



PGDBG[®] Debugger Guide

Parallel Debugging for Scientists and Engineers

Release 2011

The Portland Group[®]

While every precaution has been taken in the preparation of this document, The Portland Group® (PGI®), a wholly-owned subsidiary of STMicroelectronics, Inc., makes no warranty for the use of its products and assumes no responsibility for any errors that may appear, or for damages resulting from the use of the information contained herein. The Portland Group retains the right to make changes to this information at any time, without notice. The software described in this document is distributed under license from STMicroelectronics and/or The Portland Group and may be used or copied only in accordance with the terms of the license agreement ("EULA").

PGI Workstation, PGI Server, PGI Accelerator, PGF95, PGF90, PGFORTRAN, and PGI Unified Binary are trademarks; and PGI, PGHPE, PGF77, PGCC, PGC++, PGI Visual Fortran, PVE, PGI CDK, Cluster Development Kit, PGPROF, PGDBG, and The Portland Group are registered trademarks of The Portland Group Incorporated. Other brands and names are property of their respective owners.

No part of this document may be reproduced or transmitted in any form or by any means, for any purpose other than the purchaser's or the end user's personal use without the express written permission of STMicroelectronics and/or The Portland Group.

PGDBG® Debugger Guide

Copyright © 2010-2011 STMicroelectronics, Inc.

All rights reserved.

Printed in the United States of America

First Printing: Release 11.0, December 2010

Second Printing: Release 11.1, January 2011

Third Printing: Release 11.2, February 2011

Third Printing: Release 11.4, April 2011

Technical support: <http://www.pgroup.com/support/>

Sales: sales@pgroup.com

Web: <http://www.pgroup.com>

Contents

Preface	xvii
Intended Audience	xvii
Documentation	xvii
Compatibility and Conformance to Standards	xvii
Organization	xviii
Conventions	xix
Terminology	xx
Related Publications	xx
 1. Getting Started	 1
Definition of Terms	1
Building Applications for Debug	1
Debugging Optimized Code	2
Building for Debug on Windows	2
User Interfaces	2
Invoking PGDBG	3
PGDBG Initialization	3
Debugging	3
Selecting a Version of Java	3
 2. The Graphical User Interface	 5
Main Components	5
Source Window	6
Source and Disassembly Displays	6
Source Window Context Menu	7
Call Stack Navigation	8
Status Message Area	8
Main Toolbar	8
Buttons	8
Drop-Down Lists	9
Debug Information Tabs	9
Command Tab	10
Events tab	10

Groups Tab	11
Call Stack Tab	11
Locals Tab	12
Memory Tab	12
MPI Messages Tab	13
Procs & Threads Tab	14
Registers Tab	14
Status Tab	15
Menu Bar	16
File Menu	16
Edit Menu	17
View Menu	17
Data Menu	18
Debug Menu	19
Help Menu	20
3. Command Line Options	21
Command-Line Options Syntax	21
Command-Line Options	21
Command-Line Options for MPI Debugging	22
4. Command Language	23
Command Overview	23
Command Syntax	23
Command Modes	23
Constants	24
Symbols	24
Scope Rules	24
Register Symbols	24
Source Code Locations	24
Lexical Blocks	25
Statements	26
Events	26
Event Commands	27
Event Command Action	28
Expressions	29
Ctrl-C	30
5. Command Summary	33
Notation Used in Command Sections	33
Command Summary	34
6. Assembly-Level Debugging	47
Assembly-Level Debugging Overview	47
Assembly-Level Debugging on Microsoft Windows Systems	47
Assembly-Level Debugging with Fortran	48
Assembly-Level Debugging with C ⁺⁺	48

Assembly-Level Debugging Using the PGDBG GUI	48
Assembly-Level Debugging Using the PGDBG CLI	48
SSE Register Symbols	49
7. Source-Level Debugging	51
Debugging Fortran	51
Fortran Types	51
Arrays	51
Operators	51
Name of the Main Routine	52
Common Blocks	52
Internal Procedures	52
Modules	53
Module Procedures	53
Debugging C++	54
Calling C++ Instance Methods	54
8. Platform-Specific Features	55
Pathname Conventions	55
Debugging with Core Files	55
Signals	56
Signals Used Internally by PGDBG	57
Signals Used by Linux Libraries	57
9. Parallel Debugging Overview	59
Overview of Parallel Debugging Capability	59
Graphical Presentation of Threads and Processes	59
Basic Process and Thread Naming	59
Thread and Process Grouping and Naming	60
PGDBG Debug Modes	60
Threads-only Debugging	61
Process-only Debugging	61
Multilevel Debugging	61
Process/Thread Sets	62
Named p/t-sets	62
p/t-set Notation	62
Dynamic vs. Static p/t-sets	63
Current vs. Prefix p/t-set	63
p/t-set Commands	64
Using Process/Thread Sets in the GUI	65
p/t set Usage	67
Command Set	67
Process Level Commands	67
Thread Level Commands	67
Global Commands	69
Process and Thread Control	69
Configurable Stop Mode	70

Configurable Wait Mode	70
Status Messages	73
The PGDBG Command Prompt	73
Parallel Events	74
Parallel Statements	75
Parallel Compound/Block Statements	75
Parallel If, Else Statements	75
Parallel While Statements	76
Return Statements	76
10. Parallel Debugging with OpenMP	77
OpenMP and Multi-thread Support	77
Multi-thread and OpenMP Debugging	77
Debugging OpenMP Private Data	78
11. Parallel Debugging with MPI	81
MPI and Multi-Process Support	81
Process Control	81
Process Synchronization	82
MPI Message Queues	82
MPI Groups	82
Use halt instead of Ctrl+C	82
SSH and RSH	83
MPI Debugging on Linux	83
Invoking PGDBG for MPI Debugging	84
Using PGDBG for MPI Debugging	85
Debugging Support for MPICH-1	86
MPI Debugging on Windows	86
Installing MSMPI	86
Building with MSMPI	87
Debug MSMPI Applications Locally	87
Debug MSMPI Applications on a Cluster	87
12. Parallel Debugging of Hybrid Applications	89
PGDBG Multilevel Debug Mode	89
Multilevel Debugging	89
13. Command Reference	91
Notation Used in Command Sections	91
Process Control	92
attach	92
cont	92
debug	92
detach	92
halt	93
load	93
next	93

nexti	93
proc	93
procs	93
quit	93
rerun	93
run	94
setargs	94
step	94
stepi	94
stepout	94
sync	94
synci	95
thread	95
threads	95
wait	95
Process-Thread Sets	95
defset	95
focus	95
undefset	95
viewset	96
whichsets	96
Events	96
break	96
breaki	97
breaks	97
catch	98
clear	98
delete	98
disable	98
do	98
doi	99
enable	99
hwatch	99
hwatchboth	99
hwatchread	99
ignore	100
status	100
stop	100
stopi	100
trace	100
tracei	101
track	101
tracki	101
unbreak	101
unbreaki	101

watch	101
watchi	102
when	102
wheni	102
Program Locations	103
arrive	103
cd	103
disasm	103
edit	103
file	103
lines	104
list	104
pwd	104
stackdump	104
stacktrace	104
where	105
/	105
?	105
Printing Variables and Expressions	105
print	105
printf	106
ascii	107
bin	107
dec	107
display	107
hex	107
oct	107
string	107
undisplay	108
Symbols and Expressions	108
assign	108
call	108
declaration	109
entry	109
lval	109
rval	109
set	110
sizeof	110
type	110
Scope	110
class	110
classes	111
decls	111
down	111
enter	111

files	111
global	111
names	111
scope	111
up	111
whereis	112
which	112
Register Access	112
fp	112
pc	112
regs	112
retaddr	112
sp	112
Memory Access	113
cread	113
dread	113
dump	113
fread	114
iread	114
lread	114
mqdump	114
sread	114
Conversions	114
addr	114
function	115
line	115
Target	115
connect	115
disconnect	115
native	115
Miscellaneous	116
alias	116
directory	116
help	116
history	117
language	117
log	117
noprint	117
pgienv	117
repeat	120
script	120
setenv	120
shell	120
sleep	121
source	121

unalias	121
use	121
Index	123

Figures

2.1. Default Appearance of PGDBG GUI	5
2.2. Source Window	6
2.3. Context Menu	7
2.4. Call Frame Selector	8
2.5. Status Messages	8
2.6. Buttons on Toolbar	8
2.7. Drop-Down Lists on Toolbar	9
2.8. Command Tab	10
2.9. Events Tab	10
2.10. Groups Tab	11
2.11. Call Stack Tab	11
2.12. Locals Tab	12
2.13. Memory Tab	12
2.14. Memory Tab in Decimal Format	13
2.15. MPI Messages Tab	13
2.16. Process (Thread) Grid Tab	14
2.17. General Purpose Registers	15
2.18. Status Tab	16
9.1. Groups Tab	65
9.2. Process/Thread Group Dialog Box	66
10.1. OpenMP Private Data in PGDBG GUI	79

Tables

2.1. Colors Describing Thread State	14
4.1. PGDBG Operators	30
5.1. PGDBG Commands	34
9.1. PGDBG Debug Modes	60
9.2. p/t-set Commands	64
9.3. PGDBG Parallel Commands	67
9.4. PGDBG Stop Modes	70
9.5. PGDBG Wait Modes	71
9.6. PGDBG Wait Behavior	72
9.7. PGDBG Status Messages	73
10.1. Thread State Is Described Using Color	78
11.1. MPICH Support	86
13.1. pgienv Commands	118

Examples

9.1. Thread IDs in Threads-only Debug Mode	61
9.2. Process IDs in Process-only Debug Mode	61
9.3. Thread IDs in Multilevel Debug Mode	61
9.4. p/t-sets in Threads-only Debug Mode	62
9.5. p/t-sets in Process-only Debug Mode	63
9.6. p/t-sets in Multilevel Debug Mode	63
9.7. Defining a Dynamic p/t-set	63
9.8. Defining a Static p/t-set	63
12.1. Thread IDs in multilevel debug mode	89
13.1. Syntax examples	91

Preface

This guide describes how to use the *PGDBG* debugger to debug serial and parallel applications built with The Portland Group (PGI) Fortran, C, and C⁺⁺ compilers for X86, AMD64 and Intel 64 processor-based systems. It contains information about how to use *PGDBG*, as well as detailed reference information on commands and its graphical interface.

Intended Audience

This guide is intended for application programmers, scientists and engineers proficient in programming with the Fortran, C, and/or C⁺⁺ languages. The PGI tools are available on a variety of operating systems for the X86, AMD64, and Intel 64 hardware platforms. This guide assumes familiarity with basic operating system usage.

Documentation

PGI Documentation is installed with every release. The latest version of *PGDBG* documentation is also available at www.pgroup.com/docs.htm. See www.pgroup.com/faq/index.htm for frequently asked *PGDBG* questions and answers.

Compatibility and Conformance to Standards

The PGI compilers and tools run on a variety of systems. They produce and/or process code that conforms to the ANSI standards for FORTRAN 77, Fortran 95, Fortran 2003, C, and C⁺⁺ and includes extensions from MIL-STD-1753, VAX/VMS Fortran, IBM/VS Fortran, SGI Fortran, Cray Fortran, and K&R C. PGF77, PGFORTRAN, PGCC ANSI C, and PGCPP support parallelization extensions based on the OpenMP 3.0 standard. PGHPF supports data parallel extensions based on the High Performance Fortran (HPF) defacto standard. The PGI Fortran Reference Manual describes Fortran statements and extensions as implemented in the PGI Fortran compilers.

PGDBG supports debugging of serial, multi-threaded, parallel OpenMP, parallel MPI and multi-process multi-threaded hybrid MPI programs compiled with PGI compilers.

For further information, refer to the following:

- American National Standard Programming Language FORTRAN, ANSI X3. -1978 (1978).
- American National Standard Programming Language C, ANSI X3.159-1989.

- ISO/IEC 9899:1999, Information technology – Programming Languages – C, Geneva, 1999 (C99).
- ISO/IEC 1539:1991, Information technology – Programming Languages – Fortran, Geneva, 1991 (Fortran 90).
- ISO/IEC 1539:1997, Information technology – Programming Languages – Fortran, Geneva, 1997 (Fortran 95).
- High Performance Fortran Language Specification, Revision 1.0, Rice University, Houston, Texas (1993), <http://www.crpc.rice.edu/HPFF>.
- High Performance Fortran Language Specification, Revision 2.0, Rice University, Houston, Texas (1997), <http://www.crpc.rice.edu/HPFF>.
- OpenMP Application Program Interface, Version 2.5, May 2005, <http://www.openmp.org>.
- Programming in VAX Fortran, Version 4.0, Digital Equipment Corporation (September, 1984).
- IBM VS Fortran, IBM Corporation, Rev. GC26-4119.
- Military Standard, Fortran, DOD Supplement to American National Standard Programming Language Fortran, ANSI x.3-1978, MIL-STD-1753 (November 9, 1978).
- HPDF Standard (High Performance Debugging Forum) <http://www.ptools.org/hpdf/draft/intro.html>

Organization

The PGDBG Debugger Manual contains these thirteen chapters that describe *PGDBG*, a symbolic debugger for Fortran, C, C⁺⁺ and assembly language programs.

Chapter 1, “*Getting Started*”

contains information on how to start using the debugger, including a description of how to build a target application for debug and how to invoke *PGDBG*.

Chapter 2, “*The Graphical User Interface*”

describes how to use the *PGDBG* graphical user interface (GUI).

Chapter 3, “*Command Line Options*”

describes the *PGDBG* command-line options.

Chapter 4, “*Command Language*”

provides detailed information about the *PGDBG* command language, which can be used from the command-line user interface or from the Command tab of the graphical user interface.

Chapter 5, “*Command Summary*”

provides a brief summary table of the *PGDBG* debugger commands with a brief description of the command as well as information about the category of command use.

Chapter 6, “*Assembly-Level Debugging*”

contains information on assembly-level debugging; basic debugger operations, commands, and features that are useful for debugging assembly code; and how to access registers.

Chapter 7, “*Source-Level Debugging*”

contains information on language-specific issues related to source debugging.

Chapter 8, “*Platform-Specific Features*”

contains platform-specific information as it relates to debugging.

Chapter 9, “*Parallel Debugging Overview*”

contains an overview of the parallel debugging capabilities of *PGDBG*.

Chapter 10, “*Parallel Debugging with OpenMP*”

describes the parallel debugging capabilities of *PGDBG* and how to use them with OpenMP.

Chapter 11, “*Parallel Debugging with MPI*”

describes the parallel debugging capabilities of *PGDBG* and how to use them with MPI.

Chapter 12, “*Parallel Debugging of Hybrid Applications*”

describes the parallel debugging capabilities of *PGDBG* and how to use them with hybrid applications.

Chapter 13, “*Command Reference*”

provides reference information about each of the *PGDBG* commands, organized by area of use.

Conventions

This guide uses the following conventions:

italic

is used for emphasis.

`Constant Width`

is used for filenames, directories, arguments, options, examples, and for language statements in the text, including assembly language statements.

Bold

is used for commands.

`[item1]`

in general, square brackets indicate optional items. In this case item1 is optional. In the context of p/t-sets, square brackets are required to specify a p/t-set.

`{ item2 | item 3 }`

braces indicate that a selection is required. In this case, you must select either item2 or item3.

`filename ...`

ellipsis indicate a repetition. Zero or more of the preceding item may occur. In this example, multiple filenames are allowed.

`FORTRAN`

Fortran language statements are shown in the text of this guide using a reduced fixed point size.

`C/ C++`

C/ C++ language statements are shown in the text of this guide using a reduced fixed point size.

The PGI compilers and tools are supported on both 32-bit and 64-bit variants of the Linux, Windows, and Mac OS operating systems on a variety of x86-compatible processors. There are a wide variety of releases and distributions of each of these types of operating systems.

Terminology

If there are terms in this guide with which you are unfamiliar, PGI provides a glossary of terms which you can access at www.pgroup.com/support/definitions.htm

Related Publications

The following documents contain additional information related to the X86 architecture and the compilers and tools available from The Portland Group.

- PGI Fortran Reference Manual describes the FORTRAN 77, Fortran 90/95, Fortran 2003, and HPF statements, data types, input/output format specifiers, and additional reference material related to the use of PGI Fortran compilers.
- System V Application Binary Interface Processor Supplement by AT&T UNIX System Laboratories, Inc. (Prentice Hall, Inc.).
- FORTRAN 95 HANDBOOK, Complete ANSI/ISO Reference (The MIT Press, 1997).
- Programming in VAX Fortran, Version 4.0, Digital Equipment Corporation (September, 1984).
- IBM VS Fortran, IBM Corporation, Rev. GC26-4119.
- The C Programming Language by Kernighan and Ritchie (Prentice Hall).
- C: A Reference Manual by Samuel P. Harbison and Guy L. Steele Jr. (Prentice Hall, 1987).
- The Annotated C⁺⁺ Reference Manual by Margaret Ellis and Bjarne Stroustrup, AT&T Bell Laboratories, Inc. (Addison-Wesley Publishing Co., 1990)
- PGI Compiler User's Guide, PGI Reference Manual, PGI Release Notes, FAQ, Tutorials, <http://www.pgroup.com/>
- MPI-CH <http://www.unix.mcs.anl.gov/MPI/mpich/>
- OpenMP <http://www.openmp.org/>

Chapter 1. Getting Started

PGDBG is a symbolic debugger for Fortran, C, C++ and assembly language programs. It provides debugger features, such as execution control using breakpoints, single-stepping, and examination and modification of application variables, memory locations, and registers.

PGDBG supports debugging of certain types of parallel applications:

- Multi-threaded and OpenMP applications.
- MPI applications.
- Hybrid applications, which use multiple threads or OpenMP as well as multiple MPI processes.

Multi-threaded and OpenMP applications may be run using more threads than the available number of CPUs, and MPI applications may allocate more than one process to a cluster node. *PGDBG* supports debugging the listed types of applications regardless of how well the number of threads match the number of CPUs or how well the number of processes match the number of cluster nodes.

Definition of Terms

Throughout this manual we use the term *host* to refer to the system on which *PGDBG* executes, *target* to refer to the program being debugged, and *target machine* to refer to the system on which the target runs. For more detailed definitions of these terms, refer to the PGI glossary of terms which you can access at www.pgroup.com/support/definitions.htm.

For an introduction to terminology used to describe parallel debugging, refer to [Chapter 9, “Parallel Debugging Overview”](#).

Building Applications for Debug

To build an application for debug, compile with the `-g` option. With this option, the compiler generates information about the symbols and source files in the program and includes it in the executable file. The option `-g` also sets the compiler optimization to level zero (no optimization) unless you specify optimization options such as `-O`, `-fast`, or `-fastsse` on the command line. Optimization options take effect whether they are listed before or after `-g` on the command line.

Debugging Optimized Code

Programs built with `-g` and optimization levels higher than `-O0` can be debugged, but due to transformations made to the program during optimization, source-level debugging may not be reliable. Assembly-level debugging (e.g., accessing registers, viewing assembly code, etc.) is reliable, even with optimized code. Programs built without `-g` can be debugged; however, information about types, local variables, arguments and source file line numbers are not available. For more information on assembly-level debugging, refer to [Chapter 6, “Assembly-Level Debugging”](#).

In programs built with both `-g` and optimization levels higher than `-O0`, some optimizations may be disabled or otherwise affected by the `-g` option, possibly changing the program behavior. An alternative option, `-gopt`, can be used to build programs with full debugging information, but without modifying program optimizations. Unlike `-g`, the `-gopt` option does not set the optimization to level zero.

Building for Debug on Windows

To build an application for debug on Windows platforms, applications must be linked with the `-g` option as well as compiled with `-g`. This process results in the generation of debug information stored in a `.dwf` file and a `.pdb` file. The PGI compiler driver should always be used to link applications; except for special circumstances, the linker should not be invoked directly.

User Interfaces

PGDBG includes both a command-line interface (CLI) and a graphical user interface (GUI).

The GUI interface supports command entry through a point-and-click interface, a view of source and assembly code, a full command-line interface panel, and several other graphical elements and features. [“Command Language”](#) and [“Command Reference”](#) describe in detail how to use the *PGDBG* CLI. [“The Graphical User Interface”](#) describes how to use the *PGDBG* GUI.

Command Line Interface (CLI)

Text commands are entered one line at a time through the command-line interface. A number of command-line options can be used when launching *PGDBG*.

For information on these options and how they are interpreted, refer to [Chapter 3, “Command Line Options”](#) and [“Command Reference”](#).

Graphical User Interface

The GUI, the default user interface, supports command entry through a point-and-click interface, a view of source and assembly code, a full command-line interface panel, and several other graphical elements and features. There may be minor variations in the appearance of the *PGDBG* GUI from host to host, depending on the type of display hardware available, the settings for various defaults and the window manager used. Except for differences caused by those factors, the basic interface remains the same across all systems.

For more information on the *PGDBG* GUI, refer to [“The Graphical User Interface”](#).

Invoking PGDBG

PGDBG is invoked using the **pgdbg** command as follows:

```
% pgdbg arguments target arg1 arg2 ... argn
```

where *arguments* may be any of the command-line arguments described in [Chapter 3, “Command Line Options”](#). See [Chapter 11, “Parallel Debugging with MPI”](#) for instructions on how to launch and debug an MPI program.

The *target* parameter is the name of the program executable file being debugged. The arguments *arg1 arg2 ... argn* are the command-line arguments to the target program. Invoking *PGDBG* as described starts the *PGDBG* GUI. For users who prefer to use a command-line interface, *PGDBG* may also be invoked with the **-text** parameter.

Note

The command shell interprets any I/O redirection specified on the *PGDBG* command line. Refer to [“Process Control,” on page 81](#) for a description of how to redirect I/O using the **run** command.

PGDBG supports both 32-bit and 64-bit applications. If the PATH environment variable is set to use the 32-bit PGI tools, a 64-bit application can be debugged by invoking *PGDBG* normally. *PGDBG* automatically determines the architecture of the target program and configures itself for 64-bit debugging. Conversely, you can debug a 32-bit application by invoking *PGDBG* even when the PATH is set to invoke the 64-bit PGI tools. For more details, refer to [Chapter 3, “Command Line Options”](#).

PGDBG Initialization

Once *PGDBG* is started, it reads symbol information from the executable file, then loads the application into memory. For large applications this process can take a few moments.

An initialization file can be useful for defining common aliases, setting breakpoints, and for other startup commands. If an initialization file named `.pgdbgrc` exists in the current directory or in the home directory, as defined by the environment variable HOME, *PGDBG* opens this file and executes the commands in it.

If an initialization file is found in the current directory, then the initialization file in the home directory, if there is one, is ignored. However, a **script** command placed in the initialization file may execute the initialization file in the home directory, or execute *PGDBG* commands in any other file.

Debugging

Once *PGDBG* is invoked and the initialization file is processed, *PGDBG* is ready to process commands. Normally, a session begins by setting one or more breakpoints, using the **break**, **stop** or **trace** commands, and then issuing a **run** command followed by **cont**, **step**, **trace** or **next**.

Selecting a Version of Java

The *PGDBG* GUI depends on Java. *PGDBG* command-line mode (`pgdbg -text`) does not depend on Java. *PGDBG* requires that the Java Virtual Machine be a specific minimum version or above. By default, *PGDBG* uses the version of Java installed with your PGI software; if you chose not to install Java when installing your

PGI software, *PGDBG* looks for Java on your PATH. Both of these can be overridden by setting the PGI_JAVA environment variable to the full path of the Java executable you wish to use.

For example, on a Linux system using the bash shell, use this command to specify the location of Java:

```
$ export PGI_JAVA=/home/myuser/myjava/bin/java
```


Chapter 2. The Graphical User Interface

The default user interface used by *PGDBG* is a graphical user interface or GUI. There may be minor variations in the appearance of the *PGDBG* GUI from host to host, depending on the type of display hardware available, the settings for various defaults and the window manager used. Except for differences caused by those factors, the basic interface remains the same across all systems.

Main Components

Figure 2.1. Default Appearance of PGDBG GUI

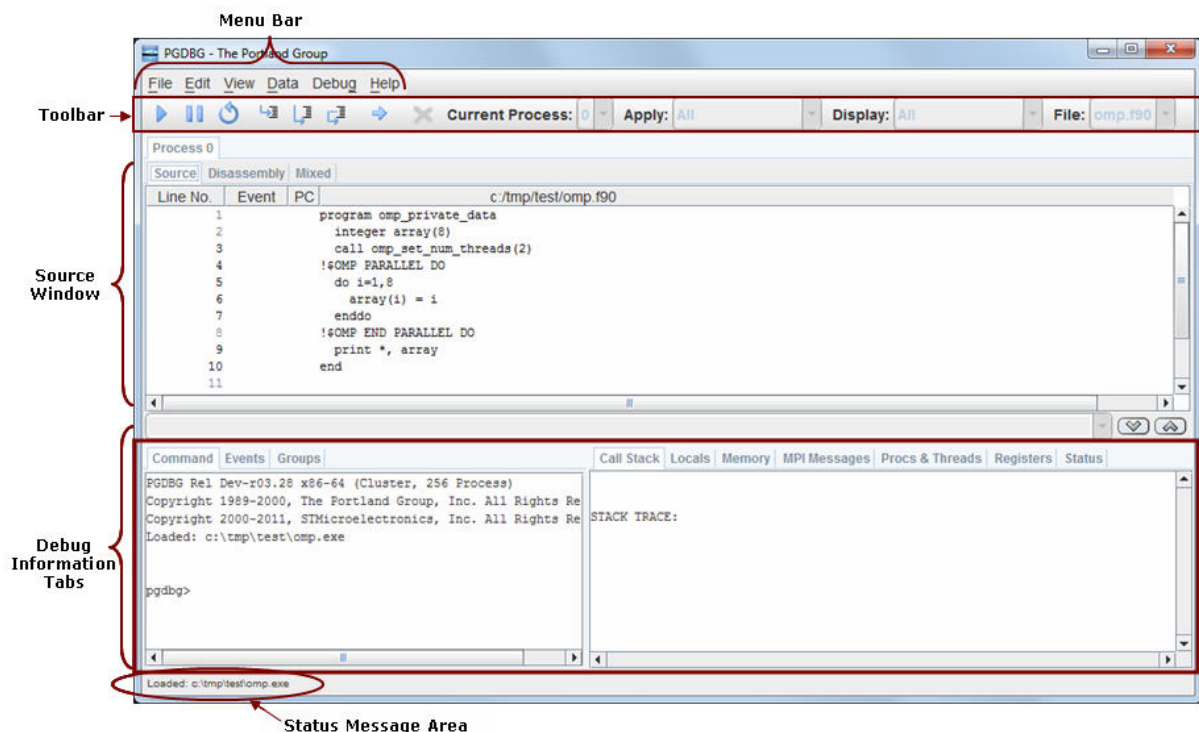


Figure 2.1, “Default Appearance of PGDBG GUI” shows the *PGDBG* GUI as it appears when *PGDBG* is invoked for the first time.

The GUI can be resized according to the conventions of the underlying window manager. Changes in window size and other settings are saved and used in subsequent invocations of *PGDBG*. To prevent changes to the default settings from being saved, uncheck the Save Settings on Exit item on the Edit menu.

The GUI is divided into five main areas, as illustrated in [Figure 2.1](#). From top to bottom these areas are the menu bar, main toolbar, source window, debug information tabs, and status message area.

The source window and all of the debug information tabs are dockable tabs. A dockable tab can be separated from the main window by either double-clicking the tab or dragging the tab off the main window. To return the tab to the main window, double-click it again or drag it back onto the main window. You can change the placement of any dockable tab by dragging it from one location to another. Right-click on a dockable tab to bring up a context menu with additional options, including closing the tab. To reopen a closed tab, use the View menu. To return the GUI to its original state, use the Edit menu's Restore Default Settings... option.

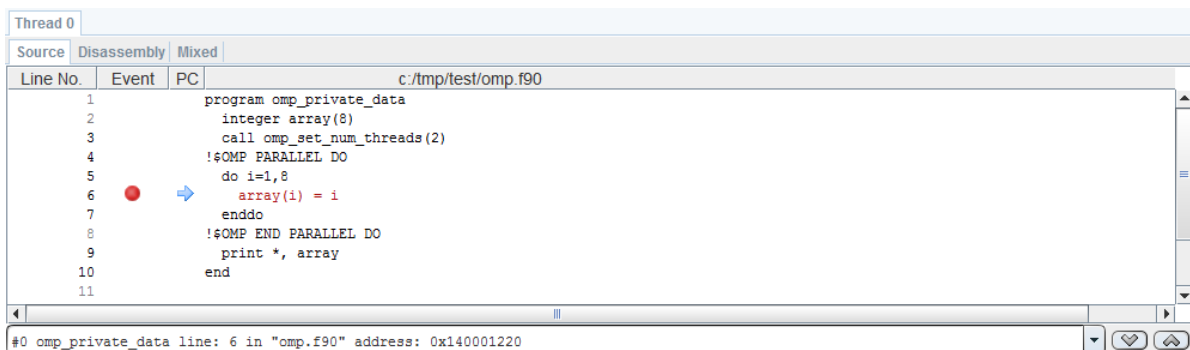
A separate window named the Program I/O window is displayed when *PGDBG* is started. Any input or output performed by the target program is entered into or displayed in this window.

The following sections explain the parts of the GUI and how they are used in a debug session.

Source Window

The source window, illustrated in [Figure 2.2](#) displays the source code for the current location. Use the source window to control the debug session, step through source files, set breakpoints, and browse source code.

Figure 2.2. Source Window



The source window contains a number of visual aids that allow you to know more about the execution of your code. The following sections describe these features.

Source and Disassembly Displays

Tabs for source, disassembly, and mixed display are contained by a tab that defines the process or thread being debugged, as illustrated in [Figure 2.2](#). When the current process or thread changes from one process or thread to another, the label on this tab will change and the contents of the display tab will be updated.

Choose between debugging at the source level, disassembly level, or with a mixture of source and disassembly. When source information is unavailable, only the disassembly tab will contain code.

The columns for line number or instruction address, debug event, program counter and location will be available in any display mode.

The line number column contains line numbers when displaying source code, instruction addresses when displaying disassembly, and a mixture of both in mixed mode. A grayed-out line number indicates a non-executable source line. Some examples of non-executable source lines are comments, non-applicable preprocessed code, some routine prologues, and some variable declarations. Breakpoints and other events cannot be set on non-executable lines.

The Event column indicates where debug events such as breakpoints or watchpoints exist. An event is indicated by a red sphere icon. Breakpoints may be set at any executable source line by left-clicking in the Event column at the desired source line. An existing breakpoint may be deleted by left-clicking on its breakpoint icon.

The PC column is the home of a blue arrow icon which marks the current location of the program counter. In other words, this arrow marks where program execution is during a debug session.

