

# PGPROF<sup>®</sup> Profiler Guide

Parallel Profiling for Scientists and Engineers

Release 2013



**The Portland Group<sup>®</sup>**

**PGI<sup>®</sup> Profiler Guide**  
Copyright © 2013 NVIDIA Corporation  
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  
Fourth Printing: Release 11.4, April 2011  
Fifth Printing: Release 12.1, January 2012  
Sixth Printing: Release 12.9, September 2012  
Seventh Printing: Release 13.1, January 2013  
Eighth Printing: Release 13.3, March 2013  
Ninth Printing: Release 13.5, May 2013  
Tenth Printing: Release 13.7, July 2013  
Eleventh Printing: Release 13.8, August 2013

Technical support: <http://www.pgroup.com/support/>  
Sales: [sales@pgroup.com](mailto:sales@pgroup.com)  
Web: <http://www.pgroup.com>

# Contents

<b>Preface</b> .....	xiii
Intended Audience .....	xiii
Supplementary Documentation .....	xiii
Compatibility and Conformance to Standards .....	xiii
Organization .....	xiv
Conventions .....	xv
Terminology .....	xvi
Related Publications .....	xvi
System Requirements .....	xvi
 <b>1. Getting Started</b> .....	 1
Basic Profiling .....	1
Methods of Collecting Performance Data .....	2
Instrumentation-based Profiling .....	2
Sample-based Profiling .....	3
Choose Profile Method .....	4
Collect Performance Data .....	4
Profiling Output File .....	4
Using System Environment Variables .....	4
Profiling with Hardware Event Counters .....	4
Profiler Invocation and Initialization .....	5
Application Tuning .....	5
Troubleshooting .....	5
Prerequisite: Java Virtual Machine .....	6
Slow Network .....	6
 <b>2. Using PGPROF</b> .....	 7
PGPROF Tabs and Icons Overview .....	8
Profile Navigation .....	9
HotSpot Navigation .....	13
Sorting Profile Data .....	13
Compiler Feedback .....	14
Special Feedback Messages .....	15

Profiling Parallel Programs .....	15
Profiling Multi-threaded Programs .....	15
Profiling MPI Programs .....	17
Scalability Comparison .....	19
Profiling Resource Utilization with Hardware Event Counters .....	20
Profiling with Hardware Event Counters (Linux Only) .....	21
Analyzing Event Counter Profiles .....	21
Profiling GPU Programs .....	22
Profiling PGI Accelerator Model Programs .....	22
Profiling CUDA Fortran Programs .....	26
<b>3. Compiler Options for Profiling .....</b>	<b>29</b>
-Mprof Syntax .....	29
Profiling Compilation Options .....	29
<b>4. Profiling with Open MPI .....</b>	<b>31</b>
Overview .....	31
Compiler Wrapper Data Files .....	31
Configure Open MPI for PGI Profiling .....	32
Modified Compiler Wrapper Data File Sample .....	33
Using PGI MPI Profiling for Open MPI .....	34
<b>5. Command Line Options .....</b>	<b>35</b>
Command Line Option Descriptions .....	35
Profiler Invocation and Startup .....	36
<b>6. Environment Variables .....</b>	<b>39</b>
System Environment Variables .....	39
<b>7. Data and Precision .....</b>	<b>41</b>
Measuring Time .....	41
Profile Data .....	41
Caveats (Precision of Profiling Results) .....	42
Accuracy of Performance Data .....	42
Clock Granularity .....	43
Source Code Correlation .....	43
<b>8. PGPROF Reference .....</b>	<b>45</b>
PGPROF User Interface Overview .....	45
PGPROF Menus .....	46
File Menu .....	46
Edit Menu .....	47
View Menu .....	48
Sort Menu .....	49
Help Menu .....	49

PGPROF Toolbar .....	50
PGPROF Statistics Table .....	51
Performance Data Views .....	51
Source Code Line Numbering .....	52
PGPROF Focus Panel .....	52
Parallelism tab .....	53
Histogram tab .....	53
Compiler Feedback tab .....	53
System Configuration tab .....	54
Accelerator Performance tab .....	54
<b>9. Command Line Interface</b> .....	<b>59</b>
Command Description Syntax .....	59
PGPROF Command Summary .....	59
Command Reference .....	60
<b>10. pgcollect Reference</b> .....	<b>65</b>
pgcollect Overview .....	65
Invoke pgcollect .....	66
Build for pgcollect .....	66
General Options .....	66
Time-Based Profiling .....	66
Time-Based Profiling Options .....	66
Event-Based Profiling .....	67
Root Privileges Requirement .....	67
Interrupted Profile Runs .....	67
Event-based Profiling Options .....	68
Defining Custom Event Specifications .....	68
PGI Accelerator Model and CUDA Fortran Profiling .....	69
Accelerator Model Profiling .....	69
CUDA Fortran Program Profiling .....	70
Performance Tip .....	70
<b>Index</b> .....	<b>71</b>



# Figures

2.1. PGPROF Overview .....	8
2.2. PGPROF Initial View .....	10
2.3. Source Code View .....	11
2.4. Assembly Level View .....	12
2.5. View Navigation Buttons .....	12
2.6. HotSpot Navigation Controls .....	13
2.7. Sort Example .....	14
2.8. Multi-Threaded Program Example .....	16
2.9. Sample MPI Profile .....	18
2.10. Sample Scalability Comparison .....	20
2.11. Profile with Hardware Event Counter .....	22
2.12. Accelerator Performance Data for Routine-Level Profiling Example .....	24
2.13. Source-Level Profiling for an Accelerator Region .....	25
2.14. Source-Level Profiling for an Accelerator Kernel .....	26
2.15. CUDA Program Profile .....	28
8.1. PGPROF User Interface .....	46
8.2. PGPROF Toolbar .....	50
8.3. Focus Panel Tabs .....	53
8.4. Accelerator Performance tab of Focus Panel .....	55
8.5. CUDA Program Profile .....	57





# Tables

2.1. PGPROF Icon Summary .....	9
9.1. PGPROF Commands .....	59



# Examples

10.1. Custom Event Example 1 .....	69
10.2. Custom Event Example 2 .....	69



# Preface

This guide describes how to use the *PGPROF* profiler to tune 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 the PGI profiling tools, as well as detailed reference information on commands and graphical interfaces.

## 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.

## Supplementary Documentation

See <http://www.pgroup.com/docs.htm> for the *PGPROF* documentation updates. Documentation delivered with *PGPROF* should be accessible on an installed system by accessing [docs/index.htm](#) in the PGI installation directory. Typically the value of the environment variable PGI is set to the PGI installation directory. See <http://www.pgroup.com/faq/index.htm> for frequently asked *PGPROF* 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, 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, PGF90, PGCC ANSI C, and PGCPP support parallelization extensions based on the OpenMP defacto 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.

*PGPROF* permits profiling of serial and parallel (multi-threaded, OpenMP and/or MPI) programs compiled with PGI compilers.

For further information, refer to the following:

- American National Standard Programming Language FORTRAN, ANSI X3. -1978 (1978).
- 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).
- American National Standard Programming Language C, ANSI X3.159-1989.
- ISO/IEC 9899:1999, Information technology – Programming Languages – C, Geneva, 1999 (C99).
- HPDF Standard (High Performance Debugging Forum) <http://www.ptools.org/hpdf/draft/intro.html>
- Fortran 2003 Standard  
(High Performance Debugging Forum) <http://www.ptools.org/hpdf/draft/intro.html>

## Organization

*The PGPROF Profiler User's Guide* contains ten chapters that describe the *PGPROF* Profiler, a tool for analyzing the performance characteristics of C, C++, F77, and F95 programs.

### Chapter 1, “*Getting Started*”

contains information on how to start using the profiler, including a description of the profiling process, information specific to certain how to profile MPI and OpenMP programs and how to profile with hardware event counters.

### Chapter 2, “*Using PGPROF*”

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

### Chapter 3, “*Compiler Options for Profiling*”

describes the compiler options available for profiling and how they are interpreted.

### Chapter 4, “*Profiling with Open MPI*”

describes how to do MPI profiling with Open MPI and PGI compilers.

### Chapter 5, “*Command Line Options*”

describes the *PGPROF* command-line options used for profiling and provides sample invocations and startup commands.

### Chapter 6, “*Environment Variables*”

contains information on environment variables that you can set to control the way profiling is performed in *PGPROF*.

**Chapter 7, “*Data and Precision*”**

contains descriptions of the profiling mechanisms that measure time, how statistics are collected, and the precision of the profiling results.

**Chapter 8, “*PGPROF Reference*”**

provides reference information about the *PGPROF* graphical user interface, including information about the menus, the toolbars, and the subwindows.

**Chapter 9, “*Command Line Interface*”**

provides information about the *PGPROF* profiler command line interface language, providing both a summary table and details about the commands. The table includes the command name, the arguments for the command, and a brief description of the command - all separated by area of use.

**Chapter 10, “*pgcollect Reference*”**

provides reference information about the **pgcollect** command. It describes the *PGPROF* command line options and how to use them to configure and control collection of application performance data.

## 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 test 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 and Windows 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](http://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, 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 User's Guide, PGI Release Notes, FAQ, Tutorials, <http://www.pgroup.com/>
- MPI-CH <http://www.unix.mcs.anl.gov/MPI/mpich/>
- OpenMP <http://www.openmp.org/>

## System Requirements

- Linux or Windows (See <http://www.pgroup.com/faq/install.htm> for supported releases)
- Intel x86 (and compatible), AMD Athlon or AMD64, or Intel 64 or Core2 processor



# Chapter 1. Getting Started

This chapter describes the *PGPROF* profiler. *PGPROF* provides a way to visualize and diagnose the performance of the components of your program. Using tables and graphs, *PGPROF* associates execution time with the source code and instructions of your program, allowing you to see where and how execution time is spent. Through resource utilization data (processor counters) and compiler feedback information, *PGPROF* also provides features to help you understand why certain parts of your program have high execution times.

You can also use the *PGPROF* profiler to profile parallel programs, including multiprocess MPI programs, multi-threaded programs such as OpenMP programs, or a combination of both. *PGPROF* provides views of the performance data for analysis of MPI communication, multiprocess and multi-thread load balancing, and scalability.

Using the Common Compiler Feedback Format (CCFF), PGI compilers save information about how your program was optimized, or why a particular optimization was not made. *PGPROF* can extract this information and associate it with source code and other performance data, allowing you to view all of this information simultaneously. *PGPROF* also supports a *feedbackonly* mode, which allows you to browse Compiler Feedback in the absence of a performance profile.

Each performance profile depends on the resources of the system where it is run. *PGPROF* provides a summary of the processor(s) and operating system(s) used by the application during any given performance experiment.

## Basic Profiling

Performance profiling can be considered a two-stage process.

- In the first stage, you collect performance data when your application runs using typical input.
- In the second stage, you analyze the performance data using *PGPROF*.

There are a variety of ways to collect performance data from your application. For basic execution-time profiling, we recommend that you use the **pgcollect** tool, which has several attributes that make it a good choice:

- You don't have to recompile or relink your application.
- Data collection overhead is low.

- It is simple to use.
- It supports multi-threaded programs.
- It supports shared objects, DLLs, and dynamic libraries.

To profile your application named `myprog`, you execute the following commands:

```
$ pgcollect myprog
$ pgprof -exe myprog
```

The information available to you when you analyze your application's performance can be significantly enhanced if you compile and link your program using the `-Minfo=ccff` option. This option saves information about the compilation of your program, compiler feedback, for use by *PGPROF*. For more information on compiler feedback, refer to [“Compiler Feedback,” on page 14](#).

For a more complete analysis, our command execution might look similar to this:

```
$ pgfortran -fast -Minfo=ccff -o myprog myprog.90
$ pgcollect myprog
$ pgprof -exe myprog
```

## Methods of Collecting Performance Data

PGI provides a number of methods for collecting performance data in addition to the basic **pgcollect** method described in the previous section. Some of these have advantages or capabilities not found in the basic **pgcollect** method. We divide these methods into two categories: instrumentation-based profiling and sample-based profiling.

### Instrumentation-based Profiling

Instrumentation-based profiling is one way to measure time spent executing the functions or source lines of your program. The compiler inserts timer calls at key points in your program and does the bookkeeping necessary to track the execution time and execution counts for routines and source lines. This method is available on all platforms on which PGI compilers are supported.

Instrumentation-based profiling:

- Provides exact call counts.
- Provides exact line/block execution counts.
- Reports time attributable to only the code in a routine.
- Reports time attributable to the code in a routine and all the routines it called.

This method requires that you recompile and relink your program using one of these compiler options:

- Use `-Mprof=func` for routine-level profiling.

Routine-level profiling can be useful in identifying which portions of code to analyze with line-level profiling.

- Use `-Mprof=lines` for source line-level profiling.

The overhead of using line-level profiling can be high, so it is more suited for fine-grained analysis of small pieces of code, rather than for analysis of large, long-running applications.

## Sample-based Profiling

Sample-based profiling uses statistical methods to determine the execution time and resource utilization of the routines, source lines, and assembly instructions of the program. Sample-based profiling is less intrusive than instrumentation-based profiling, so profiling runs take much less time. Further, in some cases it is not necessary to rebuild the program.

### Note

The basic **pgcollect** method described earlier in “[Basic Profiling](#)” is a time-based sampling method. **pgcollect** also supports event-based profiling on linux86-64.

The following sections describe both time-based and event-based sampling. For information on the differences in how instrumentation- and sample- based profiling measure time, refer to “[Measuring Time,](#)” on page 41.

### Time-based Sampling

With time-based sampling the program's current instruction address (program counter) is read, and tracked, at statistically significant intervals. Instruction addresses where a lot of time is spent during execution are read numerous times. The profiler can map these addresses to source lines and/or functions in your program, providing an easy way to navigate from the function where the most time is spent, to the line, to the assembly instruction.

You can build your program using the `-Mprof=time` compiler option for time-based sampling of single-threaded Linux programs. When using `-Mprof=time`, you are required only to re-link your program. However, unless you compile with `-Minfo=ccff`, compiler feedback will not be available.

As described previously in “[Basic Profiling](#)”, we recommend using **pgcollect** for time-based profiling.

### Event-based Sampling

As well as reading the program's instruction address, event-based sampling uses various methods to read and track the values of selected hardware counters. These counters track processor events such as data cache misses and floating point operations. You can use this information to help determine not just that time is being spent in a particular block of code, but why so much time is spent there. If there is bottleneck related to a particular resource, such as the level two data cache, these counters can help you discover where the bottleneck is occurring.

Event-based sampling requires that a performance tool named *OProfile* be co-installed with the PGI software on the Linux system.

OProfile is a performance profiling utility for Linux systems. It runs in the background collecting information at a low overhead and providing profiles of code based on processor hardware events. When installed, **pgcollect** collects this type of performance data for analysis with *PGPROF*. For more information on OProfile, see <http://oprofile.sourceforge.net/>.

Run your program using the **pgcollect** command for event-based sampling with OProfile.

## Note

---

MPI profiling is not available with **pgcollect** profiling.

## Choose Profile Method

Use the following guidelines to decide which performance data collection method to use:

- A good starting point for any performance analysis is to use time-based sampling with **pgcollect**, as described in [“Basic Profiling,” on page 1](#).
- If you are profiling an MPI application on Linux, build your application using `-Mprof=time,<mpi>`, where `<mpi>` is the supported MPI distribution that you are using, for example, `mpich1`.
- If your MPI application also uses OpenMP or multiple threads per process and you want to determine where the majority of time is spent, build with `-Mprof=func,<mpi>`. Then build that portion of the program with `-Mprof=lines,<mpi>` to isolate the performance problem.
- If you want exact execution counts, build with `-Mprof=func` or `-Mprof=lines`.
- On Linux86-64 platforms on which OProfile is installed, once you have collected a time-based profile using either instrumentation- or sample-based profiling, consider further examining the resource utilization of those portions of code where the most time is spent. You do this with event-based sampling, using the **pgcollect** command with event-based sampling options as described in [Chapter 10, “pgcollect Reference”](#).

## Collect Performance Data

To obtain the performance data required for *PGPROF*, you must run your program.

- If you use any method other than the **pgcollect** command to collect data, run your program normally using a representative input data set.
- If you use the **pgcollect** command to collect data, refer to [“Basic Profiling,” on page 1](#) for information of how to execute a profiling run of your program. For specific details on **pgcollect**, refer to [Chapter 10, “pgcollect Reference”](#).

## Profiling Output File

In all profiling methods, once the program's profiling run is complete, a file named `pgprof.out` is written to the program's working directory. This file contains the performance data used by *PGPROF* to analyze the program's performance.

## Using System Environment Variables

You can use system environment variables to change the way profiling is performed. For more information on these variables, refer to [Chapter 6, “Environment Variables”](#).

## Profiling with Hardware Event Counters

You can also profile using hardware event counters. For more specific information on this type of profiling, refer to [“Profiling Resource Utilization with Hardware Event Counters,” on page 20](#).

## Profiler Invocation and Initialization

*PGPROF* is invoked as follows:

```
% pgprof.exe [options] [datafile]
```

If invoked without any options or arguments, *PGPROF* attempts to open a data file named `pgprof.out`, and assumes that application source files are in the current directory. The program executable name, specified when the program was run, is usually stored in the profile data file. If all program-related activity occurs in a single directory, *PGPROF* needs no options.

Probably the most common method to invoke the profiler is this:

```
% pgprof -exe <execname>
```

When you use this command to launch *PGPROF*:

- If a `pgprof.out` file exists in the current directory, *PGPROF* opens it and uses `<execname>` to display the profile data.
- If no `pgprof.out` file exists in the current directory, no profile data is displayed. Further, when the user selects the menu `File | Open Profile...`, the Text Field for Executable is set with `<execname>` in the dialog.

For information on all available profiler options and how they are interpreted, refer to [Chapter 3, “Compiler Options for Profiling”](#). For information on the command line options for the Profiler, refer to [Chapter 5, “Command Line Options”](#). For sample launch commands; refer to “[Profiler Invocation and Startup](#),” on page 36.

## Application Tuning

So how do you make your program faster? The process of tuning your program ranges from simple to complex.

- In the simple case, you may be able to easily tune the application and improve performance dramatically simply by adding a compiler option when you build. The Compiler Feedback and System Configuration tabs in the *PGPROF* user interface contain information that can help identify these situations.
- In a slightly more challenging scenario, you may need to restructure part of your code to allow the compiler to optimize it more effectively. For instance, the Compiler Feedback for a given loop may provide a hint to remove a call from the loop. If the call can be moved out of the loop or inlined, the loop might be vectorized by the next compile.
- More difficult cases involve memory alignment and algorithm restructuring. These issues are beyond the scope of this manual.

## Troubleshooting

If you are having trouble during invocation or the initialization process, use the following sections for tips on what might be causing your problem.

## Prerequisite: Java Virtual Machine

The PGPROF GUI depends on the Java Virtual Machine (JVM) which is part of the Java Runtime Environment (JRE). PGPROF requires that the JRE be a specific minimum version or above.

### Linux

When PGI software is installed on Linux, the version of Java required by the profiler is also installed. PGPROF uses this version of Java by default. You can override this behavior in two ways: set your PATH to include a different version of Java; or, set the PGI\_JAVA environment variable to the full path of the Java executable. The following example uses a bash command to set PGI\_JAVA:

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

### OSX

The PGI profiler on OSX uses the version of Java installed by Apple's OSX software updater. If your system is configured such that Java is not installed in the default location, you need to set your PATH to include the Java bin directory or use the PGI\_JAVA environment variable to specify the full path to the java executable.

