145
- libraries
 - dbfork 41
 - debugging SHMEM library code 128
- limiting array display 321
- line most recently selected 233
- line number area 146
- line numbers 150
- line numbers for specifying blocks 295
- LINES_PER_SCREEN variable 201
- List and Vector Transformations figure 279
- lists of processes 146
- lists with inconsistent widths 267
- lists, iterating over 268
- LM_LICENSE_FILE environment variable 45
- lo16 assembler operator 382
- lo32 assembler operator 382
- loading
 - core file 47
 - file into TotalView 42
 - new executables 46, 47
 - programs 42
 - remote executables 48
- local hosts 43
- locations, toggling breakpoints at 340
- lockstep group 27, 240, 249
 - defined 26
 - L specifier 256
 - number of 254
 - overview 254
- Log page 71, 149
- long variable names, displaying 286
- \$long_branch assembler pseudo op 382
- Lookup Function command 125, 212, 213, 215, 216
- Lookup Variable command 119, 212, 282, 286, 289, 312

- specifying slides 322
- loop infinite, *see* infinite loop
- lower adjacent array statistic 332
- lower bounds 299
 - non default 300
 - of array slices 320
- lysm TotalView pseudo op 383

M

- M state 55
- machine instructions
 - data type 304
 - data type for 304
 - displaying 290, 307
- Mail with Daemons figure 19
- main() 125
 - stopping before entering 125
- make_actions.tcl sample macro 190, 197
- manager threads 23, 28
- manual hold and release 221
- Manual Launching of Debugger Server figure 80
- manually starting tvdsrv 84
- markers, in graphs 174
- master process, recreating slave processes 139
- master thread 115
 - OpenMP 115, 120
 - stack 117
- matching processes 244
- matching stack frames 333
- maximum array statistic 332
- mean array statistic 332
- median array statistic 332
- Memorize All command 155
- Memorize command 155
- memory
 - analyzing 161
 - data segment 162
 - displaying areas of 289
 - heap 162
 - stack 162
 - text segment 162
 - virtual stack 162
- memory locations, changing values of 296
- Memory Statistics command 161
- menus, context 146
- mesh, drawing as 178
- message passing deadlocks 109
- Message Passing
 - Interface/Chameleon Standard, *see* MPICH
- Message Passing Toolkit 109
- Message Queue command 108, 109
- message queue display 104, 113
- Message Queue Graph 108
 - diving 108
 - rearranging shape 108
 - updating 108
- message queue graph 12
- Message Queue Graph command 107
- Message Queue Graph figure 13
- Message Queue Graph window figure 107
- Message Queue Window figure 110
- Message Queue Window Showing Pending Receive Operation figure 111
- message states 107
- message tags, reserved 127
- message-passing programs 137
- messages
 - envelope information 112
 - operations 110
 - reserved tags 127
 - unexpected 112
- messages from TotalView, saving 201
- middle mouse button 145
- middle mouse dive 156
- minimum array statistic 332
- missing TID 251
- mixed state 55
- mixing arena specifiers 268
- modify watchpoints, *see* watchpoints
- modifying code behavior 355
- module data definition 311
- modules 310, 312

- debugging, Fortran 312
 - displaying Fortran data 310
- monitoring TotalView sessions 71
- More Conditions figure 7
- more processing 201
- more prompt 201
- mouse button
 - diving 145
 - left 145
 - middle 145
 - right 146
 - selecting 145
- mouse buttons, using 145
- MP_ADAPTER_USE environment variable 99
- MP_CPU_USE environment variable 99
- MP_EUIDEVELOP environment variable 112
- mpicc_r compilers 109
- MPI
 - attaching to 105
 - attaching to HP job 98
 - attaching to running job 97
 - buffer diving 111
 - communicators 109
 - library state 109
 - on HP Alpha 96
 - on HP machines 97
 - on IBM 99
 - on SGI 104
 - process diving 111
 - processes, starting 103
 - starting on HP Alpha 96
 - starting on SGI 104
 - starting processes 97, 105
 - troubleshooting 113
- MPI program
 - toolbar for 14
- MPI_Init() 101, 109
 - breakpoints and timeouts 140
- MPI_Iprobe() 112
- MPI_Recv() 112
- MPICH 92, 93
 - and SIGINT 113
 - and the TOTALVIEW environment variable 93

- attach from TotalView 94
 - attaching to 94
 - ch_lfshmem device 92, 95
 - ch_mpl device 92
 - ch_p4 device 92, 95
 - ch_shmem device 95
 - ch_smem device 92
 - configuring 92
 - Debugging Tips 139
 - diving into process 94
 - MPICH/ch_p4 140
 - mpirun command 92, 93
 - naming processes 96
 - obtaining 92
 - P4 96
 - p4pg files 96
 - starting TotalView using 92
 - using –debug 113
 - MPICH –tv command-line option 92
 - mpirun command 22, 92, 93, 98, 104, 139
 - examples 98
 - for HP MPI 98
 - options to TotalView through 139
 - passing options to 139
 - mpirun process 105
 - MPL_Init() 101
 - and breakpoints 101
 - mprun command 106
 - mpxlf_r compiler 109
 - mpxlf90_r compiler 109
 - MOD, *see* message queue display
 - multiple classes, resolving 215
 - Multiple indicator 334
 - multiple sessions 121
 - multiprocess debugging 11
 - multiprocess programming library 41
 - multiprocess programs
 - and signals 58
 - attaching to 52
 - compiling 40
 - process groups 223
 - setting and clearing breakpoints 346
 - multiprocessing 22
 - multithreaded debugging 11
 - multithreaded signals 234
- N**
- n option, of rsh command 85
 - n single process server launch command 81
 - named groups 149, 275
 - named sets 275
 - names of processes in process
 - groups 224
 - namespaces 202
 - TV:: 202
 - TV::GUI:: 202
 - naming MPICH processes 96
 - naming rules
 - for control groups 224
 - for share groups 224
 - nan filter 327
 - nanq filter 327
 - NaNs 324, 326
 - array statistic 332
 - nans filter 327
 - navigating
 - source code 216
 - ndenorm filter 327
 - nested dive 153
 - defined 291
 - window 291, 292
 - Nested Dive figure 153
 - Nested Dives figure 292
 - nested stack frame
 - running to 245
 - New Base Window
 - in Data Window 172
 - New Base Window command (Visualizer) 172
 - New Program command 42, 46, 49, 51, 53, 80, 84, 88
 - Next command 138, 227
 - "next" commands 231
 - Next Instruction command 227
 - \$nid built-in variable 374
 - ninf filter 327
 - no_stop_all command-line option 139
 - node ID 374
 - nodes, attaching from to poe 102
 - nodes, detaching 135
 - None (laminar) command 333
 - None (sort) command 330
 - nonexistent operators 270
 - non-sequential program execution 195
- O**
- O option 40
 - offsets, for machine instructions 307
 - \$oldval built-in variable 374
 - omitting array stride 321
 - omitting components in creating scope 296
 - omitting period in specifier 267
 - omitting width specifier 267
 - <opaque> data type 306
 - opaque type definitions 306
 - Open process window at
 - breakpoint check box 59
 - Open process window on signal check box 58
 - OpenMP 113, 115
 - debugging 114
 - debugging applications 113
 - master thread 115, 118, 120
 - master thread stack context 117
 - on HP Alpha 115
 - private variables 117
 - runtime library 114
 - shared variables 117, 120
 - stack parent token 120
 - THREADPRIVATE common blocks 118
 - THREADPRIVATE variables 119
 - threads 115
 - viewing shared variables 118
 - worker threads 115
 - OpenMP Shared Variable figure 118
 - OpenMP Stack Parent Token Line figure 120