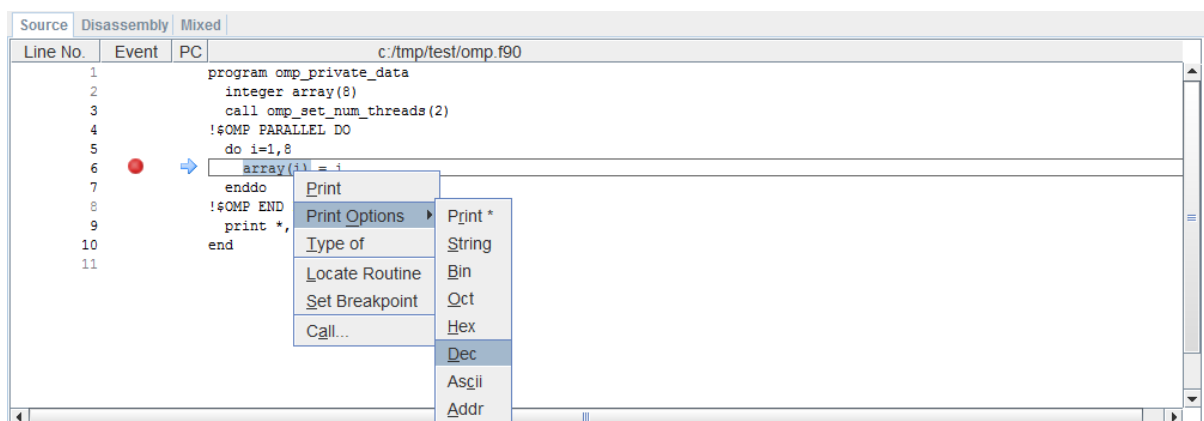
The title of the fourth column in the display windows is dependent on display mode. In the Source tab, this column will contain the name and path of the displayed source file. In the Disassembly and Mixed tabs, this column will contain the name of the disassembled function.

Source Window Context Menu

The display tabs in the source window support a context menu that provides convenient access to commonly used features. To bring up this context menu, first select a line in the source or disassembly code by clicking on it. Within the selected line, highlight a section of the text and right-click with the mouse to produce the menu. The context menu options use the selected text as input.

In the example in [Figure 2.3](#), the variable `array(i)` is highlighted and the context menu is set to print its value as a decimal integer:

Figure 2.3. Context Menu



The context menu in [Figure 2.3](#) also provides shortcuts to the [Type Of](#), [Locate Routine...](#), [Set Breakpoint...](#), and [Call...](#) menu options.

Call Stack Navigation

The call frame selector consists of the dropdown list and two navigation buttons. It sits below the display tabs in the source window.

Figure 2.4. Call Frame Selector




You can select any frame on the current call stack by using the call frame's drop-down list. The call stack can also be navigated using the up and down buttons to the right of the drop-down list. For more information on stack navigation, refer to the [up](#) and [down](#) commands.

Status Message Area

PGDBG displays debugging status messages in a messages area at the very bottom of the GUI.

Figure 2.5. Status Messages




Main Toolbar

The *PGDBG* GUI's main toolbar contains several buttons and four drop-down lists.

Buttons

Figure 2.6. Buttons on Toolbar



Most of the buttons on the main toolbar have corresponding entries on the Debug menu. The functionality invoked from the toolbar is the same as that achieved by selecting the menu item. Refer to the [“Debug Menu”](#) descriptions for details on how Run, Stop, Restart, Step Into, Step Over, Step Out, and PC work.

The Cancel button is grayed-out unless an operation is in progress that can be cancelled before it is completed.

Drop-Down Lists

Figure 2.7. Drop-Down Lists on Toolbar



As illustrated in [Figure 2.7](#), the main toolbar contains four drop-down lists. A drop-down list displays information while also offering an opportunity to change the displayed information if other choices are available. When no or one choice is available, a drop-down list is grayed-out. When more than one choice is available, the drop-down arrow in the component can be clicked to display the available choices.

Current Process or Current Thread

The first drop-down list displays the current process or current thread. The list's label changes depending on whether processes or threads are described. When more than one process or thread is available, use this drop-down list to specify which process or thread should be the current one. The current process or thread controls the contents of the source and disassembly display tabs. The function of this drop-down list is the same as that of the Procs & Threads tab in the debug information tabs.

Apply

The second drop-down list is labeled Apply. The selection in the Apply drop-down determines the set of processes and threads to which action commands are applied. Action commands are those that control program execution and include, for example, **cont**, **step**, **next**, and **break**. By default, action commands are applied to all processes and threads. When more than one process or thread exists, you have additional options in this drop-down list from which to choose. The Current Group option designates the process and thread group selected in the Groups tab, and the Current Process and Current Thread options designate the process or thread selected in the Current Process or Current Thread drop-down.

Display

The third drop-down list is labeled Display. The selection in the Display drop-down determines the set of processes and threads to which data display commands are applied. Data display commands are those that print the values of expressions and program state and include, for example, **print**, **names**, **regs** and **stack**. The options in the Display drop-down are the same as those in the Apply drop-down but can be changed independently.

File

The fourth drop-down list is labeled File. It displays the source file that contains the current target location. It can be used to select another file for viewing in the source window.

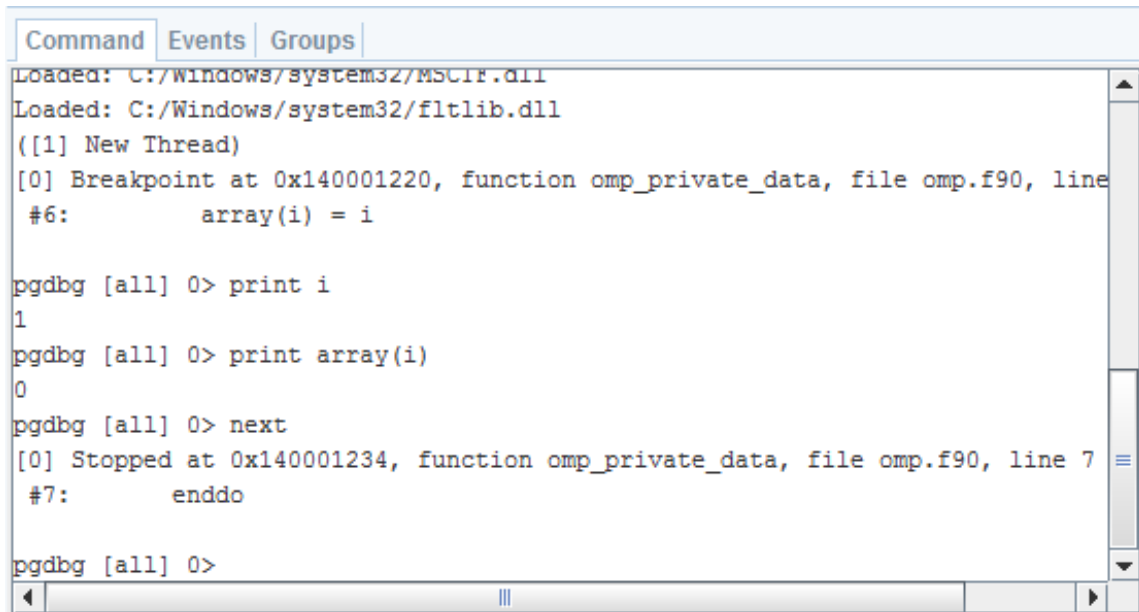
Debug Information Tabs

Debug information tabs take up the lower half of the *PGDBG* GUI. Each of these tabs provides a particular function or view of debug information. The following sections discuss the tabs as they appear from left-to-right in the GUI's default configuration.

Command Tab

The Command tab provides an interface in which to use the *PGDBG* command language. Commands entered in this panel are executed and the results are displayed there.

Figure 2.8. Command Tab

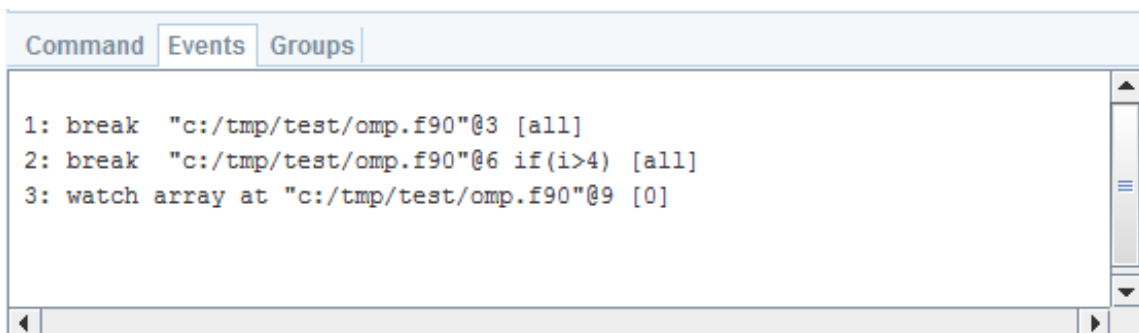


Using this tab is much like interacting with the debugger in text mode; the same list of commands is supported. For a complete list of commands, refer to [Chapter 5, “Command Summary”](#).

Events tab

The Events tab displays the current set of events held by the debugger. Events include breakpoints and watchpoints, as illustrated in [Figure 2.9](#).

Figure 2.9. Events Tab



Groups Tab

The Groups tab displays the current set of user-defined groups of processes and threads. The group selected (highlighted) in the Groups tab defines the Current Group as used by the Apply and Display drop-down lists. In [Figure 2.10](#), the ‘evens’ group is the Current Group.

Figure 2.10. Groups Tab

Command Events Groups	
Group Name	Composition
all	[*. *]
evens	[*.0, *.2]
odds	[*.1, *.3]
<div> Add... Modify... Remove... </div>	

To change the set of defined groups use the Add..., Modify..., and Remove... buttons on the Groups tab.

Note

A defined group of processes and threads is also known as a process/thread-set or p/t-set. For more information on p/t-sets, refer to “[p/t-set Notation](#)” in [Chapter 9](#), “[Parallel Debugging Overview](#)”.

Call Stack Tab

The Call Stack tab displays the current call stack. An arrow ‘=>’ indicates the current stack frame.

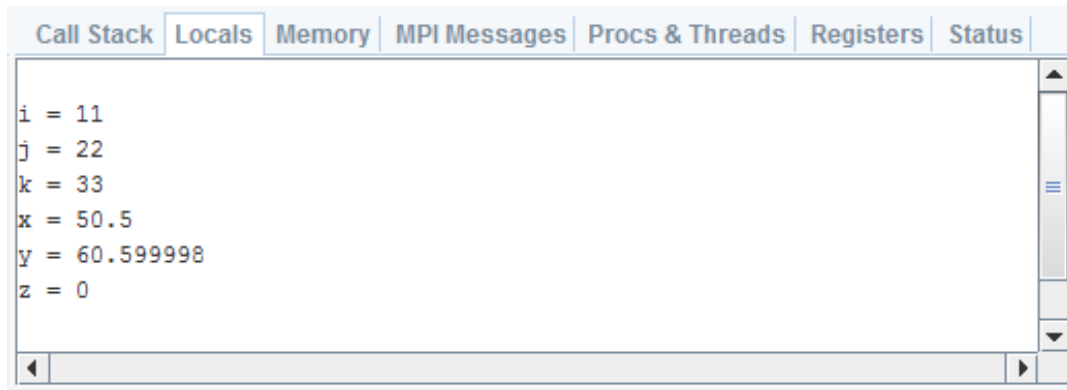
Figure 2.11. Call Stack Tab

Call Stack	Locals	Memory	MPI Messages	Procs & Threads	Registers	Status
STACK TRACE: #5 prog line: "subs.f90"@3 address: 0x140001158 #4 sub1 line: "subs.f90"@9 address: 0x140001182 #3 sub2 line: "subs.f90"@15 address: 0x1400011A2 #2 sub3 line: "subs.f90"@21 address: 0x1400011C2 #1 sub4 line: "subs.f90"@27 address: 0x1400011E2 => #0 sub5 line: "subs.f90"@33 address: 0x1400011FD						

Locals Tab

The Locals tab displays the current set of local variables and each of their values.

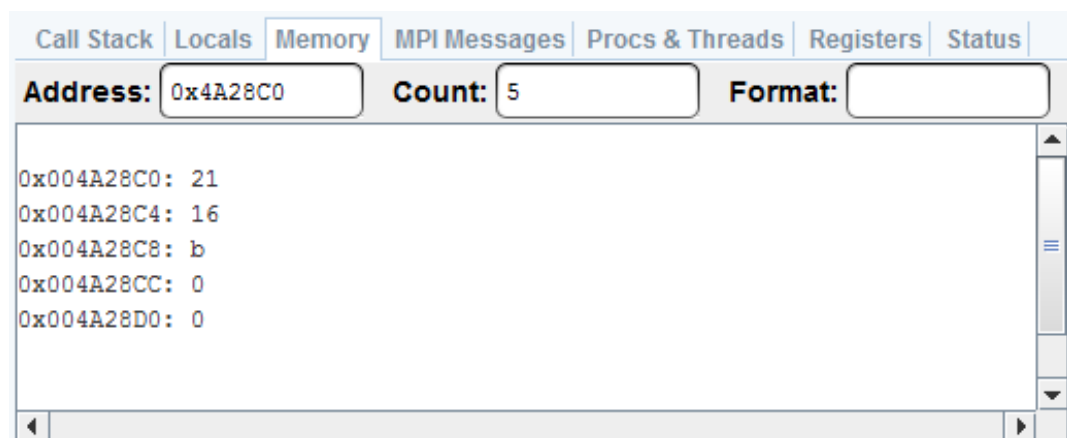
Figure 2.12. Locals Tab



Memory Tab

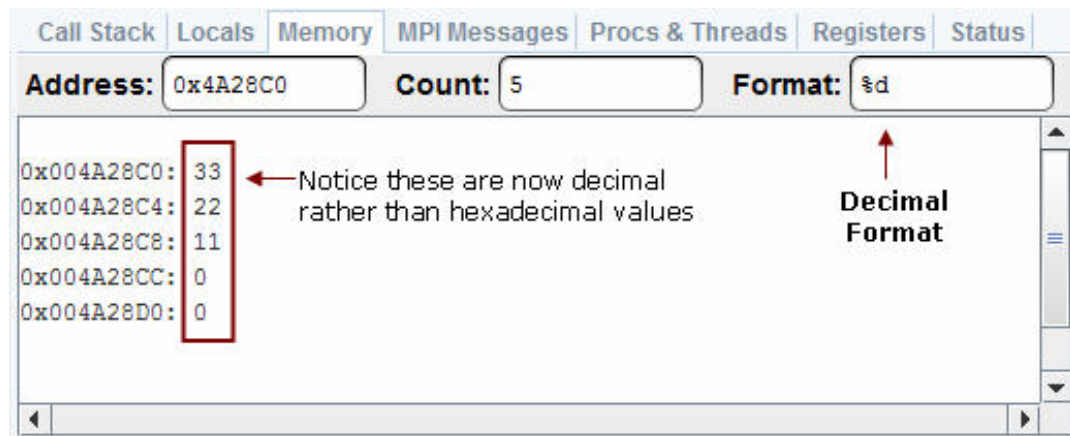
The Memory tab displays a region of memory starting with a provided Address which can be a memory address or a symbol name. One element of memory is displayed by default, but this amount can be changed via the Count field. [Figure 2.13](#) illustrates this process.

Figure 2.13. Memory Tab



The default display format for memory is hexadecimal. The display format can be changed by providing a printf-like format descriptor in the Format field. A detailed description of the supported format strings is available in [“Memory Access”](#) in [Chapter 13, “Command Reference”](#).

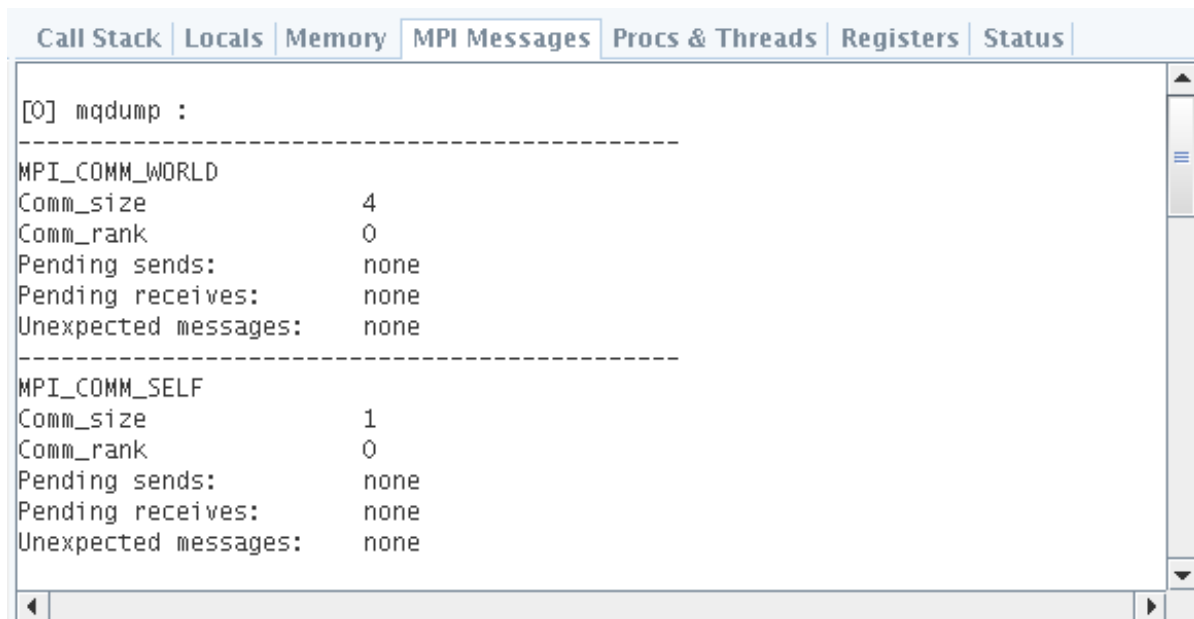
Figure 2.14. Memory Tab in Decimal Format



MPI Messages Tab

The MPI Messages tab provides a listing of the MPI message queues as by [Figure 2.15](#).

Figure 2.15. MPI Messages Tab



Message queue information applies only to MPI applications. When debugging a non-MPI application, this tab is empty. Additionally, message queue information is not supported by Microsoft MPI so this tab contains no data on Windows.

Procs & Threads Tab

The Procs & Threads tab provides a graphical display of the processes and threads in a debug session.

The Process Grid in [Figure 2.16](#) has four processes. The thicker border around process 0 indicates that it is the current process; its threads are represented pictorially. Thread 0.0, as the current thread of the current process, has the thickest border. Clicking on any process or thread in this grid changes that process or thread to be the current process or thread.

Figure 2.16. Process (Thread) Grid Tab



Use the slider at the bottom of the grid to zoom in and out.

The color of each element indicates the state of that process or thread. For a list of colors and states, refer to [Table 2.1](#).

Table 2.1. Colors Describing Thread State

Option	Description
Stopped	Red
Signaled	Blue
Running	Green
Terminated	Black

Registers Tab

The target machine's architecture determines the number and type of system registers. Registers are organized into groups based on their type and function. Each register group is displayed in its own tab contained in the Registers tab. Registers and their values are displayed in a table. Values are shown for all the threads of the currently selected process.

In [Figure 2.17](#), the General Purpose registers are shown for threads 0-3 of process 0.

Figure 2.17. General Purpose Registers

Call Stack Locals Memory MPI Messages Procs & Threads Registers Status					
GP FLAGS X87 XMM MXCSR					
Format: hex 64 Mode: scalar					
P0	T0	T1	T2	T3	
rax	0x2	0x2	0x2	0x2	
rbx	0x2E3150	0x2E3150	0x2E3150	0x2E3150	
rcx	0x0	0x2	0x4	0x6	
rdx	0x0	0x0	0x0	0x0	
rdi	0x1	0x1	0x1	0x1	
rsi	0x0	0x0	0x0	0x0	
rbp	0x12FF30	0x12FF30	0x12FF30	0x12FF30	
rsp	0x12FC80	0x250FC80	0x2D0FC80	0x350FC80	
r8	0x8	0x8	0x8	0x8	
r9	0x40	0x250FEA0	0x2D0FEA0	0x350FEA0	
r10	0x0	0x0	0x0	0x0	
r11	0x140001175	0x140001175	0x140001175	0x140001175	
r12	0x0	0x0	0x0	0x0	
r13	0x0	0x0	0x0	0x0	
r14	0x0	0x0	0x0	0x0	
r15	0x0	0x0	0x0	0x0	

The values in the registers table are updated each time the program stops. Values that change from one stopping point to the next are highlighted in yellow.

Register values can be displayed in a variety of formats. The formatting choices provided for each register group depends on the type of registers in the group and whether the current operating system is 64- or 32-bit. Use the Format drop-down list to change the displayed format.

Vector registers, such as XMM and YMM registers, can be displayed in both scalar and vector modes. Change the Mode drop-down list to switch between these two modes.

Status Tab

The Status tab provides a text summary of the status of the program being debugged. The state and location of each thread of each process is shown. In [Figure 2.18](#), each of four processes has two threads.

Figure 2.18. Status Tab

Call Stack Locals Memory MPI Messages Procs & Threads Registers Status					
0	ID	PID	STATE	SIG/CODE	LOCATION
=>	0	5792	Stopped	TRAP	prog line: "mpi.f90"@11 address: 0x140001179
	1	3432	Stopped	STOP	NtWaitForMultipleObjects address: 0x774D046A
1	ID	PID	STATE	SIG/CODE	LOCATION
=>	0	5288	Stopped	TRAP	prog line: "mpi.f90"@11 address: 0x140001179
	1	5696	Stopped	STOP	NtWaitForMultipleObjects address: 0x774D046A
2	ID	PID	STATE	SIG/CODE	LOCATION
=>	0	4772	Stopped	TRAP	prog line: "mpi.f90"@11 address: 0x140001179
	1	5228	Stopped	STOP	NtWaitForMultipleObjects address: 0x774D046A
3	ID	PID	STATE	SIG/CODE	LOCATION
=>	0	5608	Stopped	TRAP	prog line: "mpi.f90"@11 address: 0x140001179
	1	4568	Stopped	STOP	NtWaitForMultipleObjects address: 0x774D046A

Menu Bar

The main menu bar contains these menus: File, Edit, View, Data, Debug and Help. This section describes these menus and their contents.

You can navigate the menus using the mouse or the system's mouseless modifier (typically the Alt key). Use the mouseless modifier together with a menu's mnemonic, usually a single character, to select a menu and then a menu item. Menu mnemonics are indicated with an underscore. For example, the File menu appears as File which indicates that 'F' is the mnemonic.

Keyboard shortcuts, such as Ctrl+V for Edit | Paste, are available for some actions. Where a keyboard shortcut is available, it is shown in the GUI on the menu next to the menu item.

Menu items that contain an ellipsis (...) launch a dialog box to assist in performing the menu's action.

File Menu

Open Target...

Begin debugging a new program. After choosing this option, select the program (i.e., the target) to debug from the file chooser dialog. The current target, if any, is automatically closed and replaced with the new target. For more information, refer to the [debug](#) command.

Attach to Target...

Attach to a locally running process. Attach is not supported for remote processes. For more information, refer to the [attach](#) command.

Detach Target

Detach from the current attached-to process. This command does not terminate the target application if it is running. This option should only be used when Attach to Target... was used to begin the debug session. For more information, refer to the [detach](#) command.

Exit

End the current debug session and close all windows.

Edit Menu**Copy**

Copy selected text to the system's clipboard.

Paste

Paste selected text to the system's clipboard.

Search Forward...

Perform a forward string search in the currently displayed source file.

Search Backward...

Perform a backward string search in the currently displayed source file.

Search Again

Repeat the last search that was performed in the source panel.

Locate Routine...

Find a routine. If symbol and source information is available for the specified routine, the routine is displayed in the source panel.

Restore Default Settings

Restore the GUI's various settings to their initial default state illustrated in [Figure 2.1, "Default Appearance of PGDBG GUI,"](#) on page 5.

Revert to Saved Settings

Restore the GUI to the state that it was in at the start of the debug session.

Save Settings on Exit

By default, *PGDBG* saves the state (size and settings) of the GUI on exit on a per-system basis. To prevent settings from being saved from one invocation of *PGDBG* to another, uncheck this option. This option must be unchecked prior to every exit since *PGDBG* always defaults to saving the GUI state.

View Menu

Use the View menu to customize the PGDBG GUI.

Many of the items on this menu contain a check box next to the name of a tab.

- When the check box is checked, the tab is visible in the GUI.
- When the check box is not checked, the tab is hidden.

View menu items that correspond to tabs include Call Stack, Command, Events, Groups, Locals, Memory, MPI Messages, Procs & Threads, Source and Status.

Registers

The Registers menu item opens a submenu containing items representing every subtab on the Registers tab. Recall that each subtab represents a register group and the set of register groups is system and

architecture dependent. Use the Registers submenu to hide or show tabs for register groups. Use the Show Selected item to hide or show the Registers tab itself.

Font...

Use the font chooser dialog box to select the font and size used in the source window debug information tabs. The default font is named *monospace* and the default size is 12.

Show Tool Tips

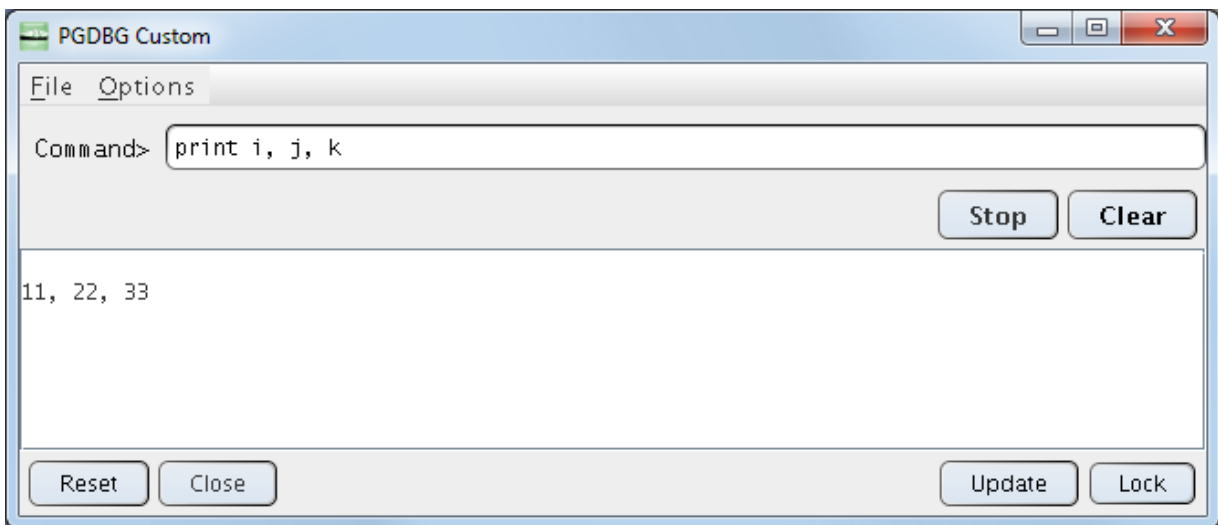
Tool tips are small temporary messages that pop up when the mouse pointer hovers over a component in the GUI. They provide additional information on the functionality of the component. Tool tips are enabled by default. Uncheck the Show Tools Tips option to prevent them from popping up.

Refresh

Update the source window and Procs & Threads tab.

Data Menu

Each Data menu item opens a Custom subwindow. The Custom subwindow provides a command field where any debugger-supported command can be entered. For example:



The Custom menu item is always enabled. It opens a Custom subwindow with a blank Command field. All the other Data menu items are enabled only when text (usually data) is selected in the display tabs. To select text, first click on a line in the Source, Disassembly or Mixed tab. Within the selected line, highlight a section of the text. With the text highlighted, open the Data menu and select the desired option.

Print

Print the value of the selected item.

Print *

Dereference and print the value of the selected item.

String

Treat the selected value as a string and print its value.

Bin

Print the value of the selected item as a base-2 integer.

Oct

Print the value of the selected item as an octal integer.

Hex

Print the value of the selected item as a hexadecimal integer.

Dec

Print the value of the selected item as a decimal integer.

Ascii

Print the ASCII value of the selected item.

Addr

Print the address of the selected item.

Type Of

Print data type information for the selected item.

Custom

Open an empty Custom subwindow. Execute any supported debugging command in its Command field.

Debug Menu

The items under this menu control the execution of the target application.

Set Breakpoint...

Set a breakpoint at the first executable source line in the specified routine.

Display Current Location

Display the current program location in the Source panel. For more information, refer to the **arrive** command.

Up

Enter the scope of the routine up one level in the call stack. For more information, refer to the **up** command.

Down

Enter the scope of the routine down one level in the call stack. For more information, refer to the **down** command.

Run

Run the target application. If the target application is currently running, this item will rerun the application from the beginning. For more information, refer to the commands **run** and **rerun**.

Set Run Arguments...

Modify the target application's runtime arguments.

Halt

Halt the running processes or threads. For more information, refer to the **halt** command.

Call...

Specify a routine to call. For more information, refer to the **call** command.

Cont

Continue execution from the current location. For more information, refer to the [cont](#) command.

Step

Continue and stop after executing one source line or one assembly-level instruction depending on whether the Source, Disassembly or Mixed tab is displayed. Step steps *into* called routines. For more information, refer to the [step](#) and [stepi](#) commands.

Next

Continue and stop after executing one source line or one assembly-level instruction depending on whether the Source, Disassembly or Mixed tab is displayed. Next steps *over* called routines. For more information, refer to the [next](#) and [nexti](#) commands.

Step Out

Continue and stop after returning to the caller of the current routine. For more information, refer to the [stepout](#) command.

Help Menu

PGDBG Help...

The *PGDBG* GUI provides an integrated help utility that includes information about every *PGDBG* command. To find a command, use one of the following tabs in the left panel: The “book” tab presents a table of contents, the “index” tab presents an index of commands, and the “magnifying glass” tab presents a search engine. Each help page, displayed on the right, may contain hyperlinks, denoted in underlined blue, to terms referenced elsewhere in the help engine. Use the arrow buttons to navigate between visited pages. Use the printer buttons to print the current help page.

About PGDBG...

This option displays a dialog box with version and copyright information on *PGDBG*. It also contains sales and support points of contact.

Chapter 3. Command Line Options

As we stated in [Chapter 1, “Getting Started”](#), *PGDBG* can interpret command-line options when present on the command line. This chapter describes these options and how they are interpreted.

Command-Line Options Syntax

The **pgdbg** command accepts several command-line options.

These options must appear on the command line before the name of the program being debugged.

Command-Line Options

The valid **pgdbg** options are these:

-attach *<pid>*

Attach to a running process with the process ID *<pid>*.

-c *<pgdbg_cmd>*

Execute the debugger command *pgdbg_cmd* before executing the commands in the startup file.

-cd *<workdir>*

Sets the working directory to the specified directory.

-core *<corefile>*

Analyze the core dump named *corefile*. [Linux only]

-dbx

Start the debugger in dbx mode, which provides a dbx-like debugger command language.

-emacs

Invoke the debugger using the Emacs GUD interface.

-help

Display a list of command-line arguments (this list).