### Windows

If an appropriately-versioned JRE is not already on your system, the PGI software installation process installs it for you. The PGI command shell and Start menu links are automatically configured to use the JRE. If you choose to skip the JRE-installation step or want to use a different version of Java to run the profiler, then set your PATH to include the Java bin directory or use the PGI\_JAVA environment variable to specify the full path to the java executable.

## Slow Network

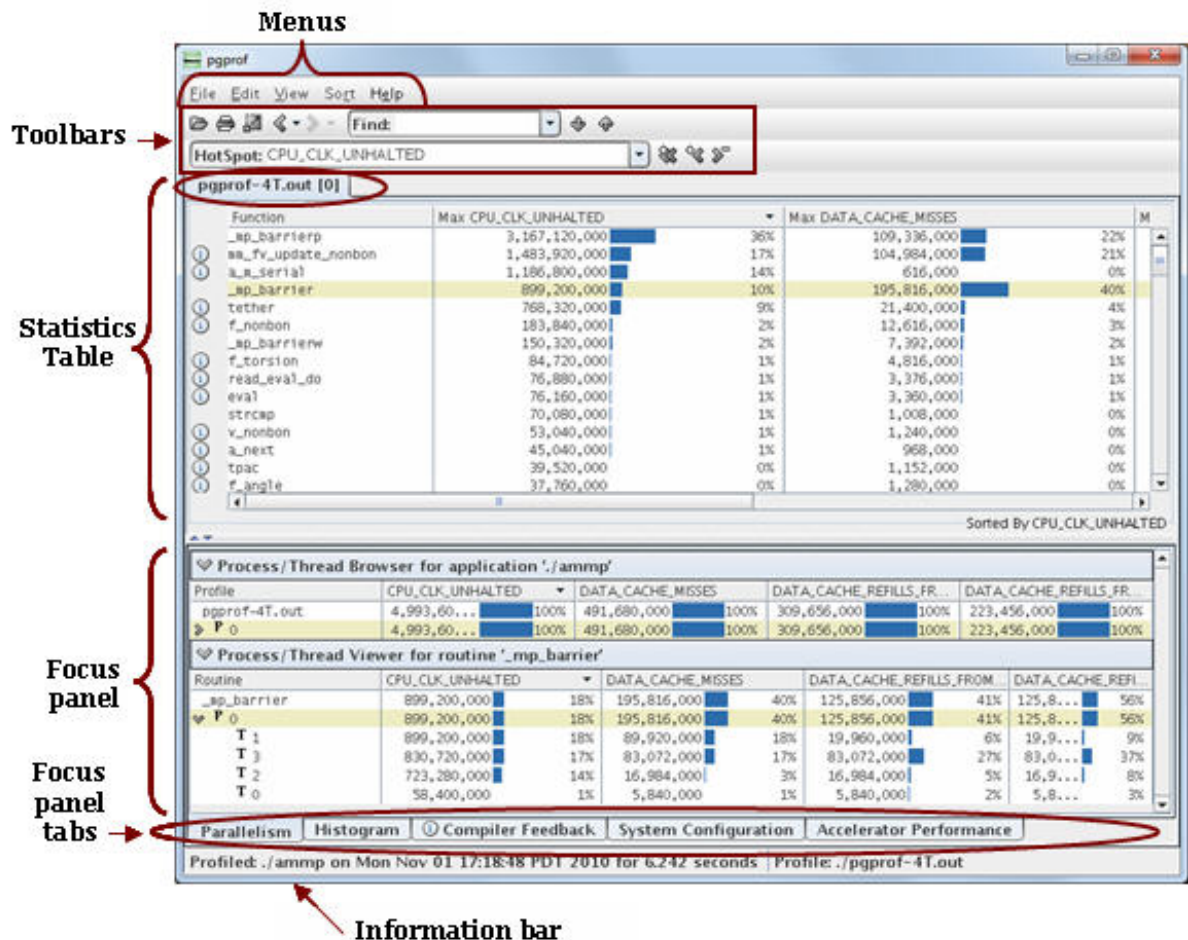
If you are viewing a profile across a slow network connection, or a connection that does not support remote display of Java GUIs, consider using the *PGPROF* command-line interface, described in [Chapter 9, “Command Line Interface,”](#) on page 59.

# Chapter 2. Using PGPROF

In [Chapter 1, “Getting Started”](#) you learned how to choose a profiling method, build your program, and execute it to collect profile data. This chapter provides a more detailed description of how to use the features of *PGPROF*, in particular:

- Profile navigation
- HotSpot navigation
- Sorting profile data
- Compiler Feedback
- Profiling parallel programs, including multi-threaded and MPI programs
- Scalability comparison
- Profiling resource utilization with hardware event counters
- Profiling accelerator programs

Figure 2.1. PGPROF Overview



PGPROF Tabs and Icons Overview

Before we describe how to navigate within *PGPROF*, it is useful to have some common terminology for the tabs and icons that you see within the application.

Closeable and Non-closeable Tabs







*PGPROF* displays both closeable and non-closeable tabs. For example, when you first invoke *PGPROF*, you see the function-level statistics table in a panel with a non-closeable tab. Then, to access profiling data specific to a given function, you double-click on the function name and a *closeable* tab opens with source code and profiling statistics for that function. This closeable tab navigation approach provides a way for you to easily view a variety of information quickly.

PGPROF Common Icons

Table 2.1 provides a summary of the common icons you see in the statistics table during profile navigation.



Table 2.1. PGPROF Icon Summary

Click this icon...	to...
	Display the corresponding assembly code for this line.
	Hide the corresponding assembly code for this line.
	Close the tab on which it is displayed.
	Display the compiler feedback for this line.
	Click to expand Focus Panel item.
	Click to hide Focus Panel item.

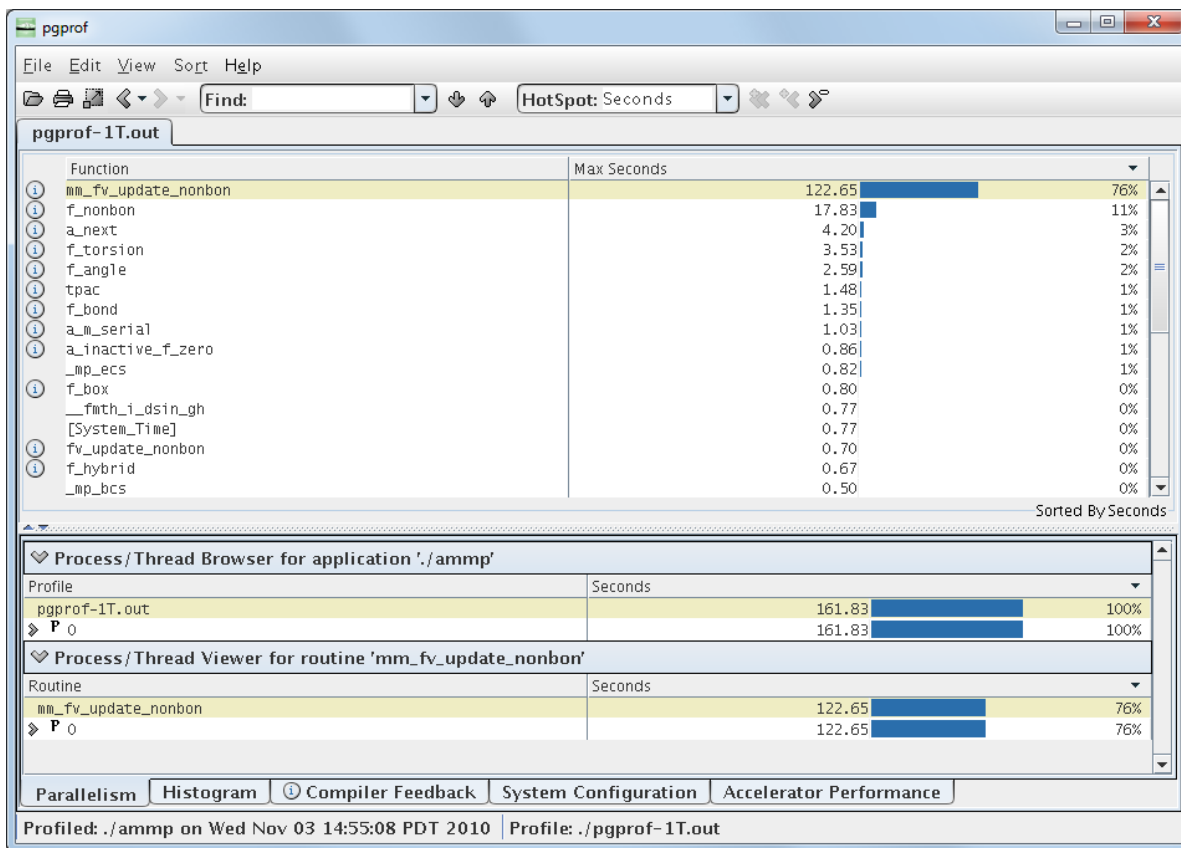
## Profile Navigation

When you first invoke *PGPROF*, it displays top-level profiling information in a non-closeable tab, as illustrated in [Figure 2.2](#).

This tab shows the Statistics Table containing a routine list in the Function column and performance data associated with each routine in the Seconds column. This list is sorted by the Seconds value, assuming there is such a value in the profile data.

By default, PGI compilers include enough symbol information in executables to allow *PGPROF* to display performance data at the source line level as well as at the routine level. However, if you compiled with the option `-Mnodwarf` or `-Mprof=func` or if you built your program using another compiler, you may only be able to access the routine-level view.

Figure 2.2. PGPROF Initial View



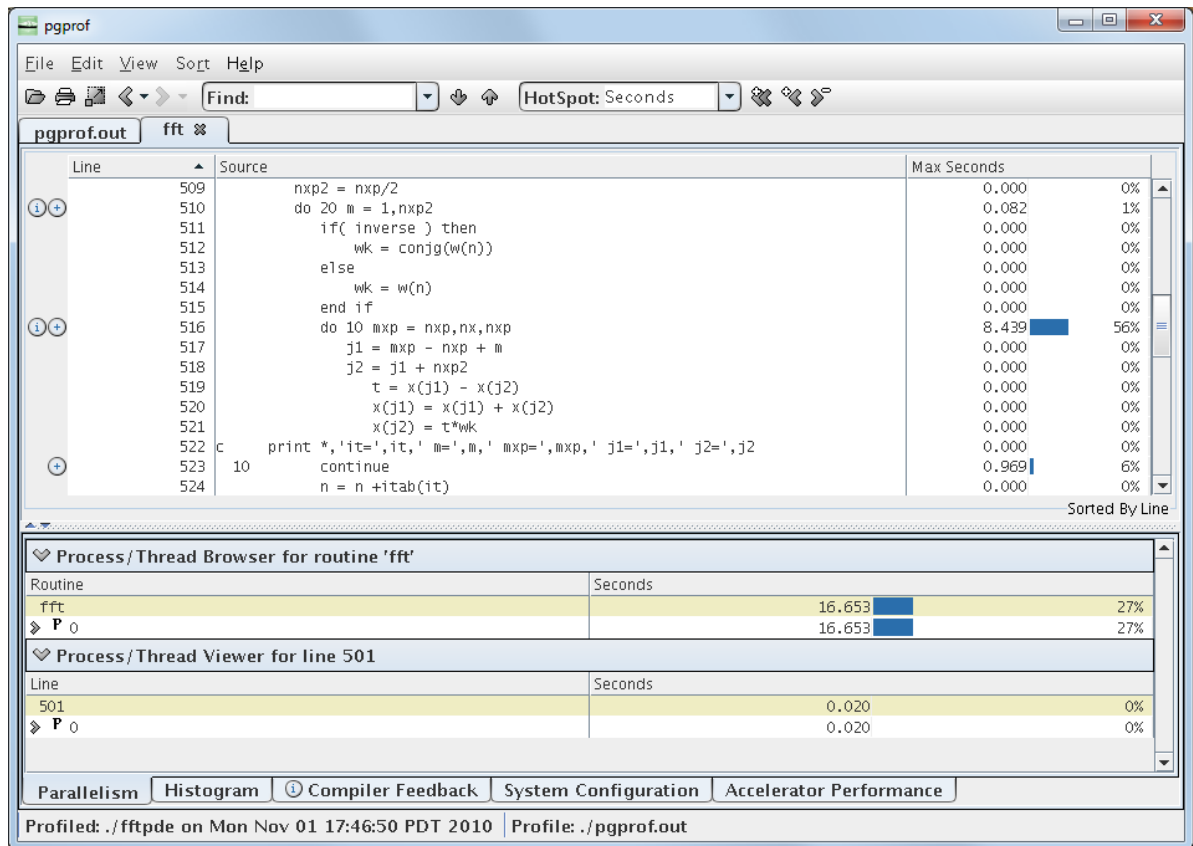
- To zoom in to the line level for a particular routine, double-click the function name.

This action opens a tab that displays profiling data specific to the given function. The tab label is the function name followed by an x icon. You use the x icon to close the tab when you no longer want to view that information.

In this tab, *PGPROF* displays the source code for that routine, together with the performance data for each line. For example, if you double-click on the function `fft`, *PGPROF* displays a new tab labelled `fft` that contains the source code for that function, as illustrated in [Figure 2.3](#).

Because your program is probably optimized, you may notice that performance data is only shown for a subset of the source lines. For example, a multi-line loop may only have line-level data for the first line of the loop.

Figure 2.3. Source Code View



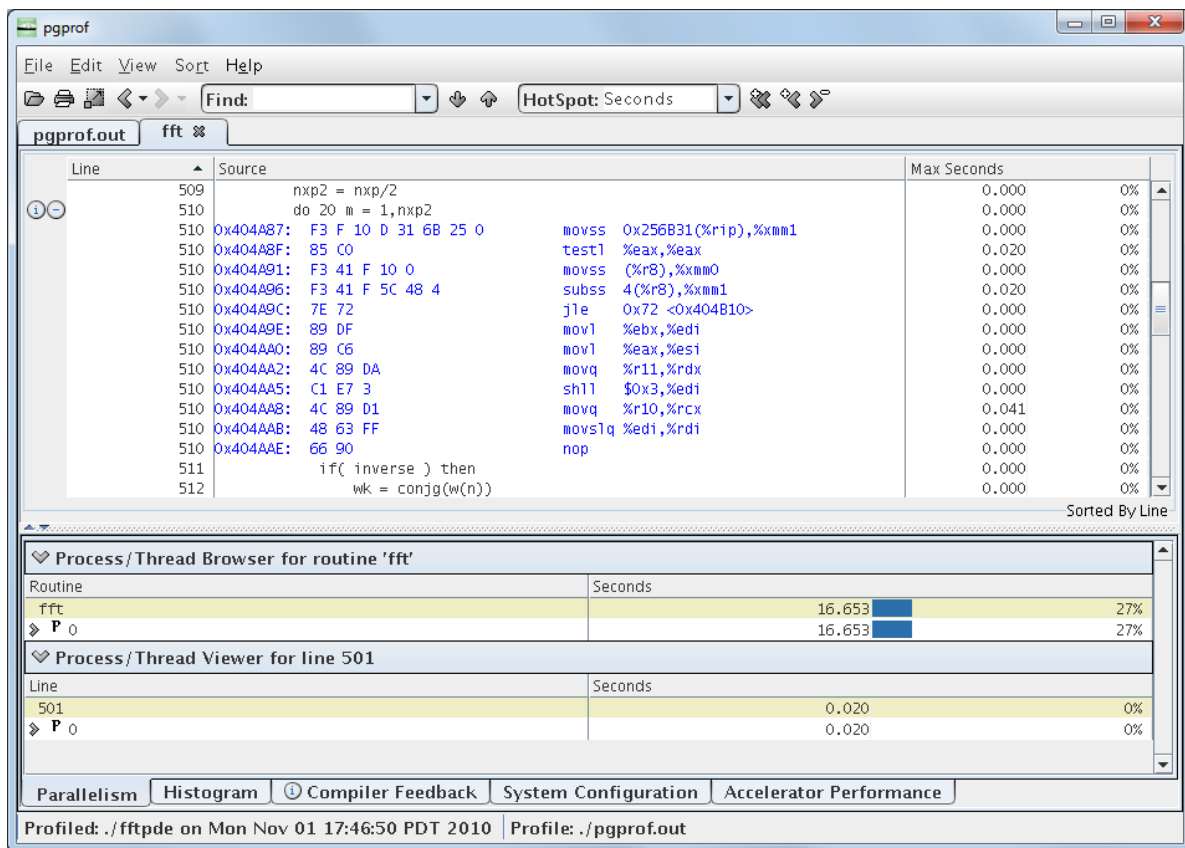
In the optimization process, the compiler may significantly reorder the assembly instructions used to implement the loop, making it impractical to associate any given instruction with a line in the loop. However, it is possible to associate all of a loop's instructions with that loop, so all of the performance data for the loop is associated with a single "line". For example, in [Figure 2.3](#), the information for the entire **do** loop at line 516 is associated with line 516.

- To zoom in to the assembly level for a particular source line, click the plus symbol (+) in the row of the Statistics Table containing that source line.

*PGPROF* displays the routine with assembly code interspersed with the source lines with which that assembly code is associated, as [Figure 2.4](#) illustrates the for loop at line 510.


*PGPROF* displays performance data associated with a specific assembly instruction in the row of the Statistics Table containing that instruction.

Figure 2.4. Assembly Level View



- To return to a previous view, use the Back button (" $\leftarrow$ ") in the Toolbar, just below the Menus.

Figure 2.5. View Navigation Buttons



The Back and Forward buttons work much like those found in web browsers, moving to previous and next views, respectively.

- To select and jump to a specific view, use the *down arrow* on each of the Forward and Back buttons.

### Note


You can have multiple function views open at a time, as illustrated in [Figure 2.4](#), where tabs for both functions `fft` and `cfft3` are displayed.

## HotSpot Navigation

The HotSpot navigation controls in the Toolbar are usually the quickest way to locate a hot spot. By *hot spot* we mean a program location that has a high value for some performance measurement such as Time, Count, and so on.

To locate the hot spot, select the desired performance measurement in the HotSpot drop-down menu in the Toolbar, then click on the "Hottest" button ("<<+"), illustrated in [Figure 2.6, "HotSpot Navigation Controls"](#), to select the highest value for that measurement in the current view.

Figure 2.6. HotSpot Navigation Controls



In addition to the HotSpot navigation controls on the toolbar, illustrated in [Figure 2.6](#), you can find the performance-critical parts of your program using the Histogram tab which shows clickable bar graphs of the performance data plotted against the address range of the program.


To find a HotSpot using the Histogram, click on the Histogram tab. In the histogram for the measurement you are interested in, click on the tallest bar. The corresponding row in the Statistics Table will be selected.

## Sorting Profile Data

*PGPROF* maintains a consistent sort order for the Statistics Table and the Histogram tab. Changing the sort order for either of these changes it for both of them. The sort order can be changed by using the Sort Menu, as described in ["Sort Menu," on page 49](#) or by clicking the column header in the Statistics Table or the row header in the Histogram tab.

The current sort order, such as sorting by the CPU Clock time, is displayed at the bottom of the Statistics Table. For example, [Figure 2.7](#) shows the message `Sort By CPU_CLK_UNHALTED` at the bottom of the Statistics Table and the Histogram.