- OpenMP THREADPRIVATE
 - Common Block Variable figure 120
 - operators
 - difference 270
 - & intersection 270
 - | union 270
 - breakpoint 270
 - error 270
 - existent 270
 - held 270
 - nonexistent 270
 - running 270
 - stopped 270
 - unheld 270
 - watchpoint 270
 - optimizations, compiling for 40
 - options
 - for visualize 180
 - in Data Window 172
 - patch_area 361
 - patch_area_length 361
 - sb 370
 - serial 88
 - setting 69
 - surface data display 179
 - Options > Auto Visualize
 - command (Visualizer) 172
 - Options command (Visualizer) 172, 174
 - Options Page 156
 - org assembler pseudo op 383
 - ORNL PVM, *see* PVM
 - "out" commands 232
 - out command
 - goal 232
 - outliers 332, 333
 - outlined routine 114, 119, 120
 - outlining, defined 114
 - output
 - assigning output to variable 200
 - from CLI 200
 - only last command executed returned 200
 - printing 200
 - returning 200
 - when not displayed 200
 - output files, setting 62
- P**
- p width specifier 257
 - p.t notation 250
 - p/t selector 247
 - p/t set browser 271
 - P/T Set Browser command 271
 - P/T Set Browser Window figure 272
 - P/T Set Browser Windows (Part 1) figure 273
 - P/T Set Browser Windows (Part 2) figure 274
 - P/T Set Control in the Process Window figure 246
 - P/T Set Control in the Tools > Evaluate Window figure 246
 - p/t sets
 - arguments to Tcl 248
 - arranged hierarchically 272
 - browser 271
 - defined 248
 - expressions 270
 - grouping 270
 - set of arenas 250
 - syntax 251
 - visualizing 271
 - p/t syntax
 - group syntax 255
 - p4 listener process 95
 - p4pg files 96
 - p4pg option 96
 - panes
 - action points list, *see* action points list pane
 - source code, *see* source code pane
 - stack frame, *see* stack frame pane
 - stack trace, *see* stack trace pane
 - panes, saving 157
 - parallel debugging tips 134
 - PARALLEL DO outlined routine 116, 117
 - Parallel Environment for AIX, *see* PE
 - parallel environments
 - execution control 206
 - Parallel Page 135
 - parallel program, defined 204
 - parallel program, restarting 139
 - parallel region 115
 - parallel tasks, starting 101
 - Parallel Virtual Machine, *see* PVM
 - parallel_attach variable 137
 - parallel_stop variables 136
 - parsing comments example 190
 - passing arguments 42
 - passing default arguments 202
 - passing environment variables to processes 70
 - Paste command 156
 - pastng 156
 - pastng between windows 156
 - pastng with middle mouse 145
 - patch space size, different than 1MB 362
 - patch space, allocating 360
 - patch_area_base option 361
 - patch_area_length option 361
 - patchng
 - function calls 357
 - programs 357
 - PATH environment variable 48, 51, 59, 60
 - pathnames, setting in procgroup file 96
 - PC Arrow Over a Stop Icon figure 346
 - PC icon 236
 - pdenorm filter 327
 - PE 99, 102, 109
 - adapter_use option 99
 - and slow processes 140
 - applications 99
 - cpu_use option 99
 - debugng tips 140
 - from command line 100
 - from poe 100
 - options to use 99