-I *<directory>*

Add *<directory>* to the list of directories that *PGDBG* uses to search for source files. You can use this option multiple times to add multiple directories to the search path.

`-jarg, <javaarg>`

Pass specified argument(s) (separated by commas) to java, e.g. `-jarg,-Xmx256m`.

`-java <jrepath>`

Add a jrepath directory to the JVM search path. Multiple '-java' options are allowed.

`-nocheckjvm`

Skip the test for the correct JVM architecture.

`-nomin`

Do not minimize the *PGDBG* console shell on startup. [Windows only]

`-s <pgdbg_script>`

Runs a **pgdbg** command script instead of `.pgdbgrc` [Linux and MacOS] or `pgdbrc` [Windows].

`-text`

Run the debugger using a command-line interface (CLI). The default is for the debugger to launch in graphical user interface (GUI) mode.

Command-Line Options for MPI Debugging

`-mpi[=<mpiexec_path>]`

Start/debug an MPI job.

`-pgserv[=<pgserv_path>]`

Specify path for pgserv, the per-node debug agent.

`-program_args`

Start/debug an MPI job. `-program_args` delimits mpiexec args.

Chapter 4. Command Language

PGDBG supports a command language that is capable of evaluating complex expressions. The command language is composed of commands, constants, symbols, locations, expressions, and statements.

You can use the command language by invoking the *PGDBG* command-line interface with the `-text` option, or in the Command tab of the *PGDBG* graphical user interface, as described in “[The Graphical User Interface](#)”.

Command Overview

Commands are named operations, which take zero or more arguments and perform some action. Commands may also return values that may be used in expressions or as arguments to other commands.

Command Syntax

Commands are entered one line at a time.

- Lines are delimited by a carriage return.
- Each line must consist of a command and its arguments, if any.
- You can place multiple commands on a single line by using the semi-colon (;) as a delimiter.

Command Modes

There are two command modes: **pgi** and **dbx**.

- The **pgi** command mode maintains the original *PGDBG* command interface.
- In **dbx** mode, the debugger uses commands compatible with the Unix-based **dbx** debugger.

PGI and **dbx** commands are available in both command modes, but some command behavior may be slightly different depending on the mode. The mode can be set when *PGDBG* is invoked by using command-line options, or while the debugger is running by using the **pgienv** command.

Constants

PGDBG supports C language style integer (hex, octal and decimal), floating point, character, and string constants.

Symbols

PGDBG uses the symbolic information contained in the executable object file to create a symbol table for the target program. The symbol table contains symbols to represent source files, subprograms (functions, and subroutines), types (including structure, union, pointer, array, and enumeration types), variables, and arguments. The *PGDBG* command-line interface is case-sensitive with respect to symbol names; a symbol name on the command line must match the name as it appears in the object file.

Scope Rules

Since several symbols in a single application may have the same name, scope rules are used to bind program identifiers to symbols in the symbol table. *PGDBG* uses the concept of a search scope for looking up identifiers. The search scope represents a routine, a source file, or global scope. When the user enters a name, *PGDBG* first tries to find the symbol in the search scope. If the symbol is not found, the containing scope (source file or global) is searched, and so forth, until either the symbol is located or the global scope is searched and the symbol is not found.

Normally, the search scope is the same as the current scope, which is the routine where execution is currently stopped. The current scope and the search scope are both set to the current routine each time execution of the target program stops. However, you can use the **enter** command to change the search scope.

A scope qualifier operator @ allows selection of out-of-scope identifiers. For example, if *f* is a routine with a local variable *i*, then:

```
f@i
```

represents the variable *i* local to *f*. Identifiers at file scope can be specified using the quoted file name with this operator, for example:

```
"xyz.c"@i
```

represents the variable *i* defined in file *xyz.c*.

Register Symbols

To provide access to the system registers, *PGDBG* maintains symbols for them. Register names generally begin with \$ to avoid conflicts with program identifiers. Each register symbol has a default type associated with it, and registers are treated like global variables of that type, except that their address may not be taken. For more information on register symbols, refer to [“SSE Register Symbols,” on page 49](#).

Source Code Locations

Some commands must refer to source code locations. Source file names must be enclosed in double quotes. Source lines are indicated by number, and may be qualified by a quoted filename using the scope qualifier operator. Further, a range of lines is indicated using the range operator ":".

Here are some examples:

<code>break 37</code>	sets a breakpoint at line 37 of the current source file.
<code>break "xyz.c"@37</code>	sets a breakpoint at line 37 of the source file xyz.c.
<code>list 3:13</code>	lists lines 3 through 13 of the current file.
<code>list "xyz.c"@3:13</code>	lists lines 3 through 13 of the source file xyz.c.

Some commands accept both line numbers and addresses as arguments. In these commands, it is not always obvious whether a numeric constant should be interpreted as a line number or an address. The description for these commands says which interpretation is used. However, *PGDBG* provides commands to convert from source line to address and vice versa. The **line** command converts an address to a line, and the **addr** command converts a line number to an address. Here are some examples:

<code>line 37</code>	means "line 37"
<code>addr 0x1000</code>	means "address 0x1000"
<code>addr {line 37}</code>	means "the address associated with line 37"
<code>line {addr 0x1000}</code>	means "the line associated with address 0x1000"

Lexical Blocks

Line numbers are used to name lexical blocks. The line number of the first instruction contained by a lexical block is used to indicate the start scope of the lexical block. In the following example, there are two variables named `var`. One is declared in function `main`, and the other is declared in the lexical block starting at line 5. The lexical block has the unique name `"lex.c"@main@5`. The variable `var` declared in `"lex.c"@main@5` has the unique name `"lex.c"@main@5@var`. The output of the **whereis** command that follows shows how these identifiers can be distinguished.

```
lex.c:
1 main()
2 {
3     int var = 0;
4     {
5         int var = 1;
6         printf("var %d\n",var);
7     }
8     printf("var %d\n",var)
9 }

pgdbg> n
Stopped at 0x8048b10, function main, file
/home/demo/pgdbg/ctest/lex.c,
line 6
#6: printf("var %d\n",var);
pgdbg> print var
1
pgdbg> which var
"lex.c"@main@5@var
pgdbg> whereis var
variable: "lex.c"@main@var
variable: "lex.c"@main@5@var
pgdbg> names "lex.c"@main@5
var = 1
```

Statements

Although *PGDBG* command-line input is processed one line at a time, statement constructs allow multiple commands per line, as well as conditional and iterative execution. The statement constructs roughly correspond to the analogous C language constructs. Statements may be of the following forms.

- *Simple Statement*: A command and its arguments. For example:

```
print i
```

- *Block Statement*: One or more statements separated by semicolons and enclosed in curly braces. Note: these may only be used as arguments to commands or as part of `if` or `while` statements. For example:

```
if(i>1) {print i; step }
```

- *If Statement*: The keyword `if`, followed by a parenthesized expression, followed by a block statement, followed by zero or more `else if` clauses, and at most one `else` clause. For example:

```
if(i>j) {print i} else if(i<j) {print j} else {print "i==j"}
```

- *While Statement*: The keyword `while`, followed by a parenthesized expression, followed by a block statement. For example:

```
while(i==0) {next}
```

Multiple statements may appear on a line separated by a semicolon. The following example sets breakpoints in routines `main` and `xyz`, continues, and prints the new current location.

```
break main; break xyz; cont; where
```

However, since the **where** command does not wait until the target application has halted, this statement displays the stack at some arbitrary execution point in the program. To control when the stack is printed, insert a **wait** command, as shown in this example:

```
break main; break xyz; cont; wait; where
```

Note

Any value returned by the last statement on a line is printed.

Statements can be parallelized across multiple threads of execution. For more information, refer to [“Parallel Statements,” on page 75](#).

Events

Breakpoints, watchpoints, and other mechanisms used to define the response to certain conditions are collectively called *events*.

- An event is defined by the conditions under which the event occurs and by the action taken when the event occurs.
- A breakpoint occurs when execution reaches a particular address.

The default action for a breakpoint is simply to halt execution and prompt the user for commands.

- A watchpoint occurs when the value of an expression changes.
- A hardware watchpoint occurs when the specified memory location is accessed or modified.

Event Commands

PGDBG supports six basic commands for defining events. Each command takes a required argument and may also take one or more optional arguments. The basic commands are **break**, **watch**, **hwatch**, **trace**, **track**, and **do**.

Event Command Descriptions

- The **break** command takes an argument specifying a breakpoint location. Execution stops when that location is reached.
- The **watch** command takes an expression argument. Execution stops and the new value is printed when the value of the expression changes.
- The **hwatch** command takes a data address argument, which can be either an identifier or a variable name. Execution stops when memory at that address is written.
- The **trace** command activates source line tracing, as specified by the arguments you supply.
- The **track** command is like **watch** except that execution continues after the new value is printed.
- The **do** command takes a list of commands as an argument. The commands are executed whenever the event occurs.

Event Command Arguments

The six event commands share a common set of optional arguments. The optional arguments provide the ability to make the event definition more specific. They are:

at line

Event occurs at indicated line.

at addr

Event occurs at indicated address.

in routine

Event occurs throughout indicated routine.

if (condition)

Event occurs only when condition is true.

do {commands}

When event occurs, execute commands.

The optional arguments may appear in any order after the required argument and should not be delimited by commas.

Event Command Examples

Here are some event definition examples:

<code>watch i at 37 if(y>1)</code>	This event definition says to stop and print the value of <code>i</code> whenever line 37 is executed and the value of <code>y</code> is greater than 1.
<code>do {print xyz} in f</code>	This event definition says that at each line in the routine <code>f</code> print the value of <code>xyz</code> .
<code>break func1 if (i==37) do {print a[37]; stack}</code>	This event definition says to print the value of <code>a[37]</code> and do a stack trace when <code>i</code> is equal to 37 in routine <code>func1</code> .

Event Command Action

It is useful to know when events take place.

- Event commands that do not explicitly define a location occur at each source line in the program. Here are some examples:

<code>do {where}</code>	prints the current location at the start of each source line.
<code>trace a.b</code>	prints the value of <code>a.b</code> each time the value has changed.
<code>track a.b</code>	prints the value of <code>a.b</code> at the start of each source line if the value has changed.

Note

Events that occur at every line can be useful, but they can make program execution very slow. Restricting an event to a particular address minimizes the impact on program execution speed, and restricting an event that occurs at every line to a single routine causes execution to be slowed only when that routine is executed.

- *PGDBG* supports instruction-level versions of several commands, such as **breaki**, **watchi**, **tracei**, **tracki**, and **doi**. The basic difference in the instruction-level version is that these commands interpret integers as addresses rather than line numbers, and events occur at each instruction rather than at each line.
- When multiple events occur at the same location, all event actions are taken before the prompt for input. Defining event actions that resume execution is allowed but discouraged, since continuing execution may prevent or defer other event actions.

For example, the following syntax creates an ambiguous situation:

```
break 37 do {continue}
break 37 do {print i}
```

With this sequence, it is not clear whether `i` will ever be printed.

- Events only occur after the **continue** and **run** commands. They are ignored by **step**, **next**, **call**, and other commands.

- Identifiers and line numbers in events are bound to the current scope when the event is defined.

For example, the following command sets a breakpoint at line 37 in the current file.

```
break 37
```

The following command tracks the value of whatever variable `i` is currently in scope.

```
track i
```

If `i` is a local variable, then it is wise to add a location modifier (`at` or `in`) to restrict the event to a scope where `i` is defined. Scope qualifiers can also specify lines or variables that are not currently in scope. Events can be parallelized across multiple threads of execution. See [“Parallel Events,” on page 74](#) for details.

Expressions

The debugger supports evaluation of expressions composed of constants, identifiers, commands that return values, and operators.

The following rules apply:

- To use a value returned by a command in an expression, the command and arguments must be enclosed in curly braces.

For example, the following command invokes the `pc` command to compute the current address, adds 8 to it, and sets a breakpoint at that address.

```
breaki {pc}+8
```

Similarly, the following command compares the start address of the current routine with the start address of routine `xyz`. It prints the value 1 if they are equal and 0 if they are not.

```
print {addr {func}}=={addr xyz}
```

- The `@` operator, introduced previously, may be used as a scope qualifier. Its precedence is the same as the C language field selection operators `."` and `"->"`.
- `PGDBG` recognizes a range operator `:"` which indicates array sub-ranges or source line ranges. The precedence of `:` is between `||` and `'='`.

Here are a few examples that use the range operator:

<code>print a[1:10]</code>	prints elements 1 through 10 of the array <code>a</code> .
<code>list 5:10</code>	lists source lines 5 through 10.
<code>list "xyz.c"@5:10</code>	lists lines 5 through 10 in file <code>xyz.c</code> .

The general format for the range operator is `[lo : hi : step]` where:

<code>lo</code>	is the array or range lower bound for this expression.
<code>hi</code>	is the array or range upper bound for this expression.
<code>step</code>	is the step size between elements.

- An expression can be evaluated across many threads of execution by using a prefix p/t-set. For more details, refer to [“Current vs. Prefix p/t-set,” on page 63](#).

[Table 4.1, “PGDBG Operators”](#) shows the C language operators that *PGDBG* supports. The *PGDBG* operator precedence is the same as in the C language.

Table 4.1. PGDBG Operators

Operator	Description	Operator	Description
*	indirection	<=	less than or equal
.	direct field selection	>=	greater than or equal
->	indirect field selection	!=	not equal
[]	C/ C++ array index	&&	logical and
()	routine call		logical or
&	address of	!	logical not
+	add		bitwise or
(type)	cast	&	bitwise and
-	subtract	~	bitwise not
/	divide	^	bitwise exclusive or
*	multiply	<<	left shift
=	assignment	>>	right shift
==	comparison	()	FORTTRAN array index
<<	left shift	%	FORTTRAN field selector
>>	right shift		

Ctrl-C

If the target application is not running, Ctrl-C can be used to interrupt long-running *PGDBG* commands. For example, a command requesting disassembly of thousands of instructions might run for a long time, and it can be interrupted by Ctrl-C. In such cases the target application is not affected.

If the target application is running, entering Ctrl-C at the *PGDBG* command prompt halts execution of the target. This is useful in cases where the target “hangs” due to an infinite loop or deadlock.

Sending Ctrl-C, also known as a SIGINT, to a program while it is in the middle of initializing its threads, by calling `omp_set_num_threads()` or entering a parallel region, may kill some of the threads if the signal is sent before each thread is fully initialized. Avoid sending SIGINT in these situations. Note that when the number of threads employed by a program is large, thread initialization may take a while.

Sending Ctrl-C to a running MPICH-1 program is not recommended. See [“Use halt instead of Ctrl+C,” on page 82](#), for details. Use the *PGDBG* **halt** command as an alternative to sending Ctrl-C to a running program. The *PGDBG* command prompt must be available in order to issue a **halt** command. The *PGDBG* command prompt is available while threads are running if **pgienv threadwait none** is set.

As described in [“Invoking PGDBG for MPI Debugging,” on page 84](#), when debugging an MPI job via the command `pgdbg -mpi . . .`, *PGDBG* spawns the job in a manner that prevents console-generated interrupts from directly reaching the MPI job launcher or any of the MPI processes. In this case, typing Ctrl-C only interrupts *PGDBG*, leaving the MPI processes running. When *PGDBG*’s thread wait mode is not set to none, you can halt the MPI job after using Ctrl-C by entering *PGDBG*’s **halt** command, even if no *PGDBG* prompt is generated.

Chapter 5. Command Summary

This chapter contains a brief summary of the *PGDBG* debugger commands. For a detailed description of each command, grouped by category of use, refer to [Chapter 13, “Command Reference”](#).

If you are viewing an online version of this manual, you can select the hyperlink under the selection category to jump to that section in the manual.

Notation Used in Command Sections

The command sections that follow use these conventions for the command names and arguments, when the command accepts one.

- Command names may be abbreviated by omitting the portion of the command name enclosed in brackets ([]).
- Argument names are chosen to indicate what kind of argument is expected.
- Arguments enclosed in brackets ([]) are optional.
- Two or more arguments separated by a vertical line (|) indicate that any one of the arguments is acceptable.
- An ellipsis (...) indicates an arbitrarily long list of arguments.
- Other punctuation (commas, quotes, etc.) should be entered as shown.

For example, the following syntax indicates that the command **list** may be abbreviated to **lis**, and that it can be invoked without any arguments or with *one* of the following arguments: an integer count, a line range, a routine name, or a line and a count.

```
lis[t] [count | lo:hi | routine | line,count]
```

Command Summary

Table 5.1. PGDBG Commands

Name	Arguments	Category
<i>ad[dr]</i>	[n line n routine var arg]	“Conversions,” on page 114
	Creates an address conversion under certain conditions.	
<i>al[ias]</i>	[name [string]]	“Miscellaneous,” on page 116
	Create or print aliases.	
<i>args</i>		“Process Control,” on page 92
	Print the current program arguments.	
<i>arri[ve]</i>		“Program Locations,” on page 103
	Print location information for the current location.	
<i>asc[ii]</i>	exp [...exp]	“Printing Variables and Expressions,” on page 105
	Evaluate and print as an ascii character.	
<i>as[ign]</i>	var=exp	“Symbols and Expressions,” on page 108
	Set variable <code>var</code> to the value of the expression <code>exp</code> .	
<i>att[ach]</i>	pid [exe]	“Process Control,” on page 92
	Attach to a running process with process ID <code>pid</code> . Use <code>exe</code> to specify the absolute path of the executable file.	
<i>bin</i>	exp [...exp]	“Printing Variables and Expressions,” on page 105
	Evaluate and print the expressions. Integer values are printed in base 2.	
<i>b[reak]</i>	[line routine] [if (condition)] [do {commands}]	“Events,” on page 96
	When arguments are specified, sets a breakpoint at the indicated line or routine. When no arguments are specified, prints the current breakpoints.	
<i>breaki</i>	[addr routine] [if (condition)] [do {commands}]	“Events,” on page 96
	When arguments are specified, sets a breakpoint at the indicated address or routine. When no arguments are specified, prints the current breakpoints.	
<i>breaks</i>		“Events,” on page 96
	Displays all the existing breakpoints	

Name	Arguments	Category
<i>call</i>	routine [(exp,...)]	“Symbols and Expressions,” on page 108
	Call the named routine.	
<i>catch</i>	[number [,number...]]	“Events,” on page 96
	With arguments, catches the signals and runs target as though signal was not sent. With no arguments, prints the list of signals being caught.	
<i>cd</i>	[dir]	“Program Locations,” on page 103
	Change to the \$HOME directory or to the specified directory dir.	
<i>clas[s]</i>	[class]	“Scope,” on page 110
	Return the current class or enter the scope of the specified class <code>class</code> .	
<i>classe[s]</i>		“Target,” on page 115
	Print the C++ class names.	
<i>clear</i>	[all routine line addr {addr}]	“Events,” on page 96
	With arguments, clears the indicated breakpoints. When no arguments are specified, this command clears all breakpoints at the current location.	
<i>con[nect]</i>	[-t name [args] -d path [args] -f file [name [args]]]	“Target,” on page 115
	Prints the current connection and the list of possible connection targets.	
<i>c[ont]</i>		“Process Control,” on page 92
	Continue execution from the current location.	
<i>cr[ead]</i>	addr	“Memory Access,” on page 113
	Fetch and return an 8-bit signed integer (character) from the specified address.	
<i>de[bug]</i>	[target [arg1 _ argn]]	“Process Control,” on page 92
	Load the specified target program with optional command-line arguments.	
<i>dec</i>	exp [,...exp]	“Printing Variables and Expressions,” on page 105
	Evaluate and print the expressions. Integer values are printed in decimal.	
<i>decl[aration]</i>	name	“Symbols and Expressions,” on page 108
	Print the declaration for the symbol based on its type according to the symbol table.	
<i>decls</i>	[routine "sourcefile" {global}]	“Scope,” on page 110
	Print the declarations of all identifiers defined in the indicated scope. If no scope is given, print the declarations for global scope.	

Name	Arguments	Category
<i>defset</i>	name [p/t-set]	“Process-Thread Sets,” on page 95
	Assign a name to a process/thread set. Define a named set.	
<i>del[ete]</i>	event-number all 0 event-number [,event-number.]	“Events,” on page 96
	Delete the event <code>event-number</code> or all events (delete 0 is the same as delete all). Multiple event numbers can be supplied if they are separated by commas.	
<i>det[ach]</i>		“Process Control,” on page 92
	Detach from the current running process.	
<i>dir[ectory]</i>	[pathname]	“Miscellaneous,” on page 116
	Add the directory pathname to the search path for source files. If no argument is specified, the currently defined directories are printed.	
<i>disab[le]</i>	event-number all	“Printing Variables and Expressions,” on page 105
	With arguments, disables the event <code>event-number</code> or all events. When no arguments are specified, prints both enabled and disabled events.	
<i>dis[asm]</i>	[count lo:hi routine addr, count]	“Program Locations,” on page 103
	Disassemble memory. If no argument is given, disassemble four instructions starting at the current address.	
<i>disc[onnect]</i>		“Events,” on page 96
	Close connection to target.	
<i>display</i>	[exp [,...exp]]	“Printing Variables and Expressions,” on page 105
	With an argument or several arguments, print expression <code>exp</code> at every breakpoint. Without arguments, list the expressions for <i>PGDBG</i> to automatically display at breakpoints.	
<i>do</i>	{commands} [at line in routine] [if (condition)]	“Events,” on page 96
	Define a <code>do</code> event. Without the optional arguments <code>at</code> or <code>in</code> , the commands are executed at each line in the program.	
<i>doi</i>	{commands} [at addr in routine] [if (condition)]	“Events,” on page 96
	Define a <code>doi</code> event. If neither the <code>at</code> or <code>in</code> argument is specified, then the commands are executed at each instruction in the program.	
<i>down</i>	[number]	“Scope,” on page 110
	Enter scope of routine down one level or <code>number</code> levels on the call stack.	
<i>dr[ead]</i>	addr	“Memory Access,” on page 113
	Fetch and return a 64 bit double from the specified address.	

Name	Arguments	Category
<i>du[mp]</i>	[addr [,count [,format]]]	“Memory Access,” on page 113
	Dumps the contents of a region of memory. The output is formatted according to a printf-like format descriptor.	
<i>edit</i>	[filename routine]	“Program Locations,” on page 103
	Edit the specified file or file containing the routine. If no argument is supplied, edit the current file starting at the current location. In the <i>PGDBG</i> GUI, command-line editors like <i>vi</i> , <i>emacs</i> , or <i>nano</i> are launched in the Program I/O Window.	
<i>enab[le]</i>	[event-number all]	“Events,” on page 96
	With arguments, this command enables the event <code>event-number</code> or all events. When no arguments are specified, prints both enabled and disabled events.	
<i>en[ter]</i>	[routine "sourcefile" global]	“Scope,” on page 110
	Set the search scope to be the indicated symbol, which may be a routine, source file or global. Using no argument is the same as using <code>enter global</code>	
<i>entr[y]</i>	[routine]	“Symbols and Expressions,” on page 108
	Return the address of the first executable statement in the program or specified routine.	
<i>fil[e]</i>	[filename]	“Program Locations,” on page 103
	Change the source file to the file <code>filename</code> and change the scope accordingly. With no argument, print the current file.	
<i>files</i>		“Scope,” on page 110
	Return the list of known source files used to create the executable file	
<i>focus</i>	[p/t-set]	“Process-Thread Sets,” on page 95
	Set the target process/thread set for commands. Subsequent commands are applied to the members of this set by default.	
<i>fp</i>		“Register Access,” on page 112
	Return the current value of the frame pointer.	
<i>fr[ead]</i>	addr	“Memory Access,” on page 113
	Fetch and print a 32-bit float from the specified address.	
<i>func[tion]</i>	[addr line]	“Conversions,” on page 114
	Return a routine symbol. If no argument is specified, return the current routine.	
<i>glob[al]</i>		“Scope,” on page 110
	Return a symbol representing global scope.	
<i>halt</i>	[command]	“Process Control,” on page 92
	Halt the running process or thread.	

Name	Arguments	Category
<i>be[lp]</i>	[command]	“Miscellaneous,” on page 116
	If no argument is specified, print a brief summary of all the commands. If a <code>command</code> name is specified, print more detailed information about the use of that command.	
<i>hex</i>	exp [...exp]	“Printing Variables and Expressions,” on page 105
	Evaluate and print expressions as hexadecimal integers.	
<i>hi[story]</i>	[num]	“Miscellaneous,” on page 116
	List the most recently executed commands. With the <code>num</code> argument, resize the history list to hold <code>num</code> commands.	
<i>hwatch</i>	addr var [if (condition)] [do {commands}]	“Events,” on page 96
	Define a hardware watchpoint.	
<i>hwatchb[otb]</i>	addr var [if (condition)] [do {commands}]	“Events,” on page 96
	Define a hardware read/write watchpoint.	
<i>hwatchr[eat]</i>	addr var [if (condition)] [do {commands}]	“Events,” on page 96
	Define a hardware read watchpoint.	
<i>ignore</i>	[number [,number...]]	“Events,” on page 96
	Ignore the specified signals and does not deliver them to the target. When no arguments are specified, prints the list of signals being ignored.	
<i>ir[eat]</i>	addr	“Memory Access,” on page 113
	Fetch and print a signed integer from the specified address.	
<i>language</i>		“Miscellaneous,” on page 116
	Print the name of the language of the current file.	
<i>lin[e]</i>	[n routine addr]	“Conversions,” on page 114
	Create a source line conversion. If no argument is given, return the current source line.	
<i>lines</i>	[routine]	“Program Locations,” on page 103
	Print the lines table for the specified routine. If no argument is specified, prints the lines table for the current routine.	
<i>lis[t]</i>	[count line,count lo:hi routine]	“Program Locations,” on page 103
	With no argument, list 10 lines centered at the current source line. If an argument is specified, list lines based on information requested.	

Name	Arguments	Category
<i>lo[ad]</i>	[prog [args]]	“Process Control,” on page 92
	Without options, print the name and arguments of the program being debugged. With arguments, invoke the debugger using the specified program and program args, if any.	
<i>log</i>	filename	“Miscellaneous,” on page 116
	Keep a log of all commands entered by the user and store it in the named file.	
<i>lr[ead]</i>	addr	“Memory Access,” on page 113
	Fetch and print an address from the specified address.	
<i>lv[al]</i>	exp	“Symbols and Expressions,” on page 108
	Return the lvalue of the expression expr.	
<i>mq[dump]</i>		“Memory Access,” on page 113
	Dump MPI message queue information for the current process.	
<i>names</i>	[routine "sourcefile" {global}]	“Scope,” on page 110
	Print the names of all identifiers defined in the indicated scope. If no scope is specified, use the search scope.	
<i>nat[ive]</i>	[command]	“Target,” on page 115
	Without arguments, print a list of the available target commands. With a command argument, send the native command directory to the target.	
<i>n[ext]</i>	[count]	“Process Control,” on page 92
	Stop after executing one or count source line(s) in the current routine.	
<i>nexti</i>	[count]	“Process Control,” on page 92
	Stop after executing one or count instruction(s) in the current routine.	
<i>nop[rint]</i>	exp	“Miscellaneous,” on page 116
	Evaluate the expression but do not print the result.	
<i>oct</i>	exp [,...exp]	“Printing Variables and Expressions,” on page 105
	Evaluate and print expressions as octal integers.	
<i>pc</i>		“Register Access,” on page 112
	Return the current program address.	
<i>pgienv</i>	[command]	“Miscellaneous,” on page 116
	Define the debugger environment. With no arguments, display the debugger settings.	
<i>p[rint]</i>	exp1 [,...expn]	“Printing Variables and Expressions,” on page 105
	Evaluate and print one or more expressions.	

Name	Arguments	Category
<i>printf</i>	"format_string", expr,...expr	“Printing Variables and Expressions,” on page 105
	Print expressions in the format indicated by the format string.	
<i>proc</i>	[id]	“Process Control,” on page 92
	Set the current process to the process identified by id. When issued with no argument, lists the location of the current thread of the current process in the current program.	
<i>procs</i>		“Process Control,” on page 92
	Print the status of all active processes, listing each process by its logical process ID.	
<i>pwd</i>		“Program Locations,” on page 103
	Print the current working directory.	
<i>q[uit]</i>		“Process Control,” on page 92
	Terminate the debugging session.	
<i>regs</i>	regs [-info] [-grp=grp1[,grp2...]] [-fmt=fmt1[,fmt2...]] [-mode=vector scalar]	“Register Access,” on page 112
	Print a formatted display of the names and values of registers. Specify the register group(s) with the <code>-grp</code> option and formatting with the <code>-fmt</code> option. Use <code>-info</code> to see a listing of available register groups and formats.	
<i>rep[eat]</i>	[first, last] [first: last:n] [num] [-num]	“Miscellaneous,” on page 116
	Repeat the execution of one or more previous history list commands.	
<i>rer[un]</i>	[arg0 arg1 ... argn] [< inputfile] [[> >& >> >>&] outputfile]	“Process Control,” on page 92
	Like the run command with one exception: if no args are specified with rerun , then no args are used when the program is launched.	
<i>ret[addr]</i>		“Register Access,” on page 112
	Return the current return address.	
<i>ru[n]</i>	[arg0 arg1 ... argn] [< inputfile] [> outputfile]	“Process Control,” on page 92
	Execute program from the beginning. If arguments arg0, arg1, and so on are specified, they are set up as the command-line arguments of the program. Otherwise, the arguments for the previous run command are used.	
<i>rv[al]</i>	expr	“Symbols and Expressions,” on page 108
	Return the rvalue of the expression expr.	
<i>sco[pe]</i>		“Scope,” on page 110
	Return a symbol for the search scope.	

Name	Arguments	Category
<i>scr[ipt]</i>	filename	“Miscellaneous,” on page 116
	Open the indicated file and execute the contents as though they were entered as commands. If you use ~ before the filename, it is expanded to the value of the environment variable HOME.	
<i>set</i>	var = exp	“Symbols and Expressions,” on page 108
	Set variable var to the value of expression.	
<i>setargs</i>	[arg1 , arg2, ... argn]	“Process Control,” on page 92
	Set program arguments to be used by the current program,	
<i>setenv</i>	name [value]	“Miscellaneous,” on page 116
	Print value of environment variable name. With a specified value, set name to value.	
<i>sh[ell]</i>	[arg0 , arg1, ... argn]	“Miscellaneous,” on page 116
	Fork a shell (defined by \$SHELL) and give it the indicated arguments (the default shell is sh). Without arguments, invokes an interactive shell, and executes until a "^D" is entered.	
<i>siz[eof]</i>	name	“Symbols and Expressions,” on page 108
	Return the size, in bytes, of the variable type name; or, if the name refers to a routine, returns the size in bytes of the subprogram.	
<i>sle[ep]</i>	[time]	“Miscellaneous,” on page 116
	Pause for time seconds. If no time is specified, pause for one second	
<i>sou[rce]</i>	filename	“Miscellaneous,” on page 116
	Open the indicated file and execute the contents as though they were entered as commands. If you use ~ before the filename, it is expanded to the value of \$HOME.	
<i>sp</i>		“Register Access,” on page 112
	Return the current stack pointer address.	
<i>sr[ead]</i>	addr	“Memory Access,” on page 113
	Fetch and print a short signed integer from the specified address	
<i>stackd[ump]</i>	[count]	“Program Locations,” on page 103
	Print a formatted dump of the stack. This command displays a hex dump of the stack frame for each active routine.	
<i>stack[trace]</i>	[count]	“Program Locations,” on page 103
	Print a stacktrace. For each active routine print the routine name, source file, line number, current address, provided that information is available.	