Figure 2.7. Sort Example



## Compiler Feedback

The PGI compilers generate a special kind of information that is saved inside the executable file so that it is available to tools, such as *PGPROF*, to help with program analysis. A compiler discovers a lot about a program during the build process. Most compilers use such information for compilation, and then discard it. However, when the `-mprof` or `-minfo=ccff` options are used, the PGI compilers save this information in the object and executable files using the *Common Compiler Feedback Format*, or CCFF.

Feedback messages provide information about what the compiler did in optimizing the code, as well as describe obstacles to optimization. Most feedback messages have associated explanations or hints that explain what the message means in more detail. Further, these messages sometimes provide suggestions for improving the performance of the program.



The information icon indicates that CCFF information is available.

In *PGPROF* you can access Compiler Feedback by clicking an information icon in the left margin of the Statistics Table. This opens the **Compiler Feedback** tab in the Focus Panel. Messages are categorized according to the type of information that they contain.

For more information on the Compiler Feedback tab, refer to [“Compiler Feedback tab,” on page 53](#).

For more information on the Common Compiler Feedback Format (CCFF), refer to the website:

[www.pgroup.com/ccff/](http://www.pgroup.com/ccff/)

## Special Feedback Messages

There are some Compiler Feedback messages that deserve some explanation, specifically, intensity messages and messages for inlined routines.

### Intensity Messages

*Computational intensity* has been defined as the number of arithmetic operations performed per memory transfer.<sup>1</sup> The key idea is this: a high compute intensity value means that the time spent on data transfer is low compared to the time spent on arithmetic; a low compute intensity value suggests that memory traffic involving data transfer may dominate the overall time used by the computer.

The PGI Compiler emphasizes floating point operations, if they are present, to calculate the compute intensity ratio within a particular loop. If floating point operations are not present, the PGI compiler uses integer operations.

In some cases it is necessary to build programs using profile-guided optimization by building with `-Mpf i` or `-Mpf o`, as described in the section *Profile-Feedback Optimization using -Mpf i/Mpf o* in the Optimizing and Parallelizing chapter of the PGI User's Guide. Profile-guided optimization can often determine loop counts and other information needed to calculate the Compute Intensity for a given statement or loop.

### Messages for Inlined Routines

Inlined functions are identified by CCFF messages. These Compiler Feedback messages for routines that have been inlined are associated with the source line where the routine is called. Further, these messages are prefixed with the routine and line number, and are indented to show the level of inlining. Currently there is not a way to view the source code of that inlined instance of the routine.

## Profiling Parallel Programs

You can use *PGPROF* to analyze the performance of parallel programs, including multi-threaded and OpenMP programs, multi-process MPI programs, and programs that are a combination of the two. *PGPROF* also provides a Scalability Analysis feature that allows you to compare two profiling runs, and thus determine how well different parts of your program scale as the number of threads or processes changes.

### Profiling Multi-threaded Programs

Multi-threaded programs that you can profile using *PGPROF* include OpenMP programs that are built with `-mp`, auto-parallelized programs that are built with `-Mconcur`, and programs that use native thread libraries such as `pthread`.

### Collecting Data from Multi-Threaded Programs

Some methods of performance data collection work better with multi-threaded programs than others. As always, the recommended approach is to use **pgcollect**, initially with time-based sampling, optionally followed by event-based sampling. Building with `-Minfo=ccff` is always a good idea when using **pgcollect**.

Alternatively, building with the compiler option `-Mprof=lines` creates a program that collects accurate multi-threaded performance profiles.

---

<sup>1</sup>R.W. Hockney and C. R. Jesshope, *Parallel Computers 2: Architecture, Programming and Algorithms* 1988

The `-Mprof=func` option works with multi-threaded programs. However, routines that contain one or more parallel regions appear to be run on a single thread; because the data collection is at the entry and exit of the routine when the parallelism is not active.

The `-Mprof=time` and `-pg` options generate programs that only collect data on a single thread.

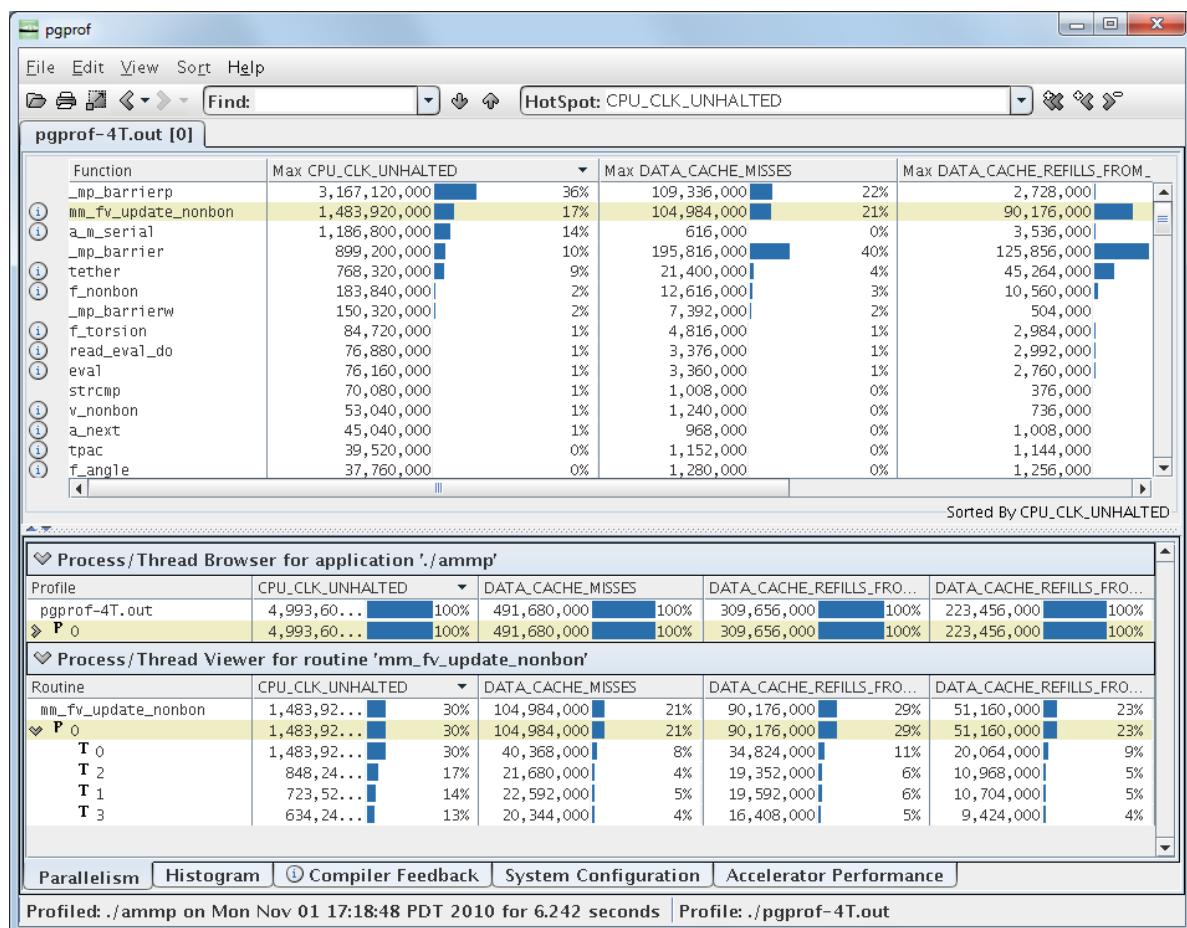
To collect data, for programs built using `-Mprof`, run your program normally. Upon successful termination, a `pgprof.out` file is created.

## Analyzing the Performance of Multi-Threaded Programs

The display of profile data for a multi-threaded program differs from that of a single-threaded program in a couple of ways:

- In the Statistics Table, the data shown is the maximum value for any single thread in the process.
- The Parallelism tab shows the thread-specific performance data for the row selected in the Statistics Table, whether the Statistics Table is in the routine-level, line-level, or assembly-level view. Click the arrow icon to the left of the P to expand the view to show all threads.

Figure 2.8. Multi-Threaded Program Example





You can use thread-specific data to determine how well-balanced the application is. Ideally, each thread would spend exactly the same amount of time on a given part of the program. If there are large disparities in the time spent by the various threads, this points to a load imbalance, where some threads are left idle while other threads are working. In this case, the resources of the system are not being used with 100% efficiency.

For example, in the program illustrated in [Figure 2.8](#), we can see that thread 0 spent 30% of the time in the routine, while thread 3 spent only 13% of the time there. Performance might improve if the work could be distributed more evenly.

## Profiling MPI Programs

You can profile MPI programs using PGI compilers and tools by building with the `-Mprof` option, which allows you to specify the implementation of MPI you want to use. Supported versions of MPI include: MPICH-1 (using the included version of MPICH-1), MPICH-2, Open MPI (Linux, Apple OS X), MVAPICH, and MS-MPI (Windows).

To collect MPI performance data, you must build your program using one of the MPI suboptions to the option `-Mprof`. These suboptions include:

For Linux:

```
-Mprof=mpich1
-Mprof=mpich2
-Mprof=mvapich1
```

For Windows:

```
-Mprof=mmpi
```

### Important

---

The MPI profiling options cannot be used alone. They may only be used in concert with another suboption of `-Mprof`, such as `lines`, `func`, or `time`.

For example, to build for time-based profiling with MPICH-1 profiling on Linux, use the following command:

```
$ pgfortran -fast -Mprof=time,mpich1 myprog.f90
```

To collect MPI profile data, run your program normally. On successful program termination, one profile data file is created for each MPI process. The master profile data file is named `pgprof.out`. The other files have names similar to `pgprof.out`, but they are numbered.

*PGPROF* MPI profiling on Linux collects counts of the number of messages and bytes sent and received. You can then use this information to analyze a program's message passing behavior.

## Open MPI Profiling

PGI provides performance profiling of MPI message passing support for Open MPI applications on Linux and Apple OS X. PGI ships a pre-built, pre-configured version of Open MPI on Apple OS X but not on Linux.

You can build your own version of Open MPI for Linux using PGI compilers; refer to [www.open-mpi.org](http://www.open-mpi.org) for information on building Open MPI. Information on configuring the Open MPI compiler wrappers for profiling is available in [Chapter 3, “Compiler Options for Profiling”](#).

To **build** an application using Open MPI, use the Open MPI compiler wrappers: `mpicc`, `mpic++`, `mpif77`, and `mpif90`. Use these wrappers with at least one of the options: `-Mprof=func`, `-Mprof=lines`, or `-Mprof=time` (Linux only). The PGI compilers do not directly support `-Mprof=openmpi`.

**Run** your program as you normally would. One or more files named `pgprof.out` is created in your working directory.

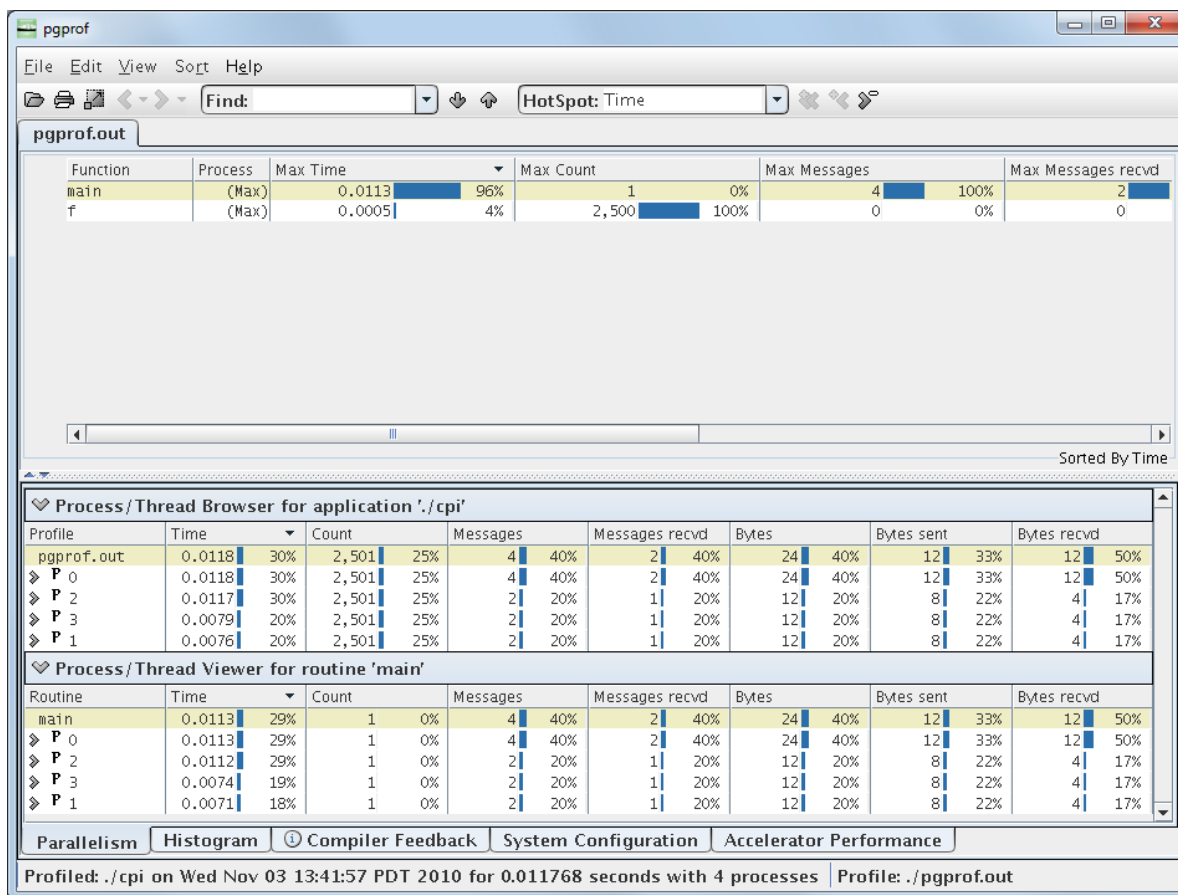
**Invoke** the profiler to see the results of your profiling run.

```
pgprof -exe your_program
```

## Analyzing the Performance of MPI Programs

[Figure 2.9](#) illustrates an MPI profile.

Figure 2.9. Sample MPI Profile



This sample shows an example MPI profile with maximum times and counts in the Statistics Table, and per process measurements in the Parallelism tab. The Parallelism tab for MPI programs is used in the same way

that it is used for multi-threaded programs, as described in [“Analyzing the Performance of Multi-Threaded Programs,” on page 16](#).

You can use the send and receive counts for messages, and the byte counts to identify potential communication bottlenecks, and use the process-specific data to find load imbalances.

## Scalability Comparison

*PGPROF* provides a Scalability Comparison feature that measures changes in the program's performance between multiple executions of an application. Generally this information is used to measure the performance of the program when it is run with a varying number of processes or threads. To use scalability comparison, first generate two or more profiles for a given application. For best results, compare profiles from the same application using the same input data with a different number of threads or processes.

Scalability is computed using the maximum time spent in each thread/process. Depending on how you profiled your program, this measurement may be displayed in the Statistics Table in a column with one of these heading titles:

Time	if you used <code>-Mprof=func</code> , <code>-Mprof=lines</code> , or <code>-Mprof=time</code>
CPU_CLK_UNHALTED	if you used <b>pgcollect</b>

### Important

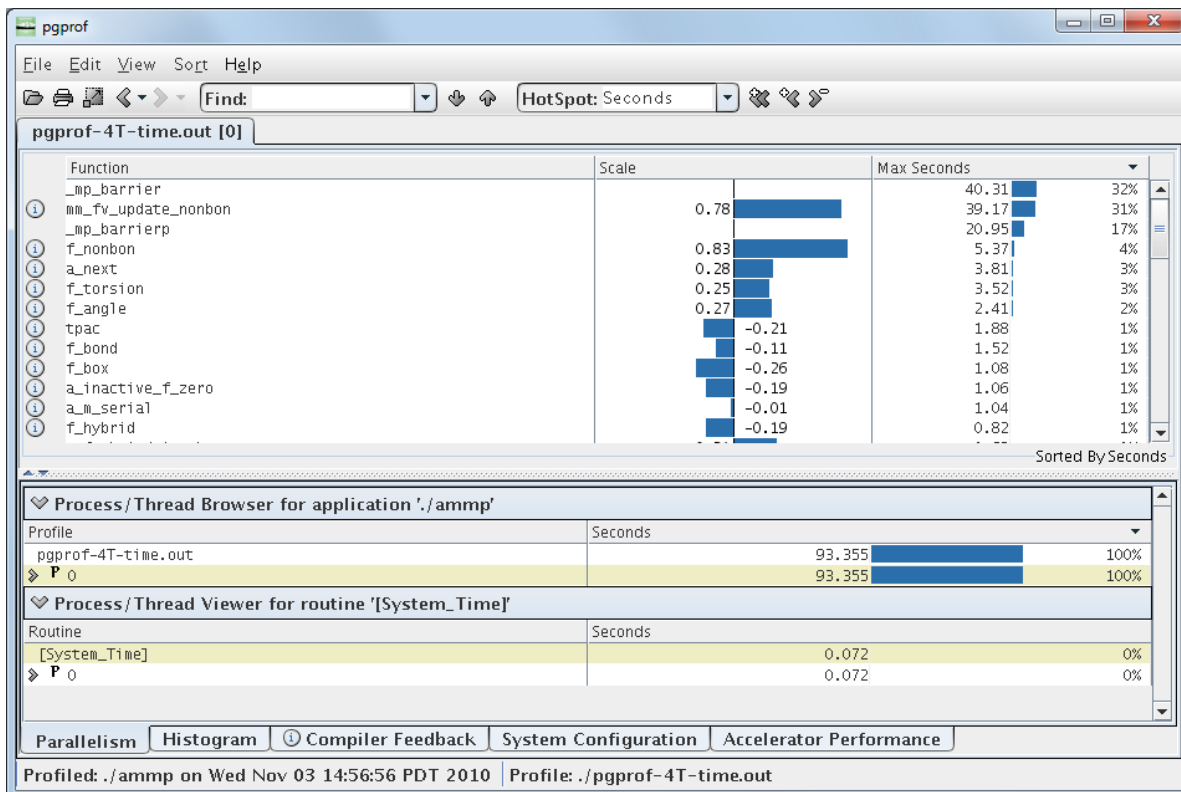
---

Profiling multi-process MPI programs with the **pgcollect** command is not supported.

The number of processes and/or threads used in each execution can be different. After generating two or more profiles, load one of them into *PGPROF*. Select the Scalability Comparison item under the File menu, described in [“File Menu,” on page 46](#), or click the Scalability Analysis button in the Toolbar. Choose a second profile for comparison. A new instance of *PGPROF* appears, with a column named `Scale` in the Statistics Table.

[Figure 2.10](#) shows the profile of a run that used four threads with Scalability Comparison to the same program run with a single thread.

Figure 2.10. Sample Scalability Comparison



Each profile entry that has timing information has a Scale value. The scale value measures how well these parts of the program scaled, or improved their performance as a result of parallelism.

- A scale value of zero indicates no change in the execution time between the two runs.
- A scale value of one means that part of the program achieved perfect scalability. For example, if a routine had a Time value of 100 seconds with one thread, and 25 seconds with four threads, it would have a Scale value of one.
- A negative value is the relative slowdown without taking the number of threads or processes into account. If a routine takes 20% more time to execute using four threads than it took using one thread, the Scale value is -0.2.
- A question mark ('?') in the Scale column indicates that *PGPROF* is unable to perform the scalability comparison for this profile entry. For example, scalability comparison may not be possible if the two profiles do not share the same executable or input data.