- switch-based communication 99
- pending messages 108
- pending receive operations 110, 111
- pending send operations 110, 112
 - configuring for 112
- pending unexpected messages 110
- performance of interpreted, and compiled expressions 359
- performance of remote debugging 73
- persist command-line option to Visualizer 167, 180
- \$pid built-in variable 374
- pid specifier, omitting 267
- pid.tid to identify thread 150
- pinf filter 327
- pipe for Visualizer 164
- piping data 157
- Plant in share group check box 348, 356
- poe
 - and mpirun 93
 - and TotalView 100
 - arguments 99
 - attaching to 102
 - interacting with 140
 - on IBM SP 94
 - placing on process list 103
 - required options to 99
 - running PE 100
 - TotalView acquires poe processes 102
- POI defined 239
- point of execution for multiprocess or multithreaded program 152
- pointer data 153
- Pointer Dive page 287
- Pointer to a Shared Variable figure 133
- pointers 153
 - as arrays 287
 - chasing 287
 - dereferencing 287
 - diving on 153
 - in Fortran 315
 - to arrays 299
- pop_at_breakpoint variable 59
- pop_on_error variable 58
- popping a window 153
- port 4142 79
- port command-line option 79
- port number for tvdsvr 43
- precision 280
 - changing 280
 - changing display 65
- predefined data types 302
- preference file 44
- Preferences
 - Action Points Page 64
 - Bulk Launch Page 65, 76
 - Bulk Launch page 79
 - Dynamic Libraries Page 65
 - Fonts Page 65
 - Formatting Page 65
 - Launch Strings Page 65, 74
 - Options Page 58
 - Parallel Page 65
 - Pointer Dive Page 65
- preferences, setting 69
- preferences6.tvd file 44
- primary thread
 - stepping failure 244
- print statements, using 5
- printing an array slice 187
- printing in an eval point 5
- private variables 115
 - in OpenMP 117
- procedures
 - debugging over a serial line 86
 - displaying 306
 - displaying declared and allocated arrays 306
- process
 - detaching 52
 - holding 245
 - state 54
 - synchronization 245
- Process > Create command 226
- Process > Detach command 53
- Process > Go command 97, 98, 101, 103, 105, 138, 226
- Process > Halt command 138, 220, 229
- Process > Hold command 222
- Process > Hold Threads command 222
- Process > Next command 227
- Process > Next Instruction command 227
- Process > Out command 242
- Process > Release Threads command 222
- Process > Run To command 227, 241
- Process > Startup command 43, 63
- Process > Startup Parameters 62
 - Arguments Page 62
 - Environment Page 70
 - Standard I/O Page 63
- Process > Startup Parameters command 62, 63
- Process > Startup Parameters Dialog Box
 - Arguments Page figure 62
- Process > Startup Parameters Dialog Box: Environment Page figure 71
- Process > Startup Parameters Dialog Box: Standard I/O Page figure 63
- Process > Step command 227
- Process > Step Instruction command 227
- Process and Thread Labels in the Process Window figure 55
- Process and Thread Switching Icons figure 12
- process as dimension in Visualizer 169
- process barrier breakpoint
 - changes when clearing 354
 - changes when setting 354
 - defined 338
 - deleting 353
 - setting 351

- process DUID 374
- process groups 26, 244, 245, 253
 - behavior 261
 - behavior at goal 244
 - displaying 224
 - stepping 243
 - synchronizing 244
- process ID 374
- process numbers are unique 204
- process states 55, 150
- process states, attached 55
- process stepping 243
- process synchronization 137
- process width specifier 251
 - omitting 267
- Process Window 4, 150
 - displaying 152
 - host name in title 147
 - raising 58
- Process Window figure 4, 151
- Process Window Tag Field Area
 - figure 152
- process/set threads
 - saving 252
- process/thread identifier 204
- process/thread notation 204
- process/thread sets 205
 - as arguments 248
 - changing focus 248
 - default 248
 - examples 253
 - implicitly defined 248
 - inconsistent widths 268
 - structure of 251
 - target 248
 - widths inconsistent 268
- process_id.thread_id 250
- process_load_callbacks variable 45
- \$processduid built-in variable 374
- processes
 - see also* automatic process acquisition
 - see also* groups
 - acquiring 94, 95, 124
 - acquiring in PVM applications 122
 - acquisition in poe 102
 - apparently hung 138
 - attaching 49, 50, 147
 - attaching to 49, 50, 102, 126
 - barrier point behavior 354
 - behavior 243
 - breakpoints shared 347
 - call tree 160
 - cleanup 127
 - copy breakpoints from master process 94
 - creating 61, 225, 227
 - creating by single-stepping 227
 - creating new 199
 - creating using Go 226
 - creating without starting 226
 - deleting 234
 - deleting related 235
 - detaching from 52
 - dimmed, in the Root Window 232
 - displaying data 153
 - diving into 50, 102
 - diving on 153
 - groups 223
 - examining 224
 - held defined 350
 - holding 221, 350, 376
 - hung 49
 - initial 204
 - killing while exiting 49
 - launching 135
 - list of 146
 - loading new executables 46
 - local 50
 - master restart 139
 - MPI 111
 - names 224
 - passing environment variables to 70
 - refreshing process info 221
 - released 350
 - releasing 221, 350, 353
 - remote 50
 - restarting 235
 - single-stepping 241
 - slave, breakpoints in 94
 - spawned 204
 - starting 226
 - state 55
 - states 55
 - status of 54
 - stepping 13, 138, 243
 - stop all related 347
 - stopped 350
 - stopped at barrier point 354
 - stopping 220, 354
 - stopping all related 57
 - stopping and deleting 355
 - stopping intrinsic 376
 - stopping spawned 94
 - switching between 11
 - synchronizing 207, 244
 - terminating 199
 - types of process groups 223
 - when stopped 243
- process-level stepping 138
- processors and threads 22
- proggroup file 96
 - using same absolute path names 96
- Program and Daemons figure 18
- Program Browser and Variables Window (Part 2) figure 285
- program control groups
 - defined 254
 - naming 224
- program counter (PC) 50, 151, 152
 - arrow icon for PC 152
 - indicator 150
 - setting 236
 - setting program counter 236
 - setting to a stopped thread 237
- program execution
 - advancing 206
 - controlling 206
- program state, changing 195
- program visualization 159
- programming TotalView 14