Name	Arguments	Category
<i>stat[us]</i>		“Events,” on page 96
	Display all the event definitions, including an event number by which the event can be identified.	
<i>s[tep]</i>	[count up]	“Process Control,” on page 92
	Step into the current routine and stop after executing one or count source line(s). If the up argument is specified, stops execution after stepping out of the current routine.	
<i>stepi</i>	[count up]	“Process Control,” on page 92
	Step into the current routine and stop after executing one or count source line(s). If the up argument is specified, stops execution after stepping out of the current routine.	
<i>stepo[ut]</i>		“Process Control,” on page 92
	Stop after returning to the caller of the current routine.	
<i>stop</i>	[at line in routine] [var] [if (condition)] [do {commands}]	“Events,” on page 96
	Set a breakpoint at the indicated routine or line. Break when the value of the indicated variable var changes.	
<i>stopi</i>	[at addr in routine] [var] [if (condition)] [do {commands}]	“Events,” on page 96
	Set a breakpoint at the indicated address or routine. Break when the value of the indicated variable var changes.	
<i>str[ing]</i>	exp [...exp]	“Printing Variables and Expressions,” on page 105
	Evaluate and print expressions as null-terminated character strings, up to a maximum of 70 characters.	
<i>sync</i>	[routine line]	“Process Control,” on page 92
	Advance the current process/thread to a specific program location, ignoring any user-defined events.	
<i>synci</i>	[routine addr]	“Process Control,” on page 92
	Advance the current process/thread to a specific program location, ignoring any user-defined events.	
<i>thread</i>	[number]	“Process Control,” on page 92
	Set the current thread to the thread identified by number; where number is a logical thread id in the current process' active thread list. When issued with no argument, thread lists the current program location of the currently active thread.	
<i>threads</i>		“Process Control,” on page 92
	Prints the status of all active threads, grouped by process.	

Name	Arguments	Category
<i>trace</i>	[at line in routine] [var routine] [if (condition)] do {commands}	“Events,” on page 96
	Activates source line tracing as specified by the arguments supplied.	
<i>tracei</i>	[at addr in routine] [var] [if (condition)] do {commands}	“Events,” on page 96
	Activates instruction tracing as specified by the arguments supplied.	
<i>track</i>	expression [at line in routine] [if (condition)] [do {commands}]	“Events,” on page 96
	Define a track event.	
<i>tracki</i>	expression [at addr in routine] [if (condition)] [do {commands}]	“Events,” on page 96
	Define an assembly-level track event.	
<i>type</i>	expr	“Symbols and Expressions,” on page 108
	Return the type of the expression.	
<i>unal[ias]</i>	name	“Miscellaneous,” on page 116
	Remove the alias definition for name, if one exists.	
<i>unb[reak]</i>	line routine all	“Events,” on page 96
	Remove a breakpoint from the statement line, the routine <code>routine</code> , or remove all breakpoints.	
<i>unbreaki</i>	addr routine all	“Events,” on page 96
	Remove a breakpoint from the address <code>addr</code> , the routine <code>routine</code> , or remove all breakpoints.	
<i>undefset</i>	[name -all]	“Process-Thread Sets,” on page 95
	Remove a previously defined process/thread set from the list of process/thread sets.	
<i>undisplay</i>	[all 0 exp]	“Printing Variables and Expressions,” on page 105
	Remove all expressions specified by previous display commands. With an argument or several arguments, remove the expression <code>exp</code> from the list of display expressions.	
<i>u[p]</i>	[number]	“Scope,” on page 110
	Move up one level or number levels on the call stack.	
<i>use</i>	[dir]	“Miscellaneous,” on page 116
	Print the current list of directories or add <code>dir</code> to the list of directories to search. If the first character in <code>pathname</code> is <code>~</code> , the value of <code>\$HOME</code> is substituted for this character.	

Name	Arguments	Category
<i>viewset</i>	name	“Process-Thread Sets,” on page 95
	List the members of a process/thread set that currently exist as active threads or list defined p/t-sets.	
<i>wait</i>	[any all none]	“Process Control,” on page 92
	Inserts explicit wait points in a command stream.	
<i>wa[tcb]</i>	expression [at line in routine] [if (condition)] [do {commands}]	“Events,” on page 96
	Define a watch event. The given expression is evaluated, and subsequently, each time the value of the expression changes, the program stops and the new value is printed.	
<i>watchi</i>	expression [at addr in routine] [if(condition)] [do {commands}]	“Events,” on page 96
	Define an assembly-level watch event.	
<i>whatis</i>	[name]	“Symbols and Expressions,” on page 108
	With no arguments, prints the declaration for the current routine. With argument name, prints the declaration for the symbol name.	
<i>when</i>	[at line in routine] [if (condition)] do {commands}	“Events,” on page 96
	Execute <i>commands</i> at every line in the program, at a specified line in the program or in the specified routine.	
<i>wheni</i>	[at addr in routine] [if(condition)] do {commands}	“Events,” on page 96
	Execute <i>commands</i> at each address in the program. If an addr is specified, the commands are executed each time the address is reached.	
<i>w[here]</i>	[count]	“Program Locations,” on page 103
	Print a stacktrace. For each active routine print the routine name, routine arguments, source file, line number, current address, provided that information is available.	
<i>whereis</i>	name	“Symbols and Expressions,” on page 108
	Print all declarations for name.	
<i>which</i>	name	“Scope,” on page 110
	Print full scope qualification of symbol name.	
<i>whichsets</i>	[p/t-set]	“Process-Thread Sets,” on page 95
	List all defined p/t-sets to which the members of a process/thread set belong.	
<i>/</i>	/ [string] /	“Program Locations,” on page 103
	Search forward for a string (<i>string</i>) of characters in the current source file	

Name	Arguments	Category
?	?[string] ?	“Program Locations,” on page 103
	Search backward for a string (<i>string</i>) of characters in the current source file.	
!	History modification	“Miscellaneous,” on page 116
	Executes a command from the command history list. The command executed depends on the information that follows the !.	
^	History modification	“Miscellaneous,” on page 116
	Quick history command substitution ^old^new^<modifier> this is equivalent to !:s/old/new/	

Chapter 6. Assembly-Level Debugging

This section provides information about *PGDBG* assembly-level debugging, including an overview and what to expect if you are using assembly-level debugging or if you did not compile your program for debugging.

Assembly-Level Debugging Overview

PGDBG supports debugging regardless of how a program was compiled. Specifically, *PGDBG* does not require that the program under debug be compiled with debugging information, such as using `-g`. It can debug code that is lacking debug information, but because it is missing information about symbols and line numbers, it can only access the program at the assembly level. *PGDBG* also supports debugging at the assembly level if debug symbols are available.

As described in [“Building Applications for Debug,” on page 1](#), the most information is available when the program is compiled using `-g` or `-gopt` with no optimization. When a program is compiled at higher levels of optimization, less information about source-level symbols and line numbers is available, even if the program was compiled with `-g` or `-gopt`. In such cases, if you want to find the source of a problem without rebuilding the program, you may need to debug at the assembly level.

If a program has been "stripped" of all symbols, either by the linker or a separate utility, then debugging is at the assembly level. *PGDBG* is only able to examine or control the program in terms of memory addresses and registers.

Assembly-Level Debugging on Microsoft Windows Systems

When applications are built without `-g` on Windows systems, the resulting binary, the `.exe` file, does not contain any symbol information. Microsoft stores symbol information in a program database, a `.pdb` file. To generate a `.pdb` file using the PGI compiler drivers, you must use `-g` during the link step. You can do this even if you did not use `-g` during the compile step. Having this `.pdb` file available provides *PGDBG* with enough symbol information to map addresses to routine names.

Assembly-Level Debugging with Fortran

To refer to Fortran symbol names when debugging at the assembly level, you must translate the names to use the naming convention that matches the calling convention in use by the compiler. For code compiled by the PGI compilers, in most cases this means translating to lower case and appending an underbar. For example, a routine that appears in the source code as "VADD" would be referred to in the debugger as "vadd_".

On 32-bit Windows systems there are alternative calling conventions. The one described above matches the convention used when the compiler is invoked with `-Miface=unix` (previously `-Munix`). For details of other 32-bit Windows calling conventions, refer to the *PGI Compiler User's Guide*.

Note

Name translation is only necessary for assembly-level debugging. When debugging at the source level, you may refer to symbol names as they appear in the source.

A special symbol, `MAIN_`, is created by PGFORTRAN to refer to the main program. PGFORTRAN generates this special symbol whether or not there is a PROGRAM statement. One way to run to the beginning of a Fortran program is to set a breakpoint on `MAIN_`, then run.

Assembly-Level Debugging with C++

C++ symbol names are "mangled" names. For the names of C++ methods, the names are modified to include not only the name as it appears in the source code, but information about the enclosing class hierarchy, argument and return types, and other information. The names are long and arcane. At the source level these names are translated by *PGDBG* to the names as they appear in the source. At the assembly level, these names are in the mangled form. Translation is not easy and not recommended. If you have no other alternative, you can find information about name mangling in the *PGI Compiler User's Guide*.

Assembly-Level Debugging Using the PGDBG GUI

This section describes some basic operations for assembly-level debugging using the *PGDBG* GUI. If you encounter the message "Can't find main function compiled -g" on startup, assembly-level debugging is required.

To get into a program in this situation, you can select the Debug | Set Breakpoint... menu option. To stop at program entry, for example, in Fortran you could enter `MAIN_` in response to the dialog query; in C or C++ you could enter `main`.

PGDBG debug information tabs that are useful in assembly-level debugging include the Call Stack, Memory, and Register tabs. Use the Disassembly tab in the source pane to view the disassembled code.

If disassembly is not automatically displayed in the source pane, use the `dis` command in either the Command tab or Data | Custom window to generate disassembly for an address or function.

Assembly-Level Debugging Using the PGDBG CLI

This section describes some basic operations for assembly-level debugging using the *PGDBG* command-line interface. When you invoke the *PGDBG* CLI and are presented with a message telling you that *PGDBG* "Can't find main function compiled -g", assembly-level debugging is required.

To get into the program, you can set a breakpoint at a named routine. To stop at program entry, for example, in Fortran you could use

```
pgdbg> break MAIN_
```

and in C/ C⁺⁺ you could use

```
pgdbg> break main
```

Some useful commands for assembly-level debugging using the *PGDBG* command-line interface include:

run

run the program from the beginning

cont

continue program execution from the current point

nexti

single-step one instruction, stepping over calls

stepi

single-step one instruction, stepping into calls

breaki

set a breakpoint at a given address

regs

display the registers

print \$<regname>

display the value of the specified register

For more information on register names, refer to [“SSE Register Symbols,” on page 49](#).

dump

dump memory locations

stacktrace

display a stack traceback

stackdump

display a traceback/dump of stack frame

SSE Register Symbols

X64 processors and x86 processors starting with Pentium III provide SSE (Streaming SIMD Enhancements) registers and a SIMD floating-point control/status register.

Each SSE register may contain four 32-bit single-precision or two 64-bit floating-point values. The *PGDBG* **regs** command reports these values individually in both hexadecimal and floating-point format. *PGDBG* provides command notation to refer to these values individually or all together.

The component values of each SSE register can be accessed using the same syntax that is used for array subscripting. Pictorially, the SSE registers can be thought of as follows:

[32-bit]

127	96 95	64 63	32 31	0
\$xmm0[3]	\$xmm0[2]	\$xmm0[1]	\$xmm0[0]	
\$xmm1[3]	\$xmm1[2]	\$xmm1[1]	\$xmm1[0]	
\$xmm2[3]	\$xmm2[2]	\$xmm2[1]	\$xmm2[0]	

To access `$xmm0[3]`, the 32-bit single-precision floating point value that occupies bits 96 – 127 of SSE register 0, use the following *PGDBG* command:

```
pgdbg> print $xmm0[3]
```

To set `$xmm2[0]` to the value of `$xmm3[2]`, use the following *PGDBG* command:

```
pgdbg> set $xmm2[3] = $xmm3[2]
```

[64-bit]

127	64 63	0
\$xmm0d[1]	\$xmm0d[0]	
\$xmm1d[1]	\$xmm1d[0]	
\$xmm2d[1]	\$xmm2d[0]	

To access the 64-bit floating point values in `xmm0`, append the character 'd' (for double precision) to the register name and subscript as usual, as illustrated in the following *pgdbg* commands:

```
pgdbg> print $xmm0d[0]
```

```
pgdbg> print $xmm0d[1]
```

In most cases, *PGDBG* detects when the target environment supports the SSE registers. In the event *PGDBG* does not allow access to SSE registers on a system that should have them, set the *PGDBG_SSE* environment variable to `on` to enable SSE support.

Chapter 7. Source-Level Debugging

This chapter describes source-level debugging, including debugging Fortran and C⁺⁺.

Debugging Fortran

Fortran Types

PGDBG displays Fortran type declarations using Fortran type names. The only exception is Fortran character types, which are treated as arrays of the C type `char`.

Arrays

Fortran array subscripts and ranges are accessed using the Fortran language syntax convention, denoting subscripts with parentheses and ranges with colons.

PGI compilers for the linux86-64 platform (AMD64 or Intel 64) support large arrays (arrays with an aggregate size greater than 2GB). You can enable large array support by compiling using these options: `-mmodel=medium -Mlarge_arrays`. *PGDBG* provides full support for large arrays and large subscripts.

PGDBG supports arrays with non-default lower bounds. Access to such arrays uses the same subscripts that are used in the target application.

PGDBG also supports adjustable arrays. Access to adjustable arrays may use the same subscripting that is used in the target application.

Operators

In general, *PGDBG* uses C language style operators in expressions and supports the Fortran array index selector `()` and the Fortran field selector `%` for derived types. However, `.eq.`, `.ne.`, and so forth are not supported. You must use the analogous C operators `==`, `!=`, and so on, instead.

Note

The precedence of operators matches the C language, which may in some cases be different than for Fortran.

See [Table 5.1, “PGDBG Commands”](#) for a complete list of operators and their definition.

Name of the Main Routine

If a **PROGRAM** statement is used, the name of the main routine is the name in the program statement. You can always use the following command to set a breakpoint at the start of the main routine.

```
break MAIN
```

Common Blocks

Each subprogram that defines a common block has a local static variable symbol to define the common. The address of the variable is the address of the common block. The type of the variable is a locally-defined structure type with fields defined for each element of the common block. The name of the variable is the common block name, if the common block has a name, or `_BLNK_` otherwise.

For each member of the common block, a local static variable is declared which represents the common block variable. Thus given declarations:

```
common /xyz/ a, b
integer a
integer b
```

then the entire common block can be printed out using,

```
print xyz
```

Individual elements can be accessed by name. For example,;

```
print a, b
```

Internal Procedures

To unambiguously reference an internal procedure, qualify its name with the name of its host using the scoping operator `@`.

For example:

```
subroutine sub1 ()
  call internal_proc ()
  contains
  subroutine internal_proc ()
    print *, "internal_proc in sub1"
  end subroutine internal_proc
end subroutine

subroutine sub2 ()
  call internal_proc ()
  contains
  subroutine internal_proc ()
    print *, "internal_proc in sub2"
  end subroutine internal_proc
end subroutine

program main
  call sub1 ()
  call sub2 ()
end program
```



```
pgdbg> whereis internal_proc
function:      "/path/ip.f90"@sub1@internal_proc
function:      "/path/ip.f90"@sub2@internal_proc

pgdbg> break sub1@internal_proc
(1)breakpoint set at: internal_proc line: "ip.f90"@5 address: 0x401E3C 1

pgdbg> break sub2@internal_proc
(2)breakpoint set at: internal_proc line: "ip.f90"@13 address: 0x401EEC 2
```

Modules

A member of a Fortran 90 module can be accessed during debugging.

```
module mod
  integer iMod
end module
subroutine useMod()
  use mod
  iMod = 1000
end subroutine
program main
  call useMod()
end program
```

- If the module is in the current scope, no qualification is required to access the module's members.

```
pgdbg> b useMod
(1)breakpoint set at: usemod line: "modv.f90"@7 address: 0x401CC4
1

Breakpoint at 0x401CC4, function usemod, file modv.f90, line 7
#7:      iMod = 1000
```

```
pgdbg> p iMod
0
```

- If the module is not in the current scope, use the scoping operator @ to qualify the member's name.

```
Breakpoint at 0x401CF0, function main, file modv.f90, line 11
#11:      call useMod()
```

```
pgdbg> p iMod
"iMod" is not defined in the current scope
```

```
pgdbg> p mod@iMod
0
```

Module Procedures

A module procedure is a subroutine contained within a module. A module procedure itself can contain internal procedures. The scoping operator @ can be used when working with these types of subprograms to prevent ambiguity.

```
module mod
  contains
  subroutine mod_proc1()
    call internal_proc()
  contains
  subroutine internal_proc()
```

```

        print *, "internal_proc in mod_proc1"
    end subroutine
end subroutine
subroutine mod_proc2()
    call internal_proc()
    contains
    subroutine internal_proc()
        print *, "internal_proc in mod_proc2"
    end subroutine
end subroutine
end module

```

```

program main
    use mod
    call mod_proc1
    call mod_proc2
end program

```

```

pgdbg> whereis internal_proc
function:      "/path/modp.f90"@mod@mod_proc1@internal_proc
function:      "/path/modp.f90"@mod@mod_proc2@internal_proc

pgdbg> break mod@mod_proc1@internal_proc
(1)breakpoint set at: internal_proc line: "modp.f90"@7 address: 0x401E3C
1
pgdbg> break mod@mod_proc2@internal_proc
(2)breakpoint set at: internal_proc line: "modp.f90"@14 address: 0x401EEC
2

```

Debugging C++

Calling C++ Instance Methods

To use the **call** command to call a C++ instance method, the object must be explicitly passed as the first parameter to the call. For example, suppose you were given the following definition of class `Person` and the appropriate implementation of its methods:

```

class Person
{
    public:
    char name[10];
    Person(char * inName);
    void print();
};

int main ()
{
    Person * pierre;
    pierre = new Person("Pierre");
    pierre->print();
    return 0;
}

```

Call the instance method `print` on object `pierre` as follows:

```
pgdbg> call Person::print(pierre)
```

Notice that `pierre` must be explicitly passed into the method because it is the `this` pointer. You can also specify the class name to remove ambiguity.

Chapter 8. Platform-Specific Features

This chapter describes the *PGDBG* features that are specific to particular platforms, such as pathname conventions, debugging with core files, and signals.

Pathname Conventions

PGDBG uses the forward slash character (/) as the path component separator on all platforms. The backslash (\) is used as the escape character in the *PGDBG* command language.

On Windows systems, a drive letter specifier may be used whenever specifying a full path, but the forward slash separator convention is still in effect. For example, to add the Windows pathname `C:\Temp\src` to the list of searched source directories, use the command:

```
pgdbg> dir C:/Temp/src
```

To set a breakpoint at line 10 of the source file specified by the relative path `sub1/main.c`, use this command:

```
pgdbg> break "sub1/main.c":10
```

Debugging with Core Files

PGDBG supports debugging of core files on Linux platforms. To invoke *PGDBG* for core file debugging, use the following options:

```
$ pgdbg -core coreFileName programName
```

Core files (or core dumps) are generated when a program encounters an exception or fault. For example, one common exception is the segmentation violation, which can be caused by referencing an invalid memory address. The memory and register states of the program are written into a core file so that they can be examined by a debugger.

The shell environment in which the application runs must be set up to allow core file creation. On many systems, the default user setting `ulimit` does not allow core file creation.

Check the `ulimit` as follows:

For sh/bash users:

```
$ ulimit -c
```

For csh/tcsh users:

```
% limit coredumpsize
```

If the core file size limit is zero or something too small for the application, it can be set to unlimited as follows:

For sh/bash users:

```
$ ulimit -c unlimited
```

For csh/tcsh users:

```
% limit coredumpsize unlimited
```

See the Linux shell documentation for more details. Some versions of Linux provide system-wide limits on core file creation.

The core file is normally written into the current directory of the faulting application. It is usually named `core` or `core.pid` where *pid* is the process ID of the faulting thread. If the shell environment is set correctly and a core file is not generated in the expected location, the system core dump policy may require configuration by a system administrator.

Different versions of Linux handle core dumping slightly differently. The state of all process threads are written to the core file in most modern implementations of Linux. In some new versions of Linux, if more than one thread faults, then each thread's state is written to separate core files using the `core.pid` file naming convention previously described. In older versions of Linux, only one faulting thread is written to the core file.

If a program uses dynamically shared objects (i.e., shared libraries named `lib*.so`), as most programs on Linux do, then accurate core file debugging requires that the program be debugged on the system where the core file was created. Otherwise, slight differences in the version of a shared library or the dynamic linker can cause erroneous information to be presented by the debugger. Sometimes a core file can be debugged successfully on a different system, particularly on more modern Linux systems, but you should take care when attempting this.

When debugging core files, *PGDBG*:

- Supports all non-control commands.
- Performs any command that does not cause the program to run.
- Generates an error message in *PGDBG* for any command that causes the program to run.
- May provide the status of multiple threads, depending on the type of core file created.

PGDBG does not support multi-process core file debugging.

Signals

PGDBG intercepts all signals sent to any of the threads in a multi-threaded program and passes them on according to that signal's disposition as maintained by *PGDBG* (see the **catch** and **ignore** commands), except for signals that cannot be intercepted or signals used internally by *PGDBG*.

Signals Used Internally by PGDBG

SIGTRAP and SIGSTOP are used by Linux for communication of application events to *PGDBG*. Management of these signals is internal to *PGDBG*. Changing the disposition of these signals in *PGDBG* (via **catch** and **ignore**) results in undefined behavior.

Signals Used by Linux Libraries

Some Linux thread libraries use SIGRT1 and SIGRT3 to communicate among threads internally. Other Linux thread libraries, on systems that do not have support for real-time signals in the kernel, use SIGUSR1 and SIGUSR2. Changing the disposition of these signals in *PGDBG* (via **catch** and **ignore**) result in undefined behavior.

Target applications compiled with the options `-pg` or `-Mprof=time` generate numerous SIGPROF signals. Although SIGPROF can be handled by *PGDBG* via the **ignore** command, debugging of applications built for sample-based profiling is not recommended.

Chapter 9. Parallel Debugging Overview

This chapter provides an overview of how to use *PGDBG* to debug parallel applications. It includes important definitions and background information on how *PGDBG* represents processes and threads.

Overview of Parallel Debugging Capability

PGDBG is a parallel application debugger capable of debugging multi-process MPI applications, multi-thread and OpenMP applications, and hybrid multi-thread/multi-process applications that use MPI to communicate between multi-threaded or OpenMP processes.

For specific information on multi-thread and OpenMP debugging, refer to [Chapter 10, “Parallel Debugging with OpenMP”](#).

For specific information on multi-process MPI debugging, refer to [Chapter 11, “Parallel Debugging with MPI”](#).

Graphical Presentation of Threads and Processes

PGDBG graphical user interface components that provide support for parallelism are described in detail in [“The Graphical User Interface”](#).

Basic Process and Thread Naming

Because *PGDBG* can debug multi-threaded applications, multi-process applications, and hybrid multi-threaded/multi-process applications, it provides a convention for uniquely identifying each thread in each process. This section gives a brief overview of this naming convention and how it is used to provide adequate background for the subsequent sections. A more detailed discussion of this convention, including advanced techniques for applying it, is provided in [“Thread and Process Grouping and Naming,” on page 60](#).

PGDBG identifies threads in an OpenMP application using the OpenMP thread IDs. Otherwise, *PGDBG* assigns arbitrary IDs to threads, starting at zero and incrementing in order of thread creation.

PGDBG identifies processes in an MPI application using MPI rank (in communicator `MPI_COMM_WORLD`). Otherwise, *PGDBG* assigns arbitrary IDs to processes; starting at zero and incrementing in order of process creation. Process IDs are unique across all active processes.

In a multi-threaded/multi-process application, each thread can be uniquely identified across all processes by prefixing its thread ID with the process ID of its parent process. For example, thread 1.4 identifies the thread with ID 4 in the process with ID 1.

An OpenMP application logically runs as a collection of threads with a single process, process 0, as the parent process. In this context, a thread is uniquely identified by its thread ID. The process ID prefix is implicit and optional. For more information on debugging threads, refer to [“Threads-only Debugging,” on page 61](#).

An MPI program logically runs as a collection of processes, each made up of a single thread of execution. Thread 0 is implicit to each MPI process. A process ID uniquely identifies a particular process, and thread ID is implicit and optional. For more information on process debugging, refer to [“Process-only Debugging,” on page 61](#).

A hybrid, or multilevel, MPI/OpenMP program requires the use of both process and thread IDs to uniquely identify a particular thread. For more information on multilevel debugging, refer to [“Multilevel Debugging,” on page 61](#).

A serial program runs as a single thread of execution, thread 0, belonging to a single process, process 0. The use of thread IDs and process IDs is allowed but unnecessary.

Thread and Process Grouping and Naming

This section describes how to name a single thread, how to group threads and processes into sets, and how to apply *PGDBG* commands to groups of processes and threads.

PGDBG Debug Modes

PGDBG can operate in four debug modes. The mode determines a short form for uniquely naming threads and processes. The debug mode is set automatically or by the **pgienv mode** command.

Table 9.1. PGDBG Debug Modes

Debug Mode	Program Characterization
Serial	A single thread of execution
Threads-only	A single process, multiple threads of execution
Process-only	Multiple processes, each process made up of a single thread of execution
Multilevel	Multiple processes, at least one process employing multiple threads of execution

PGDBG initially operates in serial mode reflecting a single thread of execution. Thread IDs can be ignored in serial debug mode since there is only a single thread of execution.

The *PGDBG* prompt displays the ID of the current thread according to the current debug mode. For a description of the *PGDBG* prompt, refer to [“The PGDBG Command Prompt,” on page 73](#).

The debug mode can be changed at any time during a debug session.

To change debug mode manually, use the **pgienv** command.

```
pgienv mode [serial|thread|process|multilevel]
```

Threads-only Debugging

Enter threads-only mode to debug a program with a single multi-threaded process. As a convenience the process ID portion can be omitted. *PGDBG* automatically enters threads-only debug mode from serial debug mode when it detects and attaches to new threads.

Example 9.1. Thread IDs in Threads-only Debug Mode

1	Thread 1 of process 0 (*. 1)
*	All threads of process 0 (*. *)
0. 7	Thread 7 of process 0 (multilevel names are valid in threads-only mode)

In threads-only debug mode, status and error messages are prefixed with thread IDs depending on context.

Process-only Debugging

Enter process-only mode to debug an application consisting of single-threaded processes. As a convenience, the thread ID portion can be omitted. *PGDBG* automatically enters process-only debug mode from serial debug mode when multiple processes are detected.

Example 9.2. Process IDs in Process-only Debug Mode

0	All threads of process 0 (0.*)
*	All threads of all processes (*.*)
1. 0	Thread 0 of process 1 (multilevel names are valid in process-only mode)

In process-only debug mode, status and error messages are prefixed with process IDs depending on context.

Multilevel Debugging

The name of a thread in multilevel debug mode is the thread ID prefixed with its parent process ID. This forms a unique name for each thread across all processes. This naming scheme is valid in all debug modes. *PGDBG* changes automatically to multilevel debug mode when at least one MPI process creates multiple threads.

Example 9.3. Thread IDs in Multilevel Debug Mode

0. 1	Thread 1 of process 0
0. *	All threads of process 0
*	All threads of all processes

In multilevel debug, mode status and error messages are prefixed with process/thread IDs depending on context.

Process/Thread Sets

You use a process/thread set (p/t-set) to restrict a debugger command to apply to just a particular set of threads. A p/t-set is a set of threads drawn from all threads of all processes in the target program. Use p/t-set notation, described in “p/t-set Notation”, to define a p/t-set.

Named p/t-sets

In the following sections, you will notice frequent references to three named p/t-sets:

- The *target p/t-set* is the set of processes and threads to which a debugger command is applied. The target p/t-set is initially defined by the debugger to be the set [all] which describes all threads of all processes.
- A *prefix p/t-set* is defined when p/t-set notation is used to prefix a debugger command. For the prefixed command, the target p/t-set is the prefix p/t-set.
- The *current p/t-set* is the p/t set currently set in the *PGDBG* environment. You can use the **focus** command to define the current p/t-set. Unless a prefix p/t-set overrides it, the current p/t set is used as the target p/t-set.

p/t-set Notation

The following rules describe how to use and construct p/t-sets:

Use a prefix p/t-set with a simple command:

```
[p/t-set prefix] command parm0, parm1, ...
```

Use a prefix p/t-set with a compound command:

```
[p/t-set prefix] simple-command [:simple-command ...]
```

p/t-id:

```
{integer|*}.{integer|*}
```

Use *p/t-id* optional notation when process-only or threads-only debugging is in effect. For more information, refer to the **pgienv** command.

p/t-range:

```
p/t-id:p/t-id
```

p/t-list:

```
{p/t-id|p/t-range} [, {p/t-id|p/t-range} ...]
```

p/t-set

```
[[]]{p/t-list|set-name}]
```

Example 9.4. p/t-sets in Threads-only Debug Mode

[0,4:6]	Threads 0,4,5, and 6
[*]	All threads
[*.1]	Thread 1. Multilevel notation is valid in threads-only mode
[*.*]	All threads

Example 9.5. p/t-sets in Process-only Debug Mode

<code>[0,2:3]</code>	Processes 0, 2, and 3 (equivalent to <code>[0.*,2:3.*]</code>)
<code>[*]</code>	All processes (equivalent to <code>[*.*]</code>)
<code>[0]</code>	Process 0 (equivalent to <code>[0.*]</code>)
<code>[*.0]</code>	Process 0. Multilevel syntax is valid in process-only mode.
<code>[0:2.*]</code>	Processes 0, 1, and 2. Multilevel syntax is valid in process-only debug mode.

Example 9.6. p/t-sets in Multilevel Debug Mode

<code>[0.1,0.3,0.5]</code>	Thread 1,3, and 5 of process 0
<code>[0.*]</code>	All threads of process 0
<code>[1.1:3]</code>	Thread 1,2, and 3 of process 1
<code>[1:2.1]</code>	Thread 1 of processes 1 and 2
<code>[clients]</code>	All threads defined by named set clients
<code>[1]</code>	Incomplete; invalid in multilevel debug mode

Dynamic vs. Static p/t-sets

The **defset** command can be used to define both dynamic and static p/t-sets. The members of a dynamic p/t-set are those active threads described by the p/t-set at the time that the p/t-set is used. By default, a p/t-set is dynamic. Threads and processes are created and destroyed as the target program runs and, therefore, membership in a dynamic set varies as the target program executes.