## Profiling Resource Utilization with Hardware Event Counters

### Important

Profiling with hardware counters is available only on Linux.

Modern x86 and x64 processors provide low-level hardware counters that can be used to track the resource utilization of a program. Tracking this information can be useful in tuning program performance because it

allows you to go beyond just knowing where the program is spending the most time and examine why it is spending time there.

Linux systems do not provide hardware counter support by default. These systems must have the OProfile package installed.

## Profiling with Hardware Event Counters (Linux Only)

*PGPROF* supports hardware counter data collection through the execution of the program under the control of the **pgcollect** command.

Collection of profile data using **pgcollect** may be done on any linux86 or linux86\_64 system where Oprofile is installed. OProfile is included as an install-time option with most Linux distributions; it may also be downloaded from <http://oprofile.sourceforge.net/>.

No special build options are required to enable event-based profiling with **pgcollect**, although building with the option `-Minfo=ccff` may provide useful compiler feedback.

For specific information on using *PGPROF* with hardware event counters, refer to [Chapter 10, “pgcollect Reference,”](#) on page 65.

## Analyzing Event Counter Profiles

If you executed your program under the control of **pgcollect**, then you can profile up to four event counters and view them in *PGPROF*. For brief descriptions of what each hardware counter measures, use

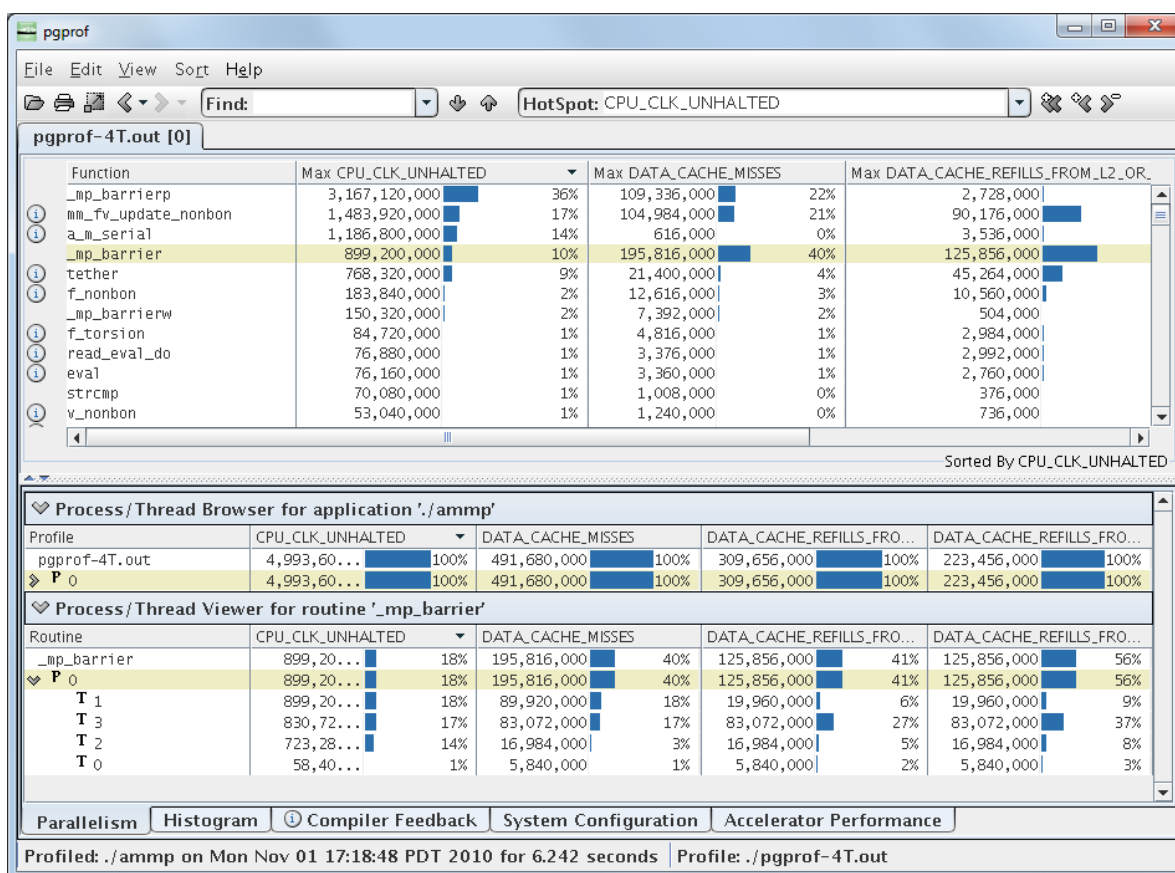
```
pgcollect --list-events
```

For more detailed information, see the processor vendor's documentation.

[Figure 2.11](#) shows a profile of four event counters: CPU\_CLK\_UNHALTED, DATA\_CACHE\_MISSES, DATA\_CACHE\_REFILLS\_FROM\_L2, DATA\_CACHE\_REFILLS\_FROM\_SYSTEM.

In this example, the routine using the most time is also getting many cache misses. Investigating the memory access behavior in that routine, and looking at the Compiler Feedback, may offer some clues for improving its performance.

Figure 2.11. Profile with Hardware Event Counter



## Profiling GPU Programs

You can use *PGPROF* to analyze the performance of GPU programs. GPU performance data is included in the profile, *pgprof.out*, when a GPU program is run using **pgcollect**. PGI provides two methods of programming GPUs: The PGI Accelerator model, which uses programs and directives to tell the compiler how to generate GPU code, and CUDA Fortran, which is used to program the GPU more directly.

The next section describes how to use **pgcollect** with PGI Accelerator model programs, and the subsequent section describes using it with CUDA Fortran programs.

## Profiling PGI Accelerator Model Programs

For the Accelerator model, the profiling procedure is the same as for host-only programs, except that *PGPROF* provides an Accelerator Performance tab that allows you to review profiling information provided by the accelerator. You do not need to build or run with any special options to collect accelerator performance data.

Here is an example of the commands you might use in a simple accelerator profiling session:

```
$ pgfortran -ta=nvidia -o myprog myprog.f90
$ pgcollect -time ./myprog
$ pgprof -exe ./myprog
```

## Note

You can build your program to print GPU performance data to the standard output by using the `time` suboption to the target accelerator option `-ta`. For example, you can use this command:

```
$ pgfortran -ta=nvidia,time myprog.f90
```

The `time` suboption has no effect on `pgcollect` or *PGPROF* profiling.

For more information on using PGI compilers to build programs for accelerators and on related terminology, refer to Chapter 7, “*Using an Accelerator*,” in the *PGI User's Guide*.

For more information on **pgcollect**, refer to [Chapter 10, “\*pgcollect Reference\*”](#).

## Analyzing Accelerator Performance Data

This section provides a basic description of how to examine accelerator performance data using *PGPROF*, including function-level analysis, region-level analysis and kernel-level analysis. A comprehensive guide to tuning accelerator programs is beyond the scope of this manual.

### Function-Level Analysis

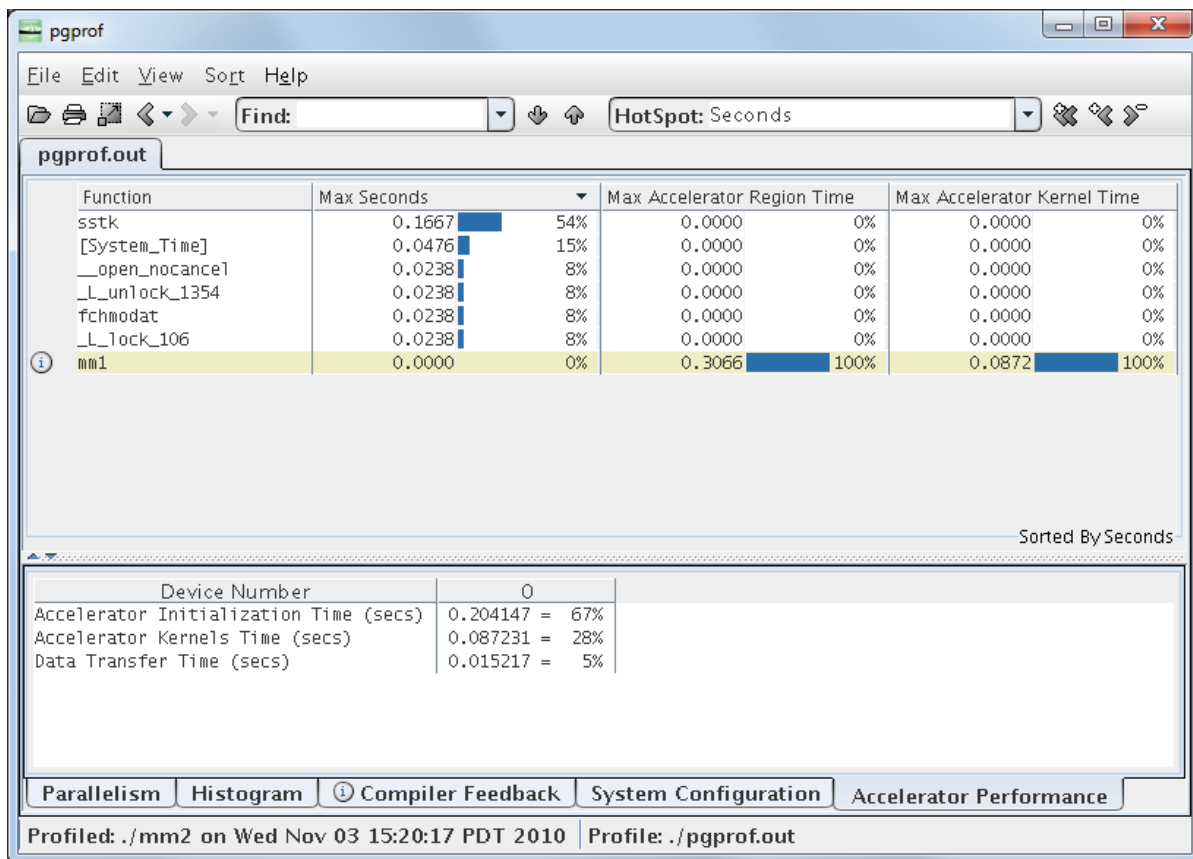
When you invoke *PGPROF* on the profile of an accelerator program, the initial view displays a function list showing host times in the Seconds column and accelerator times in the Accelerator Region Time column and Accelerator Kernel Time column. [Figure 2.12](#) illustrates a routine-level view with the routine `jacobi` selected and the Accelerator Performance tab chosen in the Focus Panel.

One of the first things to look at in tuning an accelerator program is whether the Data Transfer Time is large relative to the Accelerator Kernels Time. In the example illustrated in [Figure 2.12](#), the Accelerator Kernels Time of 4.134521 seconds is much larger than the Data Transfer Time of 0.132602 seconds, so we have efficient use of the accelerator.

If data transfer time is taking a significant portion of the total time, you would want to investigate if transfer time could be reduced using *data regions*, described in Chapter 7, *Using an Accelerator*, of the PGI User's Guide.

If data transfer time is relatively high and you have already considered data regions, you might want to examine the Compiler Feedback. You must compile with `-Minfo=ccff` to be able to do this. Check if the compiler is generating `copyin/copyout` operations that use slices of your arrays. If so, you may want to override the compiler to `copyin/copyout` the entire array.

Figure 2.12. Accelerator Performance Data for Routine-Level Profiling Example



For more information on compiler feedback, refer to [“Compiler Feedback,” on page 14](#).

## Region-Level Analysis

As with host-only profiles, you can drill down to the source code level by double-clicking on the routine name in the Function column. For an accelerator program, the display centers on the accelerator region directive for the longest-executing region. The Accelerator Performance tab shows a breakdown of timing statistics for the region and the accelerator kernels it contains.

### Note

A routine can contain more than one accelerator region.

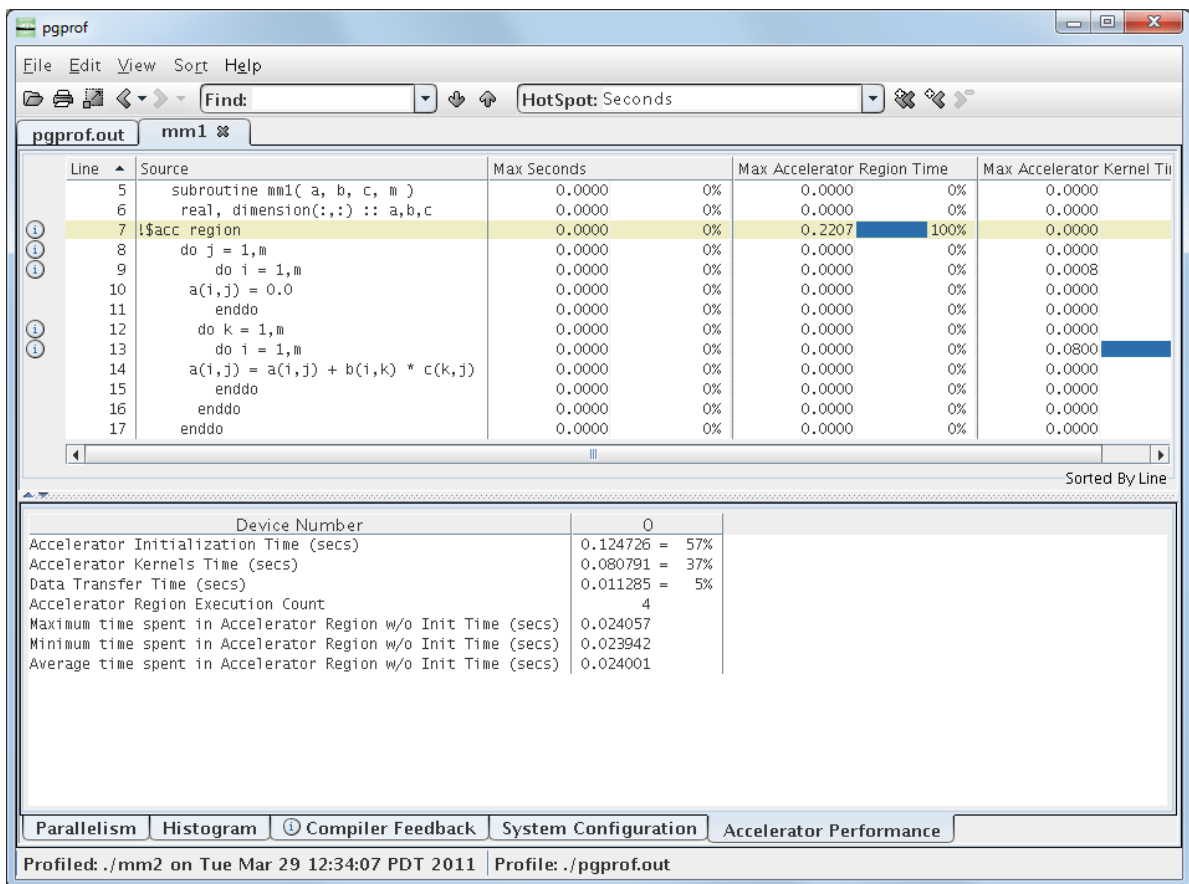
[Figure 2.13](#) shows an example of a source-level view with an accelerator region directive selected.

### Note

In this illustration, if you want to see the Seconds column, you could scroll to the right in the Statistics Table.



Figure 2.13. Source-Level Profiling for an Accelerator Region



## Kernel-Level Analysis

Since an accelerator region can contain multiple distinct kernels, you may want to examine performance data for an individual kernel. You do this by selecting the first source line of the kernel.

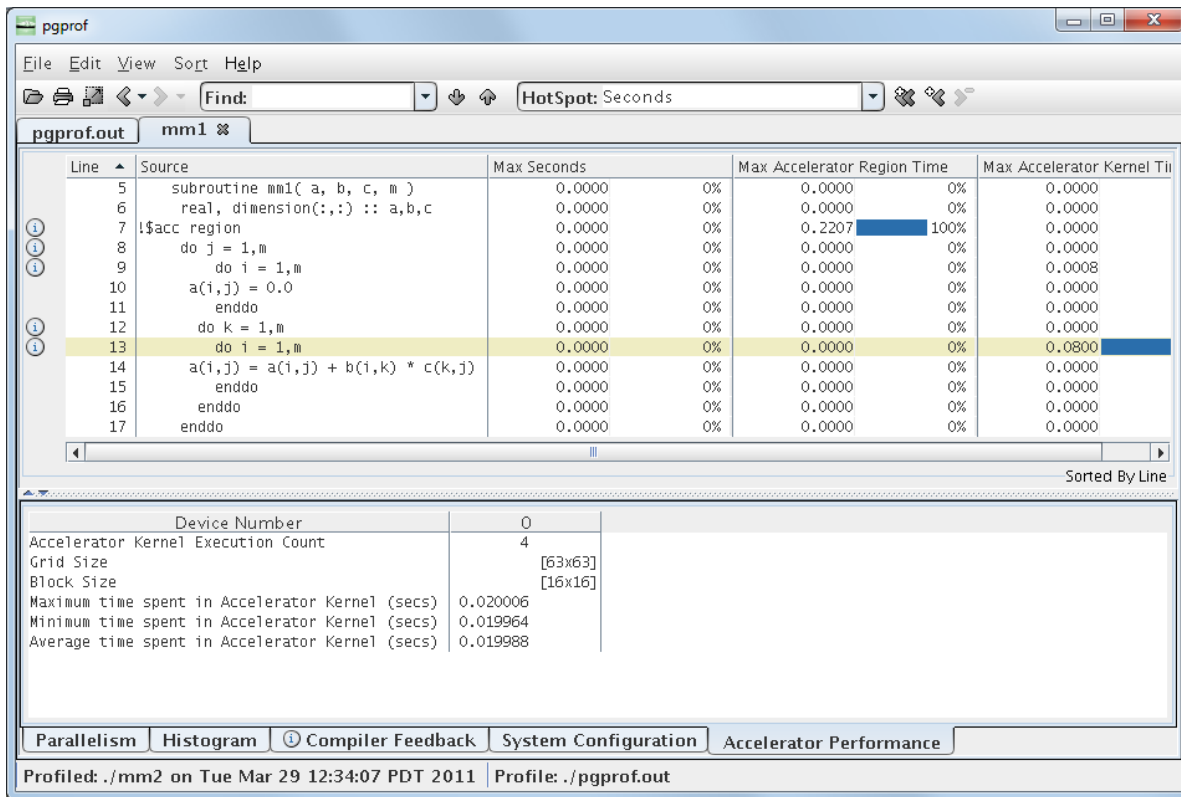
In the source-level view, the first line of a kernel has data listed in the Accelerator Kernel Time column.

To navigate to the longest-executing kernel:

1. Select Accelerator Kernel Time in the HotSpot selector in the upper-right portion of the user interface.
2. Click the double left arrow (<<+) located next to the HotSpot selector.

In [Figure 2.14](#) the selected line in the main Statistics Table has a value only in the Accelerator Kernel Time. The Accelerator Performance tab displays all the details for the Accelerator Kernel performance data.

Figure 2.14. Source-Level Profiling for an Accelerator Kernel



For more information on tuning accelerator programs, refer to Chapter 7, *Using an Accelerator*, of the PGI User's Guide.

## Profiling CUDA Fortran Programs

For CUDA Fortran, **pgcollect** provides an option `-cuda` that enables collection of performance data on the CUDA device. Analysis of this performance data is much the same as for accelerator model programs, as described in the previous section, except that the data is collected from counters on the device and in the CUDA driver.