- programs
 - compiling 3, 40
 - compiling using `-g` 40
 - correcting 358
 - deleting 234
 - killing 234
 - loading by process ID 47
 - not compiled with `-g` 40
 - patching 6, 357
 - restarting 235
- prompt and width specifier 259
- PROMPT variable 203
- Properties command 140, 338, 343, 346, 351, 356
- properties, of action points 5
- prototypes for temp files 77
- prun command 103
- pthread ID 205
- threads, *see* threads
- pushing focus 249
- PVM
 - acquiring processes 122
 - attaching procedure 126
 - attaching to tasks 126
 - automatic process acquisition 124
 - cleanup of tvdsvr 127
 - creating symbolic link to tvdsvr 122
 - daemons 127
 - debugging 121
 - message tags 127
 - multiple instances not allowed by single user 121
 - multiple sessions 121
 - running with DPVM 122
 - same architecture 126
 - search path 123
 - starting actions 124
 - tasker 124
 - tasker event 125
 - tasks 121, 122
 - TotalView as tasker 121
 - TotalView limitations 121
 - tvdsvr 125
 - Update Command 126
 - pvm command 122, 124
 - PVM groups, unrelated to process groups 122
 - PVM Tasks and Configuration Window figure 127
 - PVM Tasks command 126
 - pvm_joining() 127
 - pvm_spawn() 122, 125
 - pvmgs process 122, 127
 - terminated 127
- Q**
 - QSW RMS applications 103
 - attaching to 104
 - debugging 103
 - starting 103
 - quad assembler pseudo op 383
 - Quadrics RMS 103
 - quartiles array statistic 332
 - queue state 107
 - quitting TotalView 46
- R**
 - R state 55, 56
 - raising process window 58
 - rank for Visualizer 165
 - ranks 107
 - recursive functions 232
 - single-stepping 231
 - redirecting
 - stdin 63
 - stdout 63
 - redive icon 154, 291
 - registers
 - editing 238
 - interpreting 237
 - relatives, attaching to 52
 - Release command 222
 - release state 222
 - Release Threads command 222
 - reloading
 - breakpoints 101
 - remembering window positions 155
 - `-remote` command-line option 43
 - Remote Debug Server Launch preferences 74
 - remote debugging 73
 - see also* PVM applications
 - launching tvdsvr 73
 - performance 73
 - remote executables, loading 48
 - remote host
 - debugging on 48
 - remote hosts 43
 - remote login 100
 - `-remote` option 43, 48
 - remote shell command, changing 84
 - removing breakpoints 146
 - remsh command 84
 - used in server launches 80
 - replacing default arguments 202
 - researching directories 61
 - reserved message tags 127
 - Reset command 214, 216
 - Reset command (Visualizer) 179
 - resetting command-line arguments 62
 - resetting the program counter 236
 - Resizing (and Sometimes Its Consequences) figure 155
 - Resolving Ambiguous Function Names Dialog Box figure 215
 - resolving ambiguous names 214
 - resolving multiple classes 215
 - resolving multiple static functions 214
 - Restart Checkpoint command 235
 - Restart command 235
 - restarting
 - parallel programs 139
 - program execution 198
 - programs 235
 - restoring focus 249
 - results, assigning output to variables 200
 - resuming
 - executing thread 236
 - execution 221, 226
 - processes with a signal 233
 - retaining the dive stack 154

- returning to original source
 - location 214
 - reusing windows 153
 - .rhosts file 84
 - right angle bracket (>) 153
 - right arrow is program counter 50
 - right mouse button 146
 - RMS applications 103
 - attaching to 104
 - starting 103
 - Root Window 11, 146
 - Attached Page 103, 146, 232, 233
 - dimmed information 232
 - Groups page 14, 149, 224
 - Log page 71, 149
 - selecting a process 153
 - starting CLI from 196
 - state indicator 55
 - Unattached Page 11, 49, 50, 55, 56, 102, 147
 - Unattached page 50, 94
 - Root Window Attached Page
 - figure 147
 - Root Window figure 11
 - Root Window Groups Page figure 149
 - Root Window Log Page figure 72, 149
 - Root Window Showing Process and Thread Status figure 86
 - Root Window Showing Remote figure 148
 - Root Window Unattached Page figure 148
 - Root Window: Group Page figure 225
 - Root Window: Unattached Page figure 95
 - Root Window's Group Page figure 14
 - Rotating and Querying figure 174
 - rotating surface 179
 - rounding modes 238
 - routine visualization 159
 - routines, diving on 153
 - routines, selecting 150
 - rsh command 84, 100
 - rules for scoping 295
 - Run To command 5, 138
 - "run to" commands 231, 244
 - running CLI commands 44
 - running groups 135
 - running operator 270
 - running state 55
- S**
- s command-line option 44, 197
 - S share group specifier 255
 - S state 56
 - S width specifier 257
 - Sample OpenMP Debugging
 - Session figure 116
 - sample programs
 - make_actions.tcl 197
 - Sample Visualizer Data Windows figure 173
 - Sample Visualizer Windows figure 171
 - sane command argument 196
 - Satisfaction group items
 - pull-down 353
 - satisfaction set 353
 - satisfied barrier 353
 - Save All (action points) command 370
 - Save All command 370
 - Save Pane command 157
 - saved action points 45
 - saving
 - action points 370
 - TotalView messages 201
 - window contents 157
 - sb option 370
 - scaling a surface 179
 - scaling data window 176
 - scope pull-down 246
 - scoping 294
 - ambiguous 296
 - as a tree 295
 - omitting components 296
 - rules 295
 - scrolling 145
 - output 201
 - undoing 216
 - Search Path command 48, 51, 59, 61, 102
 - search order 59
 - search paths
 - for initialization 44
 - order 59
 - setting 59, 123
 - search_path variable 61
 - search_port command-line option 79
 - searching 212
 - case-sensitive 212
 - for source code 215
 - functions 213
 - locating closest match 212
 - source code 213
 - wrapping to front or back 212
 - Searching, *see* Edit > Find, View > Lookup Function, View Lookup Variable
 - searching, variable not found 212
 - segments
 - data 162
 - text 162
 - select button 145
 - Select Directory Dialog Box figure 61
 - selected line, running to 245
 - selecting
 - different stack frame 150
 - routines 150
 - source code, by line 237
 - source line 227
 - text 156
 - sending signals to program 59
 - serial command-line option 87
 - serial line
 - baud rate 88
 - debugging over a 86
 - radio button 89
 - starting TotalView 88
 - serial option 88
 - server launch 74
 - command 75
 - enabling 75
 - replacement character %C 80