Example 9.7. Defining a Dynamic p/t-set

<code>defset clients [*..1:3]</code>	Defines a named set <code>clients</code> whose members are threads 1, 2, and 3 of all processes that are currently active when <code>clients</code> is used. Membership in <code>clients</code> changes as processes are created and destroyed.
--------------------------------------	---

Membership in a static set is fixed at definition time. The members of a static p/t-set are those threads described by that p/t-set when it is defined. Use a `!` to specify a static set.

Example 9.8. Defining a Static p/t-set

<code>defset clients [!*.1:3]</code>	Defines a named set <code>clients</code> whose members are threads 1, 2, and 3 of those processes that are currently active at the time of the definition.
--------------------------------------	--

Note

p/t-sets defined with `defset` are not mode-dependent and are valid in any debug mode.

Current vs. Prefix p/t-set

The current p/t-set is set by the **focus** command. The current p/t-set is described by the debugger prompt and depends on debug mode. For a description of the *PGDBG* prompt, refer to [“The PGDBG Command Prompt,”](#)

on page 73. You can use a p/t-set to prefix a command that overrides the current p/t-set. The prefix p/t-set becomes the target p/t-set for the command. The target p/t-set defines the set of threads that will be affected by a command.

- In the following command line, the target p/t-set is the current p/t-set:

```
pgdbg [all] 0.0> cont
Continue all threads in all processes
```

- In contrast, a prefix p/t-set is used in the following command so that the target p/t-set is the prefix p/t-set, shown in this example in bold:

```
pgdbg [all] 0.0> [0.1:2] cont
Continue threads 1 and 2 of process 0 only
```

In both of the above examples, the current p/t-set is the debugger-defined set [all]. In the first case, [all] is the target p/t-set. In the second case, the prefix p/t-set overrides [all] and becomes the target p/t-set. The **continue** command is applied to all active threads in the target p/t-set. Also, using a prefix p/t-set does not change the current p/t-set.

p/t-set Commands

You can use the following commands to collect threads and processes into logical groups.

- Use **defset** and **undefset** to manage a list of named p/t-sets.
- Use **focus** to set the current p/t-set.
- Use **viewset** to view the active members described by a particular p/t-set, or to list all the defined p/t-sets.
- Use **whichsets** to describe the p/t-sets to which a particular process/thread belongs.

Table 9.2. p/t-set Commands

Command	Description
defset	Define a named process/thread set. This set can later be referred to by name. A list of named sets is stored by <i>PGDBG</i> .
focus	Set the target process/thread set for commands. Subsequent commands are applied to the members of this set by default.
undefset	Undefine a previously defined process/thread set. The set is removed from the list. The debugger-defined p/t-set [all] cannot be removed.
viewset	List the members of a process/thread set that currently exist as active threads, or list all the defined p/t-sets.
whichsets	List all defined p/t-sets to which the members of a process/thread set belong.

Examples of the p/t-set commands in the previous table follow.

Use **defset** to define the p/t-set *initial* to contain only thread 0:

```
pgdbg [all] 0> defset initial [0]
"initial" [0] : [0]
```

Use the **focus** command to change the current p/t-set to *initial*:

```
pgdbg [all] 0> focus [initial]
[initial] : [0]
[0]
```

Advance the thread using the current p/t-set, which is `initial`:

```
pgdbg [initial] 0> next
```

The **whichsets** command shows that thread 0 is a member of two defined p/t-sets:

```
pgdbg [initial] 0> whichsets [initial]
Thread 0 belongs to:
all
initial
```

The **viewset** command displays all threads that are active and are members of defined p/t-sets:

```
pgdbg [initial] 0> viewset
"all" [*.*] : [0.0,0.1,0.2,0.3]
"initial" [0] : [0]
```

You can use the **focus** command to set the current p/t-set back to `[all]`:

```
pgdbg [initial] 0> focus [all]
[all] : [0.0,0.1,0.2,0.3]
[*.*]
```

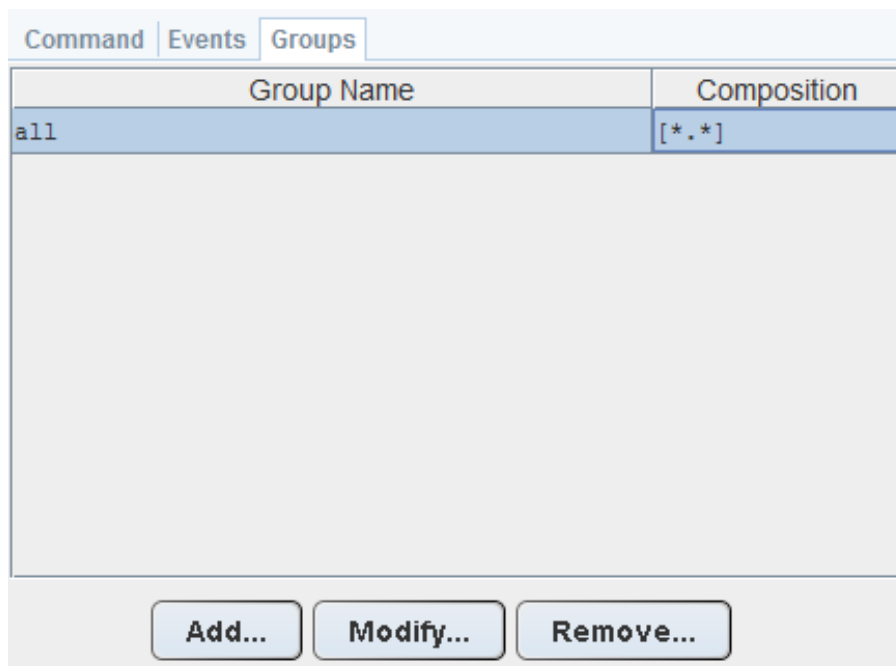
The **undefset** command undefines the initial p/t-set:

```
pgdbg [all] 0> undefset initial
p/t-set name "initial" deleted.
```

Using Process/Thread Sets in the GUI

The previous examples illustrate how to manage named p/t-sets using the command-line interface. A similar capability is available in the *PGDBG* GUI. “[Groups Tab](#),” on [page 11](#) provides an overview of the Groups tab.

Figure 9.1. Groups Tab



The Groups tab contains a table with two columns: a Group Name column and a p/t-set Composition column. The entries in the Composition column are the same p/t-sets used in the command-line interface.

Using this tab you can create, select, modify and remove p/t sets.

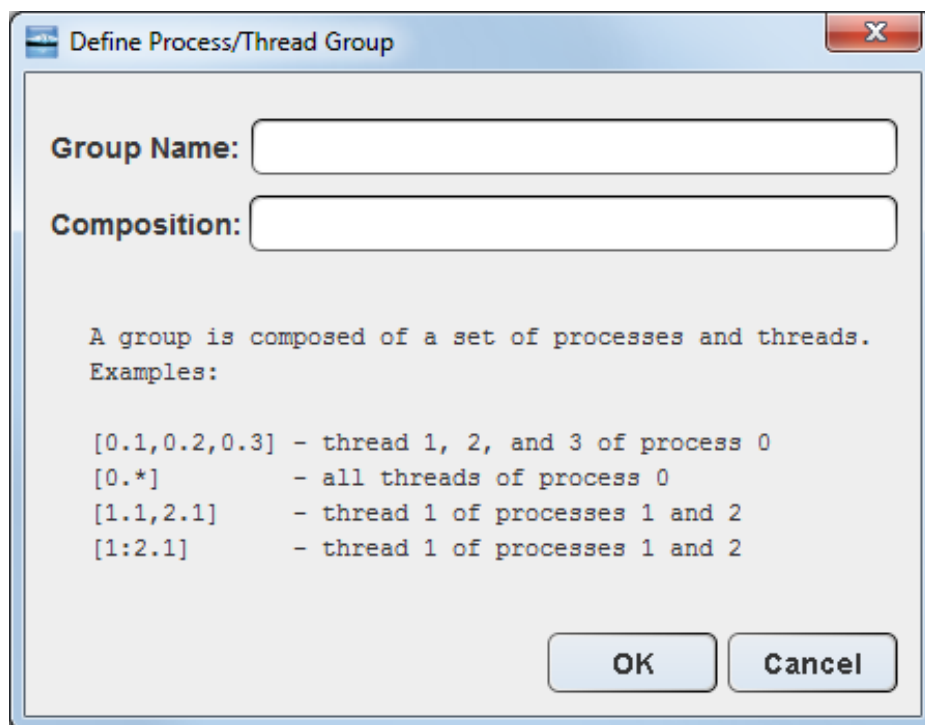
Create a p/t set

To create a p/t set in the Groups tab:

1. Click the Add button. This opens a dialog box similar to the one in [Figure 9.2](#).
2. Enter the name of the p/t-set in the Group Name field and enter the p/t-set in the Composition field.
3. Click OK to add the p/t-set.

The new p/t-set appears in the Groups table. Clicking the Cancel button or closing the dialog box aborts the operation.

Figure 9.2. Process/Thread Group Dialog Box



Select a p/t set

To select a p/t-set, click the desired p/t-set in the table. The selected p/t-set defines the Current Group used in the Apply and Display drop-down lists on the main toolbar.

Modify a p/t set

To modify an existing p/t-set, select the desired group in the Group table and click the Modify... button. A dialog box similar to that in [Figure 9.2](#) appears, except that the Group Name and Composition fields contain

the selected group's name and p/t-set respectively. You can edit the information in these fields and click OK to save the changes.

Remove a p/t set

To remove an existing p/t-set, select the desired item in the Groups Table and click the Remove... button. *PGDBG* displays a dialog box asking for confirmation of the removal request.

p/t set Usage

When Current Group is selected in either the Apply or Display drop-down lists on the main toolbar, the currently selected p/t-set in the Groups tab defines the Current Group.

Command Set

For the purpose of parallel debugging, the *PGDBG* command set is divided into three disjoint subsets according to how each command reacts to the current p/t-set. Process level and thread level commands can be parallelized. Global commands cannot be parallelized.

Table 9.3. PGDBG Parallel Commands

Commands	Action
Process Level Commands	Parallel by current p/t-set or prefix p/t-set
Thread Level Commands	Parallel by prefix p/t-set only; current p/t-set is ignored.
Global Commands	Non-parallel commands

Process Level Commands

The process level commands are the *PGDBG* control commands.

The *PGDBG* control commands apply to the active members of the current p/t-set by default. A prefix set can be used to override the current p/t-set. The target p/t-set is the prefix p/t-set if present.

cont	next	step	stepout	synci
halt	nexti	stepi	sync	wait

Apply the **next** command to threads 1 and 2 of process 0:

```
pgdbg [all] 0.0> focus [0.1:2]
pgdbg [0.1:2] 0.0> next
```

Apply the **next** command to thread 3 of process 0 using a prefix p/t-set:

```
pgdbg [all] 0.0> [0.3] n
```

Thread Level Commands

The following commands are not concerned with the current p/t-set. When no p/t-set prefix is used, these commands execute in the context of the current thread of the current process by default. That is, thread level

commands ignore the current p/t-set. Thread level commands can be applied to multiple threads by using a prefix p/t-set. When a prefix p/t-set is used, the commands in this section are executed in the context of each active thread described by the prefix p/t-set. The target p/t-set is the prefix p/t-set if present, or the current thread (not the current p/t-set) if no prefix p/t-set exists.

The thread level commands are:

addr	do	hwatch	print	stack
ascii	doi	iread	regs	stackdump
assign	dread	line	retaddr	string
bin	dump	lines	rval	track
break*	entry	lval	scope	tracki
cread	fp	noprint	set	watch
dec	fread	oct	sizeof	watchi
decl	func	pc	sp	whatis
disasm	hex	pf	sread	where

* breakpoints and variants (**stop**, **stopi**, **break**, **breaki**): if no prefix p/t-set is specified, [all] is used (overriding current p/t-set).

The following actions occur when a prefix p/t-set is used:

- The threads described by the prefix are sorted per process by thread ID in increasing order.
- The processes are sorted by process ID in increasing order, and duplicates are removed.
- The command is then applied to the threads in the resulting list in order.

Without a prefix p/t-set, the **print** command executes in the context of the current thread of the current process, thread 0.0, printing rank 0:

```
pgdbg [all] 0.0> print myrank
0
```

With a prefix p/t-set, the thread members of the prefix are sorted and duplicates are removed. The **print** command iterates over the resulting list:

```
pgdbg [all] 0.0> [2:3.*,1:2.*] print myrank
[1.0] print myrank:
1
[2.0] print myrank:
2
[2.1] print myrank:
2
[2.2] print myrank:
2
```



```
[3.0] print myrank:
3
[3.2] print myrank:
3
[3.1] print myrank:
3
```

Global Commands

The rest of the *PGDBG* commands ignore threads and processes, or are defined globally for all threads across all processes. The current p/t-set and prefix p/t-set (if any) are ignored.

The following is a list of commands that are defined globally.

?	defset	funcs	quit	threads
/	delete	help	repeat	unalias
alias	directory	history	rerun	unbreak
arrive	disable	ignore	run	undefset
breaks	display	log	script	use
call	edit	pgienv	shell	viewset
catch	enable	proc	source	wait
cd	files	procs	status	whereis
debug	focus	pwd	thread	whichsets

Process and Thread Control

PGDBG supports thread and process control everywhere in the program. Threads and processes can be advanced in groups anywhere in the program.

The *PGDBG* control commands are:

cont	next	step	stepout	synci
halt	nexti	stepi	sync	wait

To describe those threads to be advanced, set the current p/t-set or use a prefix p/t-set.

A thread inherits the control operation of the current thread when it is created. If the current thread single-steps over an `_mp_init` call (found at the beginning of every OpenMP parallel region) using the **next** command, then all threads created by `_mp_init` step into the parallel region as if by the **next** command.

A process inherits the control operation of the current process when it is created. So if the current process returns from a call to `MPI_Init` under the control of a **cont** command, the new process does the same.

Configurable Stop Mode

PGDBG supports configuration of how threads and processes stop in relation to one another. *PGDBG* defines two **pgienv** environment variables, **threadstop** and **procstop**, for this purpose. *PGDBG* defines two stop modes, synchronous (sync) and asynchronous (async).

Table 9.4. PGDBG Stop Modes

Command	Result
sync	Synchronous stop mode; when one thread stops at a breakpoint (event), all other threads are stopped soon after.
async	Asynchronous stop mode; each thread runs independently of the other threads. One thread stopping does not affect the behavior of another.

Thread stop mode is set using the **pgienv** command as follows:

```
pgienv threadstop [sync|async]
```

Process stop mode is set using the **pgienv** command as follows:

```
pgienv procstop [sync|async]
```

PGDBG defines the default to be asynchronous for both thread and process stop modes. When debugging an OpenMP program, *PGDBG* automatically enters synchronous thread stop mode in serial regions, and asynchronous thread stop mode in parallel regions.

The **pgienv** environment variables **threadstopconfig** and **procstopconfig** can be set to automatic (auto) or user defined (user) to enable or disable this behavior:

```
pgienv threadstopconfig [auto|user]
pgienv procstopconfig [auto|user]
```

Selecting the user-defined stop mode prevents the debugger from changing stop modes automatically. Automatic stop configuration is the default for both threads and processes.

Configurable Wait Mode

Wait mode describes when *PGDBG* accepts the next command. The wait mode is defined in terms of the execution state of the program. Wait mode describes to the debugger which processes/threads must be stopped before it will accept the next command.

In certain situations, it is desirable to be able to enter commands while the program is running and not stopped at an event. The *PGDBG* prompt does not appear until all processes/threads are stopped. However, a prompt may be available before all processes/threads have stopped. Pressing <enter> at the command line

brings up a prompt if it is available. The availability of the prompt is determined by the current wait mode and any pending **wait** commands.

PGDBG accepts a compound statement at each prompt. Each compound statement is a sequence of semicolon-separated commands, which are processed immediately in order.

The wait mode describes when to accept the next compound statement. *PGDBG* supports three wait modes, which can be applied to processes and/or threads.

Table 9.5. PGDBG Wait Modes

Command	Result
all	The prompt is available only after all threads have stopped since the last control command.
any	The prompt is available only after at least one thread has stopped since the last control command.
none	The prompt is available immediately after a control command is issued.

- Thread wait mode describes which threads *PGDBG* waits for before accepting new commands.

Thread wait mode is set using the **pgienv** command as follows:

```
pgienv threadwait [any|all|none]
```

- Process wait mode describes which processes *PGDBG* waits for before accepting new commands.

Process wait mode is set using the **pgienv** command as follows:

```
pgienv procwait [any|all|none]
```

If process wait mode is set to none, then thread wait mode is ignored.

The *PGDBG* CLI defaults to:

```
threadwait all
procwait any
```

If the target program goes MPI parallel, then **procwait** is changed to none automatically by *PGDBG*.

If the target program goes thread parallel, then **threadwait** is changed to none automatically by *PGDBG*. The **pgienv** environment variable `threadwaitconfig` can be set to automatic (auto) or user defined (user) to enable or disable this behavior.

```
pgienv threadwaitconfig [auto|user]
```

Selecting the user defined wait mode prevents the debugger from changing wait modes automatically. Automatic wait mode is the default thread wait mode.

The *PGDBG* GUI defaults to:

```
threadwait none
```

```
procwait none
```

Setting the wait mode may be necessary when invoking the *PGDBG* GUI using the `-s` (script file) option. This step ensures that the necessary threads are stopped before the **next** command is processed).

PGDBG also provides a **wait** command that can be used to insert explicit wait points in a command stream. **wait** uses the target p/t-set by default, which can be set to wait for any combination of processes/threads. You can use the **wait** command to insert wait points between the commands of a compound command.

The `threadwait` and `procwait` **pgienv** variables can be used to configure the behavior of **wait**. For more information, refer to **pgienv** usage in [“Configurable Wait Mode,” on page 70](#).

[Table 9.6, “PGDBG Wait Behavior”](#) describes the behavior of **wait**.

Suppose *s* is the target p/t-set. In the table,

- *P* is the set of all processes described by *s*.
- *p* is a single process.
- *T* is the set of all threads described by *s*.
- *t* is a single thread.

Table 9.6. PGDBG Wait Behavior

Command	threadwait	procwait	Wait Set
wait	all any none	all	Wait for T
wait	all	any none	Wait for all threads in at least one p in P
wait	any none	any none	Wait for all t in T for at least one p in P
wait all	all any none	all	Wait for T
wait all	all	any none	Wait for all threads of at least one p in P
wait all	any none	any none	Wait for all t in T for at least one p in P
wait any	all	all	Wait for at least one thread for each process p in P
wait any	all any none	any none	Wait for at least one t in T
wait any	any none	all	Wait for at least one thread in T for each process p in P

Command	threadwait	procwait	Wait Set
wait none	all any none	all any none	Wait for no threads

Status Messages

PGDBG can produce a variety of status messages during a debug session. This feature can be useful in the CLI if the graphical aids provided by the GUI are unavailable. Use the `pgienv` command to enable or disable the types of status messages produced by setting the verbose environment variable to an integer-valued bit mask:

```
pgienv verbose <bitmask>
```

The values for the bit mask, listed in the following table, control the type of status messages desired.

Table 9.7. PGDBG Status Messages

Value	Type	Information
0x0	Standard	Disable all messages.
0x1	Standard	Report status information on current process/thread only. A message is printed when the current thread stops and when threads and processes are created and destroyed. Standard messaging is the default and cannot be disabled.
0x2	Thread	Report status information on all threads of current processes. A message is reported each time a thread stops. If process messaging is also enabled, then a message is reported for each thread across all processes. Otherwise, messages are reported for threads of the current process only.
0x4	Process	Report status information on all processes. A message is reported each time a process stops. If thread messaging is also enabled, then a message is reported for each thread across all processes. Otherwise, messages are reported for the current thread only of each process.
0x8	SMP	Report SMP events. A message is printed when a process enters or exits a parallel region, or when the threads synchronize. The <i>PGDBG</i> OpenMP handler must be enabled.
0x16	Parallel	Report process-parallel events (default).
0x32	Symbolic debug information	Report any errors encountered while processing symbolic debug information (e.g. ELF, DWARF2).

The PGDBG Command Prompt

The *PGDBG* command prompt reflects the current debug mode, as described in [“PGDBG Debug Modes,” on page 60](#).

In serial debug mode, the *PGDBG* prompt looks like this:

```
pgdbg>
```

In threads-only debug mode, *PGDBG* displays the current p/t-set in square brackets followed by the ID of the current thread:

```
pgdbg [all] 0>
Current thread is 0
```

In process-only debug mode, *PGDBG* displays the current p/t-set in square brackets followed by the ID of the current process:

```
pgdbg [all] 0>
Current process is 0
```

In multilevel debug mode, *PGDBG* displays the current p/t-set in square brackets followed by the ID of the current thread prefixed by the id of its parent process:

```
pgdbg [all] 1.0>
Current thread 1.0
```

The **pgienv promptlen** variable can be set to control the number of characters devoted to printing the current p/t-set at the prompt.

Parallel Events

This section describes how to use a p/t-set to define an event across multiple threads and processes. Events, such as breakpoints and watchpoints, are user-defined events. User-defined events are thread-level commands, described in “[Thread Level Commands](#),” on page 67.

Breakpoints, by default, are set across all threads of all processes. A prefix p/t-set can be used to set breakpoints on specific processes and threads. For example:

```
i) pgdbg [all] 0> b 15
ii) pgdbg [all] 0> [all] b 15
iii) pgdbg [all] 0> [0.1:3] b 15
```

(i) and (ii) are equivalent. (iii) sets a breakpoint only in threads 1,2,3 of process 0.

By default, all other user events are set for the current thread only. A prefix p/t-set can be used to set user events on specific processes and threads. For example:

```
i) pgdbg [all] 0> watch glob
ii) pgdbg [all] 0> [*] watch glob
```

(i) sets a watchpoint for glob on thread 0 only. (ii) sets a watchpoint for glob on all threads that are currently active.

When a process or thread is created, it inherits all of the breakpoints defined for the parent process or thread. All other events must be defined explicitly after the process or thread is created. All processes must be stopped to add, enable, or disable a user event.

Events may contain if and do clauses. For example:

```
pgdbg [all] 0> [*] break func if (glob!=0) do {set f = 0}
```

The breakpoint fires only if glob is non-zero. The do clause is executed if the breakpoint fires. The if and do clauses execute in the context of a single thread. The conditional in the if clause and the body of the do execute

in the context of a single thread, the thread that triggered the event. The conditional definition as above can be restated as follows:

```
[0] if (glob!=0) {[0] set f = 0}
[1] if (glob!=0) {[1] set f = 0}
...
```

When thread 1 hits `func`, `glob` is evaluated in the context of thread 1. If `glob` evaluates to non-zero, `f` is bound in the context of thread 1 and its value is set to 0.

Control commands can be used in `do` clauses, however they only apply to the current thread and are only well defined as the last command in the `do` clause. For example:

```
pgdbg [all] 0> [*] break func if (glob!=0) do {set f = 0; c}
```

If the **wait** command appears in a `do` clause, the current thread is added to the wait set of the current process. For example:

```
pgdbg [all] 0> [*] break func if (glob!=0) do {set f = 0; c; wait}
```

If conditionals and `do` bodies cannot be parallelized with prefix `p/t`-sets. For example, the following command is illegal:

```
pgdbg [all] 0> break func if (glob!=0) do {[*] set f = 0} ILLEGAL
```

The body of a `do` statement cannot be parallelized.

Parallel Statements

This section describes how to use a `p/t`-set to define a statement that executes for multiple threads and processes.

Parallel Compound/Block Statements

Each command in a compound statement is executed in order. The target `p/t`-set is applied to all statements in a compound statement. The following two examples (i) and (ii) are equivalent:

```
i) pgdbg [all] 0>[*] break main; cont; wait; print f@11@i
ii) pgdbg [all] 0>[*] break main; [*]cont; [*]wait; [*]print f@11@i
```

Use the **wait** command if subsequent commands require threads to be stopped, as the **print** command in the example does.

The `threadwait` and `procwait` environment variables do not affect how commands within a compound statement are processed. These **pgienv** environment variables describe to *PGDBG* under what conditions (runstate of program) it should accept the next (compound) statement.

Parallel If, Else Statements

A prefix `p/t`-set can be used to parallelize an `if` statement. An `if` statement executes in the context of the current thread by default. The following example:

```
pgdbg [all] 0> [*] if (i==1) {break func; c; wait} else {sync func2}
```

is equivalent to the following pseudo-code:

```
for the subset of [*] where (i==1)
break func; c; wait; for the subset of [*] where (i!=1) sync func2
```

Parallel While Statements

A prefix p/t-set can be used to parallelize a while statement. A while statement executes in the context of the current thread by default. The following example:

```
pgdbg [all] 0> [*] while (i<10) {n; wait; print i}
```

is equivalent to the following pseudo-code:

```
loop:
if the subset of [*] is the empty set
goto done
endif
for the subset [s] of [*] where (i<10)
[s]n; [s]wait; [s]print i;
endfor
goto loop
```

The while statement terminates when either the subset of the target p/t-set matching the while condition is the empty set, or a return statement is executed in the body of the while.

Return Statements

The return statement is defined only in serial context since it cannot return multiple values. When return is used in a parallel statement, it returns the last value evaluated.

Chapter 10. Parallel Debugging with OpenMP

This chapter provides information on how to debug OpenMP applications. Before reading this chapter, review the information in [Chapter 9, “Parallel Debugging Overview”](#).

OpenMP and Multi-thread Support

PGDBG provides full control of threads in parallel regions. Commands can be applied to all threads, a single thread, or a group of threads. Thread identification in *PGDBG* uses the native thread numbering scheme for OpenMP applications; for other types of multi-threaded applications thread numbering is arbitrary. OpenMP PRIVATE data can be accessed accurately for each thread. *PGDBG* provides understandable status displays regarding per-thread state and location.

Advanced features provide for configurable thread stop modes and wait modes, allowing debugger operation that is concurrent with application execution.

Multi-thread and OpenMP Debugging

PGDBG automatically attaches to new threads as they are created during program execution. *PGDBG* reports when a new thread is created and the thread ID of the new thread is printed.

```
([1] New Thread)
```

The system ID of the freshly created thread is available through the **threads** command. You can use the **procs** command to display information about the parent process.

PGDBG maintains a conceptual current thread. When using the *PGDBG* CLI, the current thread is chosen by using the **thread** command.

```
pgdbg [all] 2> thread 3  
pgdbg [all] 3>
```

When using the *PGDBG* GUI, the current thread can be selected using the Current Thread drop-down list or by clicking in the Thread Grid. A subset of *PGDBG* commands known as thread-level commands apply only to the current thread. See [“Thread Level Commands,” on page 67](#), for more information.

The **threads** command lists all threads currently employed by an active program. It displays each thread's unique thread ID, system ID (OS process ID), execution state (running, stopped, signaled, exited, or killed), signal information and reason for stopping, and the current location if stopped or signaled. An arrow (=>) indicates the current thread. The process ID of the parent is printed in the top left corner. The **threads** command does not change the current thread.

```
pgdbg [all] 3> threads
0 ID PID STATE SIGNAL LOCATION
=> 3 18399 Stopped SIGTRAP main line: 31 in "omp.c" address: 0x80490ab
  2 18398 Stopped SIGTRAP main line: 32 in "omp.c" address: 0x80490cf
  1 18397 Stopped SIGTRAP main line: 31 in "omp.c" address: 0x80490ab
  0 18395 Stopped SIGTRAP f line: 5 in "omp.c" address: 0x8048fa0
```

In the GUI, thread state is represented by a color in the process/thread grid.

Table 10.1. Thread State Is Described Using Color

Thread State	Description	Color
Stopped	The thread is stopped at a breakpoint, or was directed to stop by <i>PGDBG</i> .	Red
Signaled	The thread is stopped due to delivery of a signal.	Blue
Running	The thread is running.	Green
Exited or Killed	The thread has been killed or has exited.	Black

Debugging OpenMP Private Data

PGDBG supports debugging of OpenMP private data for all supported languages. When an object is declared private in the context of an OpenMP parallel region, it essentially means that each thread team has its own copy of the object. This capability is shown in the following Fortran and C/ C++ examples, where the loop index variable *i* is private by default.

FORTTRAN example:

```
program omp_private_data
  integer array(8)
  call omp_set_num_threads(2)
!$OMP PARALLEL DO
  do i=1,8
    array(i) = i
  enddo
!$OMP END PARALLEL DO
  print *, array
end
```

C/ C++ example:

```
#include <omp.h>
int main ()
{
  int i;
  int array[8];
  omp_set_num_threads(2);
#pragma omp parallel
```

```

{
#pragma omp for
  for (i = 0; i < 8; ++i) {
    array[i] = i;
  }
}
for (i = 0; i < 8; ++i) {
  printf("array[%d] = %d\n",i, array[i]);
}
}

```

Compile the examples with a PGI compiler. The display of OpenMP private data in the resulting executables as debugged by *PGDBG* is as follows:

```

pgdbg [all] 0> [*] print i
[0] print i:
1
[1] print i:
5

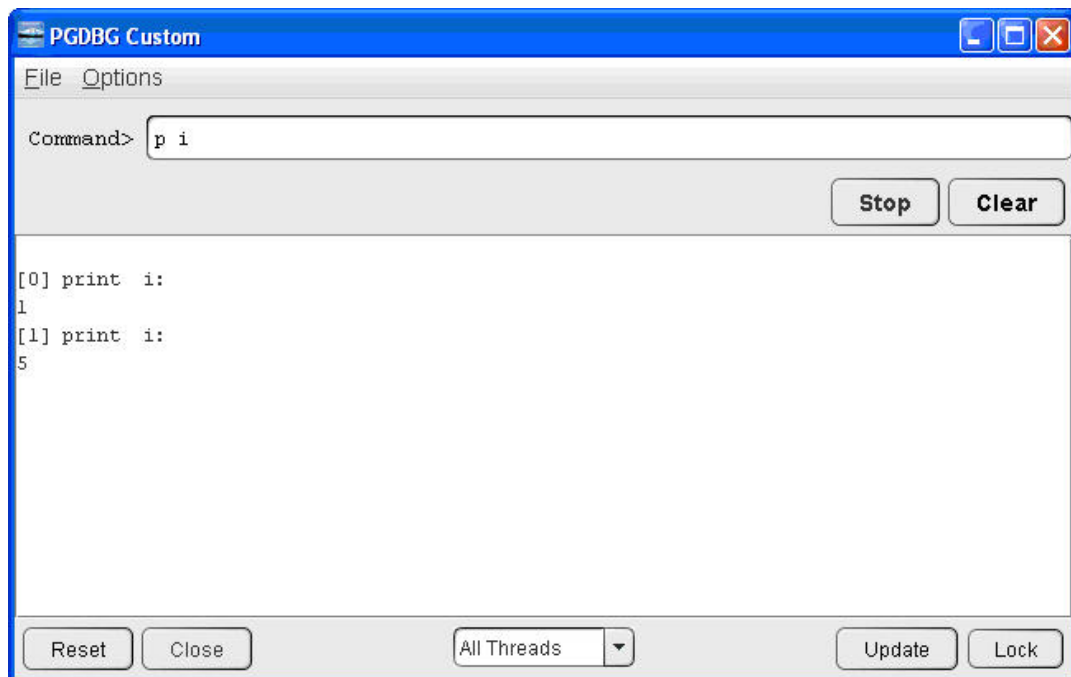
```

The example specifies `[*]` for the p/t-set to execute the **print** command on all threads. [Figure 10.1](#) shows the values for `i` in the *PGDBG* GUI using a Custom Window.