If you are profiling a program that uses CUDA Fortran kernels running on a GPU, **pgcollect -cuda** collects performance data from CUDA-enabled GPU and includes it in the profile output for the program. The syntax for this command option is:

```
-cuda[=gmem|branch|cfg:<cfgpath>|cc13|cc20|list]
```

The sub-options modify the behavior of **pgcollect -cuda** as described here:

- branch** Collect branching and warp statistics.
- cc13** Use counters for compute capability 1.3. [default]
- ccnm** Use counters for compute capability *n.m*.  
Use `pgcollect -help` to see which compute capabilities your system supports.
- cfg:<cfgpath>** Specify `<cfgpath>` as CUDA profile config file.

<code>gmem</code>	Collect global memory access statistics.
<code>list</code>	List cuda event names available for use in profile config file.

## Performance Profiling with Pre-defined Counter Configurations

The `-gmem` and `-branch` suboptions initiate profiling with predefined sets of performance counters to measure specific areas of GPU resource utilization.

- `-gmem` measures accesses to global memory.
- `-branch` tracks divergent branches and thread warp statistics.

Some of the counters used for `-gmem` and `-branch` differ depending on the version (compute capability) of the GPU you are using. To ensure that you use the counters available on your GPU, you must specify the compute capability you want to use. You can do this in two ways:

- On the **pgcollect** command line. For example, to specify compute capability 1.3, you can use:

```
pgcollect -cuda=branch,cc13 myprog
```

- In a special file in your home directory. The home directory is specified by the environment variable `HOME`.

The name of the file depends on your OS:

- On Windows, the name of the file is `mypgirc`.
- On Linux and Mac OS, the name of the file is `.mypgirc`.

In this file you put a line that indicates compute capability 1.3 or 2.0:

```
COMPUTECAP=13
or
COMPUTECAP=20
```

### Note

Placing this line in this file also affects the compiler defaults with respect to compute capability.

## Performance Profiling with User-defined Counter Configurations

You have the ability to specify which counters to use in data collection. To do this, you create a profile configuration file with any filename. You can do this using this command:

```
pgcollect -cuda=list
```

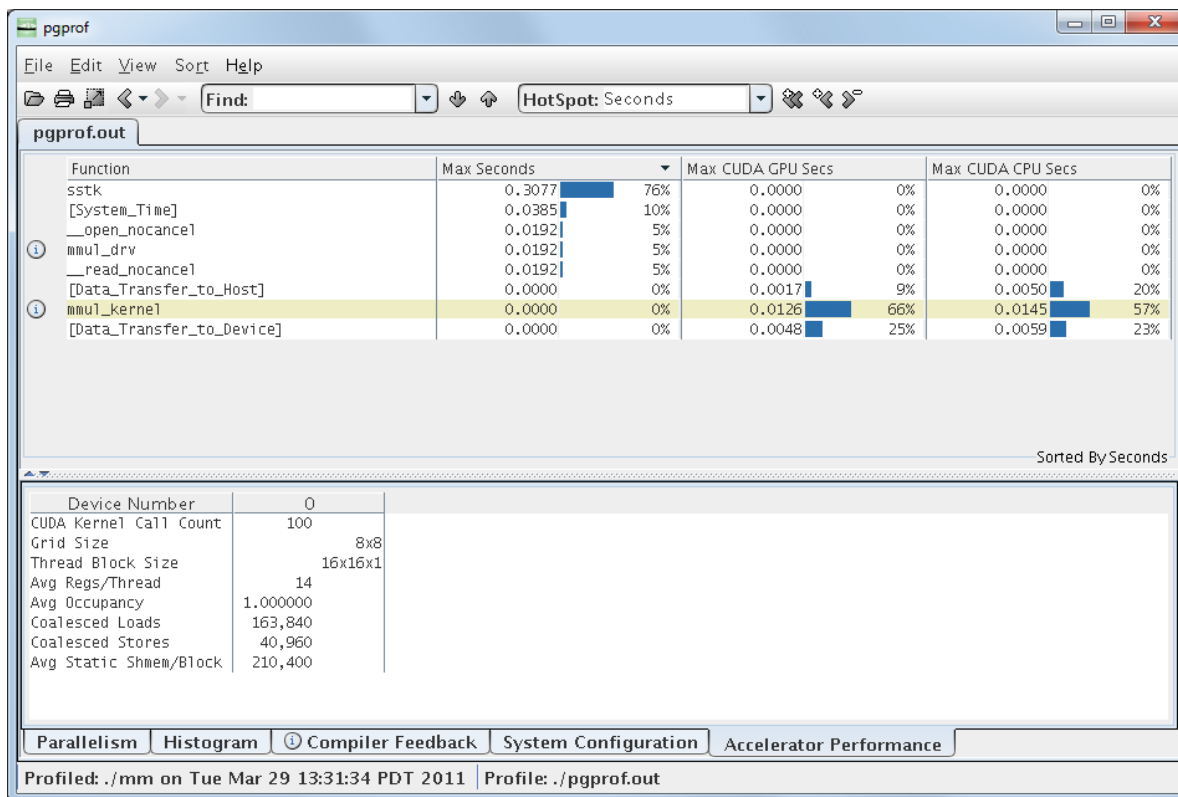
To specify the counters to use, place a list of NVIDIA counters in your file, listing one counter per line. In general, the number of counters you can list is limited to four, although with compute capability 2.0 you may be able to use more, depending on the counters selected. In addition, you may always list certain data collection options that do not depend on hardware counters, such as these:

<code>gridsize</code>	<code>stasmemperblock</code>
<code>threadblocksize</code>	<code>regperthread</code>
<code>dynsmemperblock</code>	<code>memtransfersize</code>

To get a full list of the counters available, use this command:

```
pgcollect -cuda=list
```

Figure 2.15. CUDA Program Profile



In Figure 2.15:

- The columns labeled Max CUDA GPU Secs and Max CUDA CPU Secs show times captured by the CUDA driver.
- The Max Seconds column contains timings for host-only code.
- Pseudo-function names [Data\_Transfer\_To\_Host] and [Data\_Transfer\_To\_Device] show the transfer times to and from the GPU.
- The Accelerator Performance Tab shows counter values collected from the GPU.

# Chapter 3. Compiler Options for Profiling

This chapter describes the PGI compiler options that are used to control profiling and how they are interpreted.

## -Mprof Syntax

You can use the following compiler options to control data collection. Most of these options are related to `-Mprof`, for which the syntax is:

```
-Mprof{=option[,option, ...]}
```

You use `-Mprof` to set performance profiling options. Use of these options causes the resulting executable to create a performance profile that can be viewed and analyzed with the *PGPROF* performance profiler.

### Note

---

If you are going to use `pgcollect` to gather performance data, you don't need to compile or link with `-Mprof`.

## Profiling Compilation Options

In the descriptions that follow, instrumentation-based profiling implies compiler-generated source instrumentation. profiling implies the use of instrumented wrappers for MPI library routines.

### `-Minfo=ccff`

Generate compiler feedback information and store it in object and executable files for later access by performance tools. Use `-Minfo=ccff` when collecting performance data using **pgcollect**. All `-Mprof` options except `-Mprof=dwarf` imply `-Minfo=ccff`.

### `-Mprof=dwarf`

Generate a subset of DWARF symbol information adequate for viewing source line information with most performance profilers.

In the PGI compilers `-Mprof=dwarf` is on by default. You can use the `-Mnodwarf` option to disable it. Source-level information is not available if you profile a program built with `-Mnodwarf`.

- `-Mprof=func`  
Perform routine-level instrumentation-based profiling.
- `-Mprof=lines`  
Perform instrumentation-based line-level profiling.
- `-Mprof=mpich1`  
Perform profiling for MPICH-1. Implies `-Mmpi=mpich1`.
- `-Mprof=mpich2`  
Perform profiling for MPICH-2. Implies `-Mmpi=mpich2`.
- `-Mprof=msmpi`  
[Windows] Perform profiling for Microsoft MSMPI. Implies option `-Mmpi=msmpi`.
- `-Mprof=mvapich1`  
[Linux] Perform profiling for MVAPICH-1. Implies `-Mmpi=mvapich1`.
- `-Mprof=sgimpi`  
Perform profiling for SGI MPI. Implies option `-Mmpi=sgimpi`.

### Note

---

This option is required even if you compile and link using the `mpicc` or `mpif90` compiler wrappers.

- `-Mprof=time`  
[Linux] Generate a profile using time-based assembly-level statistical sampling. This is equivalent to using the `-pg` option, except the profile is saved in a file named `pgprof.out` rather than in `gmon.out`.
- `-pg`  
[Linux] Enable gprof-style (sample-based) profiling. Running an executable compiled with this option produces a `gmon.out` profile file which contains routine, line, and assembly-level profiling data.

# Chapter 4. Profiling with Open MPI

This chapter describes how to do MPI profiling with Open MPI and PGI compilers.

The explanations and examples in this chapter require PGI 12.6 or later and assume the use of Open MPI 1.6 or later on linux86\_64, although this information may apply to other releases as well. Further, the Open MPI installation must be configured and built for use with PGI compilers.

PGI supports profiling with Open MPI on Linux and OS X.

## Overview

The Open MPI documentation *strongly* recommends using the MPI compiler wrappers (e.g. `mpicc`, `mpic++`, `mpif77`, `mpif90`) to build Open MPI applications, rather than using direct invocation of the compiler via `pgcc` or `pgfortran`. Compiler options are passed straight through the MPI compiler wrappers to the compilers.

- **For Open MPI profiling on OS X**, the PGI distribution includes a preconfigured version of Open MPI ready for use with the PGI compilers. To do MPI profiling in this case, simply build your MPI program using the Open MPI compiler wrappers and one of the PGI profile instrumentation options: `-Mprof=time`, `-Mprof=func`, or `-Mprof=lines`. You do not need to follow any of the subsequent instructions in this chapter.
- **For Open MPI profiling on Linux**, your Open MPI installation requires some special configuration to support MPI profiling with the PGI compilers and tool. You accomplish this by editing the compiler wrapper "data" files.

## Compiler Wrapper Data Files

In `your_openmpi_dir/share/openmpi` there are some configuration files that you can use to control how the MPI compiler wrappers handle various compiler options:

```
mpicc-wrapper-data.txt
mpic++-wrapper-data.txt
mpif77-wrapper-data.txt
mpif90-wrapper-data.txt
```

You must modify these files to support MPI profiling with Open MPI and PGI compilers. PGI cannot provide pre-configured files for Linux since each Open MPI installation and/or PGI installation may be different.

## Configure Open MPI for PGI Profiling

To configure Open MPI for PGI profiling, you edit the compiler wrapper data files.

### Important

---

Before you begin, make backup copies of your original wrapper data files.

Make these modifications:

1. Add the line `compiler_args=` before any other configuration lines.
2. Copy the entire data block in the sample file *twice*. You now have three identical data blocks in your wrapper file.

You need a data block for each of these compiler options:

- The default with no profiling, as it was originally.
  - `-Mprof=func|lines`
  - `-Mprof=time`
3. In the second data block, modify the `compiler_args=` and the `compiler_flags=` lines. The PGI profiling options are shown just to the right of the equal sign. The compiler flags you select immediately follow the equal sign, with a space between each flag.

Your lines should look *similar* to these:

```
compiler_args=-Mprof=func;-Mprof=lines
...
compiler_flags=
...
```

4. In the third data block, modify the `compiler_args=` and the `compiler_flags=` lines. The PGI profiling options are shown just to the right of the equal sign. The compiler flags in this data block should include: `-W0, -profile, lines` at the beginning of the list of flags you select.

Your lines should look *similar* to these:

```
compiler_args=-Mprof=time
...
compiler_flags=-W0,-profile,lines
...
```

5. In both the second and third data blocks, modify the `libs=` line so that `-lpgnod_prof_openmpi` comes just before `-lmpi`.

### Note

---

Do not modify any other lib values.

The new `libs=` line looks similar to this:

```
libs=-lpgnod_prof_openmpi -lmpi -lopen-rte -lopen-pal -lrt
-lldl -Wl,--export-dynamic -lnsl -lutil -lpthread -ldl
```



## Note

Do not assume that the `libs=` line in this example contains the correct libraries for your Open MPI installation. This `libs=` line is just an example. Your compiler wrapper file may contain different libraries, and it is important that you use those libraries.

## Modified Compiler Wrapper Data File Sample

When you complete your modifications, your new wrapper data file has three data blocks that look *similar* to these. The lines you modified are in bold.

```

compiler_args=
project=Open MPI
project_short=OMPI
version=1.2.8
language=C
compiler_env=CC
compiler_flags_env=CFLAGS
compiler=pgcc
extra_includes=
preprocessor_flags=-D_REENTRANT
compiler_flags=
linker_flags=
libs=-lmpi -lopen-rte -lopen-pal -lrt -ldl -Wl,--export-dynamic -lnsl -lutil
-lpthread -ldl
required_file=
includedir=${includedir}
libdir=${libdir}

compiler_args=-Mprof=func;-Mprof=lines
project=Open MPI
project_short=OMPI
version=1.2.8
language=C
compiler_env=CC
compiler_flags_env=CFLAGS
compiler=pgcc
extra_includes=
preprocessor_flags=-D_REENTRANT
compiler_flags=
linker_flags=
libs=-lpgnod_prof_openmpi
-lmpi -lopen-rte -lopen-pal -lrt -ldl -Wl,--export-dynamic -lnsl
-lutil -lpthread -ldl
required_file=
includedir=${includedir}
libdir=${libdir}

compiler_args=-Mprof=time
project=Open MPI
project_short=OMPI
version=1.2.8
language=C
compiler_env=CC
compiler_flags_env=CFLAGS
compiler=pgcc
extra_includes=
preprocessor_flags=-D_REENTRANT

```

```

compiler_flags=-W0,-profile,lines
linker_flags=
libs=-lpgnod_prof_openmpi
-lmpi -lopen-rte -lopen-pal -lrt -ldl -Wl,--export-dynamic -lnsl
-lutil -lpthread -ldl
required_file=
includedir=${includedir}
libdir=${libdir}

```

### Note

---

Do not assume that the data blocks in this example contain the correct data for your Open MPI installation. These data blocks are just examples. Your compiler wrapper file may contain different data, and it is important that you use the data in your original wrappers.

## Using PGI MPI Profiling for Open MPI

Once you have successfully configured the compiler wrapper data files for PGI profiling options, you can profile your program.

- For a single-threaded MPI program, use `-Mprof=time`.
- For a multi-threaded program in which each MPI process uses multiple threads, use either `-Mprof=func` for function-level profiling or `-Mprof=lines` for line-level profiling. Be aware that `-Mprof=lines` can introduce significant overhead, so long-running programs can take much longer to complete.

1. Build your program. For example:

```
$ mpif90 -fast -Mprof=time -o myprog myprog.f90
```

2. Run your program:

```
$ mpiexec -n 4 myprog
```

This step should create multiple `pgprof*.out*` files, one for each process.

3. Examine your profile data

```
$ pgprof -exe myprog
```

# Chapter 5. Command Line Options

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

## Command Line Option Descriptions

The following list describes the options and how *PGPROF* interprets them.

### datafile

A single datafile name may be specified on the command line. For profiled MPI applications, the specified datafile should be that of the initial MPI process. Access to the profile data for all MPI processes is available in that case, and data may be filtered to allow inspection of the data from a subset of the processes.

The default datafile name is `pgprof.out`. If no datafile argument is used, *PGPROF* attempts to use `pgprof.out` in the current directory.

### -exe <filename>

Set the executable to filename. The default filename is `a.out`.

### -feedbackonly (Linux only)

Only browse source code and Compiler Feedback information. Do not load any performance data from profile runs.

### -help

Prints a list of available command-line arguments.

### -I <srcpath>

Specify the source file search path.

*PGPROF* always looks for a program source file in the current directory first. If it does not find the source file in the current directory, it consults the search path specified in `srcpath`.

The `srcpath` argument is a string containing one or more directories separated by a path separator. The path separator is platform dependent: on Linux and Mac OS, it is a colon ( : ), and on Windows it is a semicolon ( ; ). Directories in the path are then searched in order from left-to-right. When a directory with a filename that matches a source file is found, that directory is used.

Here is an example for Linux and Mac OS. In this example, the profiler first looks for source files in the current directory, then in the `../src` directory, followed by the STEPS directory.

```
-I ../src:STEPS
```

Here is the same example for Windows:

```
-I ../src:STEPS
```

For more information, see the Open Profile... item in the description of the [“File Menu,” on page 46](#).

`-jarg, arg1[, arg2,..., argn]`

Pass specified arguments, separated by commas, to java. For example, the following option passes the argument `-Xmx256m` to java.

```
-jarg, -Xmx256m
```

This option is provided for troubleshooting purposes and is expected to rarely be used. If you do use this option, be certain not to forget the comma between the option and the first argument.

`-scale “file(s)”`

Compare scalability of datafile with one or more files. A list of files may be specified by enclosing the list within quotes and separating each filename with a space. For example:

```
-scale one.out two.out
```

This example compares the profiles `one.out` and `two.out` with datafile (or `pgprof.out` by default). If only one file is specified quotes are not required.

For sample based profiles (e.g., `gmon.out`) specified with this option, *PGPROF* assumes that all profile data was generated by the same executable. For information on how to specify multiple executables in a sample-based scalability comparison, see the Scalability Comparison... item in the description of the [“File Menu,” on page 46](#).

`-text`

Use the *PGPROF* Command-Line Interface (CLI).

`-V`

Print version information.

## Profiler Invocation and Startup

Let’s take a look at some common ways to invoke the profiler, describing what each launch command means.

### % **pgprof**

- If a `pgprof.out` file exists in the current directory, *PGPROF* tries to open it.
- If an executable name can be determined from the `pgprof.out` file, the GUI is populated according to profile data, if valid.
- If an executable name can NOT be determined from the `pgprof.out` file, then a dialog is opened on top of the main window with the following message:

Can't determine executable for file 'pgprof.out' Please use 'File | Open Profile...' menu to specify one

- If no `pgprof.out` file exists in the current directory, the GUI is not populated and no dialog appears.

#### % **pgprof -exe <execname>**

- If a `pgprof.out` file exists in the current directory, *PGPROF* tries to open it and use `<execname>`. Further, the GUI is populated according to profile data, if valid.
- If no `pgprof.out` file exists in the current directory, the GUI is not populated and no dialog appears. Further, when the user selects the menu `File | Open Profile...`, then the Text Field for Executable is set with `<execname>` in the dialog.

#### % **pgprof -exe <execname> <profilename>**

*PGPROF* tries to open the profile `<profilename>` using `<execname>` for the executable name. Further, the GUI is populated according to profile data, if valid.

#### % **pgprof -feedbackonly**

- If an `a.out` file exists in the current directory, *PGPROF* tries to open it. Further, if `a.out` is an executable with valid DWARF/ELF/CCFF information, then *PGPROF* populates the GUI. You are then in *feedback-only mode*.
- If no `a.out` file exists in the current directory, the GUI is not populated and no dialog appears. Further, when the user selects the menu `File | Open Profile...`, then the Feedback only checkbox is selected in the dialog.

#### % **pgprof -exe <execname> -feedbackonly**

*PGPROF* tries to open the executable `<execname>`. Further, if the executable `<execname>` is valid with DWARF/ELF/CCFF info, then *PGPROF* populates the GUI. You are then in *feedback-only mode*.

#### % **pgprof -exe <execname> -feedbackonly <profilename>**

#### Note

---

`<profilename>` is ignored without warning

*PGPROF* tries to open the executable `<execname>`. Further, if the executable contains valid DWARF/ELF/CCFF information, then *PGPROF* populates the GUI. You are then in *feedback-only mode*.