- server on each processor 19
- server option 79
 - not secure 79
- server_launch_enabled variable 75, 79, 84
- server_launch_string variable 75
- server_launch_timeout variable 76
- service threads 23, 28
- Set Barrier command 351
- set expressions 270
- set indicator, uses dot 250, 271
- Set PC command 237
- Set Signal Handling Mode
 - command 123
- set_pw command-line option 84
- set_pw single process server
 - launch command 81
- set_pws bulk server launch
 - command 82
- setting
 - barrier breakpoint 351
 - breakpoints 101, 146, 189, 241, 340, 346
 - breakpoints while running 340
 - command arguments 62
 - command line arguments 61, 62
 - environment variables 70
 - evaluation points 146, 356
 - groups 261
 - input and output files 62
 - options 69
 - preferences 69
 - search paths 59, 123
 - thread specific breakpoints 374
- setting up, debug session 39
- setting up, parallel debug session 91
- setting up, remote debug session 73
- setting X resources 69
- SGROUP variable 262
- shading graph 178
- shape arrays, deferred types 314
- Share > Halt command 220
- share groups 26, 223, 232, 254, 353
 - defined 26
 - determining 225
 - determining members of 225
 - discussion 223
 - naming 224
 - overview 254
 - S specifier 255
- SHARE_ACTION_POINT variable 343, 347, 348
- shared library, specifying name in
 - scope 295
- shared memory library code, *see* SHMEM library code
 - debugging
- shared variables 115
 - in OpenMP 117
 - OpenMP 117, 120
 - procedure for displaying 117
- sharing action points 348
- shell, example of invoking CLI
 - program 197
- SHLIB_PATH environment
 - variable 45
- SHMEM library code debugging 128
- SHMEM Sample Session figure 129
- showing areas of memory 289
- SIGALRM 140
- SIGFPE errors (on SGI) 57
- SIGINT signal 113
- signal handling mode 57
- signal/resignal loop 59
- signalHandlingMode option 56
- signals
 - affected by hardware registers 57
 - clearing 234
 - continuing execution with 233
 - default handling behavior 57
 - discarding 59
 - error option 59
 - handler routine 56
 - handling 56
- handling in PVM applications 123, 124
- handling in TotalView 56
- handling mode 57
- ignore option 59
- resend option 59
- sending continuation signal 233
- SIGALRM 140
- SIGTERM 123, 124
- stop option 59
- stops all related processes 57
- that caused core dump 53
- Signals command 57
- SIGSTOP
 - used by TotalView 57
 - when detaching 52
- SIGTERM signal 123, 124
 - stops process 123
 - terminates threads on SGI 115
- SIGTRAP, used by TotalView 57
- single process server launch 73, 74, 80
- single process server launch
 - command
 - %D 81
 - %L 81
 - %P 81
 - %R 81
 - %verbosity 81
 - callback_option 81
 - n 81
 - set_pw 81
 - working_directory 81
- single-stepping 229, 241
 - commands 229
 - in a nested stack frame 245
 - into function calls 230
 - not allowed for a parallel
 - region 115
 - on primary thread only 242
 - operating system
 - dependencies 231, 234
 - over function calls 230
 - recursive functions 231
 - skipping elements 321
 - slash in group specifier 256

- sleeping state 56
- Sliced UPC Array figure 131
- slices 8, 323
 - defining 320
 - descriptions 322
 - displaying one element 323
 - examples 320, 321
 - in sorts 330
 - lower bound 320
 - of arrays 319
 - operations using 315
 - stride elements 320
 - upper bound 320
 - with the variable command 322
- smart stepping, defined 242
- SMP machines 92
- sockets 86
- Sort > Ascending command 330
- Sort > Descending command 330
- Sort > None command 330
- Sorted Variable Window figure 330
- sorting array data 330
- Source As > Assembler 216
- Source As > Both 216, 237
- Source As > Both command 237
- Source As > Source 216
- source code
 - examining 216
 - finding 213, 215
 - navigating 216
- Source command 216
- source file, specifying name in scope 296
- source lines
 - ambiguous 227
 - editing 218
 - searching 227
 - selecting 227
- Source Pane 150
- source-level breakpoints 340
- space allocation
 - dynamic 360
 - static 360, 361
- spawned processes 204
- stopping 94
- specifier combinations 257
- specifiers
 - and dfocus 258
 - and prompt changes 259
 - examples 257
- specifying groups 255
- specifying search directories 61
- splitting up work 19
- stack
 - master thread 117
 - trace, examining 284
 - unwinding 237
- stack context of the OpenMP
 - master thread 117
- stack frame 286
 - current 216
 - examining 284
 - matching 333
 - pane 150
 - selecting different 150
- Stack Frame Pane 6, 150, 290
- stack memory 162
- stack parent token 120
 - diving 120
- Stack Trace Pane 150, 152
 - displaying source 153
- stack virtual memory 162
- standard deviation array statistic 332
- Standard I/O Page 63
- standard input, and launching
 - tvdsvr 85
- start(), stopping within 125
- start_pes() SHMEM command 128
- starting
 - CLI 41, 42, 196
 - groups 226
 - parallel tasks 101
 - TotalView 4, 41, 42, 53, 100
 - tvdsvr 43, 73, 79, 125
 - tvdsvr manually 84
- starting program under CLI
 - control 198
- Startup and Initialization
 - Sequence figure 44
- Startup command 43
- startup file 44
- Startup Parameters
 - Environment page 70
- Startup Parameters command 62, 63
 - Arguments Page 62
 - Standard I/O Page 63
- state characters 56
- states
 - and status 55
 - initializing 44
 - of processes and threads 55
 - process and thread 55
 - unattached process 56
- static constructor code 227
- static functions, resolving
 - multiple 214
- static patch space allocation 360, 361
- statically linked, stopping in
 - start() 125
- statistics for arrays 331
- status
 - and state 55
 - of processes 54
 - of threads 54
- status registers
 - examining 237
 - interpreting 237
- stdin, redirect to file 62
- stdout, redirect to file 62
- Step 1: A Program Starts figure 29
- Step 2: Forking a Process figure 30
- Step 3: Exec'ing a Process figure 31
- Step 5: Creating a Second Version figure 31
- Step 6: Creating a Remote Process figure 32
- Step 7: A Thread is Created figure 33
- Step command 138, 227
- "step" commands 230
- step command 4
- Step Instruction command 227