Note

All Threads is selected in the Context drop-down list to display the value on both threads.

Figure 10.1. OpenMP Private Data in PGDBG GUI



Chapter 11. Parallel Debugging with MPI

PGDBG is a parallel application debugger capable of debugging multi-process MPI applications and hybrid multi-thread/multi-process applications that use MPI to communicate between multi-threaded or OpenMP processes. This section begins with a general overview of how to use *PGDBG* to debug parallel MPI applications before detailing how to launch MPI applications under debug using the various supported platforms and versions of MPI.

For information on compiling a program using MPI, refer to “Using MPI” in the *PGI Compiler User’s Guide*.

MPI and Multi-Process Support

PGDBG can debug MPI applications running on the local system or distributed across a cluster. MPI applications must be started under debugger control. Process identification uses the MPI rank within `MPI_COMM_WORLD`.

An MPI program under debug cannot be restarted. The program must be reloaded as though it were a new program. You must exit the debugger and start a new debug session.

MPI debugging is supported on Linux, Windows, and Mac OSX.

Process Control

Here are some general things to consider when debugging an MPI program:

- Use p/t-sets to focus on a set of processes. Be mindful of process dependencies.
- For a running process to receive a message, the sending process must be allowed to run.
- Process synchronization points, such as `MPI_Barrier`, do not return until all processes have hit the sync point.
- `MPI_Finalize` acts as an implicit barrier except when using MPICH-1 where Process 0 returns while Processes 1 through n-1 exit.

You can apply a control command, such as **cont** or **step**, to a stopped process while other processes are running. A control command applied to a running process is applied to the stopped threads of that process and is ignored by its running threads.

PGDBG automatically switches to process wait mode `none` as soon as it attaches to its first MPI process. See the **pgienv** command and “Configurable Wait Mode,” on page 70 for details.

Process Synchronization

Use the *PGDBG* **sync** command to synchronize a set of processes to a particular point in the program. The following command runs all processes to `MPI_Finalize`:

```
pgdbg [all] 0.0> sync MPI_Finalize
```

The following command runs all threads of process 0 and process 1 to `MPI_Finalize`:

```
pgdbg [all] 0.0> [0:1.*] sync MPI_Finalize
```

A **sync** command only successfully syncs the target processes if the sync address is well defined for each member of the target process set, and all process dependencies are satisfied. If these conditions are not met, a member could wait forever for a message. The debugger cannot predict if a text address is in the path of an executing process.

MPI Message Queues

PGDBG can dump MPI message queues. When using the CLI, use the **mqdump** command, described in “Memory Access,” on page 113. When using the GUI, the message queues are displayed in the MPI Messages debug information tab.

The following error message may appear in the MPI Messages tab or when invoking **mqdump**:

```
ERROR: MPI Message Queue library not found.
Try setting 'PGDBG_MQS_LIB_OVERRIDE' environment variable
or set via the PGDBG command: pgienv mqslib <path>.
```

If this message is displayed, then the `PGDBG_MQS_LIB_OVERRIDE` environment variable should be set to the absolute path of `libtvmppich.so` or another shared object that is compatible with the version of MPI being used. The default path can also be overridden via the **mqslib** variant of the **pgienv** command.

Microsoft MPI does not currently provide support for dumping message queues.

MPI Groups

PGDBG identifies each process by its `MPI_COMM_WORLD` rank. In general, *PGDBG* currently ignores MPI groups.

Use halt instead of Ctrl+C

Entering `Ctrl+C` from the *PGDBG* command line can be used to halt all running processes. However, this is not the preferred method to use while debugging an MPICH-1 program. *PGDBG* automatically switches to process wait mode `none` (**pgienv procwait none**) as soon as it attaches to its first MPI process.

Setting **pgienv procwait none** allows commands to be entered while there are running processes, which allows the use of the **halt** command to stop running processes without the use of Ctrl+C.

Note

halt cannot interrupt a **wait** command. Ctrl+C must be used for this.

In MPI debugging, **wait** should be used with care.

SSH and RSH

By default, *PGDBG* uses rsh for communication between remote *PGDBG* components. *PGDBG* can also use ssh for secure environments. The environment variable `PGRSH` should be set to `ssh` or `rsh`, to indicate the desired communication method.

If you opt to use `ssh` as the mechanism for launching the remote components of *PGDBG*, you may want to do some additional configuration. The default configuration of `ssh` can result in a password prompt for each remote cluster node on which the debugger runs. Check with your network administrator to make sure that you comply with your local security policies when configuring `ssh`. The following steps provide one way to configure SSH to eliminate this prompt.

```
$ ssh-keygen -t dsa
$ eval `ssh-agent -s`
$ ssh-add
<make sure that $HOME is not group-writable>
$ cd $HOME/.ssh
$ cp id_dsa.pub authorized_keys
```

Then for each cluster node you use in debugging, use:

```
$ ssh <host>
```

A few things that are important related to this example are these:

- The **ssh-keygen** command prompts for a passphrase that is used to authenticate to the `ssh-agent` during future sessions. The passphrase can be anything you choose.
- Once you answer the prompts to make the initial connection, subsequent connections should not require further prompting.
- The **ssh-agent -s** command is correct for the `sh` or `bash` shells. For `csh` shells, use **ssh-agent -c**.

After logging out and logging back in, the `ssh-agent` must be restarted and reauthorized. For example, in a `bash` shell, this is accomplished as follows:

```
$ eval `ssh-agent -s`
$ ssh-add
```

You must enter the passphrase that was initially given to `ssh-add` to authenticate to the `ssh-agent`.

For further information, consult your `ssh` documentation.

MPI Debugging on Linux

When installed as part of the PGI Cluster Development Kit (CDK) on Linux platforms, *PGDBG* supports multi-process MPI debugging. The *PGI CDK* contains versions of MPICH, MPICH2, and MVAPICH pre-configured to

support debugging cluster applications with PGDBG. Versions of MPI not included in the *PGI CDK* must be configured to support *PGDBG*; refer to the *PGI Installation Guide* or www.pgroup.com/support/faq.htm for more information.

Invoking PGDBG for MPI Debugging

The command used to start MPI debugging under MPICH-1 using the *PGDBG* GUI is:

```
% mpirun -np nprocs -dbg=pgdbg executable [ arg1,...argn ]
```

To use MPI debugging in text mode, be certain that the `DISPLAY` variable is undefined in the shell that is invoking `mpirun`. If this variable is set, you can undefine it by using one of the following commands:

For `sh/bash` users, use this command:

```
$ unset DISPLAY
```

For `csh/tcsh` users, use this command:

```
% unsetenv DISPLAY
```

To launch and debug an MPICH-2, MVAPICH, or HPMPPI job, use this command:

```
% pgdbg [-text] -mpi[:<launcher>] <mpiexec_args> [ -program_args arg1,...argn ]
```

The default setting for `<launcher>` in `-mpi:<launcher>` is `mpiexec`.

Note

If the path for `<launcher>` is not part of the `PATH` environment variable, then you must specify the full path to the `<launcher>` command.

```
% pgdbg [-text] -mpi:mpiexec <mpiexec_args> [ -program_args arg1,...argn ]
```

- For HPMPPI, `<launcher>` is `mpirun`, so use `-mpi:mpirun`
- For MVAPICH, `<launcher>` is `mpirun_rsh`, so use `-mpi:mpirun_rsh`

When debugging an MPI job by invoking *PGDBG* with the `-mpi` option, each process is stopped before the first assembly instruction in the program. Continuing execution using **step** or **next** is not appropriate; instead, use the **cont** command.

Another way to invoke the *PGDBG* GUI for debugging an MPI job applies only to the CDK version of MPICH-2:

```
% mpiexec -np nprocs -pgi executable [ arg1,...argn ]
```

Note

You cannot restart an MPI application from within *PGDBG*. You must exit the debugger and start a new debug session.

Newer versions of the Linux kernel support a security feature that allows shared objects to be loaded at randomized addresses. Older versions of *PGDBG* assumed that a shared object loaded by an MPI application would be loaded at the same address for each process of an MPI job; thus incorrect mapping of symbols in shared objects could occur when this mode is enabled.

PGDBG supports debugging of MPI jobs running on Linux kernels when this address randomization mode is enabled. However, when this mode is enabled, the current implementation of *PGDBG* does not share symbol table information associated with shared objects that are loaded by each process of an MPI job, which increases memory usage by *PGDBG*. Therefore, PGI recommends that this kernel mode be disabled on Linux clusters where *PGDBG* is used to debug MPI applications.

You can disable randomization mode by executing the following command as root on each node of the cluster:

```
sysctl -w kernel.randomize_va_space=0
```

Note

PGDBG emits a warning whenever it detects that it is being invoked on a multi-process MPI job when this kernel mode is enabled.

Using PGDBG for MPI Debugging

PGDBG automatically attaches to new MPI processes as they are created by the running MPI application. *PGDBG* displays an informational message as it attaches to the freshly created processes.

```
([1] New Process)
```

The MPI global rank is printed with the message. You can use the **procs** command to list the host and the PID of each process by rank. The current process is indicated by an arrow (**=>**). You can use the **proc** command to change the current process by process ID.

```
pgdbg [all] 0.0> proc 1; procs
Process 1: Thread 0 Stopped at 0x804a0e2, function main, file MPI.c, line 30
#30: aft=time(&aft);
  ID  IPID  STATE  THREADS  HOST
   0   24765  Stopped    1      local
=> 1   17890  Stopped    1      red2.wil.st.com
```

The execution state of a process is described in terms of the execution state of its component threads. See [Table 10.1, “Thread State Is Described Using Color,” on page 78](#) for a description of how thread state is represented in the GUI.

The *PGDBG* command prompt displays the current process and the current thread. In the above example, the current process was changed to process 1 by the **proc 1** command and the current thread of process 1 is 0; this is written as 1.0:

```
pgdbg [all] 1.0>
```

See [“Process and Thread Control,” on page 69](#) for a complete description of the prompt format.

The following rules apply during a *PGDBG* debug session:

- *PGDBG* maintains a conceptual current process and current thread.
- Each active process has a thread set of size ≥ 1 .
- The current thread is a member of the thread set of the current process.

Certain commands, when executed, apply only to the current process or the current thread. See [“Process Level Commands,” on page 67](#) and [“Thread Level Commands,” on page 67](#) for more information.

The PGI license keys restrict the total number of MPI processes that can be debugged. In addition, there are internal limits on the number of threads per process that can be debugged.

Debugging Support for MPICH-1

With the CDK version of MPICH-1, *PGDBG* supports redirecting stdin, stdout, and stderr with the following MPICH switches:

Table 11.1. MPICH Support

Command	Output
<code>-stdout <file></code>	Redirect standard output to <file>
<code>-stdin <file></code>	Redirect standard input from <file>
<code>-stderr <file></code>	Redirect standard error to <file>

PGDBG also provides support for the following MPICH switches:

Command	Output
<code>-nolocal</code>	<i>PGDBG</i> runs locally, but no MPI processes run locally
<code>-all-local</code>	<i>PGDBG</i> runs locally, all MPI processes run locally

For information about how to configure an arbitrary installation of MPICH to use *PGDBG*, see the *PGDBG* online FAQ at www.pgroup.com/support/faq.htm.

When *PGDBG* is invoked via **mpirun** the following *PGDBG* command-line arguments are not accessible. A workaround is listed for each.

Argument	Workaround
<code>-dbx</code>	Include 'pgienv dbx on' in .pgdbgrc file.
<code>-s startup</code>	Use .pgdbgrc default script file and the script command.
<code>-c "command"</code>	Use .pgdbgrc default script file and the script command.
<code>-text</code>	Clear your DISPLAY environment variable before invoking mpirun.
<code>-t <target></code>	Add to the beginning of the PATH environment variable a path to the appropriate <i>PGDBG</i> .

MPI Debugging on Windows

PGDBG supports Microsoft's version of MPI called MSMPI. *PGDBG* can debug MSMPI programs running locally or on a distributed system. This section provides general information about building with and debugging MSMPI applications

Installing MSMPI

The MSMPI headers and libraries are available in the Microsoft HPC Pack SDK. This install package is available for download directly from Microsoft. You must install the HPC Pack SDK before you can build, run or debug MSMPI applications.

Building with MSMPI

To build an application using the MSMPI libraries, use the option `-Mmpi=msmpi`. This compiler flag inserts options into the compile and link lines to pick up the MSMPI headers and libraries. For `-Mmpi=msmpi` to work properly, the `CCP_SDK` environment variable must be set. This variable is typically set when the Microsoft HPC Pack SDK is installed.

Debug MSMPI Applications Locally

MSMPI applications can be run and debugged locally. In other words, an HPC Server cluster is not required to take advantage of MSMPI.

To invoke the *PGDBG* debugger to debug an MSMPI application locally, use the `pgdbg -mpi` option:

```
pgdbg -mpi[:<path>] <mpiexec_args> [-program_args arg1,...argn]
```

The location of `mpiexec` should be part of your `PATH` environment variable. Otherwise, you should specify the pathname for **mpiexec** as `<path>` in `-mpi[:<path>]`.

In this example, to debug an MSMPI application named `prog` using four processes running on the host system, use a command like:

```
PGI$ pgdbg -mpi -n 4 prog.exe
```

Debug MSMPI Applications on a Cluster

PGDBG provides support for debugging MSMPI applications on Windows HPC Server 2008 clusters. A PGI CDK license is required to enable *PGDBG*'s distributed debugging capabilities.

Microsoft's cluster management software uses a job management application to launch and manage executables on the head and cluster nodes. To begin distributed debugging on a cluster, invoke `pgdbg` with both the `-pgserv` and `-mpi` options:

```
pgdbg -pgserv:<path_to_pgserv.exe> -mpi[:<job submit command>]
```

The `-pgserv` option causes the *PGDBG* remote debug agent, called **pgserv**, to be copied into the current working directory when debugging is launched. This action ensures that **pgserv** can be found on all the nodes.

The **job submit** command references Microsoft's HPC Job Manager.

Note

The current working directory must be designated as a shared directory across all nodes of the cluster. All nodes must have access to this directory in order for distributed execution and debugging to succeed.

In this example, to debug an MSMPI application named `prog` using four processes running on a Windows cluster, use a command like:

```
PGI$ pgdbg -pgserv -mpi:job.exe
      submit /numprocessors:4 /workdir:\\head-node\sharedir mpiexec prog.exe
```


Chapter 12. Parallel Debugging of Hybrid Applications

PGDBG supports debugging hybrid multi-thread/multi-process applications that use MPI to communicate between multi-threaded or OpenMP processes. Multi-threaded and OpenMP applications may be run using more threads than the available number of CPUs, and MPI applications may allocate more than one process to a cluster node. *PGDBG* supports debugging the supported types of applications regardless of how well the requested number of threads matches the number of CPUs or how well the requested number of processes matches the number of cluster nodes.

PGDBG Multilevel Debug Mode

As described in “[PGDBG Debug Modes](#),” on page 60, *PGDBG* can operate in four debug modes. The mode determines a short form for uniquely naming threads and processes.

The debug mode is set automatically or by the **pgienv** command.

When *PGDBG* detects multilevel debugging, it sets the debug mode to multilevel. To manually set the debug mode to multilevel, use the **pgienv** command:

```
pgdbg> pgienv mode multilevel
```

Multilevel Debugging

The name of a thread in multilevel debug mode is the thread ID prefixed with its parent process ID. This forms a unique name for each thread across all processes. This naming scheme is valid in all debug modes. *PGDBG* changes automatically to multilevel debug mode from process-only debug mode or threads-only debug mode when at least one MPI process creates multiple threads.

Example 12.1. Thread IDs in multilevel debug mode

0.1	Thread 1 of process 0
0.*	All threads of process 0
*	All threads of all processes

In multilevel debug, mode status and error messages are prefixed with process/thread IDs depending on context. Further, in multilevel debug mode, *PGDBG* displays the current p/t-set in square brackets followed by the ID of the current thread prefixed by the id of its parent process:

```
pgdbg [all] 1.0>  
Current thread 1.0
```

For more information on p/t sets, refer to [“Process/Thread Sets,” on page 62](#).

Chapter 13. Command Reference

This chapter describes the *PGDBG* command set in detail, grouping the commands by these categories:

Conversions	Miscellaneous	Process-Thread Sets	Scope
Events	Printing Variables and Expressions	Program Locations	Symbols and Expressions
Memory Access	Process Control	Register Access	Target

For an alphabetical listing of all the commands, with a brief description of each, refer to “[Command Summary](#),” on page 34 in “*Command Summary*”.

Notation Used in Command Sections

The command sections that follow use these conventions for the command names and arguments, when the command accepts one.

- Command names may be abbreviated by omitting the portion of the command name enclosed in brackets ([]).
- Argument names are italicized.
- Argument names are chosen to indicate what kind of argument is expected.
- Arguments enclosed in brackets([]) are optional.
- Two or more arguments separated by a vertical line (|) indicate that any one of the arguments is acceptable.
- An ellipsis (...) indicates an arbitrarily long list of arguments.
- Other punctuation, such as commas and quotes, must be entered as shown.

Example 13.1. Syntax examples

Example 1:

```
lis[t] [count | lo:hi | routine | line,count]
```

This syntax indicates that the command **list** may be abbreviated to **lis**, and that it can be invoked without any arguments or with *one* of the following: an integer count, a line range, a routine name, or a line and a count.

Example 2:

```
att[ach] pid [exe]
```

This syntax indicates that the command **attach** may be abbreviated to **att**, and, when invoked, must have a process ID argument, *pid*. Optionally you can specify an executable file, *exe*.

Process Control

The following commands control the execution of the target program. *PGDBG* lets you easily group and control multiple threads and processes. For more details, refer to “[Basic Process and Thread Naming](#),” on [page 59](#).

attach

```
att[ach] pid [exe]
```

Attach to a running process with process ID *pid*. Use *exe* to specify the absolute path of the executable file. For example, `attach 1234` attempts to attach to a running process whose process ID is 1234. You may enter something like `attach 1234 /home/demo/a.out` to attach to a process ID 1234 called `/home/demo/a.out`.

PGDBG attempts to infer the arguments of the attached target application. If *PGDBG* fails to infer the argument list, then the program behavior is undefined if the **run** or **rerun** command is executed on the attached process.

The `stdio` channel of the attached process remains at the terminal from which the program was originally invoked.

The **attach** command is not supported for MPI programs.

cont

```
c[ont]
```

Continue execution from the current location.

debug

```
de[bug] [target [ arg1... argn]]
```

Load the specified target program with optional command-line arguments.

detach

```
det[ach]
```

Detach from the current running process.

halt

```
halt [command]
```

Halt the running process or thread.

load

```
lo[ad] [program [args]]
```

Without arguments, **load** prints the name and arguments of the program being debugged. With arguments, **load** loads the specified *program* for debugging. Provide program arguments as needed.

next

```
n[ext] [count]
```

Stop after executing one source line in the current routine. This command steps over called routines. The *count* argument stops execution only after executing *count* source lines.

nexti

```
nexti [count]
```

Stop after executing one instruction in the current routine. This command steps over called routines. The *count* argument stops execution only after executing *count* instructions.

proc

```
proc [id]
```

Set the current process to the process identified by *id*. When issued with no argument, **proc** lists the location of the current thread of the current process in the current program. For information on how processes are numbered, refer to [“MPI Debugging on Linux,” on page 83](#).

procs

```
procs
```

Print the status of all active processes, listing each process by its logical process ID.

quit

```
q[uit]
```

Terminate the debugging session.

rerun

```
rer[un] [arg0  
arg1 ... argn] [< inputfile ] [ [ > | >& | >> | >>& ] outputfile ]
```

The **rerun** command is the same as **run** with one exception: if no args are specified with **rerun**, then no args are used when the program is launched.

run

```
ru[n] [arg0 arg1
... argn] [< inputfile ] [ [ > | >& | >> | >>& ] outputfile ]
```

Execute the program from the beginning. If arguments *arg0*, *arg1*, and so on are specified, they are set up as the command-line arguments of the program. Otherwise, the arguments for the previous **run** command are used. Standard input and standard output for the target program can be redirected using < or > and an input or output filename.

setargs

```
setargs [arg1, arg2, ... argn]
```

Set program arguments for use by the **run** command. The **rerun** command does not use the arguments specified by **setargs**.

step

```
s[tep] [count | count]
```

Stop after executing one source line. This command steps into called routines. The *count* argument stops execution after executing *count* source lines. The *up* argument stops execution after stepping out of the current routine (see [stepout](#)).

stepi

```
stepi [count | up]
```

Stop after executing one instruction. This command steps into called routines. The *count* argument stops execution after executing *count* instructions. The *up* argument stops the execution after stepping out of the current routine (see [stepout](#)).

stepout

```
stepo[ut]
```

Stop after returning to the caller of the current routine. This command sets a breakpoint at the current return address and continues execution to that point. For this command to work correctly, it must be possible to compute the value of the return address. Some routines, particularly terminal (i.e. leaf) routines at higher optimization levels, may not set up a stack frame. Executing **stepout** from such a routine causes the breakpoint to be set in the caller of the most recent routine that set up a stack frame. This command halts execution immediately upon return to the calling routine.

sync

```
sy[nc] line | func
```

Advance to the specified source location, either the specified *line* or the first line in the specified function *func*, ignoring any user-defined events.

synci

```
synci addr | func
```

Advance to the specified address *addr*, or to the first address in the specified function *func*, ignoring any user-defined events.

thread

```
thread [number]
```

Set the current thread to the thread identified by *number*; where *number* is a logical thread id in the current process' active thread list. When issued with no argument, **thread** lists the current program location of the currently active thread.

threads

```
threads
```

Print the status of all active threads. Threads are grouped by process. Each process is listed by its logical process id. Each thread is listed by its logical thread id.

wait

```
wait [any | all | none]
```

Return the *PGDBG* prompt only after specific processes or threads stop.

Process-Thread Sets

The following commands deal with defining and managing process thread sets. See [“Process/Thread Sets,” on page 62](#), for a detailed discussion of process-thread sets.

defset

```
defset name [p/t-set]
```

Assign a *name* to a process/thread set. In other words, define a named set of processes/threads. This set can then be referred to by its *name*. A list of named sets is stored by *PGDBG*.

focus

```
focus [p/t-set]
```

Set the target process/thread set for *PGDBG* commands. Subsequent commands are applied to the members of this set by default.

undefset

```
undefset [name | -all]
```

Remove a previously defined process/thread set from the list of process/thread sets. The debugger-defined p/t-set **[all]** cannot be removed.

viewset

```
viewset [name]
```

List the active members of the named process/thread set. If no process/thread set is given, list the active members of all defined process/thread sets.

whichsets

```
whichsets [p/t-set]
```

List all defined p/t-sets to which the members of a process/thread set belong. If no process/thread set is specified, the target process/thread set is used.

Events

The following commands deal with defining and managing events.

break

```
b[reak]
b[reak] line [if condition] [do {commands}]
b[reak] routine [if(condition)] [do {commands}]
```

When no arguments are specified, the **break** command prints the current breakpoints. Otherwise, set a breakpoint at the indicated line or routine. If a routine is specified, and the routine was compiled for debugging, then the breakpoint is set at the start of the first statement in the routine (after the routine's prologue code). If the routine was not compiled for debugging, then the breakpoint is set at the first instruction of the routine, prior to any prologue code. This command interprets integer constants as line numbers. To set a breakpoint at an address, use the **addr** command to convert the constant to an address, or use the **breaki** command.

When a condition is specified with *if*, the breakpoint occurs only when the specified condition is true. If *do* is specified with a command or several commands as an argument, the command or commands are executed when the breakpoint occurs.

The following table provides examples of using **break** to set breakpoints at various locations.

This break command...	Sets breakpoints...
<code>break 37</code>	at line 37 in the current file
<code>break "xyz.c"@37</code>	at line 37 in the file <code>xyz.c</code>
<code>break main</code>	at the first executable line of routine <code>main</code>
<code>break {addr 0xf0400608}</code>	at address <code>0xf0400608</code>
<code>break {line}</code>	at the current line
<code>break {pc}</code>	at the current address

The following more sophisticated command stops when routine `xyz` is entered only if the argument `n` is greater than 10.

```
break xyz if(xyz@n > 10)
```

The next command prints the value of `n` and performs a stack trace every time line 100 in the current file is reached.

```
break 100 do {print n; stack}
```

breaki

```
breaki
breaki routine [if (condition)] [do {commands}]
breaki addr [if (condition)] [do {commands}]
```

When no arguments are specified, the **breaki** command prints the current breakpoints. Otherwise, this command sets a breakpoint at the indicated address *addr* or *routine*.

- If a routine is specified, the breakpoint is set at the first address of the routine. This means that when the program stops at this type of breakpoint the prologue code which sets up the stack frame will not yet have been executed. As a result, values of stack arguments may not yet be correct.
- Integer constants are interpreted as addresses.
- To specify a line, use the **lines** command to convert the constant to a line number, or use the **break** command.
- The *if* and *do* arguments are interpreted in the same way as for the **break** command.

The following table provides examples of setting breakpoints using **breaki**.

This breaki command...	Sets breakpoints...
<code>breaki 0xf0400608</code>	at address 0xf0400608
<code>breaki {line 37}</code>	at line 37 in the current file
<code>breaki "xyz.c"@37</code>	at line 37 in the file <code>xyz.c</code>
<code>breaki main</code>	at the first executable address of routine <code>main</code>
<code>breaki {line}</code>	at the current line
<code>breaki {pc}</code>	at the current address

In the following slightly more complex example, when `n` is greater than 3, the following command stops and prints the new value of `n` at address 0x6480:

```
breaki 0x6480 if(n>3) do {print "n=", n}
```

breaks

```
breaks
```

Display all the existing breakpoints.

catch

```
catch [sig:sig] [sig [, sig...]]
```

When no arguments are specified, the **catch** command prints the list of signals being caught. With the *sig:sig* argument, this command catches the specified range of signals. With a list of signals, catch the signals with the specified number(s). When signals are caught, *PGDBG* intercepts the signal and does not deliver it to the target application. The target runs as though the signal was never sent.

clear

```
clear [ all | routine | line | {addr addr}]
```

Clear one or more breakpoints. Use the *all* argument to clear all breakpoints. Use the *routine* argument to clear all breakpoints from the first statement in the specified routine. Use the *line* number argument to clear all breakpoints from the specified line number in the current source file. Use the *addr* argument, clear breakpoints from the specified address *addr*.

When no arguments are specified, the **clear** command clears all breakpoints at the current location.

delete

```
del[ete] [event-number | 0 | all | event-number [, event-number...]]
```

Use the **delete** command without arguments to list all defined events by their event-number.

Use the **delete** command with arguments to delete events. Delete all events with *all* or delete just the event with the specified *event-number*. Providing the argument 0, that is, using **delete 0**, is the same as using **delete all**.

disable

```
disab[le] [event-number | all ]
```

When no arguments are specified, the **disable** command prints both enabled and disabled events.

With arguments, this command disables the event specified by *event-number* or *all* events. Disabling an event definition suppresses actions associated with the event, but leaves the event defined so that it can be used later. (see [enable](#).)

do

```
do {commands} [if (condition)]
do {commands} at line [if (condition)]
do {commands} in routine [if (condition)]
```

Define a **do** event. This command is similar to [watch](#) except that instead of defining an expression, it defines a list of commands to be executed. Without the optional arguments *at* or *in*, the commands are executed at each line in the program.

Use *at* with a *line* number to specify the commands to be executed each time that line is reached. Use *in* with a *routine* to specify the commands to be executed at each line in the routine. The optional *if* argument has the same meaning that it has in the [watch](#). If a condition is specified, the **do** commands are executed only when the condition is true.

doi

```
doi {commands} [if (condition)]
doi {commands} at addr [if (condition)]
doi {commands} in routine [if (condition)]
```

Define a **doi** event. This command is similar to **watchi** except that instead of defining an expression, **doi** defines a list of commands to be executed. If an address *addr* is specified, then the commands are executed each time that the specified address is reached. If a *routine* is specified, then the commands are executed at each instruction in the routine. If neither an address nor a routine is specified, then the commands are executed at each instruction in the program. The optional *if* argument has the same meaning that it has in the **do** and **watch** commands. If a condition is specified, the **doi** commands are executed only when the condition is true.

enable

```
enab[le] [event-number | all ]
```

Without arguments, the **enable** command prints both enabled and disabled events.

With arguments, this command enables the event *event-number* or *all* events.

hwatch

```
hwatch addr | var [if (condition)] [do {commands}]
```

Define a hardware watchpoint. This command uses hardware support to create a watchpoint for a particular address or variable. The event is triggered by hardware when the byte at the given address is written. This command is only supported on systems that provide the necessary hardware and software support.

Note

Only one hardware watchpoint can be defined at a time.

When the optional *if* argument is specified, the event action is only triggered if the expression is true. When the optional *do* argument is specified, then the commands are executed when the event occurs.

hwatchboth

```
hwatchb[oth] addr | var [if (condition)] [do {commands}]
```

Define a hardware read/write watchpoint. This event is triggered by hardware when the byte at the given address or variable is either read or written. As with **hwatch**, system hardware and software support must exist for this command to be supported. The optional *if* and *do* arguments have the same meaning as for the **hwatch** command.

hwatchread

```
hwatchb[oth] addr | var [if (condition)] [do {commands}]
```

Define a hardware read watchpoint. This event is triggered by hardware when the byte at the given address or variable is read. As with **hwatch**, system hardware and software support must exist for this command to be supported. The optional *if* and *do* arguments have the same meaning as for the **hwatch** command.

ignore

```
ignore [sig:sig] [sig [, sig...]]
```

Without arguments, the **ignore** command prints the list of signals being ignored. With the *sig:sig* argument this command ignores the specified range of signals. With a list of signals the command ignores signals with the specified number.

When a particular signal number is ignored, signals with that number sent to the target application are not intercepted by *PGDBG*; rather, the signals are delivered to the target.

For information on intercepting signals, refer to [catch](#).

status

```
stat[us]
```

Display all the event definitions, including an event number by which each event can be identified.

stop

```
stop var
stop at line [if (condition)][do {commands}]
stop in routine [if (condition)][do {commands}]
stop if (condition)
```

Break when the value of the indicated variable *var* changes. Use the *at* argument and a *line* to set a breakpoint at a line number. Use the *in* argument and a *routine* name to set a breakpoint at the first statement of the specified routine. With the *if* argument is used, the debugger stops when the condition is true.

stopi

```
stopi var
stopi at address [if (condition)][do {commands}]
stopi in routine [if (condition)][do {commands}]
stopi if (condition)
```

Break when the value of the indicated variable *var* changes. Set a breakpoint at the indicated address or routine. Use the *at* argument and an *address* to specify an address at which to stop. Use the *in* argument and a *routine* name to specify the first address of the specified routine at which to stop. When the *if* argument is used, the debugger stops when the condition is true.

trace

```
trace var [if (condition)][do {commands}]
trace routine [if (condition)][do {commands}]
trace at line [if (condition)][do {commands}]
trace in routine [if (condition)][do {commands}]
trace inclass class [if (condition)][do {commands}]
```

Use *var* to activate tracing when the value of *var* changes. Use *routine* to activate tracing when the subprogram *routine* is called. Use *at* to display the specified *line* each time it is executed. Use *in* to display the current line while in the specified *routine*. Use *inclass* to display the current line while in each member function of the specified *class*. If a condition is specified, tracing is only enabled if the condition evaluates to true. The *do* argument defines a list of commands to execute at each trace point.