# Chapter 6. Environment Variables

This chapter describes the system environment variables that you can set to change the way profiling is performed.

## System Environment Variables

As you learned in [“Basic Profiling,” on page 1](#), a profiled program collects call counts and/or time data. When the program terminates, a profile data file is generated. Depending on the profiling method used, this data file is called `pgprof.out` or `gmon.out`.

You can set the following system environment variables to change the way profiling is performed:

- `GMON_ARCS` – Use this environment variable to set the maximum number of arcs (caller/callee pairs).

The default is 4096. This option only applies to `gprof` style profiling, that is, programs compiled with the `-pg` option.

- `PGPROF_DEPTH` – Use this environment variable to change the maximum routine call depth for `PGPROF` profiled programs.

The default is 4096 and is applied to programs compiled with any of the following options: `-Mprof=func`, `-Mprof=lines`, or `-Mprof=time`.

- `PGPROF_EVENTS` – Use this environment variable to specify hardware (event) counters from which to collect data.

This variable is applied to programs executed with the **pgcollect** command using one of the event-based profiling options. The use of hardware (event) counters is discussed in further detail in [“Profiling Resource Utilization with Hardware Event Counters,” on page 20](#).

- `PGPROF_NAME` – Use this environment variable to change the name of the output file intended for `PGPROF`.

The default is `pgprof.out`. This option is only applied to programs compiled with any of the following options: `-Mprof=[func | lines | MPI | time]`. If a program is compiled with the `-pg` option, then the output file is always called `gmon.out`.





# Chapter 7. Data and Precision

This chapter contains descriptions of the profiling mechanism that measures time, how statistics are collected, and the precision of the profiling results.

## Measuring Time

The sample-based profiling mechanism collects total CPU time for programs that are compiled with the options `-pg` and `-Mprof=time`, or executed with **pgcollect -time**, as described in [“Sample-based Profiling,” on page 3](#). The profiling mechanism collects cycle counts for programs run under the control of **pgcollect** or executed with **pgcollect** event-based sampling. *PGPROF* automatically converts CPU cycles into CPU time.

Programs compiled for instrumentation-based profiling with `-Mprof=lines` or `-Mprof=func` employ a virtual timer for measuring the elapsed time of each running process/thread. This data collection method employs a single timer that starts at zero (0) and is incremented at a fixed rate while the active program is being profiled. For multiprocessor programs, there is a timer on each processor, and the profiler's summary data (minimum, maximum and per processor) is based on each processor's time executing in a function. How the timer is incremented and at what frequency depends on the target machine. The timer is read from within the data collection functions and is used to accumulate `COST` and `TIME` values for each line, function, and the total execution time. The line level data is based on source lines; however, in some cases, there may be multiple statements on a line and the profiler shows data for each statement.

### NOTE

---

For instrumentation-based profiling, information provided for longer running functions are more accurate than for functions that only execute for a short time relative to the overhead of the individual timer calls. Refer to [“Caveats \(Precision of Profiling Results\),” on page 42](#) for more information about profiler accuracy.

## Profile Data

The following statistics are collected and may be displayed by the *PGPROF* profiler.

### BYTES

For MPI profiles only. This is the number of message bytes sent and received.

### BYTES RECEIVED

For MPI profiles only. This is the number of bytes received in a data transfer.

### BYTES SENT

For MPI profiles only. This is the number of bytes sent.

### CALLS

The number of times a function is called.

### COST

The sum of the differences between the timer value entering and exiting a function. This includes time spent on behalf of the current function in all children whether profiled or not. *PGPROF* can provide cost information when you compile your program with either the `-Mprof=cost` or the `-Mprof=lines` option. For more information, refer to [“Basic Profiling,” on page 1](#).

### COUNT

The number of times a line or function is executed.

### LINE NUMBER

For line mode, this is the line number for that line. For function mode, this is the line number of the first line of the function. *PGPROF* sometimes generates multiple statements for a single source line; thus multiple profiling entries might appear for a single source line. To distinguish them, *PGPROF* uses the notation: *lineNo.statementNo*

### MESSAGES

For MPI profiles only. This is the number of messages sent and received by the function or line.

### RECEIVES

For MPI profiles only. This is the number of messages received by the function or line.

### SENDS

For MPI profiles only. This is the number of messages sent by the function or line.

### TIME

The time spent only within the function or executing the line. The *TIME* does not include time spent in functions called from this function or line. *TIME* may be displayed in seconds or as a percent of the total time.

## Caveats (Precision of Profiling Results)

### Accuracy of Performance Data

The collection of performance data always introduces some overhead, or intrusion, that can affect the behavior of the application being monitored. How this overhead affects the accuracy of the performance data depends on the performance monitoring method chosen, system software and hardware attributes, the load on the system during data collection, and the idiosyncrasies of the profiled application. Although the *PGPROF* implementation attempts to minimize intrusion and maximize accuracy, it would be unwise to assume the data is beyond question.

## Clock Granularity

Many target machines provide a clock resolution of only 20 to 100 ticks per second. Under these circumstances, a routine must consume at least a few seconds of CPU time to generate meaningful line level times.

## Source Code Correlation

At higher optimization levels, and especially with highly vectorized code, significant code reorganization may occur within functions. The *PGPROF* profiler allows line profiling at any optimization level. In some cases, the correlation between source and data may at times appear inconsistent. Compiling at a lower optimization level or examining the assembly language source may help you interpret the data in these cases.



# Chapter 8. PGPROF Reference

This chapter provides a reference guide to the features of the *PGPROF* performance profiler.

For information about how to invoke *PGPROF*, refer to [“Profiler Invocation and Initialization,” on page 5](#).

For information about using the *PGPROF* text-based command-line interface, refer to [Chapter 3, “Compiler Options for Profiling”](#).

For information about how to choose a profiling method, build your program, and execute it to collect profile data, refer to [Chapter 1, “Getting Started”](#).

## PGPROF User Interface Overview

On startup, *PGPROF* attempts to load the profile datafile specified on the command line or the default, `pgprof.out`. If no file is found, a file chooser dialog box is displayed. Choose a profile datafile from the list or select Cancel.

When a profile datafile is opened, *PGPROF* populates the user interface, as illustrated and labeled in [Figure 8.1](#).

### Menu Bar

Contains these menus: File, Edit, View, Sort, and Help.

### Toolbar

Provides navigation shortcuts and controls for frequently performed operations.

### Statistics Table

Displays profile summary information for each profile entry. Information can be displayed at up to three levels - routine, line, or assembly - depending on the type of profile data collected, how the program was built, and whether the *PGPROF* source file search path has been set to include the program source directories. The initial view is the routine level view.

### Focus Panel

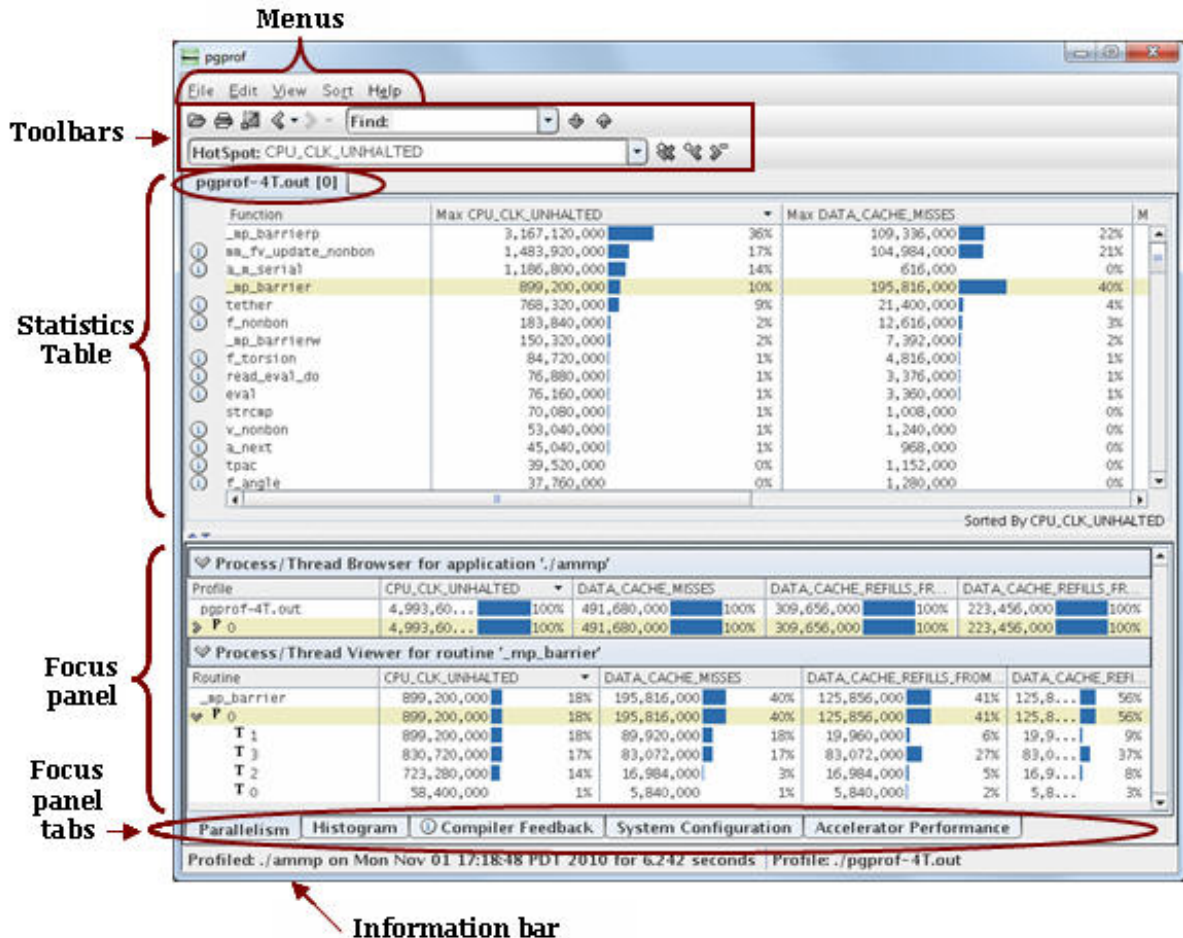
Consists of tabbed panes labeled Parallelism, Histogram, Compiler Feedback, System Configuration, and Accelerator Performance.

### Information Bar

Displays the profile summary information such as the name of the executable, the time and date of the profile run, execution time, number of processes, if more than one, and the datafile name.

The following sections describe each of these components in more detail.

Figure 8.1. PGPROF User Interface



## PGPROF Menus

PGPROF had the following menus: File, Edit, View, Sort, and Help. This section describes each menu in detail. Keyboard shortcuts, when available, are listed next to menu items.

### File Menu

The File menu contains the following items:

- **New Window** (control N) – Select this option to create a copy of the current profiler window on your screen.
- **Open Profile...** – Select this option to begin analyzing a different profile. When you see the dialog box, fill in or browse to the information requested about the profile data file (default `pgprof.out`), the

executable file, and the location of the source files. When you click OK, a new profile session is started using the information specified in the dialog box.

If the Source Path is the only parameter that is changed from current session parameters, then the current session uses the new Source Path to search for sources.

- **Set Source Directory...** – Select this option to add or remove a directory in the source file search path.
- **Scalability Comparison...** – Select this option to open another profile for scalability comparison. As you did for the `Open Profile...` option described above, provide information about the profile data file, the executable file, and the location of the source files. Notice that the new profile contains a Scale column in its Statistics table.

### Note

---

Another method to open profiles for scalability comparison is by using the `-scale` command-line option explained in [“Profiler Invocation and Initialization,” on page 5](#).

For more information on scalability, refer to [“Scalability Comparison,” on page 19](#).

- **Print...** – Select this option to make a hard copy of the current profile data. The profiler processes data from the Statistics table and sends the output to a printer. A printer dialog box appears.

You can select a printer using the Name drop-down list under Print Service. Alternately, click the Print To File check box to send the output to a file. Other print options may be available; however, they are dependent on the specific printer and the Java Runtime Environment (JRE).

- **Print to File...** – Option, output is not sent to printer, but is formatted as an editable text file. After selecting this menu item, a Save File dialog box appears. Enter or choose an output file in the dialog box. Click Cancel to abort the print operation.
- **Close...** – Select this option to close the current profiling session. This option is enabled only when more than one profile is open.
- **Exit...** – Select this option to end the profiling session and exit the profiler.

## Edit Menu

Use the Edit menu to launch a text search in the Statistics Table, and to restore, revert or save user preference settings. This menu contains the following items:

- **Search Forward...** – Displays a dialog box that prompts for the text to be located. Once the text is entered and the OK button selected, PGPROF searches forward to the next occurrence of the text in the function list, source code, or assembly code displayed in the Statistics Table. Matching text is displayed in red. A search can also be invoked using the Find text box on the main toolbar.
- **Search Backward...** – Displays a dialog box that prompts for the text to be located. Once the text is entered and the OK button selected, PGPROF searches backward to the previous occurrence of the text in the function list, source code, or assembly code displayed in the Statistics Table. Matching text is displayed in red.
- **Search Again** – Use this option to repeat the last search.

- **Clear Search** – Use this option to clear the search and turn the color of all matching text back to black.
- **Restore Default Settings...** – Use this option to restore the configuration of the user interface to the original default settings.
- **Revert to Saved Settings...** – Use this option to restore the configuration of the GUI to the previously saved settings. For more information, refer to the See the Save Settings on Exit option.
- **Save Settings on Exit...** – When this check box is selected, PGPROF saves the current GUI configuration settings on exit. These settings include the size of the main window, position of the horizontal dividers, the bar chart colors, the selected font, the tool tips preference, and the options selected in the View menu. When PGPROF is started again, these saved settings are used. To prevent saving these settings on exit, clear this check box. On Linux and Mac OS, settings are saved on a per-user basis. On Windows, settings are saved on a per-user per-system basis.

### Note

---

You can also use the Find: box in the toolbar to invoke the *PGPROF* search facility.

## View Menu

Use the View menu to change the configuration of the *PGPROF* user interface. This menu contains the following items:

- **Select Columns...** - Invokes a dialog box that allows you to select which columns of the Statistics Table are to be displayed, and how to display the data in the columns.

The choices for how to display the data are: Value, Percent, Bar, or All, though not all of these choices are available for all columns.

- **Select Graph Colors...** – This menu option opens a color chooser dialog box and a bar chart preview panel.

The preview panel contains the bar chart bar colors, and the three bar chart attributes.

- The bar chart bars can be 'gradient filled', meaning that the color of the bar gradually transitions from the Bar Start Color to the Bar End Color. To have solid colored bars without gradient fill, which is the default, simply set both of these colors to the same color.
- The Filled Text Color attribute represents the text color inside the filled portion of the bar chart.
- The Unfilled Text Color attribute represents the text color outside the filled portion of the bar chart.
- The Background Color attribute represents the color of the unfilled portion of the bar chart.
- The Reset button allows you to reset the selected bar chart or attribute to its previously selected color.
- The OK button accepts your changes and closes the dialog box.

### Note

---

Closing the dialog box is the same as choosing OK.

To modify a bar chart or attribute color:



1. Click the radio button.
2. Choose a color from the Swatches, HSB, or RGB pane.
3. Click the OK button to accept the changes and close the dialog box.

*PGPROF* saves color selections for subsequent runs unless the Save Settings on Exit box is unchecked, as described later in this section.

- **Font...** – This menu option opens the Fonts dialog box.

You can change the font and/or font size using this dialog's drop-down lists. As you change the font, you can preview the changes in the Sample Text pane.

To change the font you must click the OK button.

Click Cancel or close the dialog box to abort any changes.

- **Show Tool Tips** - Select this check box to enable tool tips. *Tool tips* are small temporary messages that pop-up when the mouse pointer is positioned over a component, such as a button, in the user interface. Tool tips provide a summary or hint about what a particular component does. Clear this check box to turn tool tips off.

## Sort Menu

Use the Sort menu to change the metric used to sort profile entries. The current sort order is displayed at the bottom of the Statistics Table and the Histogram tab.

The default sorting metric is time for function-level profiling and source line number for line-level profiling. The sort is performed in descending order, from highest to lowest value, except when sorting by filename, function name, or line number. Filename, function name, and line number sorting is performed in ascending order; lowest to highest value. Sorting is explained in greater detail in [“Sorting Profile Data,” on page 13](#).

## Help Menu

The Help menu contains the following items:

- PGPROF Help... – This option invokes *PGPROF*'s integrated help utility. The help utility includes an HTML version of this manual. To find a help topic, use one of the tabs in the left panel:
  - The *book* tab presents a table of contents.
  - The *index* tab presents an index of commands.
  - 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.

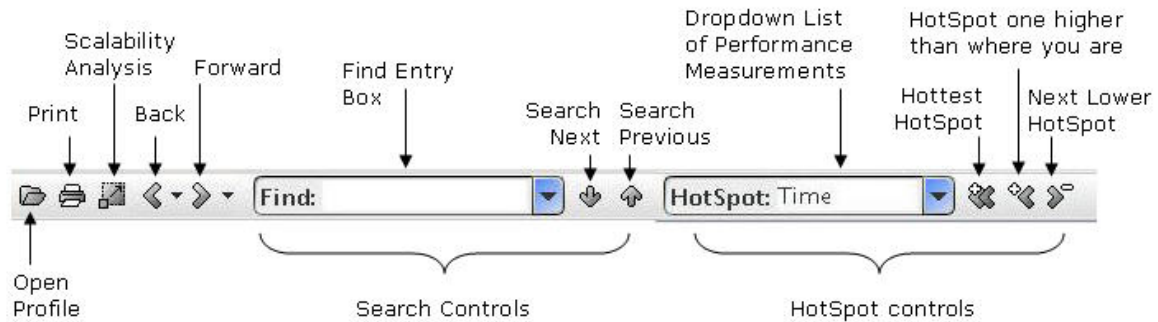
Use the printer buttons to print the current help page.

- About PGPROF... – This option opens a dialog box with version and contact information for *PGPROF*.

## PGPROF Toolbar

As illustrated in [Figure 8.2](#), the *PGPROF* toolbar provides navigation shortcuts and controls for frequently performed operations.

Figure 8.2. PGPROF Toolbar



The toolbar includes these buttons and controls:

- **Open Profile** button – clicking this button is the same as selecting File | Open Profile... from the menu bar.
- **Print** button – clicking this button is the same as selecting File | Print... from the menu bar.
- **Scalability Analysis** button – clicking this button is the same as selecting File | Scalability Comparison... from the menu bar.
- **Forward** and **Back** buttons – click these buttons to navigate forward and back to previous and subsequent views, respectively.

Use the down-arrow to display the full list of views, and to select a view to jump to. These lists use a notation to describe the profile views as follows:

```
profile_data_file@source_file@routine@line@address
```

The address field is omitted for line-level views, and both the line and address fields are omitted for routine-level views. For example, the following item in a list would describe a view that uses profile data from `pgprof.out`, and is displaying line 370 in the routine named `solver` in source file `main.f`.

```
pgprof.out@main.f@solver@370
```

- **Search** controls – use these to locate information. The controls include:
  - A text box labeled *Find:*. Entering a search string here and hitting Enter is the same as using the dialog box invoked from the Edit | Search Forward... menu bar item.
  - Two buttons labeled with down and up arrows, respectively. These buttons provide Search Next and Search Previous operations, similar to Edit | Search Again. Search Next searches for the next occurrence of the last search string below the current location, and Search Previous searches for the next occurrence above.