- stepping
 - see also* single-stepping
 - apparently hung 138
 - at process width 243
 - at thread width 244
 - goals 243
 - into 230
 - multiple statements on a line 230
 - over 230
 - primary thread can fail 244
 - process group 243
 - processes 138
 - Run (to selection) Group command 138
 - smart 242
 - target program 206
 - thread group 243
 - threads 268
 - using a numeric argument in CLI 230
- stepping a group 242
- stepping a process 243
- stepping commands 227
- stepping processes and threads 13
- STL variables, displaying 286
- \$stop assembler pseudo op 382
- Stop Before Going Parallel Question Dialog Box figure 136
- \$stop built-in function 376
- Stop control group on error check box 59
- Stop control group on error signal option 57
- stop execution 5
- STOP icon 146, 241, 340, 345
 - for breakpoints 146, 340
- stop, defined in a multiprocess environment 206
- STOP_ALL variable 343, 347
- \$stopall built-in function 376
- Stopped Execution of Compiled Expressions figure 360
- stopped operator 270
- stopped process 354
 - stopped state 55
 - unattached process 56
 - stopped thread 28
 - stopping
 - all related processes 57
 - groups 135
 - processes 220, 355
 - spawned processes 94
 - threads 220
 - \$stopprocess assembler pseudo op 382
 - \$stopprocess built-in function 376
 - \$stopthread built-in function 376
 - stride 320
 - default value of 321
 - elements 320
 - in array slices 320
 - omitting 321
 - string assembler pseudo op 383
 - <string> data type 304
 - structs
 - see also* structures
 - defined using typedefs 301
 - how displayed 300
 - structures 300
 - see also* structs
 - editing types 298
 - laminating 334
 - stty sane command 196
 - subroutines, displaying 153
 - suffixes
 - of processes in process groups 224
 - sum array statistic 332
 - Suppress All command 344
 - suppressing action points 344
 - surface
 - coloring 179
 - display 179
 - in directory window 172
 - rotating 179
 - scaling 179
 - translating 180
 - zooming 180
 - Surface command (Visualizer) 172
 - Surface Data Window 177
 - display 177
 - Surface Options Dialog Box figure 178
 - Surface visualization window 171
 - surface window, creating 172
 - suspended windows 372
 - switch-based communication for PE 99
 - switch-based communications 99
 - symbol lookup 295
 - and context 295
 - symbol name representation 294
 - symbol scoping, defined 295
 - symbol specification, omitting components 296
 - symbol table debugging information 40
 - symbolic addresses, displaying assembler as 217
 - Symbolically command 217
 - synchronizing execution 221
 - synchronizing processes 207, 244, 245
 - system PID 205
 - system TID 205
 - system variables, *see* CLI variables
 - sysuid 150, 205
 - \$sysuid built-in variable 374

T

 - T state 55, 56
 - t width specifier 257
 - tag field 345
 - tag field area 150
 - target process/thread set 206, 248
 - target program stepping 206
 - target, changing 249
 - tasker event 125
 - tasks
 - attaching to 126
 - diving into 126
 - PVM 121
 - starting 101
 - Tcl
 - and CLI 193

- and the CLI 14
- books for learning xviii
- CLI and thread lists 194
- version based upon 193
- Tcl and CLI relationship 195
- TCP/IP address, used when starting 43
- TCP/IP sockets 86
- temp file prototypes 77
- terminating processes 199
- testing when a value changes 363
- text
 - editing 156
 - locating closest match 212
 - saving window contents 157
 - selecting 156
- text assembler pseudo op 383
- text editor, default 214
- text segment 162
- text segment memory 162
- third party visualizer 164
- Thread > Continuation Signal command 52, 233
- Thread > Continuation Signal Dialog Box figure 52, 234
- Thread > Go command 226
- Thread > Hold command 222
- Thread > Set PC command 237
- thread as dimension in Visualizer 169
- thread group 245
 - stepping 243
- thread groups 26, 244, 253
 - behavior 261
 - behavior at goal 244
- thread ID 150, 205
 - system 374
 - TotalView 374
- thread local storage 118
 - variables stored in different locations 118
- thread numbers are unique 204
- Thread Objects command 317
- Thread Objects Page on an IBM AIX machine figure 318
- thread objects, displaying 317
- Thread of Interest 225
 - thread of interest 250, 252
 - defined 220, 250
 - Thread Pane 150
 - thread state 55
 - thread stepping 268
 - platforms where allowed 244
 - thread width specifier 251
 - omitting 267
 - THREADPRIVATE common block procedure for viewing variables in 119
 - THREADPRIVATE variables 119
- threads
 - call tree 160
 - creating 20
 - dimmed, in the Root Window 232
 - displaying source 153
 - diving on 150, 153
 - finding window for 150
 - holding 221, 245, 352
 - ID format 150
 - listing 150
 - manager 23
 - not available on all systems 26
 - opening window for 150
 - releasing 221, 350, 352
 - resuming executing 236
 - service 23
 - setting breakpoints in 374
 - single-stepping 241
 - stack trace 150
 - state 54
 - states 55
 - status of 54
 - stepping 13
 - stopping 220
 - switching between 11
 - systid 150
 - tid 150
 - user 23
 - workers 23, 25
- Threads figure 20, 23
- threads model 20
- thread-specific breakpoints 374
- Three Dimensional Array Sliced to Two Dimensions figure 167
- Three Dimensional Surface Visualizer Data Display figure 178
- tid 150, 205
- \$tid built-in variable 374
- TID missing in arena 251
- timeouts
 - avoid unwanted 140
 - during initialization 101
 - for connection 76
 - TotalView setting 100
- TOI defined 220
 - again 239
- Tool > P/T Set Browser command 271
- toolbar
 - controls 246
 - using 220, 246
 - width controls 246
- Toolbar figure 219
- Toolbar with Pulldown figure 14
- Tools > Call Tree command 159
- Tools > Call Tree Dialog Box figure 160
- Tools > Command Line command 42, 196
- Tools > Create Checkpoint command 235
- Tools > Evaluate command 165, 170, 371, 373
- Tools > Evaluate Dialog Box figure 372, 373
- Tools > Fortran Modules command 312
- Tools > Memory Statistics command 161
- Tools > Memory Usage Window figure 161, 163
- Tools > Message Queue command 108, 109
- Tools > Message Queue Graph command 107
- Tools > PVM Tasks command 126
- Tools > Restart Checkpoint command 235
- Tools > Statistics command 331