Use the command **pgienv speed** to set the time in seconds between trace points. Use the **clear** command to remove tracing for a line or routine.

tracei

```
tracei var [if (condition)][do {commands}]
tracei at addr [if (condition)][do {commands}]
tracei in routine [if (condition)][do {commands}]
tracei inclass class [if (condition)][do {commands}]
```

Activate tracing at the instruction level. Use *var* to activate tracing when the value of *var* changes. Use *at* to display the instruction at *addr* each time it is executed. Use *in* to display memory instructions while in the subprogram *routine*. Use *inclass* to display memory instructions while in each member function of the specified *class*. If a condition is specified, tracing is only enabled if the condition evaluates to true. The *do* argument defines a list of commands to execute at each trace point.

Use the command **pgienv speed** to set the time in seconds between trace points. Use the **clear** command to remove tracing for a line or routine.

track

```
track expression [at line | in func] [if (condition)][do {commands}]
```

Define a track event. This command is equivalent to **watch** except that execution resumes after the new value of the expression is printed.

tracki

```
tracki expression [at addr | in func] [if (condition)][do {commands}]
```

Define an assembly-level track event. This command is equivalent to **watchi** except that execution resumes after the new value of the expression is printed.

unbreak

```
unb[reak] line | routine | all
```

Remove a breakpoint from the specified *line* or *routine*, or remove *all* breakpoints.

unbreaki

```
unbreaki addr | routine | all
```

Remove a breakpoint from the specified address *addr* or *routine*, or remove *all* breakpoints.

watch

```
wa[tch] expression
wa[tch] expression [if (condition)][do {commands}]
wa[tch] expression at line [if (condition)][do {commands}]
wa[tch] expression in routine [if (condition)][do {commands}]
```

Define a watch event. The given expression is evaluated, and subsequently, each time the value of the expression changes, the program stops and the new value of the expression is printed. If a *line* is specified,

the expression is only evaluated at that line. If a *routine* is specified, the expression is evaluated at each line in the routine. If no location is specified, the expression is evaluated at each line in the program. If a *condition* is specified, the expression is evaluated only when the condition is true. If *commands* are specified using *do*, they are executed whenever the expression is evaluated and its value changes.

The watched expression may contain local variables, although this is not recommended unless a routine or address is specified to ensure that the variable is only evaluated when it is in the current scope.

NOTE

Using watchpoints indiscriminately can dramatically slow program execution.

Using the *at* and *in* arguments speeds up execution by reducing the amount of single-stepping and expression evaluation that must be performed to watch the expression. For example:

```
watch i at 40
```

may not slow program execution noticeably, while

```
watch i
```

does slow execution considerably.

watchi

```
watchi expression
watchi expression [if (condition)][do {commands}]
watchi expression at addr [if (condition)][do {commands}]
watchi expression in routine [if (condition)][do {commands}]
```

Define an assembly-level watch event. This command functions similarly to the **watch** command with two exceptions: 1) the argument interprets integers as addresses rather than line numbers and 2) the *expression* is evaluated at every instruction rather than at every line.

This command is useful when line number information is limited, which may occur when debug information is not available or assembly must be debugged. Using **watchi** causes programs to execute more slowly than **watch**.

when

```
when do {commands} [if (condition)]
when at line do {commands} [if (condition)]
when in routine do {commands} [if (condition)]
```

Execute *commands* at every line in the program, at a specified *line* in the program, or in the specified *routine*. If an optional *condition* is specified, commands are executed only when the *condition* evaluates to true.

wheni

```
wheni do {commands} [if (condition)]
wheni at addr do {commands} [if (condition)]
wheni in routine do {commands} [if (condition)]
```

Execute *commands* at each address in the program. If an address *addr* is specified, the commands are executed each time the address is reached. If a *routine* is specified, the commands are executed at each line in the routine. If an optional *condition* is specified, commands are executed whenever the *condition* evaluates to true.

Program Locations

This section describes *PGDBG* program location commands.

arrive

```
arri[vel]
```

Print location information for the current location.

cd

```
cd [dir]
```

Change directories to the \$HOME directory or to the specified directory *dir*.

disasm

```
dis[asm] [ count | lo:hi | routine | addr, count ]
```

Disassemble memory.

If no argument is given, disassemble four instructions starting at the current address. If an integer *count* is given, disassemble *count* instructions starting at the current address. If an address range (*lo:hi*) is given, disassemble the memory in the range. If a *routine* is given, disassemble the entire routine. If the routine was compiled for debugging and source code is available, the source code is interleaved with the disassembly. If an address *addr* and a *count* are both given, disassemble *count* instructions starting at the provided address.

edit

```
edit [filename | routine]
```

Use the editor specified by the environment variable \$EDITOR to edit a file.

If no argument is supplied, edit the current file starting at the current location. To edit a specific file, provide the *filename* argument. To edit the file containing the subprogram *routine*, specify the routine name.

In the *PGDBG* GUI, command-line editors like *vi*, *emacs*, or *nano* are launched in the Program I/O Window. On Windows platforms, arguments to the editor may need to be quoted to account for spaces in pathnames.

file

```
file [filename]
```

Change the source file to the file *filename* and change the scope accordingly. With no argument, print the current file.

lines

```
lines [routine]
```

Print the lines table for the specified *routine*. With no argument, prints the lines table for the current routine.

list

```
lis[t] [ count | line,num | lo:hi | routine[,num] ]
```

Provide a source listing.

By default, **list** displays ten lines of source centered at the current source line. If a *count* is given, list the specified *count* number of lines. If a *line* and *count* are both given, start the listing of *count* lines at *line*. If a line range (*lo:hi*) is given, list the indicated source lines in the current source file. If a *routine* name is given, list the source code for the indicated routine. If a *number* is specified with *routine*, list the first *number* lines of the source code for the indicated routine.

```
list [dbx mode]
```

The **list** command works somewhat differently when *PGDBG* is in dbx mode.

```
lis[t] [ line | first,last | routine | file ]
```

By default, list displays ten lines of source centered at the current source line. If a *line* is provided, the source at that line is displayed. If a range of line numbers is provided (*first,last*), lines from the first specified line to the last specified line are displayed. If a *routine* is provided, the display listing begins in that routine. If a *file* name is provided, the display listing begins in that file. File names must be quoted.

pwd

```
pwd
```

Print the current working directory.

stackdump

```
stackd[ump] [count]
```

Print a formatted dump of the stack. This command displays a hex dump of the stack frame for each active routine. This command is an assembly-level version of the [stacktrace](#) command. If a *count* is specified, display a maximum of *count* stack frames.

stacktrace

```
stack[trace] [count]
```

Print the call stack. Print the available information for each active routine, included the routine name, source file, line number, and current address. This command also prints the names and values of any arguments, when available. If a *count* is specified, display a maximum of *count* stack frames. The **stacktrace** and [where](#) commands are equivalent.

where

```
w[here] [count]
```

Print the call stack. Print the available information for each active routine, included the routine name, source file, line number, and current address. This command also prints the names and values of any arguments, when available. If a *count* is specified, display a maximum of *count* stack frames. The **where** and **stacktrace** commands are equivalent.

/

```
/
/string/
```

Search forward for a *string* of characters in the current source file. With a specified string, search for the next occurrence of *string* in the current source file.

?

```
?
?string?
```

Search backward for a *string* of characters in the current source file. Without arguments, search for the previous occurrence of *string* in the current source file.

Printing Variables and Expressions

This section describes *PGDBG* commands used for printing and setting variables. The primary print commands are **print** and **printf**, described at the beginning of this section. The rest of the commands for printing provide alternate methods for printing.

print

```
p[rint] exp1 [...expn]
```

Evaluate and print one or more expressions. This command is invoked to print the result of each line of command input. Values are printed in a format appropriate to their type. For values of structure type, each field name and value is printed. Character pointers are printed as a hex address followed by the character string.

Character string constants print out literally using a comma-separated list. For example:

```
pgdbg> print "The value of i is ", i
```

Prints this:

```
"The value of i is", 37
```

The array sub-range operator (:) prints a range of an array. The following examples print elements 0 through 9 of the array *a*:

C/ C++ example 1:

```
pgdbg> print a[0:9]
a[0:4]: 0 1 2 3 4
a[5:9]: 5 6 7 8 9
```

FORTRAN example 1:

```
pgdbg> print a(0:9)
a(0:4): 0 1 2 3 4
a(5:9): 5 6 7 8 9
```

Notice that the output is formatted and annotated with index information. *PGDBG* formats array output into columns. For each row, the first column prints an index expression which summarizes the elements printed in that row. Elements associated with each index expression are then printed in order. This is especially useful when printing slices of large multidimensional arrays.

PGDBG also supports array expression strides. Below are examples for C/ C++ and FORTRAN.

C/ C++ example 2:

```
pgdbg> print a[0:9:2]
a[0:8] 0 2 4 6 8
```

FORTRAN example 2:

```
pgdbg> print a(0:9:2)
a(0:8): 0 2 4 6 8
```

The print statement may be used to display members of derived types in FORTRAN or structures in C/ C++. Here are examples.

C/ C++ example 3:

```
typedef struct tt {
    int a[10];
}TT;
TT d = {0,1,2,3,4,5,6,7,8,9};
TT * p = &d;
```

```
pgdbg> print d.a[0:9:2]
d.a[0:8:2]: 0 2 4 6 8
pgdbg> print p->a[0:9:2]
p->a[0:7:2]: 0 2 4 6
p->a[8]: 8
```

FORTRAN example 3:

```
type tt
integer, dimension(0:9) :: a
end type
type (tt) :: d
data d%a / 0, 1, 2, 3, 4, 5, 6, 7, 8, 9 /
```

```
pgdbg> print d%a(0:9:2)
d%a(0:8:2): 0 2 4 6 8
```

printf

```
printf "format_string", expr,...expr
```

Print expressions in the format indicated by the format string. This command behaves like the C library function printf. For example:

```
pgdbg> printf "f[%d]=%G",i,f[i]
f[3]=3.14
```

The **pgienv stringlen** command sets the maximum number of characters that print with a **print** command. For example, the char declaration below:

```
char *c="a whole bunch of chars over 1000 chars long....";
```

By default, the **print c** command prints only the first 512 (default value of stringlen) bytes. Printing of C strings is usually terminated by the terminating null character. This limit is a safeguard against unterminated C strings.

ascii

```
asc[ii] exp [,...exp]
```

Evaluate and print *exp* as an ASCII character. Control characters are prefixed with the '^' character; for example, 3 prints as ^c. Otherwise, values that cannot be printed as characters are printed as integer values prefixed by '\'. For example, 250 is printed as \250.

bin

```
bin exp [,...exp]
```

Evaluate and print the expressions. Integer values are printed in base2.

dec

```
dec exp [,...exp]
```

Evaluate and print the expressions. Integer values are printed in decimal.

display

```
display [ exp [,...exp] ]
```

Without arguments, list the expressions for *PGDBG* to automatically display at breakpoints. With one or more arguments, print expression *exp* at every breakpoint. For more information, refer to the **undisplay** command.

hex

```
hex exp [,...exp]
```

Evaluate and print expressions as hexadecimal integers.

oct

```
oct exp [,...exp]
```

Evaluate and print expressions as octal integers.

string

```
str[ing] exp [,...exp]
```

Evaluate and print expressions as null-terminated character strings. This command prints a maximum of 70 characters.

undisplay

```
undisplay 0 | all | exp [...exp]
```

Remove all expressions specified by previous display commands. With an argument or several arguments, remove the expression *exp* from the list of display expressions.

Symbols and Expressions

This section describes the commands that deal with symbols and expressions.

assign

```
as[sign] var = exp
```

Set variable *var* to the value of the expression *exp*. The variable can be any valid identifier accessed properly for the current scope. For example, given a C variable declared `'int * i'`, you can use the following command to assign the value 9999 to it.

```
assign *i = 9999
```

call

```
call routine [(exp,...)]
```

Call the named *routine*. C argument passing conventions are used. Breakpoints encountered during execution of the routine are ignored. Fortran functions and subroutines can be called, but the argument values are passed according to C conventions. *PGDBG* may not always be able to access the return value of a Fortran function if the return value is an array. In the example below, *PGDBG* calls the routine `foo` with four arguments:

```
pgdbg> call foo(1,2,3,4)
```

If a signal is caught during execution of the called routine, *PGDBG* stops the execution and asks if you want to cancel the **call** command. For example, suppose a command is issued to call `foo` as shown above, and for some reason a signal is sent to the process while it is executing the call to `foo`. In this case, *PGDBG* prints the following prompt:

```
PGDBG Message: Thread [0] was signalled while executing a function
reachable from the most recent PGDBG command line call to foo. Would you
like to cancel this command line call? Answering yes will revert the register
state of Thread [0] back to the state it had prior to the last call to foo
from the command line. Answering no will leave Thread [0] stopped in the call
to foo from the command line.
Please enter 'y' or 'n' > y
Command line call to foo cancelled
```

Answering yes to this question returns the register state of each thread back to the state they had before invoking the **call** command. Answering no to this question leaves each thread at the point they were at when the signal occurred.

Note

Answering no to this question and continuing execution of the called routine may produce unpredictable results.

declaration

```
decl[aration] name
```

Print the declaration for the symbol *name* based on its type according to the symbol table. The symbol must be a variable, argument, enumeration constant, routine, structure, union, enum, or typedef tag.

For example, given the C declarations:

```
int i, iar[10];
struct abc {int a; char b[4]; struct
abc *c;}val;
```

the **decl** command provides the following output:

```
pgdbg> decl I
int i
```

```
pgdbg> decl iar
int iar[10]
```

```
pgdbg> decl val
struct abc val
```

```
pgdbg> decl abc
struct abc {
  int a;
  char b[4];
  struct abc *c;
};
```

entry

```
entr[y] [routine]
```

Return the address of the first executable statement in the program or specified *routine*. This is the first address after the routine's prologue code.

lval

```
lv[al] expr
```

Return the lvalue of the expression *expr*. The lvalue of an expression is the value it would have if it appeared on the left hand side of an assignment statement. Roughly speaking, an lvalue is a location to which a value can be assigned. This may be an address, a stack offset, or a register.

rval

```
rv[al] expr
```

Return the rvalue of the expression *expr*. The rvalue of an expression is the value it would have if it appeared on the right hand side of an assignment statement. The type of the expression may be any scalar, pointer, structure, or function type.

set

```
set var=expression
```

Set variable *var* to the value of *expression*. The variable can be any valid identifier accessed properly for the current scope. For example, given a C variable declared `int * i`, the following command could be used to assign the value 9999 to it.

```
pgdbg> set *i = 9999
```

sizeof

```
siz[EOF] name
```

Return the size, in bytes, of the variable type *name*. If *name* refers to a routine, **sizeof** returns the size in bytes of the subprogram.

type

```
type expr
```

Return the type of the expression *expr*. The expression may contain structure reference operators (`.`, and `->`), dereference (`*`), and array index (`[]`) expressions. For example, given the C declarations:

```
int i, iar[10];
struct abc {int a; char b[4];
struct abc *c;}val;
```

the **type** command provides the following output:

```
pgdbg> type i
int
pgdbg> type iar
int [10]
pgdbg> type val
struct abc
pgdbg> type val.a
int
```

```
pgdbg> type val.abc->b[2]
char
```

```
pgdbg> whatis
whatis name
```

With no arguments, print the declaration for the current routine.

With the name argument *name*, print the declaration for the symbol *name*.

Scope

The following commands deal with program scope. See “[Scope Rules](#)”, for a discussion of scope meaning and conventions.

class

```
class[s [class]
```

Without arguments, **class** returns the current class. With a *class* argument, enter the scope of class *class*.

classes

```
classse[s]
```

Print the C++ class names.

decls

```
decls [routine | "sourcefile" | {global} ]
```

Print the declarations of all identifiers defined in the indicated scope. If no scope is given, print the declarations for the current search scope.

down

```
down [number]
```

Enter the scope of the routine down one level or *number* levels on the call stack.

enter

```
en[ter] [routine | "sourcefile" | global ]
```

Set the search scope to be the indicated scope, which may be a *routine*, *sourcefile* or *global*. Using **enter** with no argument is the same as using **enter global**.

files

```
files
```

Return the list of known source files used to create the executable file.

global

```
glob[al]
```

Return a symbol representing global scope. This command is useful in combination with the scope operator @ to specify symbols with global scope.

names

```
names [routine | "sourcefile" | global ]
```

Print the names of all identifiers defined in the indicated scope. If no scope is specified, use the search scope.

scope

```
sco[pe]
```

Return a symbol for the search scope. The search scope is set to the current routine each time program execution stops. It may also be set using the **enter** command. The search scope is always searched first for symbols.

up

```
up [number]
```

Enter the scope of the routine up one level or *number* levels from the current routine on the call stack.

whereis

```
whereis name
```

Print all declarations for *name*.

which

```
which name
```

Print the full scope qualification of symbol *name*.

Register Access

System registers can be accessed by name. For details on referring to registers in *PGDBG*, refer to “[SSE Register Symbols](#),” on page 49.

fp

```
fp
```

Return the current value of the frame pointer.

pc

```
pc
```

Return the current program address.

regs

```
regs
regs -info
regs -grp=grp1[,grp2...]
regs -fmt=fmt1[,fmt2...]
regs -mode=scalar|vector
```

Print the names and values of registers. By default, **regs** prints the General Purpose registers. Use the `-grp` option to specify one or more register groups, the `-fmt` option to specify one or more display formats, and `-mode` to specify scalar or vector mode. Use the `-info` option to display the register groups on the current system and the display formats available for each group. All optional arguments with the exception of `-info` can be used with the others.

retaddr

```
ret[addr]
```

Return the current return address.

sp

```
sp
```

Return the current value of the stack pointer.

Memory Access

The following commands display the contents of arbitrary memory locations. For each of these commands, the *addr* argument may be a variable or identifier.

cread

```
cr[ead]addr
```

Fetch and return an 8-bit signed integer (character) from the specified address.

dread

```
dr[ead]addr
```

Fetch and return a 64-bit double from the specified address.

dump

```
du[mp] address[, count[,format-string]]
```

This command dumps the contents of a region of memory. The output is formatted according to a descriptor. Starting at the indicated address, values are fetched from memory and displayed according to the format descriptor. This process is repeated *count* times.

Interpretation of the format descriptor is similar to that used by **printf**. Format specifiers are preceded by %.

The recognized format descriptors are for decimal, octal, hex, or unsigned:

```
%d, %D, %o, %O, %x, %X, %u, %U
```

Default size is machine dependent. The size of the item read can be modified by either inserting 'h', or 'l' before the format character to indicate half word or long word. For example, if your machine's default size is 32-bit, then %hd represents a 16-bit quantity. Alternatively, a 1, 2, or 4 after the format character can be used to specify the number of bytes to read.

```
%c
```

Fetch and print a character.

```
%C
```

Fetch and print a float (lower case) or double (upper case) value using **printf** f, e, or g format.

```
%f, %F, %e, %E, %g, %G
```

Fetch and print a null terminated string.

```
%s
```

Interpret the next object as a pointer to an item specified by the following format characters. The pointed-to item is fetched and displayed.

```
%p<format-chars>
```

Pointer to int. Prints the address of the pointer, the value of the pointer, and the contents of the pointed-to address, which is printed using hexadecimal format.

```
%px
```

Fetch an instruction and disassemble it.

```
%i
```

Display address about to be dumped.

```
%w, %W
```

Display nothing but advance or decrement current address by n bytes.

```
%z<n>, %Z<n>, %z<-n>, %Z<-n>
```

Display nothing but advance current address as needed to align modulo n .

```
%a<n>, %A<n>
```

Display nothing but advance current address as needed to align modulo n .

fread

```
fr[ead]addr
```

Fetch and print a 32-bit float from the specified address.

iread

```
ir[ead] addr
```

Fetch and print a signed integer from the specified address.

lread

```
lr[ead] addr
```

Fetch and print an address from the specified address.

mqdump

```
mq[dump]
```

Dump MPI message queue information for the current process. For more information on **mqdump**, refer to [“MPI Message Queues,” on page 82](#).

sread

```
sr[ead]addr
```

Fetch and print a short signed integer from the specified address.

Conversions

The commands in this section are useful for converting between different kinds of values. These commands accept a variety of arguments, and return a value of a particular kind.

addr

```
ad[dr] [n | line n | routine | var | arg ]
```

Create an address conversion under these conditions:

- If an integer is given, return an address with the same value.
- If a line is given, return the address corresponding to the start of that line.
- If a routine is given, return the first address of the routine.
- If a variable or argument is given, return the address where that variable or argument is stored.

For example,

```
breaki {line {addr 0x22f0}}
```

function

```
func[tion] [[addr...] | [line...] ]
```

Return a routine symbol. If no argument is specified, return the current routine. If an address is given, return the routine containing *addr*. An integer argument is interpreted as an address. If a *line* is specified, return the routine containing that line.

line

```
lin[e] [ n | routine | addr ]
```

Create a source line conversion. If no argument is given, return the current source line. If an integer *n* is given, return it as a line number. If a *routine* is given, return the first line of the routine. If an address is given, return the line containing that address.

For example, the following command returns the line number of the specified address:

```
line {addr 0x22f0}
```

Target

The following commands are applicable to system architectures for which multiple debugging environment targets are available. The commands in this section do not apply to the x86 or x86-64 environments.

connect

```
con[nect]
con[nect] -t target [args]
con[nect] -d path [args]
con[nect] -f file
con[nect] -f file name [args]
```

Without arguments, connect prints the current connection and the list of possible connection targets. Use *-t* to connect to a specific target. Use *-d* to connect to a target specified by *path*. Use *-f* to print a list of possible targets as contained in a *file*, or to connect to a target selected by *name* from the list defined in *file*. Pass configuration arguments to the target as appropriate.

disconnect

```
disc[onnect]
```

Close connection to the current target.

native

```
nati[ve] [command]
```

Without arguments **native** prints the list of available target commands. Given a *command* argument, **native** sends *command* directly to the target.

Miscellaneous

The following commands provide shortcuts, mechanisms for querying, customizing and managing the *PGDBG* environment, and access to operating system features.

alias

```
al[ias] [ name [string] ]
```

Create or print aliases.

- If no arguments are given print all the currently defined aliases.
- If just a *name* is given, print the alias for that name.
- If both a *name* and *string* are given, make *name* an alias for *string*. Subsequently, whenever *name* is encountered it is replaced by *string*.

Although *string* may be an arbitrary string, *name* must not contain any space characters.

For example, the following statement creates an alias for xyz.

```
alias xyz print "x= ",x,"y= ",y,"z= ",z;
cont
```

Now whenever xyz is typed, *PGDBG* responds as though the following command was typed:

```
print "x= ",x,"y= ",y,"z= ",z;
cont
```

directory

```
dir[ectory] [pathname]
```

Add the directory *pathname* to the search path for source files.

If no argument is specified, the currently defined directories are printed. This command assists in finding source code that may have been moved or is otherwise not found by the default *PGDBG* search mechanisms.

For example, the following statement adds the directory `morestuff` to the list of directories to be searched.

```
dir morestuff
```

Now, source files stored in `morestuff` are accessible to *PGDBG*.

If the first character in *pathname* is `~`, then `$HOME` replaces that character.

help

```
help [command]
```

If no argument is specified, print a brief summary of all the commands. If a *command* is specified, print more detailed information about the use of that command.

history

```
history [num]
```

List the most recently executed commands. With the *num* argument, resize the history list to hold *num* commands.

History allows several characters for command substitution:

!! [modifier]	Execute the previous command.
! num [modifier]	Execute command number <i>num</i> .
!-num [modifier]	Execute the command that is <i>num</i> commands from the most current command
!string [modifier]	Execute the most recent command starting with string.
!?string? [modifier]	Execute the most recent command containing string.
^	Command substitution. For example, <code>^old^new^<modifier></code> is equivalent to <code>!:s/old/new/</code> .

There are two possible history modifiers. To substitute the value *new* for the value *old* use:

```
:s/old/new/
```

To print the command without executing it use:

```
:p
```

Use the **pgienv** history command to toggle whether or not the history record number is displayed. The default value is on.

language

```
language
```

Print the name of the language of the current file.

log

```
log filename
```

Keep a log of all commands entered by the user and store it in the named file. This command may be used in conjunction with the **script** command to record and replay debug sessions.

noprint

```
noprint] exp
```

Evaluate the expression but do not print the result.

pgienv

```
pgienv [command]
```

Define the debugger environment. With no arguments, display the debugger settings.

Table 13.1. pgienv Commands

Use this command...	To do this...
help pgienv	Provide help on pgienv
pgienv	Display the debugger settings
pgienv dbx on	Set the debugger to use dbx style commands
pgienv dbx off	Set the debugger to use PGI style commands
pgienv history on	Display the history record number with prompt
pgienv history off	Do not display the history number with prompt
pgienv exe none	Ignore executable's symbolic debug information
pgienv exe symtab	Digest executable's native symbol table (typeless)
pgienv exe demand	Digest executable's symbolic debug information incrementally on command
pgienv exe force	Digest executable's symbolic debug information when executable is loaded
pgienv solibs none	Ignore symbolic debug information from shared libraries
pgienv solibs symtab	Digest native symbol table (typeless) from each shared library
pgienv solibs demand	Digest symbolic debug information from shared libraries incrementally on demand
pgienv solibs force	Digest symbolic debug information from each shared library at load time
pgienv mode serial	Single thread of execution (implicit use of p/t-sets)
pgienv mode thread	Debug multiple threads (condensed p/t-set syntax)
pgienv mode process	Debug multiple processes (condensed p/t-set syntax)
pgienv mode multilevel	Debug multiple processes and multiple threads
pgienv omp [on/off]	Enable/Disable the <i>PGDBG</i> OpenMP event handler. This option is disabled by default. The <i>PGDBG</i> OpenMP event handler, when enabled, sets breakpoints at the beginning and end of each parallel region. Breakpoints are also set at each thread synchronization point. The handler coordinates threads across parallel constructs to maintain source level debugging. This option, when enabled, may significantly slow down program performance. Enabling this option is recommended for localized debugging of a particular parallel region only.
pgienv prompt <name>	Set the command-line prompt to <name>
pgienv promptlen <num>	Set maximum size of p/t-set portion of prompt
pgienv speed <secs>	Set the time in seconds <secs> between trace points

Use this command...	To do this...
<code>pgienv stringlen <num></code>	Set the maximum # of chars printed for <code>^char *'s'</code>
<code>pgienv termwidth <num></code>	Set the character width of the display terminal.
<code>pgienv logfile <name></code>	Close logfile (if any) and open new logfile <code><name></code>
<code>pgienv threadstop sync</code>	When one thread stops, the rest are halted in place
<code>pgienv threadstop async</code>	Threads stop independently (asynchronously)
<code>pgienv procstop sync</code>	When one process stops, the rest are halted in place
<code>pgienv procstop async</code>	Processes stop independently (asynchronously)
<code>pgienv threadstopconfig auto</code>	For each process, debugger sets thread stopping mode to 'sync' in serial regions, and 'async' in parallel regions
<code>pgienv threadstopconfig user</code>	Thread stopping mode is user defined and remains unchanged by the debugger.
<code>pgienv procstopconfig auto</code>	Not currently used.
<code>pgienv procstopconfig user</code>	Process stop mode is user defined and remains unchanged by the debugger.
<code>pgienv threadwait none</code>	Prompt available immediately; do not wait for running threads
<code>pgienv threadwait any</code>	Prompt available when at least one thread stops
<code>pgienv threadwait all</code>	Prompt available only after all threads have stopped
<code>pgienv procwait none</code>	Prompt available immediately; do not wait for running processes
<code>pgienv procwait any</code>	Prompt available when at least a single process stops
<code>pgienv procwait all</code>	Prompt available only after all processes have stopped
<code>pgienv threadwaitconfig auto</code>	For each process, the debugger sets the thread wait mode to 'all' in serial regions and 'none' in parallel regions. (default)
<code>pgienv threadwaitconfig user</code>	The thread wait mode is user-defined and remains unchanged by the debugger.
<code>pgienv mqslib default</code>	Set MPI message queue debug library by inspecting executable.
<code>pgienv mqslib <path></code>	Determine MPI message queue debug library to <code><path></code> .

Use this command...	To do this...
<code>pgienv verbose <bitmask></code>	<p>Choose which debug status messages to report. Accepts an integer valued bit mask of the following values:</p> <ul style="list-style-type: none"> • 0x0 - Disable all messages. • 0x1 - Standard messaging (default). Report status information on current process/thread only. • 0x2 - Thread messaging. Report status information on all threads of (current) processes. • 0x4 - Process messaging. Report status information on all processes. • 0x8 - OpenMP messaging (default). Report OpenMP events. • 0x10 - Parallel messaging (default). Report parallel events. • 0x20 - Symbolic debug information. Report any errors encountered while processing symbolic debug information (e.g. STABS, DWARF). Pass 0x0 to disable all messages. • Pass 0x0 to disable all messages.

repeat

```
rep[eat] [first, last]
rep[eat] [first:last:n]
rep[eat] [num ]
rep[eat] [-num ]
```

Repeat the execution of one or more previous [history](#) list commands. With the *num* (i.e., to num), re-execute the last *num* commands. With the *first* and *last* arguments, re-execute commands number *first* to *last* (optionally *n* times).

script

```
scr[ipt] filename
```

Open the indicated file and execute the contents as though they were entered as commands. Use ~ before the filename in place of the environment variable \$HOME.

setenv

```
setenv name | name value
```

Print the value of the environment variable *name*. With a specified *value*, set *name* to *value*.

shell

```
shell [arg0, arg1,... argn]
```

Fork a shell and give it the indicated arguments . The default shell type is sh or defined by \$Shell. If no arguments are specified, an interactive shell is invoked, and executes until a Ctrl+D is entered.

sleep

```
sleep [time]
```

Pause for one second or *time* seconds.

source

```
source filename
```

Open the indicated file and execute the contents as though they were entered as commands. Use ~ before the filename in place of the environment variable \$HOME.

unalias

```
unalias name
```

Remove the alias definition for *name*, if one exists.

use

```
use [dir]
```

Print the current list of directories or add *dir* to the list of directories to search. The character ~ or environment variable \$HOME can be used interchangeably.