- **HotSpot Navigation** controls – use these to navigate to the most significant measurements taken in the profiling run. The controls include:
  - A drop-down menu labeled *HotSpot:*, which you use to select the specific performance measurement of interest.
  - Three navigation buttons, containing *Forward* and *Back* icons with associated plus (+) and minus (-) signs.

When the profile is first displayed, the Statistics Table selects the row for the routine with the highest measured Time as though you had clicked on that row. To navigate to the row with the next-highest Time, you click on the button labeled with the *Forward* icon and the minus (-) sign, denoting the next Time HotSpot lower than the current one. Once you have navigated to this second HotSpot, the *Back* HotSpot buttons are activated, allowing you to navigate to the hottest HotSpot using the "<<" button, or to the next higher Time, using the "<" button.

You can use the HotSpot drop-down menu to change the measurement used to identify the HotSpots. The default selection in the HotSpot menu is *Time*, assuming that Time is one of the available measurements. You can click on the down-arrow in the drop-down menu to select any other metric listed in the menu, then click the "Hottest" button to navigate to the row showing the routine with the highest measured value for that metric.

## PGPROF Statistics Table

This section describes the *PGPROF* Statistics Table. The Statistics Table displays an overview of the performance data, and correlates it with the associated source code or assembly instructions. This is where you should start when analyzing performance data with *PGPROF*.

The Statistics Table displays information at up to three levels, depending on the type of profile data collected, how the program was built, and whether the *PGPROF* source file search path has been set to include the program source directories.

### Performance Data Views

The Statistics Table allows you to zoom in and out on the components of your program by providing several views: the routine-level view, the line-level view, and the assembly-level view.

- The initial view when you invoke *PGPROF* is the routine-level view.
- To navigate to the line level from the routine level, double-click on the Statistics Table row corresponding to the function of interest. If the program was built so that it does not contain line location information, then this action results in an assembly-level display.
- To navigate to the assembly code level from the line level, click the assembly code icon, the plus (+) symbol, on the Statistics Table row that corresponds to the source line of interest.

#### Note

---

You can use the View | Select Columns... menu option to select the data shown in the Statistics Table.

## Routine-level view

The routine-level view shows a list of the functions or subprograms in your application, with the performance data for that routine in the same row of the table. In addition, if there is any compiler feedback information for the routine, a round button containing the letter 'i' is at the far left of the row. Clicking that button populates the Compiler Feedback tab with the compiler feedback relating to that routine.

## Line-level View

You access the line-level view of a routine by clicking that routine's row in the routine-level view. *PGPROF* opens a new tab showing the line-level information for the routine. The tab label is the routine name and the tab contains an x which allows you to close the tab when you are done viewing the source code. The Statistics Table in the new tab shows the source code for the selected function, with performance data and Compiler Feedback buttons as with the routine-level view.

## Assembly-level View

You access the assembly-level view of a source line or routine by clicking the assembly code icon, the plus (+) symbol, on the Statistics Table row that corresponds to the row of interest in the line-level view. The table changes to show the assembly code, interspersed with the source lines that were compiled to generate the code.

## Source Code Line Numbering

In the optimization process, the compiler may reorder the assembly instructions such that they can no longer be associated with a single line. Therefore, for optimized code, a *source line* may actually be a code block consisting of multiple source lines. This occurrence is common, and expected, and should not interfere with the tuning process when using *PGPROF*.

*PGPROF* sometimes shows multiple rows in the Statistics Table for a single source line. The line numbers for such lines are shown in the Statistics Table using the notation

```
line.statement
```

There are several situations where this line numbering can occur:

- When there is more than one statement in a source line, as in a C/C++ program where one line contains multiple statements, separated by semicolons (;).
- When the compiler generates multiple alternative implementations of a loop. The compiler may create alternate versions to handle differences in the data and how it is stored in memory.
- When there is a complicated or conditional loop setup.

For these cases, it is generally safe to sum the times and counts of all the lines. However, take care, particularly with call counts, not to double-count measurements.

## PGPROF Focus Panel

The Focus Panel consists of a number of tabs that allow you to select more detailed views of your profile data.

Figure 8.3. Focus Panel Tabs



### Parallelism tab

This tab displays a table with detailed profile information organized by processes and threads. Profile information for the entire application is labeled 'Profile' while information for the currently-selected routine is labeled 'Routine.' Information is listed by process. Each process can be expanded to reveal profiling information by each thread in that process. To expand a process into its threads, click on the '>' icon on the left of the 'P' icon.

### Histogram tab

This tab displays a histogram of one or more profiled data items.

- Each bar graph corresponds to one of the performance measurements.
- Each vertical bar corresponds to a profile entry, that is, performance data associated with a program location.
- The bars are sorted in the order specified in the Sort menu, described in [“Sort Menu,” on page 49](#), and the current sorting metric is labeled in the lower-right hand corner of the table itself.
- Clicking on a bar displays information for the corresponding profile item in the Statistics Table.
- Double-clicking on a bar drills down into the profile for the portion of the program corresponding to the bar.
- Selected bars are highlighted in yellow.

### Compiler Feedback tab

This tab displays information provided by the compiler regarding the characteristics of a selected piece of the program, including optimization and parallelization information, obstacles to optimization or parallelization, and hints about how to improve the performance of that portion of the code. Such information is available at the line level and the routine level.

If Compiler Feedback information is available, round, blue buttons, containing a lower-case 'i', are displayed on the left side of the Statistics Table. To access the information, click on one of these *info* buttons.

The information is separated into categories of information about these items:

- A source line
- Variables referenced inside a routine
- Routines referenced inside another routine
- How a file was compiled

Each category is represented by a wide bar that functions like a button. Clicking the bar expands the display to show the information in that category. If no information is available in a given category, that category is not listed.

This information is only available if the program was compiled and also linked using either the `-Mprof` or the `-Minfo=cff` option. In some cases it is necessary to build programs using profile-guided optimization by building with `-Mpfi` or `-Mpfo`, as described in the section *Profile-Feedback Optimization using -Mpfi/Mpfo* in the Optimizing and Parallelizing chapter of the PGI User's Guide. Profile-guided optimization can often determine loop counts and other information needed to calculate the Compute Intensity for a given statement or loop.

## System Configuration tab

This tab displays System and Accelerator tabs containing information about the system on which the profile run was executed.

### System Tab

Can include information such as process(es), process manufacturer, processor model, processor, the program's OS target, cores per socket, total cores, processor frequency, CUDA driver version, and NVRM version.

#### Note

---

The *Program's OS Target* is the operating system platform that the executable was built for. Although the processor may be a 64-bit processor, the executable may target a 32-bit platform.

#### Tip

---

If you need further explanations for any of these items, refer to vendor processor documentation.

### Accelerator tab

Contains information about the GPU(s) that are on the system on which the profile was run.

If there is no GPU on this system, the Accelerator tab is empty.

For each GPU, also known as a device, the Accelerator tab can include information such as the device name, device revision number, global memory set, number of multiprocessors, number of cores, concurrent copy and execution, total constant memory, total shared memory per block, registers per block, warp size, maximum threads per block, maximum block dimensions, maximum grid dimensions, maximum memory pitch, texture alignment, and clock rate.

#### Tip

---

If you need further explanations for any of these items, refer to vendor GPU documentation.

## Accelerator Performance tab

This tab displays profiling information collected by **pgcollect** on for programs built using CUDA Fortran or the PGI Accelerator Model. For more information on **pgcollect**, refer to [Chapter 10, "pgcollect Reference"](#).

### Accelerator Model Profiles

The profiling information is relative either to an Accelerator Region or to an Accelerator Kernel.

### Accelerator Region

An accelerator region is a region of code that has been executed on the accelerator device. An accelerator region might transfer data between the host and the accelerator device. Further, an accelerator region can be split into several accelerator kernels.

### Accelerator Kernel

An accelerator kernel is a compute intensive, highly parallel portion of code executed on an accelerator device. Each compiler-generated kernel is code executed by a block of threads mapped into a grid of blocks.

Figure 8.4 illustrates one possible display for the Accelerator Performance tab, one that is relative to the Accelerator Kernel:

Figure 8.4. Accelerator Performance tab of Focus Panel

Device Number	0	
Accelerator Kernel Execution Count	4	
Grid Size		[63x63]
Block Size		[16x16]
Maximum time spent in Accelerator Kernel (secs)	0.000216	
Minimum time spent in Accelerator Kernel (secs)	0.000206	
Average time spent in Accelerator Kernel (secs)	0.000209	

Parallelism	Histogram	① Compiler Feedback	System Configuration	Accelerator Performance
-------------	-----------	---------------------	----------------------	-------------------------

Profiled: ./mm2 on Tue Mar 29 12:34:07 PDT 2011	Profile: ./pgprof.out
---	-----------------------

PGPROF displays two Accelerator events in the Statistic table:

- *Accelerator Region Time* – the time, in seconds, spent in the Accelerator region
- *Accelerator Kernel Time* – the time, in seconds, spent in the Accelerator kernel.

When a user selects a line for which one of these events is non-zero, the table in the Accelerator Performance tab contains details about that event. The information displayed depends on the selection.

If a user selects a line in which both events are non-zero, then the Accelerator Performance tab displays only Accelerator Initialization Time, Accelerator Region Time, and Accelerator Kernel Time.

### Accelerator Region Timing Information

Time is reported in seconds. When you select a non-zero Accelerator Region Timing item, you see the following information in the Accelerator Performance tab:

- *Accelerator Initialization Time* – time spent in accelerator initialization for the selected region.
- *Accelerator Kernel Time*– time spent in compute kernel(s) for the selected region.
- *Data Transfer Time*– time spent in data transfer between host and accelerator memory.
- *Accelerator Execution Count*– execution count for the selected region.
- *Maximum time spent in accelerator region (w/o init)*– the maximum time spent in a single execution of selected region.
- *Minimum time spent in accelerator region (w/o init)*– the minimum time spent in a single execution of selected region.

- *Average time spent in accelerator region (w/o init)* – the average time spent per execution of selected region.

### Note

---

The table does not contain values that are not relevant, such as zero values or values that cannot be computed. For example, in a routine-level profile, a routine can execute multiple accelerator regions. In this instance, only time spent in Initialization, in the Region, and in the Kernel can be accurately computed so other values are not displayed in the Accelerator Performance tab.

### Accelerator Kernel Timing Information

Time is reported in seconds. When you select a non-zero Accelerator Kernel Timing item, you see the following information in the Accelerator Performance tab:

- *Kernel Execution Count* – execution count for the selected kernel.
- *Grid Size* – the size, in 1D [X] or 2D [XxY], of the grid used to execute blocks of threads for the selected kernel.
- *Block Size* – the size, in 1D [X], 2D [XxY] or 3D [XxYxZ], of the thread blocks for the selected kernel.
- *Maximum time spent in accelerator kernel* – the maximum time spent in a single execution of selected kernel.
- *Minimum time spent in accelerator kernel* – the minimum time spent in a single execution of selected kernel.
- *Average time spent in accelerator kernel* – the average time spent per execution of selected kernel.

### Note

---

When there are multiple invocations of the same kernel in which the grid-size and/or block-size changes, the size information displayed in the Accelerator Performance tab is expressed as a range. For example, if the same kernel could be executed with a 2D-block of size [2,64] and a 2D-block of size [4,32], then the size displayed in Accelerator Performance tab is the range: [2-4, 32-64].

### CUDA Fortran Profiles

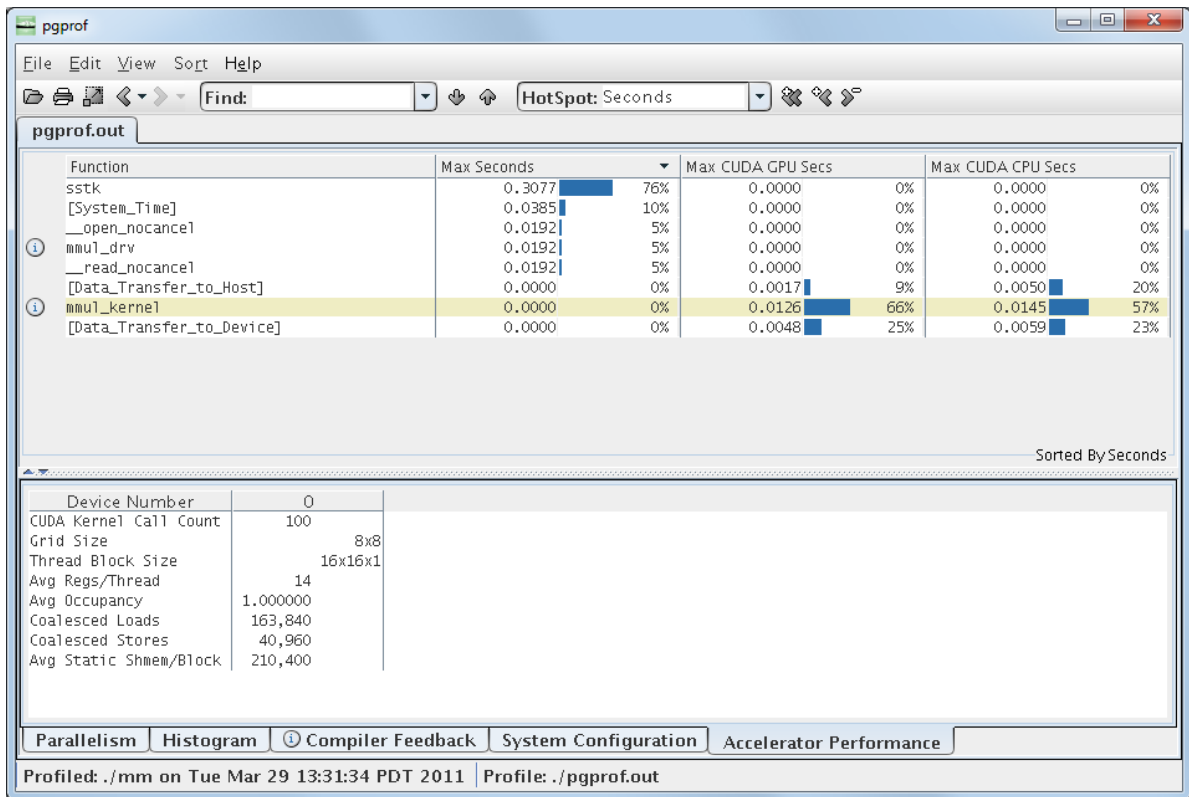
Profiles generated by **pgcollect** for CUDA Fortran programs capture data from GPU performance counters. The specific counters available for a given GPU depend on the GPU's compute capability.

In [Figure 8.5](#):

- The columns labeled Max CUDA GPU Secs and Max CUDA CPU Secs show times captured by the CUDA driver.
- The Max Seconds column contains timings for host-only code.
- Pseudo-function names [Data\_Transfer\_To\_Host] and [Data\_Transfer\_To\_Device] show the transfer times to and from the GPU.
- The Accelerator Performance Tab shows counter values collected from the GPU.



Figure 8.5. CUDA Program Profile





# Chapter 9. Command Line Interface

The command line interface (CLI) for non-GUI versions of the *PGPROF* profiler is a simple command language. This command language is available in the profiler through the `-text` option. The language is composed of commands and arguments separated by white space. A **pgprof>** prompt is issued unless input is being redirected.

This chapter describes *PGPROF*'s command line interface, providing both a summary and then more details about the commands.

## Command Description Syntax

This chapter describes the profiler's command set.

- Command names are printed in bold and may be abbreviated as indicated.
- Arguments enclosed by brackets ('[*...*']') are optional.
- Separating two or more arguments by '|' indicates that any one is acceptable.
- Argument names in italics are chosen to indicate what kind of argument is expected.
- Argument names that are not in italics are keywords and should be entered as they appear.

## PGPROF Command Summary

[Table 9.1](#) summarizes the commands for use in the CLI version of *PGPROF*, providing the applicable arguments and a brief description of the use of each command. The section that follows the table provides more details about each command.

Table 9.1. PGPROF Commands

Name	Arguments	Usage
a[sm]	routine [[>] filename]	Display the instruction and line level data together with the source and assembly for the specified routine.
c[cff]	file[@function] [line_num]	Display compiler feedback for the specified file, function, or source line

Name	Arguments	Usage
d[isplay]	[display options]   all   none	Specify display information.
he[lp]	[command]	Provide brief command synopsis.
h[istory]	[size]	Display the history list, which stores previous commands in a manner similar to that available with csh or dbx.
l[ines]	function [[>] filename]	Display the line level data together with the source for the specified function.
lo[ad]	[datafile]	Load a new dataset. With no arguments reloads the current dataset.
m[erge]	datafile	Merge the profile data from the named datafile into the current loaded dataset.
pro[cess]	processor_num	For multi-process profiles, specify the processor number of the data to display.
p[rint]	[[>] filename]	Display the currently selected function data.
q[uit]		Exit the profiler.
sel[ect]	calls   timecall   time   cost   cover   all [[>] cutoff]	Display data for a selected subset of the functions.
so[rt]	[by] [max   avg   min   proc   thread] calls   cover   timecall   time   cost   name   msgs   msgs_sent   msgs_rcv   bytes   bytes_sent   bytes_rcv   visits   file]	Function level data is displayed as a sorted list.
src[dir]	directory	Set the source file search path.
s[tat]	[no]min [no]avg [no]max [no]procl [no]thread [no]all]	Set which process fields to display (or not to display when using the arguments beginning with “no”)
th[read]	thread_num	Specify a thread for a multi-threaded process profile.
t[imes]	raw   pct	Specify whether time-related values should be displayed as raw numbers or as percentages. The default is pct.
! (history)	!! num   -num   string	Repeat recent commands

## Command Reference

This section provides more details about the commands in the previous Command Summary Table.

### asm

```
a[sm] routine [[>] filename]
```

Display the instruction and line level data together with the source and assembly for the specified routine. If the filename argument is present, the output is placed in the named file. The '>' means redirect output, and is optional. This command is only available on platforms that support assembly-level profiling.

## ccff

```
c[ccff] file[@function] [line_number]
```

Display compiler feedback for the specified file, function, or source line. PGI compilers can produce information in the Common Compiler Feedback Format (CCFF) that provides details about the compiler's analysis and optimization of your program. Often this information can illuminate ways in which to further optimize a program.

The CCFF information is produced by default when using the `-Mprof` compiler option, but if you are profiling with the **pgcollect** command, you must build your program with the `'-Minfo=ccff'` compiler option to produce this information.

## display

```
d[isplay] [display options] | all | none
```

Specify display information. This includes information on minimum values, maximum values, average values, or per processor/thread data. Below is a list of possible display options:

```
[no]calls [no]cover [no]time [no]timecall [no]cost [no]proc [no]thread [no]msgs [no]msgs_sent  
[no]msgs_recv [no]bytes [no]bytes_sent [no]name [no]file [no]line [no]lineno [no]visits [no]scale  
[no]stmtno
```

## help

```
he[lp] [command]
```

Provide brief command synopsis. If the command argument is present, only information for that command is displayed. The character "?" may be used as an alias for help.

## history

```
h[istory] [size]
```

Display the history list, which stores previous commands in a manner similar to that available with `csh` or `dbx`. The optional size argument specifies the number of lines to store in the history list.

## lines

```
l[ines] function [[>] filename]
```

Display the line level data together with the source for the specified function. If the filename argument is present, the output is placed in the named file. The '>' means redirect output, and is optional.