- Tools > Thread Objects
 - command 317
- Tools > Variable Browser
 - command 283
- Tools > Visualize command 10, 169, 336
- Tools > Watchpoint command 10, 365, 368
- Tools > Watchpoint Dialog Box
 - figure 365
- TotalView
 - and MPICH 92
 - as PVM tasker 121
 - core files 41
 - initializing 43
 - interactions with Visualizer 164
 - programming 14
 - quitting 46
 - relationship to CLI 194
 - starting 4, 41, 42, 53, 100
 - starting on remote hosts 43
 - starting the CLI within 196
 - Visualizer configuration 165
- TotalView Assembler Language 380
- TotalView assembler operators
 - hi16 382
 - hi32 382
 - lo16 382
 - lo32 382
- TotalView assembler pseudo ops
 - \$debug 382
 - \$hold 382
 - \$holdprocess 382
 - \$holdprocessstopall 382
 - \$holdstopall 382
 - \$holdthread 382
 - \$holdthreadstop 382
 - \$holdthreadstopall 382
 - \$holdthreadstopprocess 382
 - \$long_branch 382
 - \$stop 382
 - \$stopall 382
 - \$stopprocess 382
 - \$stopthread 382
 - align 382
 - ascii 382
 - asciz 382
 - bss 383
 - byte 383
 - comm 383
 - data 383
 - def 383
 - double 383
 - equiv 383
 - fill 383
 - float 383
 - global 383
 - half 383
 - lcomm 383
 - lysm 383
 - org 383
 - quad 383
 - string 383
 - text 383
 - word 383
 - zero 383
- totalview command 41, 44, 53, 96, 100, 104
 - for HP MPI 98
 - starting on a serial line 88
- TotalView data types
 - <address> 302
 - <char> 302
 - <character> 302
 - <code> 302, 304
 - <complex*16> 302
 - <complex*8> 302
 - <complex> 302
 - <double precision> 302
 - <double> 302
 - <extended> 303
 - <float> 303
 - <int> 303
 - <integer*1> 303
 - <integer*2> 303
 - <integer*4> 303
 - <integer*8> 303
 - <integer> 303
 - <logical*1> 303
 - <logical*2> 303
 - <logical*4> 303
 - <logical*8> 303
 - <logical> 303
 - <long long> 303
 - <long> 303
 - <opaque> 306
 - <real*16> 303
 - <real*4> 303
 - <real*8> 303
 - <real> 303
 - <short> 303
 - <string> 303, 304
 - <void> 303, 304
- TotalView Debugger Server, *see* tvdsrv
- TotalView Debugging Session
 - Over a Serial Line figure 87
- TOTALVIEW environment variable 93, 139
- TotalView program
 - quitting 46
- totalview subdirectory, *see* .totalview subdirectory
- TotalView Visualizer Connection
 - figure 165
- TotalView Visualizer Relationships
 - figure 164
- TotalView Visualizer
 - see* Visualizer
- TotalView windows
 - action point List pane 152
 - editing cursor 156
- totalviewcli command 41, 42, 44, 53, 104, 196, 198
 - remote 43
 - starting on a serial line 88
- translating a surface 180
- translating data window 176
- transposing axis 175
- TRAP_FPE environment variable
 - on SGI 57
- troubleshooting xxii
 - MPI 113
- tv command-line option 92
- TV::namespace 202
- TV::GUI::namespace 202
- TVD.breakpoints file 370
- TVDB_patch_base_address
 - object 361
- tvdb_patch_space.s 362

tvdrc file, *see* .tvdrc initialization file
 tvdsvr 43, 48, 73, 74, 76, 86, 87, 359
 attaching to 126
 –callback command-line option 84
 cleanup by PVM 127
 editing command line for poe 102
 fails in MPI environment 113
 launch problems 76, 78
 launching 80
 launching, arguments 85
 manually starting 84
 –port command-line option 79
 –search_port command-line option 79
 –server command-line option 79
 –set_pw command-line option 84
 starting 79
 starting for serial line 87
 starting manually 79, 84
 symbolic link from PVM directory 122
 with PVM 125
 tvdsvr command
 starting 73
 timeout while launching 76, 78
 use with PVM applications 122
 TVDSVRLAUNCHCMD
 environment variable 80
 Two Computers Working on One Problem figure 19
 Two Dimensional Surface
 Visualizer Data Display figure 177
 Two More Variable Window figure 9
 Two Variable Windows figure 9
 two-dimensional graphs 173
 type casting 297
 examples 304
 type strings
 built-in 302

 editing 297
 for opaque types 306
 supported for Fortran 298
 type transformation variable 280
 typedefs
 defining structs 301
 how displayed 300
 types
 user defined type 314
 types supported for C language 298

U

UDT 314
 UDWP, *see* watchpoints
 UID, UNIX 79
 Unattached Page 11, 49, 55, 56, 94, 102, 147
 Unattached page 50
 Unattached Page figure 50
 unattached process states 56
 summary 56
 undive icon 154, 214, 291
 Undive/Dive Controls figure 214
 undiving, from windows 292
 unexpected messages 108, 112
 unheld operator 270
 union operator 270
 unions 300
 how displayed 301
 Uniprocessor figure 18
 unique process numbers 204
 unique thread numbers 204
 unsuppressing action points 344
 unwinding the stack 237
 UPC Laminated Variable figure 133
 UPC Variable Window Showing Nodes figure 131
 Update command 103, 221, 233
 updating groups 275
 updating visualization displays 169
 upper adjacent array statistic 333
 upper bounds 299
 of array slices 320
 USEd information 312

user defined data type 314
 user mode 23
 user threads 23
 User Threads and Service Threads figure 24
 User, Service, and Manager Threads figure 24
 Using an Expression to Change a Value figure 297
 Using Assembler figure 381

V

value field 371
 values
 changing 156
 editing 7
 Variable Browser command 283
 variable scoping 294
 Variable Window
 closing 290
 displaying 281
 duplicating 292
 in recursion, manually refocus 286
 laminated display 333
 stale in pane header 286
 tracking addresses 286
 updates to 286
 Variable Window figure 168
 Variable Window for a Global Variable figure 282
 Variable Window for Area of Memory figure 289
 Variable Window for small_array figure 323
 Variable Window with Machine Instructions figure 290
 variables
 assigning p/t set to 252
 at different addresses 334
 CGROUP 254, 262
 changing the value 296
 changing values of 296
 data format 281
 display width 280
 displaying all globals 283
 displaying contents 153