Index

Symbols

- .pdb file, 47
- .pgdbgrc file
 - initialization, 3
- \$EDITOR, 103
- 32-bit Windows, 48
- g option, -gopt option, 47

A

- add
 - directory pathname, 116
- addr
 - command, 114
- address
 - 32-bit float, 114
 - 64-bit double, 113
 - conversion, 114
 - current, 112
 - current program, 112
 - fetch, 114
 - print, 19, 114
 - print integer, 114
 - print short integer, 114
 - read double, 113
 - read integer, 113
 - return, 109
 - set breakpoint, 100
 - short signed integer, 114
 - signed integer, 113, 114
- alias
 - command, 116
 - create, 116
 - print, 116

- remove, 121
- Application
 - terminate target, 17
- arguments
 - intepretation, 25
 - print name and value, 105
 - print names, 104
 - print values, 104, 105
 - target program, 3
- Arguments
 - set, 94
- arrays
 - Fortran, 51
 - large, 51
 - ranges, 51
 - subscripts, 51
- arrive
 - command, 103
- ascii
 - command, 107
 - print, 19, 107
- assembly-level
 - debug with C++, 48
 - debug with Fortran, 48
 - debug with PGDBG GUI, 48
- assign
 - command, 108
- async command, 70
- Attach
 - command, 92
 - running process, 16
- Audience Description, xvii

B

- bin
 - command, 107
- Binary
 - print, 18, 107
- blocks
 - common, 52
 - Fortran, 52
 - lexical, 25
 - statements, 26
- break
 - command, 27, 96, 96
 - conditional, 100

- on variable change, 100, 100
- breaki
 - command, 49, 97, 97
- breakpoints
 - at address, 49
 - clear, 98
 - clear all, 98
 - display all, 97
 - display existing, 97
 - print, 96, 97
 - print current, 96, 97
 - remove, 101, 101
 - remove all, 101, 101
 - remove from address, 101
 - set, 19, 96, 97, 100, 100
 - set at address, 97
 - variable, 100, 100
- breaks
 - command, 97
- breaks command, 97
- Buttons
 - toolbar, 8

C

- C++, 48
 - Instance Methods, 54
 - symbol names, 48
- call
 - command, 54, 108
 - routine, 19
 - stack, 19
- calling conventions, 48
 - Fortran, 48
- Call Stacks
 - display, 11
 - tab, 11
- cancel
 - call command, 108
- catch
 - command, 98
- catch command, 98
- cd
 - command, 103
- change
 - directories, 103
- Class

- command, 110
- Classes
 - command, 111
- clear
 - breakpoints, 98
 - command, 98, 98
- code
 - source locations, 24
- command
 - argument interpretation, 25
 - blocks, 26
 - categories, 91
 - conditional execution, 102
 - constants, 24
 - control, 69
 - events, 27
 - Invoke PGDBG, 3
 - log, 117
 - modes, 23
 - notation, 33
 - PGDBG, 23
 - PGDBG set, 91
 - print use, 116
 - prompt, 73
 - recently executed, 117
 - set, 67
 - Summary Table, 33
 - symbols, 24
 - syntax, 23
- Command
 - tab, 10
- command line
 - PGDBG options, 21, 21, 22
- Command-Line Options
 - syntax, 21
- Commands
 - execute, 19
- common blocks, 52
- Configure
 - stop mode, 70
 - wait mode, 71
- Conformance to Standards, xvii
- constants, 24
- cont command, 49, 92
- Continue
 - cont command, 92
- execution, 20, 20, 20, 20
- control-B, 17
- control-C, 30
 - MPI use, 30
 - thread initialization issues, 30
- control-D, 19
- control-E, 17
- control-F, 17
- control-G, 20
- control-H, 19
- control-L, 18
- control-N, 20
- control-O, 20
- control-P, 18
- control-R, 19
- control-S, 20
- control-U, 19
- conventions
 - calling, 48
 - calling conventions, 48
 - in text, xix
- conversions, 114
- convert
 - address, 114
 - address to line, 25
 - line to address, 25
- Copy, 17
- Copyright
 - display, 20
- core files
 - generation, 55
 - location, 56
 - name, 56
 - set size limit, 56
- cread
 - command, 113
- create
 - aliases, 116
- D**
- Data
 - print type, 19
- Data Menu, 18
- Data menu
 - Addr, 19
 - ascii, 19
- Bin, 18
- Custom, 19, 19
- decimal, 19
- Hex, 19
- Oct, 19
- Print, 18
- Print *, 18
- String, 18
- Type of, 19
- dbx
 - command mode, 23
- Debig
 - threads, 14
- debug
 - assemble-level with C++, 48
 - assemble-level with Fortran, 48
 - assemble-level with PGDBG GUI, 48
 - assembly-level, 47
 - assembly-level commands, 49
 - assembly-level menu options, 48
 - C++, 54
 - command, 92
 - command-line interface, 48
 - Fortran source, 51
 - g option, 47
 - modes, 60
 - MPI, 81
 - multilevel, 89
 - name translation, 48
 - on Microsoft Windows systems, 47
 - on windows, 2
 - parallel, 59, 67
 - PGDBG features, 1, 1
 - using memory addresses, 47
 - using registers, 47
 - with core files, 55
 - with -Munix, 48
- Debug
 - commands, 10
 - events, 10
 - groups, 11, 13, 13, 15
 - memory, 12
 - menu, 19
 - processes, 14

- program status, 16
 - Debug Information Tabs
 - Call Stack tab, 11
 - Command tab, 10
 - Events tab, 10
 - Groups tab, 11, 13, 13, 15
 - Locals tab, 12
 - Memory tab, 12
 - Process(Thread) Grid, 14
 - Status tab, 16
 - Debug Menu, 19
 - Debug menu, 20, 20
 - Call, 19
 - Cont, 20
 - Display Current Location, 19
 - Down, 19
 - Halt, 19
 - Run, 19
 - Set Run Arguments, 19
 - Step, 20
 - Up, 19
 - debug mode
 - multilevel, 89
 - process-only, 61
 - serial, 60
 - threads-only, 61
 - dec
 - command, 107
 - decimal
 - print, 19, 107
 - declaration command, 109
 - declarations
 - print, 111
 - symbol, 109
 - decls
 - command, 111
 - define
 - command list to execute, 98, 99
 - debugger environment, 118
 - do event, 98
 - doi event, 99
 - event, 100
 - instruction-level track event, 101
 - instruction-level watch event, 102
 - read/write watchpoint, 99
 - read watchpoint, 99
 - track event, 101
 - watchpoint, 99
 - defset
 - command, 64, 95
 - delete
 - command, 98
 - event number, 98
 - Detach
 - command, 92
 - end debug session, 16
 - directory
 - add pathname, 116
 - add to search list, 121
 - change, 103
 - command, 116
 - working, 104
 - disable
 - command, 98
 - event number, 98
 - tool tips, 18
 - disasm command, 103
 - disassemble
 - Memory, 103
 - display
 - breakpoints, 97, 97
 - command, 107
 - debugger settings, 118
 - event definition, 100
 - event definitions, 100
 - expressions, 107
 - OpenMP private data, 79
 - program location, 19
 - registers, 49
 - routine scope, 19, 19
 - unique thread ID, 78
 - do
 - command, 27, 98
 - Documentation
 - accessing, xvii
 - location, xvii
 - doi
 - command, 99
 - Down
 - command, 111
 - menu item, 19
 - dread
 - command, 113
 - dump
 - command, 49, 113
 - memory contents, 113
 - MPI message queue, 114
 - Dynamic p/t-set, 63
- ## E
- Edit
 - file, 103
 - menu, 17, 17, 17, 17, 17
 - specify editor, 103
 - edit
 - command, 103
 - file, 103
 - enable
 - command, 99
 - tool tips, 18
 - enter
 - command, 111
 - entry
 - command, 109
 - Environment
 - debugger, 118
 - define, 118
 - Environment variables
 - threadstoconfig, 71
 - Environment variables
 - \$EDITOR, 103
 - HOME, 3
 - name, 120
 - PATH, 3
 - PGI_JAVA, 4
 - set, 120
 - evaluate
 - without printing, 117
 - Events, 26, 96
 - at address, 27
 - at line, 27
 - commands, 27
 - conditional, 27
 - definitions, 100
 - delete, 98
 - disable, 98, 99
 - enable, 99
 - hardware triggered, 99, 99, 99

- in routine, 27
- multiple at same location, 28
- parallel, 74
- print, 98, 99
- program speed, 28
- status, 100
- tab, 10
- track, 101
- tracki, 101
- watch, 101
- watchi, 102

Execute

- command, 19, 102, 103
- conditional, 102, 103
- continue, 20
- rerun command, 94
- run command, 94
- single line, 20, 20

Expressions, 29

- evaluate, 117
- lvalue, 109
- print, 105
- print formatted, 106
- print with pgienv, 106
- rvalue, 109
- type, 110

F

file command, 103

File Menu, 16

Files

- .exe, 47
- .pdb, 47
- .pgdbgrc, 3
- Attach to Target menu, 16
- change, 103
- change source file, 103
- command, 111
- DetachTarget menu, 16
- edit, 103
- execute contents, 120, 121
- Exit menu item, 17
- initialization hierarchy, 3
- menu, 16
- open for debug, 16
- Open Target menu, 16

- source file list, 111
- source list, 111

focus command, 64

Fonts

- change, 18
- default in debugger, 18
- select, 18

fork

- shell, 120

Fortran

- debugging, 51
- symbol names, 48

Fortran 90 modules, 53

fp

- command, 112

frame pointer, 112

- value, 112

Frames

- call stack, 8
- selecting, 8

fread

- command, 114

function

- command, 115

G

Global

- commands, PGDBG, 111

Global commands, 69

grid

- color meaning, 14
- refresh, 18

Groups

- debug, 11

Groupss

- tab, 13, 13, 15

GUI

- PGDBG, 5

H

halt

- command, 83, 93
- control-C, 30
- running processes, 19
- running threads, 19

Hardware

- read/write watchpoint, 99
- read watchpoint, 99, 99
- watchpoint, 99

Help

- About PGDBG menu item, 20
- menu, 20
- on PGDBG commands, 20
- PGDBG menu item, 20

help

- command, 116

Help Menu, 20

Hex

- print, 19

hex

- command, 107

hexadecimal

- print, 107

Hide

- tabs for register groups, 17

history

- command, 117
- modifiers, 117
- repeat command, 120
- resize list, 117

HOME

- environment variable, 3

Host

- defined, 1

HPF, xvii

HPMPI

- debug, 84

hwatchboth command, 99

hwatch command, 27, 99

hwatchread command, 99, 99

hybrid applications

- parallel debugging, 89

I

ID

- process, 93

identifiers

- declarations, 111

if else

- parallel statements, 75

if statement, 26

ignore

- command, 100
 - signals, 100
- ignore command, 100
- Initialization
 - PGDBG, 2
- Initialize
 - PGDBG, 3
 - PGDBG file, 3
- instance
 - methods, 54
- instruction
 - tracing, 101
- integer
 - print as binary, 107
 - print as decimal, 107
 - print as hexadecimal, 107
 - print as octal, 107
- internal
 - procedures, 52, 52
- interrupt
 - control-C, 30
- Invocation
 - PGDBG, 2
- Invoke
 - PGDBG, 3
 - PGDBG for MPI debug, 84
- iread
 - command, 114

J

- Java
 - specify location, 4
 - version selection, 3
- JVM
 - Java and PGDBG, 3, 4

L

- language
 - command, 117
- Lexical blocks, 25
- line command, 115
- lines
 - command, 104
- Lines
 - table, print, 104
- list

- command, 104
 - source lines, 104
- Locals
 - tab, 12
- Locate
 - routine, 17
 - string, 105, 105
- Locate Routine, 17
- location
 - menu item, 19
- Location
 - change, 103
 - current, 103
 - program, 19
- log
 - all commands, 117
 - command, 117, 117
- lread
 - command, 114
- lval
 - command, 109
- lvalue
 - defined, 109

M

- Main routine
 - name, 52
- Manual organization, xviii
- Memory
 - access commands, 113
 - disassemble, 103
 - display addresses, 12
 - dump, 113
 - tab, 12
- menu, 7
- Menu items
 - About PGDBG, 20
 - Addr, 19
 - ASCII, 19
 - Attach to Target, 16
 - bin, 18
 - binary, 18
 - Call, 19
 - Cont, 20
 - Copy, 17
 - Custom, 19

- custom, 19
- Dec, 19
- decimal, 19
- Detach Target, 16
- Display Current Location, 19
- Down, 19
- Exit, 17
- Font, 18
- Halt, 19
- Hex, 19
- hexadecimal, 19
- Locate Routine, 17
- Next, 20
- Oct, 19
- octal, 19
- Open Target, 16
- Paste, 17
- PGDBG Help, 20
- print, 18
- print *, 18
- Refresh, 18
- Registers, 17
- Restore Default Settings, 17
- Revert to Saved Settings, 17
- Run, 19
- Save Settings on Exit, 17
- Search Again, 17
- Search Backward, 17
- Search Forward, 17
- Set Breakpoint, 19
- Set Run Arguments, 19
- Show Tool Tips, 18
- Step, 20
- Step Out, 20
- string, 18
- type, 19
- Up, 19
- Menus
 - assembly-level options, 48
 - context, 7
 - file, 16
 - Help, 20
- Messages
 - MPI, 82
 - MPI queue, 114
 - queues, 82

- status, 8, 73
- Microsoft Windows
 - debug, 47
- Miscellaneous commands, 116
- Modes
 - stop, 70
 - wait, 70
- modules
 - debug access, 53
 - Fortran 90, 53
 - procedures, 53
- MPI
 - debug considerations, 81
 - Debugging, 81
 - debugging options, 22
 - debug multi-process, 81
 - global rank, 85
 - groups, 82
 - listener processes, 82
 - local process, 85
 - message queue dump, 114
 - message queues, 82
 - MPICH-1, 86
 - multi-process debug, 83
 - parallel debug, 81
 - process, local, 85
- MPI_COMM_WORLD, 82
- MPICH
 - support, 86
- MPICH-2
 - debug, 84
- mqdump
 - command, 114
- MSMPI
 - debug, 84
- multilevel
 - debugging, 89
 - error messages, 90
 - mode status, 90
- multilevel debugging, 61
- MVAPICH
 - debug, 84

N

- Names
 - command, 111

- declarations, 112
- identifiers, 111
- print declarations, 112
- registers, 112
- remove alias, 121
- translation, 48
- Next
 - command, 20
- next
 - command, 93
- nexti
 - command, 49, 93
- noprint
 - command, 117
- O**
- oct
 - command, 107
- Octal
 - print, 19, 107
- Open
 - submenu containing Registers tab, 17
- OpenMP, xvii
 - parallel debug, 77
 - private data debug, 78
- Operators
 - @, 29, 52, 111
 - in expressions, 51
 - range, 29
 - scope, 52
 - scope qualifier @, 24

- Optimize
 - code, 2
 - g use, 2
- Options
 - command line, 21, 21, 21, 22
 - g, 2, 47
 - gopt, 47
 - menu, 17
 - Munix, 48
 - O0, 2

P

- p/t-sets, 62
- commands, 64

- create, 66, 66
- current, 62, 63
- define dynamic, 63
- define static, 63
- dynamic vs static, 63
- Editor, 66
- ignore, 67
- modify, 66
- multilevel debug mode, 63
- multiple threads and processes, 74
- notation, 62
- override current, 67
- prefix, 62, 64
- process-only debug mode, 63
- remove, 67
- select, 66
- target, 62
- thread-only debug mode, 62
- undefine, 65
- Parallel
 - debug commands, 67
 - debugging, 59
 - debugging, overview, 59
 - debug hybrid apps, 89
 - debug with MPI, 81
 - events, 74
 - regions, stepi command, 94
 - statements, 75
 - statements, return, 76
- Paste, 17
- PATH
 - environment variable, 3
- pathname
 - add to search path, 116
- pause, 121
- pc
 - command, 112
- PGDBG
 - Assembly-level debugging, 47
 - C++ debugging, 54
 - Command-Line Arguments, 21, 21
 - Command-Line MPI Debugging, 22
 - Command prompt, 73
 - Commands, 23, 91

- Commands Summary, 33
- Conversions, 114
- Debugger, 1, 1
- Debug modes, 60, 89
- Default GUI appearance, 5
- Events, 26, 96
- Expressions, 29
- Fortran arrays, 51
- Fortran Common Blocks, 52
- Fortran debugging, 51
- Graphical user interface, 2, 5
- Initialization, 2
- initialize, 3
- Internal Procedures, 52
- Invocation, 2
- Main Window, 5, 5
- Memory access, 113
- Miscellaneous commands, 116
- Name of main routine, 52
- Operators, 30, 51
- Printing and setting variables, 105
- Process commands, 67
- Process control commands, 92
- Program locations, 103
- Register access, 112
- Register symbols, 24
- Scope, 110
- Scope rules, 24
- Source code locations, 24
- start session, 3
- Statements, 26
- Status messages, 73
- Symbols and expressions, 108
- Thread commands, 67
- Wait modes, 71
- PGDBG Commands
 - addr, 114
 - alias, 116
 - arrive, 103
 - ascii, 107
 - assign, 108
 - attach, 34, 92
 - bin, 107
 - break, 96
 - break command, 96
 - breaki, 49, 97
 - breaki command, 97
 - breaks, 97
 - breaks command, 97
 - call, 108
 - catch, 98
 - catch command, 98
 - cd, 103
 - class, 110
 - classes, 111
 - clear, 98
 - clear command, 98
 - cont, 49, 92
 - cread, 113
 - debug, 92
 - dec, 107
 - declaration, 109
 - decls, 111
 - defset, 64
 - defset command, 95
 - delete, 98
 - detach, 92
 - directory, 116
 - disable, 98
 - disasm, 103
 - display, 107
 - do, 98
 - doi, 99
 - down, 111
 - dread, 113
 - dump, 49, 113
 - edit, 103
 - enable, 99
 - enter, 111
 - entry, 109
 - file, 103
 - files, 111
 - focus, 64
 - focus command, 95
 - fp, 112
 - fread, 114
 - function, 115
 - global, 111
 - halt, 93
 - help, 116
 - hex, 107
 - history, 117
 - hwatch, 99
 - hwatchboth, 99
 - hwatchread, 99, 99
 - ignore, 100
 - iread, 114
 - language, 117
 - line, 115
 - lines, 104
 - list, 104
 - log, 117, 117
 - lread, 114
 - lval, 109
 - mqdump, 114
 - names, 111
 - next, 93
 - nexti, 49, 93
 - noprint, 117
 - oct, 107
 - pc, 112
 - pgienv, 118
 - print, 49, 105
 - printf, 106
 - proc, 93
 - procs, 93
 - pwd, 104
 - quit, 93
 - regs, 49, 112
 - repeat, 120, 120
 - rerun, 94
 - retaddr, 112
 - run, 49, 94
 - rval, 109
 - scope, 111
 - script, 120, 120
 - search backward, 105
 - search forward, 105
 - set, 110
 - setargs, 94
 - setenv, 120, 120
 - shell, 120, 120
 - sizeof, 110
 - sleep, 121, 121
 - source, 121, 121
 - sp, 112
 - sread, 114
 - stackdump, 49, 104

- stacktrace, 49, 104
- status, 100
- step, 94
- stepi, 49, 94
- stepout, 94
- stop, 100
- stopi, 100
- string, 107
- sync command, 94, 95
- synci command, 94, 95
- thread command, 95
- threads command, 95
- trace, 100
- tracei, 101
- track, 101
- tracki, 101
- type, 110
- unalias, 121, 121
- unbreak, 101
- unbreaki, 101
- undefset, 65
- undefset command, 95
- undisplay, 108
- up, 112
- use, 121, 121
- viewset
 - viewset command, 64
- viewset command, 96
- wait command, 95
- watch, 101
- watchi, 102
- when, 102
- wheni, 103
- where, 105
- whereis, 112
- which, 112
- whichsets, 64
- whichsets command, 96
- PGDBG control commands, 69
- PGDBG GUI
 - assembly-level debugging, 48
- PGDBG Signals, 57
- pgi
 - command mode, 23
- PGI_JAVA
 - environment variable, 4
- pgienv, 100, 101, 118
 - command, 118
- pgienv command arguments, 118
- Print
 - active threads, 95
 - address, 19, 114
 - aliases, 116
 - all registers, 112
 - arg values and names, 104
 - ascii, 19, 107
 - binary, 18, 107
 - breakpoints, 96, 96, 97, 97
 - command, 49, 105
 - command info, 116
 - command summary, 116
 - current, 103
 - current file, 103
 - current location, 103
 - current working directory, 104
 - data type, 19
 - data value, 18, 18
 - dec, 19
 - decimal, 107
 - defined aliases, 116
 - defined directories, 116
 - directory list, 121
 - environment variable name, 120
 - events, 98, 99
 - expressions, 105, 107, 107
 - formatted stack dump, 104
 - formatted expressions, 106
 - formatted register names, 112
 - hex, 19
 - hexadecimal, 107
 - identifier declarations, 111
 - identifier names, 111
 - ignored signals, 100, 100
 - integer address, 114
 - language name, 117
 - lines table, 104
 - list of signals ignored, 100
 - location, 103
 - name declarations, 112
 - noprint, 117
 - octal, 19, 107
 - procs command, 93
- register info, 112
- register value, 49
- scope qualification, 112
- scope qualified symbol name, 112
- short integer address, 114
- signals, 98, 98
- stack dump, 104
- stacktrace, 104, 104, 105
- string, 18
- strings, 107
- symbol declaration, 109
- values, 105
- values as change, 101
- watched event values, 101
- printf command, 106
- proc
 - command, 93
- procedures
 - Fortran 90 modules, 53
 - internal, 52, 52
- process
 - assign name, 95
 - IDs, 61
 - proc command, 93
 - process/thread set, 62
 - process and thread control, 69
 - process level commands, 67
 - process-only debugging, 61
 - stop mode, 70
 - wait mode, 71
- Process/Thread
 - element color, 14
- Processes
 - MPI rank, 60
 - parallel debugging, 59
 - print, 93
 - Process-parallel debugging, 85
 - Process-thread sets, 95
- Process grid tab, 14
- process set
 - list members, 96
 - membership, 96
 - remove, 95
 - set target, 95
- procwait, 71
- program location

- arrive, 19
- sync command, 94, 95
- synci command, 94, 95
- thread command, 95
- Programs
 - status, 16
- prologue code, 109
- prompt
 - return, 95
- pwd
 - command, 104
- Q**
- quit
 - command, 93
- R**
- read
 - watchpoint, 99, 99
- record session, 117, 117
- Refresh
 - Process/Thread Grid, 18
 - windows, 18
- Registers
 - access, 112
 - formatted names, 112
 - print info, 112
 - symbols, 49
 - view mmenu, 17
- register symbols, 24
- regs
 - command, 49, 112
- Related Publications, xx
- remove
 - alias definition, 121
 - all expressions, 108
 - breakpoint, 101, 101
 - expression from display list, 108
- Repeat
 - command, 120
 - search, 17
- replay debug session, 117, 117
- Rerun
 - target application, 19
- rerun command, 94
- Restore

- default settings, 17
- retaddr
 - command, 112
- return
 - address, 109
 - lvalue, 109
 - routine, 115
 - rvalue, 109
 - size of var type name, 110
 - statement, 76
 - type of expression, 110
- Revert
 - saved settings, 17
- Routines
 - breakpoint, 19
 - call, 108
 - clear breakpoints, 98
 - disassemble, 103
 - display in source panel, 17
 - edit, 103
 - enter scope, 112
 - first line, 115
 - instruction tracing, 101
 - list source code, 104
 - locate, 17
 - main name, 52
 - print lines table, 104
 - print name, 104, 105, 105
 - request, 19
 - return, 115
 - scope, 19, 19, 111
 - set breakpoint, 100
 - size of, 110
 - source line tracing, 100
 - step, 94
 - stepi, 94
 - step into, 20
 - stepout command, 94
 - step out of, 20
 - step over, 20
 - symbol, 115
- rsh communication, 83
- Run
 - arguments, 19
 - target application, 19
- run command, 49, 94

- Runtime
 - arguments, 19
- rval
 - command, 109
- rvalue
 - defined, 109
- S**
- Sales
 - contact information, 20
- Save
 - GUI settings, 17
- scope, 110
 - change, 103
 - class, 110
 - classes, 111
 - command, 111
 - current, 24
 - enter, 112
 - global, 111, 111
 - identifiers defined, 111
 - operator, 52
 - print identifier names, 111
 - print symbol name qualification, 112
 - qualifier operator, 24
 - routine, 19, 19, 111
 - rules, 24
 - search, 24, 111
 - set, 111
 - start, 25
 - up one level, 112
- script command, 120
- Search
 - backward, 105
 - command, 105
 - for strings, 105, 105
 - forward, 105
 - keyword, 17, 17
 - last keyword, 17
 - path, 116
 - scope, 24, 111
- Search Again, 17
- Search Backward, 17
 - command, 105
- Search Forward, 17

- Search Forward command, 105
- Select
 - frame, 8
- Sessions
 - end debug, 16, 17
 - PGDBG, 3
 - terminate, 93
- Set
 - breakpoints, 19, 19
 - command, 110
 - search scope, 111
 - variable value, 108, 110
- setargs
 - command, 94
- setenv command, 120
- Settings
 - display for debugger, 118
 - restore, 17, 17
 - Restore Default Settings menu item, 17
 - revert, 17
 - Revert to Saved Settings menu item, 17
 - saved, 17
 - save GUI state, 17
 - Save Settings on Exit menu item, 17
- shell
 - command, 120
 - invoke, 120
- Show
 - tabs for register groups, 17
 - tool tips, 18
- signals, 56, 57
 - ignore, 100
 - ignored, 100, 100
 - interrupt, 98
 - Linux Libraries, 57
 - list, 98
 - PGDBG, 57
 - print, 98
 - Print, 100
 - SIGPROF, 57
- size
 - variable, 110
- sizeof
 - command, 110
- sleep command, 121
- Source
 - current, 115
 - line conversion, 115
 - list lines, 104
- source code
 - locations, 24
- source command, 121
- source file
 - change, 103
- source line
 - conversion, 115
- source line tracing, 100
- Source Window, 6, 8, 9
 - Call frame, 8
 - Context Menu, 7
- sp
 - command, 112
- sread
 - command, 114
- SSE Register Symbols, 49
- ssh communication, 83
- stack
 - display frames, 104
 - frame, 104
 - frames, display, 104
 - frames, display hex dump, 104
 - pointer, 112
 - pointer value, 112
 - print dump, 104
 - print stacktrace, 105
 - print trace, 104, 105
- stackdump
 - command, 49, 104
- stack frames
 - display, 105
- stacktrace
 - command, 49, 104
- Start
 - debug session, 16
 - PGDBG debugger, 1, 1
- statements
 - block, 26
 - compound, 75
 - constructs, 26
 - execution order, 75
 - if, 26
 - parallel, 75
 - parallel if else, 75
 - parallel while, 76
 - PGDBG, 26
 - return, 76
 - simple, 26
 - while, 26
- static p/t-set, 63
- Status
 - message area, 8
 - program, 16
- status
 - command, 100
 - events, 100
 - messages, 73
- status message area, 8
- Statuss
 - tab, 16
- Step
 - into routines, 20
 - out of routine, 20
 - over routines, 20
- step command, 94
- stepi
 - command, 49, 94
- Step into
 - called routines, 20
- stepout
 - command, 94
- Step Out, 20
- Step over
 - called routines, 20
- Stop
 - after return to caller, 94
 - at value change, 101, 102
 - configure mode, 70
 - execution, 94
 - modes, 70
- stop
 - command, 100
- stopi
 - command, 100
- string
 - command, 107

- Strings
 - locate, 105, 105
 - print, 18, 107
 - subroutines
 - nested, 52
 - Support
 - information, 20
 - symbol
 - declarations, 109
 - name qualification, 112
 - symbol names
 - C++, 48
 - Fortran, 48
 - Symbols, 24
 - global scope, 111
 - MAIN_, 48
 - print declaration, 109
 - register, 24
 - routine, 115
 - scope-qualified name, 112
 - search scope, 111
 - SSE register, 49
 - Symbols and Expressions, 108
 - sync
 - command, 94
 - sync command, 70, 82
 - synci command, 94, 95
- T**
- Tables
 - routine lines, 104
 - Tabs
 - Call Stack, 11
 - Command, 10
 - Events, 10
 - Groups, 11, 13, 13, 15
 - Locals, 12
 - Memory, 12
 - Process Grid, 14
 - Status, 16
 - Thread Grid, 14
 - Target
 - Application
 - run or rerun, 19
 - arguments to, 3
 - runtime arguments, 19
 - usage, 3
 - Terminology
 - PGDBG, 1
 - Terms, 1
 - text mode debug, 84
 - Thread Grid tab, 14
 - Thread level commands, 68
 - Threads
 - assign name, 95
 - command, 95
 - commands, 67
 - grouping, 60
 - IDs in multilevel debug mode, 89
 - location, 95
 - logical id, 95
 - naming, 60
 - naming convention, 59
 - naming scheme, 89
 - OpenMP, 59
 - parallel debugging, 59
 - process/thread set, 62
 - stop mode, 70
 - threads-only debugging, 61
 - wait mode, 71
 - threads
 - command, 95
 - Threads, configure, 71
 - threads command, 78
 - thread set
 - list members, 96
 - membership, 96
 - remove, 95
 - set target, 95
 - threadstoconfig environment
 - variable, 71
 - threadwait, 71
 - Toolbar
 - buttons, 8
 - trace
 - command, 27, 100
 - conditional, 100
 - source, 100
 - subprogram routines, 100
 - tracei
 - command, 101
 - conditional, 101
 - source, 101
 - subprogram routines, 101
 - track
 - command, 27, 101
 - event, 101
 - tracki
 - command, 101
 - event, 101
 - type
 - command, 110
- U**
- unalias command, 121
 - unbreak command, 101
 - unbreaki
 - command, 101
 - undefset command, 65, 95
 - undisplay
 - command, 108
 - up
 - command, 112
 - menu item, 19
 - use command, 121
 - Utilities
 - help, 20
- V**
- Variables
 - breakpoint, 100, 100
 - display local, 12
 - instruction tracing, 101
 - set value, 108, 110
 - trace changes, 100, 101
 - Versions
 - display, 20
 - View
 - Font menu item, 18
 - Refresh, 18
 - Registers menu item, 17
 - Show Tool Tips menu item, 18
 - View Menu, 17
 - viewset command, 96
- W**
- wait command, 72, 95
 - wait mode, 70, 71

- process, 71
- thread, 71
- watch
 - command, 101
 - event, 101
- watch command, 27
- watchi
 - command, 102
 - event, 102
- Watchpoints
 - define, 99, 99
 - hardware, 99
 - hardware read, 99, 99
- when command, 102
- wheni command, 103
- where
 - command, 105
- whereis
 - command, 112
- which command, 112
- whichsets command, 64, 96
- while
 - parallel statements, 76
- while statement, 26
- Window
 - source, 6, 9
- Windows
 - build for debug, 2
 - PGDBG main, 5, 5
 - refresh, 18
- working directory
 - print, 104
- write
 - watchpoint, 99