- displaying long names 286
- diving 153
- GROUP 262
- in modules 312
- in Stack Frame Pane 7
- intrinsic, *see* built-in functions
- laminated display 333
- locating 212
- precision 280
- previewing size and precision 281
- setting command output to 200
- SGROUP 262
- stored in different locations 118
- ttf 280
- watching for value changes 10
- WGROUP 262
- verbosity bulk server launch command 82
- verbosity level 105
- verbosity single process server launch command 81
- View > Assembler > By Address command 217
- View > Assembler > Symbolically command 217
- View > Dive Anew command 283
- View > Dive In All command 293
- View > Dive Thread command 317
- View > Dive Thread New command 317
- View > Graph command 171
- View > Graph command (Visualizer) 172
- View > Laminate > None command 333
- View > Laminate > Process command 333
- View > Laminate > Thread command 333
- View > Laminate Thread command 119
- View > Lookup Function command 125, 212, 213, 215, 216
- View > Lookup Function Dialog Box figure 214, 215
- View > Lookup Variable command 119, 212, 282, 286, 289, 312
 - specifying slices 322
- View > Lookup Variable Dialog Box figure 213
- View > Reset command 214, 216
- View > Reset command (Visualizer) 176, 179
- View > Sort > Ascending command 330
- View > Sort > Descending command 330
- View > Sort > None command 330
- View > Source As > Assembler command 216
- View > Source As > Both command 216, 237
- View > Source As > Source command 216
- View > Surface command (Visualizer) 171, 172
- View > Variable command 117
- viewing assembler 217
- virtual stack memory 162
- visualization
 - deleting a dataset 172
 - translating a surface 180
 - zooming a surface 180
- \$visualize 377
- Visualize command 10, 167, 169, 336
- visualize command 180
- visualize, *see* \$visualize
- Visualizer 170, 335
 - autolaunch options, changing 165
 - choosing method for displaying data 168
 - configuring 165
 - configuring launch 165
 - creating graph window 172
 - creating surface window 172
 - data sets to visualize 167
 - data types 167
 - data window 170, 172
 - data window manipulation commands 176
 - dataset defined 167
 - dataset numeric identifier 167
 - dataset parameters 179
 - deleting datasets 172
 - dimensions 169
 - directory window 170, 171
 - display not automatically updated 169
 - exiting from 172
 - file command-line option 167, 180
 - graphs, display 173, 174
 - graphs, manipulating 176
 - interactions with TotalView 164
 - laminated data panes 169
 - launch command, changing shell 167
 - launch from command line 180
 - launch options 165
 - method 168
 - new or existing dataset 167
 - number of arrays 167
 - persist command-line option 167, 180
 - pipe 164
 - rank 165
 - relationship to TotalView 164
 - rotating 179
 - scaling a surface 179
 - selecting datasets 171
 - shell launch command 166
 - slices 167
 - surface data display options 179
 - Surface Data Window 177
 - third party 164
 - using casts 170
 - windows, types of 170

- visualizer
 - closing connection to 166
 - customized command for 165
- Visualizer Graph Data Window
 - figure 175
- visualizing
 - data 163, 172
 - data sets from a file 180
 - from variable window 168
 - in expressions using \$visualize 169
- <void> data type 304
- W**
- W state 55
- W width specifier 257
- W workers group specifiers 256
- Waiting for Command to
 - Complete window 138
- Waiting to Complete Message
 - Box figure 372
- warn_step_throw variable 58
- watching memory 366
- Watchpoint command 10, 365, 368
- watchpoint operator 270
- Watchpoint Properties dialog box 366
- watchpoint state 55
- watchpoints 10, 363
 - \$newval 368
 - \$oldval 368
 - alignment 369
 - conditional 363, 368
 - copying data 368
 - creating 365
 - defined 207, 338
 - disabling 366
 - diving into 366
 - enabling 366
 - evaluated, not compiled 370
 - evaluating an expression 363
 - example of triggering when value goes negative 369
 - length compared to \$oldval or \$newval 369
 - lists of 152
 - lowest address triggered 367
 - modifying a memory location 363
 - monitoring adjacent locations 367
 - multiple 367
 - not saved 370
 - PC position 367
 - platform differences 364
 - problem with stack variables 366
 - supported platforms 363
 - testing a threshold 363
 - testing when a value changes 363
 - triggering 363, 367
 - watching memory 366
- WGROU variable 262
- When a job goes parallel or calls exec() radio buttons 136
- When a job goes parallel radio buttons 137
- When Done, Stop radio buttons 352
- When Hit, Stop radio buttons 352
- width pulldown 246
- width relationships 252
- width specifier 250
 - omitting 267
- Width Specifiers figure 252
- Window > Duplicate Base
 - command 154, 292
- Window > Duplicate command 154, 292
- Window > Memorize All
 - command 155
- Window > Memorize command 155
- Window > Update command 103, 221, 233
- window contents, saving 157
- window, copying 154
- windows 290
 - closing 154, 290
 - copying between 156
 - data 172
 - Data Window (Visualizer) 173
 - Directory Window 171
 - event log 71
 - graph data 173
 - pasting between 156
 - popping 153
 - resizing 155
 - Surface Data Window 177
 - suspended 372
- Windows > Update command (PVM) 126
- word assembler pseudo op 383
- worker threads 23, 115
- workers group 27, 245
 - defined 26
 - overview 254
- workers group specifier 256
- working directory 61
- working independently 18
- working_directory bulk server launch command 82
- working_directory single process server launch command 81
- X**
- X resources setting 69
- Xdefaults file, *see* .Xdefaults file
- xterm, launching tvdsrv from 85
- Z**
- Z state 56
- zero assembler pseudo op 383
- zero count array statistic 333
- zombie state 56
- zone coloring 179
- zone maps 177
- zooming a surface 180
- zooming data window 176
- Zooming, Rotating, About an Axis figure 181