## load

```
lo[ad] [datafile]
```

Load a new dataset. With no arguments reloads the current dataset. A single argument is interpreted as a new data file. With two arguments, the first is interpreted as the program and the second as the data file.

## merge

```
m[erge] datafile
```

Merge the profile data from the named datafile into the current loaded dataset. The datafile must be in standard pgprof.out format, and must have been generated by the same executable file as the original dataset (no datafiles are modified.)

## process

```
pro[cess] processor_num
```

For multi-process profiles, specify the processor number of the data to display.

## print

```
p[rint] [[>] filename]
```

Display the currently selected function data. If the filename argument is present, the output is placed in the named file. The '>' means redirect output, and is optional.

## quit

```
q[uit]
```

Exit the profiler.

## select

```
sel[ect] calls | timecall | time | cost | cover | all [[>] cutoff]
```

Display data for a selected subset of the functions. This command is used to set the selection key and establish a cutoff percentage or value. The cutoff value must be a positive integer, and for time related fields is interpreted as a percentage. The '>' means greater than, and is optional. The default is all.

## sort

```
so[rt] [by] [max | avg | min | proc | thread] calls | cover | timecall | time |  
cost | name | msgs | msgs_sent | msgs_recv | bytes | bytes_sent | bytes_recv |  
visits | file]
```

Function level data is displayed as a sorted list. This command establishes the basis for sorting. The default is max time.

## srcdir

```
src[dir] directory
```

Set the source file search path.

## stat

```
s[tat] [no]min|[no]avg|[no]max|[no]proc|[no]thread|[no]all]
```

Set which process fields to display (or not to display when using the arguments beginning with “no”).

## thread

```
th[read] thread_num
```

Specify a thread for a multi-threaded process profile.

## times

```
t[imes] raw | pct
```

Specify whether time-related values should be displayed as raw numbers or as percentages. The default is pct.

## ! (history)

```
!!
```

Repeat previous command.

```
! num
```

Repeat previous command numbered num in the history list.

```
!-num
```

Repeat the num-th previous command numbered num in the history list.

```
! string
```

Repeat most recent command starting with string from the history list.





# Chapter 10. pgcollect Reference

The **pgcollect** command is a development tool used to collect performance data for analysis using the **pgprof** performance profiler. This chapter describes how to use **pgcollect**.

## pgcollect Overview

**pgcollect** runs the specified program with the supplied arguments. While the program runs, **pgcollect** gathers performance statistics. When the program exits, the data that is gathered is written to a file. You can then use this file in the *PGPROF* performance profiler to analyze and tune the performance of the program.

The **pgcollect** command supports two distinct methods of performance data collection:

### Time-based sampling

Creates a time-based profile that correlates execution time to code, showing the amount of time spent in each routine, each source line, and each assembly instruction in the program. For more information on time-based profiling, refer to [“Time-based Sampling,” on page 3](#).

### Event-based sampling

Supported only on linux86-64 systems, creates an event-based profile that correlates hardware events to program source code. In this method, **pgcollect** uses hardware event counters supported by the processor to gather resource utilization data, such as cache misses.

### Note

---

This method requires co-installation of the open source performance tool *OProfile*.

For more information on event-based profiles, refer to [“Event-based Sampling,” on page 3](#).

Both forms of the **pgcollect** command gather performance data that can be correlated to individual threads, including OpenMP threads, as well as to shared objects, dynamic libraries, and DLLs.

For current availability of **pgcollect** and **pgcollect** features on a given platform, refer to the PGI Release Notes.

## Invoke pgcollect

The command you use to invoke **pgcollect** depends on the type of profile you wish to create.

Use the following command to invoke **pgcollect** for time-based sampling:

```
pgcollect [-time] program [program_args]
```

Use the following command to invoke **pgcollect** for event-based sampling available on Linux86-64:

```
pgcollect [<event_options>] program_or_script [program_or_script_args]
```

`program` or `program_or_script` are either the filename of the program to be profiled, or the name of a script that invokes the program. When applicable, you can provide arguments for the specified program or script: *program\_args* or *program\_or\_script\_args*.

The following sections describe the **pgcollect** command-line options in more detail.

## Build for pgcollect

If your program was built with PGI compilers, you do not need to use any special options to use **pgcollect**. However, if your programs are built using the `-Minfo=ccff` option, then *PGPROF* can correlate compiler feedback and optimization hints with the source code and performance data.

If you built your program using a non-PGI compiler, consider building with debugging information so you can view source-level performance data. Be aware, however, that building with debugging information may change the performance of your program.

## General Options

This section describes options that apply to all forms of the **pgcollect** command. For options specific to controlling time-based or event-based profiling, refer to “[Time-Based Profiling Options](#)” and “[Event-based Profiling Options](#)” respectively.

**-V**

Display the version of **pgcollect** being run.

**-help**

Show **pgcollect** usage and switches.

## Time-Based Profiling

Time-based profiling runs the program using time-based sampling. This form of **pgcollect** uses operating system facilities for sampling the program counter at 10-millisecond intervals.

### Time-Based Profiling Options

**-time**

Provide time-based sampling only. The sampling interval is 10 milliseconds. This option is the default.

When using **pgcollect** for time-based sampling, you can have multiple instances of **pgcollect** running simultaneously, but doing so is not recommended, since this will probably skew your performance results.

## Event-Based Profiling

You can use the **pgcollect** command on linux86-64 to drive an OProfile session. Event-based profiling provides several predefined data collection options that gather data from commonly used counters.

For event-based sampling, the only required argument is the `program_or_script`, which is either the filename of the program to be profiled, or the name of a script that invokes the program. Using a script can be useful if you want to produce an aggregated profile of several invocations of the program using different data sets. In this situation, use the `-exe` option, which allows the data collection phase to determine which program is being profiled.

When applicable, you can provide arguments for the specified program or script.

Since OProfile provides only system-wide profiling, when you invoke **pgcollect** it provides a locking mechanism that allows only one invocation to be active at a time.

### Note

The **pgcollect** locking mechanism is external to OProfile and does not prevent other profile runs from invoking **opcontrol** through other mechanisms.

## Root Privileges Requirement

When using **pgcollect** for event-based profiling, you control the OProfile kernel driver and the sample collection daemon via the OProfile command **opcontrol**. This control requires root privileges for management operations. Thus, invocations to **opcontrol** performed by **pgcollect** are executed via the **sudo** command.

When using **pgcollect**, you control the OProfile kernel driver and the sample collection daemon via the OProfile command **opcontrol**. This control requires root privileges for management operations. Thus, invocations to **opcontrol**, which are performed when **pgcollect** is used, are executed via the **sudo** command.

One technique that requires minimal updates to the `/etc/sudoers` files is to assume that all users in a group are allowed to execute **opcontrol** with group privileges. For example, you could make the following changes to `/etc/sudoers` to permit all members of the group 'sw' to run **opcontrol** with root privileges.

```
# User alias specification
User_Alias SW = %sw
...
SW ALL=NOPASSWD: /usr/bin/opcontrol
```

## Interrupted Profile Runs

**pgcollect** shuts down the OProfile daemon when interrupted. However, if the script is terminated with **SIGKILL**, you must execute the following:

```
pgcollect -shutdown
```

Executing this command is important because if the OProfile daemon is left running, disk space on the root file system eventually is exhausted.

## Event-based Profiling Options

### -check-events

Do not execute a profiling run, just check the event settings specified on the command line.

### -exe <exename>

Specify the program to be profiled. Use the `-exe` option only when the *program\_or\_script* argument is a script that invokes the program.

### -list-events

List profiling events supported by the system.

### -shutdown

Shut down the profiling interface. You only need to use this option in rare cases when a profiling run was interrupted and OProfile was not shut down properly.

## Predefined Performance Data Collection Options

### -allcache

Profile instruction, data, and branch cache misses

### -dcache

Profile various sources of data cache misses

### -imisses

Profile instruction cache-related misses.

### -hwtime <millisecs>

Provide time-based sampling only. Specify the sampling interval in milliseconds.

## User-Defined Performance Data Collection Options

### -es-function <name>

Set profile events via a shell function.

### -event <spec>

Manually add an event profile specification. An event profile specification is an **opcontrol** `--event` argument; that is, the event profile specification provided on the command line is appended to `--event=` and passed as an argument to **opcontrol**.

### -post-function <name>

Execute a shell function after profiling is complete.

## Defining Custom Event Specifications

The `pgcollect --event=EVENTSPEC` options are accumulated and used to specify events to be measured. For more information about these events, refer to the **opcontrol** man page.

x64 processors provide numerous event counters that measure the usage of a variety of processor resources. Not all processors support the same set of counters. To see which counters are supported on a given system, use the following command:

```
pgcollect -list-events
```

The output of this command also provides information on event masks (the hex value in the event specification) and minimum overflow values.

Here are two examples of shell functions providing event specifications to pgcollect. These functions would be implemented in a `.pgoprun` file:

### Example 10.1. Custom Event Example 1

This function specifies the events needed to calculate cycles per instruction (`CPU_CLK_UNHALTED` / `RETIRED_INSTRUCTIONS`). The fewer cycles used per instruction, the more efficient a program is.

```

cpi_data () {
    event[${#event[@]}]--event=CPU_CLK_UNHALTED:500000:0x00:0:1
    event[${#event[@]}]--event=RETIRED_INSTRUCTIONS:500000:0x00:0:1
}

```

To use these events, invoke **pgcollect** with the following arguments:

```
-es-function cpi_data
```

### Example 10.2. Custom Event Example 2

#### Example 2

This function specifies events needed to determine memory bandwidth:

```

mem_bw_data () {
    event[${#event[@]}]--event=CPU_CLK_UNHALTED:500000:0x00:0:1
    event[${#event[@]}]--event=SYSTEM_READ_RESPONSES:500000:0x07:0:1
    event[${#event[@]}]--event=QUADWORD_WRITE_TRANSFERS:500000:0x00:0:1
    event[${#event[@]}]--event=DRAM_ACCESSES:500000:0x07}:0:1
}

```

To use these events, invoke **pgcollect** with the following arguments:

```
-es-function mem_bw_data
```

## PGI Accelerator Model and CUDA Fortran Profiling

If you are profiling a program that uses the PGI Accelerator model or CUDA Fortran, **pgcollect** automatically collects information for you.

### Accelerator Model Profiling

**pgcollect** automatically collects and includes performance information for the PGI Accelerator model programs in the profile output for the program.

#### Note

Inclusion of the accelerator performance information in the program's profile output occurs for both time-based sampling and, on Linux, for event-based sampling.

## CUDA Fortran Program Profiling

If you are profiling a program that uses CUDA Fortran kernels running on a GPU, **pgcollect -cuda** collects performance data from CUDA-enabled GPUs and includes it in the profile output for the program. The syntax for this command option is:

```
-cuda [=gmem | branch | cfg:<cfgpath> | cc13 | cc20 | list]
```

The sub-options modify the behavior of **pgcollect -cuda** as described here:

**branch**

Collect branching and warp statistics.

**cc13**

Use counters for compute capability 1.3. [default]

**ccnm**

Use counters for compute capability *nm*.

### Tip

---

Use `pgcollect -help` to see which compute capabilities your system supports.

**cfg:<cfgpath>**

Specify <cfgpath> as CUDA profile config file.

**gmem**

Collect global memory access statistics.

**list**

List CUDA event names available for use in profile config file.

## Performance Tip

On some Linux systems, initialization of the CUDA driver for accelerator hardware that is in a power-save state can take a significant amount of time. You can avoid this delay in one of these ways:

- Run the `pgcudainit` program in the background, which keeps the GPU powered on and significantly reduces initialization time for subsequent programs. For more information on this approach, refer to Chapter 7 of the *PGI User's Guide*.
- Use the **pgcollect** option `-cudainit` to eliminate much of the initialization overhead and to provide a more accurate profile.

```
pgcollect -time -cudainit myaccelprog
```

### Note

---

In release 10.5, the option `-cudainit` was called `-accinit`. These two options have exactly the same functionality.

# Index

## Symbols

- allcache
  - pgcollect, 68
- check-events
  - pgcollect, 68
- cuda, 70
  - branch, 70
  - cc13, 70
  - cc20, 70
  - gmem, 70
  - list, 70
- dcache
  - pgcollect, 68
- es-function
  - pgcollect, 68
- event
  - pgcollect, 68
- exe
  - pgcollect, 68
- help
  - pgcollect, 66
- hwtime
  - pgcollect, 68
- imisses
  - pgcollect, 68
- list-events
  - pgcollect, 68
- Minfo
  - ccff, 29
- Mprof
  - dwarf, 29
  - func, 30
  - hwcts, 4

- lines, 30
- mpich1, 30
- mpich2, 30
- msmpi, 30, 30
- mvapich, 30
- time, 4, 30
- pg, 30
- post-function
  - pgcollect, 68
- shutdown
  - pgcollect, 68
- time
  - pgcollect, 66
- V
  - pgcollect, 66

## A

- Accelerator
  - kernel, 55
  - kernel time, 55
  - region, 55
  - region time, 55
- Analyze
  - performance data, MPI programs, 18
  - performance data, multi-threaded programs, 16
- Application
  - tuning, 5
- asm
  - PGPROF command, 60
- Attribute
  - modify color, 48
- Audience Description, xiii

## B

- Background
  - color, 48
- Bytes
  - profile data, 41
  - received, 42
  - sent, 42

## C

- call
  - routine or function, 42

- Caveats, 42
- CCFE, 1, 14
  - Minfo, 29
  - PGPROF command, 61
- clear
  - search, 48
- Clock
  - granularity, 43
  - resolution, 43
- Close
  - PGPROF, 47
- Collect
  - performance data, multi-threaded programs, 15
- Colors
  - background, 48
  - filled text, 48
  - modify attribute, 48
  - set in GUI, 48
  - unfilled text, 48
- command line
  - PGPROF options, 35
- Compare
  - scalability, 19, 47
- Compiler
  - feedback, 14
  - Feedback Format, 1
- Configure
  - es-function pgcollect option, 68
  - event pgcollect option, 68
  - post-function pgcollect option, 68
- Conformance to Standards, xiii
- conventions
  - in text, xv
- Cost
  - information in profile, 42
- Count
  - line or function execution, 42

## D

- Data
  - analysis, MPI programs, 18
  - analysis, multi-threaded programs, 16
  - collection, 4

- collection, multi-threaded programs, 15
- precision, 41
- profile, 41
- sort profile, 13
- transfer time, 55
- view performance, 51
- datafile
  - PGPROF command line option, 35
- Default
  - restore settings, 48
- Display
  - columns, 48
- display
  - PGPROF command, 61
- Documentation
  - accessing, xiii
  - location, xiii
- dwarf
  - Mprof, 29

**E**

- Edit
  - menu, PGPROF, 47
- Environment
  - variables, PGPROF, 39
- Environment variables
  - GMON\_ARCS, 39
  - PGPROF, 39
  - PGPROF\_DEPTH, 39
  - PGPROF\_EVENTS, 39
  - PGPROF\_NAME, 39
  - system, 39
  - using, 4
- Event-based profiling, 67
- Event counters
  - hardware, 21
- Events
  - counters, 4
  - counters, Linux, 21
  - custom specification, pgcollect, 68
  - profile, 68
  - profile specification, 68
  - settings check, 68
- exe

- PGPROF command line option, 35
- Execute
  - shell function after profile, 68
- Exit
  - PGPROF, 47

**F**

- Feedback
  - CCFE, 1
  - compiler, 14
  - messages, 15
  - tab, Focus Panel, 53
- feedbackonly
  - PGPROF command line option, 35
- Files
  - print profile data, 47, 47
  - profile output, 4
- Focus
  - Accelerator tab
    - Accelerator, 54
  - Compiler Feedback tab
    - Compiler, 53
  - histogram tab
    - Histogram, 53
  - panel in PGPROF, 45, 52
  - parallelism tab
    - Parallelism, 53
  - System Information tab
    - System, 54
- Fonts
  - set, 49

**G**

- GMON\_ARCS, 39
- gmon.out, 30, 30
- GUI
  - set colors, 48

**H**

- Hardware
  - counters, profile with, 20
  - event counters, 4, 21
- Help
  - menu, 49
  - pgcollect, 66
  - PGPROF command, 61

- PGPROF command line option, 35, 36
- profiler usage, 66
- Histogram tab
  - PGPROF, 53
- history
  - PGPROF command, 61, 63
- HotSpot
  - controls, 13
  - navigation, 51
- HPF, xiii

**I**

- I
  - PGPROF command line option, 35
- Information
  - profile summary, 46
- Inlining
  - routines, 15
- Intensity
  - computational, 15
  - messages, 15
- Invoke
  - pgcollect, 66, 66
  - profiler, 5

**J**

- jarg
  - PGPROF command line option, 36
- Java
  - jarg, 36

**K**

- Kernels
  - accelerator, 55
  - accelerator time, 55

**L**

- Launch
  - profiler, 5
- Lines
  - Mprof option, 30
  - numbering, 42, 52
  - profiling and optimization, 43
- lines
  - PGPROF command, 61



- Linux
  - profiles, 4
- load
  - PGPROF command, 61
- M**
- Manual organization, xiv
- Measure
  - time, 41
- Menu items
  - Backward search, 47
  - Clear Search, 48
  - Close, 47
  - Exit, 47
  - Font, 49
  - Forward search, 47
  - Help, PGPROF, 49
  - new Window, 46
  - Open Profile, 46
  - PGPROF Help, 49
  - Print, 47
  - Print to File, 47, 47
  - Restore Default Settings, 48
  - Revert to Saved Settings, 48
  - Save Settings on Exit, 48
  - Scalability comparison, 47
  - Search Again, 47
  - Select Columns, 48
  - Select Graph Colors, 48
  - Show Tool Tips, 49
- Menus
  - File, PGPROF, 46
  - PGPROF, 45, 46
  - View, 48
- merge
  - PGPROF command, 62
- Messages
  - feedback, 15
  - inlined, 15
  - intensity, 15
  - number received, 42
  - number sent, 42
  - received, 42, 42
  - sent, 42, 42
- MPI
  - data analysis, 18
  - performance analysis, 18
- MPICH-1
  - Mprof option, 30
- MPICH-2
  - Mprof option, 30
- MS-MPI
  - Mprof option, 30, 30
- Multi-threaded
  - data analysis, 16
  - data collection, 15
  - profiling, 15
- MVAPICH
  - Mprof option, 30
- N**
- Navigation
  - HotSpot, 13, 51
  - jump to profile, 12
  - previous PGPROF view, 12
  - Profile, 9
  - select profile, 12
- Network
  - slow connections, 6
- O**
- OpenMP, xiii
- OProfile
  - pgcollect, 67, 67
- Options
  - Minfo=ccff, 29
  - Mprof=dwarf, 29
  - Mprof=func, 30
  - Mprof=lines, 30
  - Mprof=mpich1, 30
  - Mprof=mpich2, 30
  - Mprof=msmpi, 30, 30
  - Mprof=mvapich, 30
  - Mprof=time, 30
  - pg, 30
- P**
- Parallel
  - program profiles, 15
- Performance
  - allcache pgcollect option, 68
  - assembly-level view, 52
  - collect data, 4
  - data views, 51
  - dcache pgcollect option, 68
  - displayed, 51
  - dmissses pgcollect option, 68
  - line-level view, 52
  - MPI programs, 18
  - multi-threaded programs, 15
  - routine-level view, 52
- pgcollect, 65
  - allcache option, 68
  - check-events option, 68
  - cuda=branch, 70
  - cuda=cc13, 70
  - cuda=cc20, 70
  - cuda=cfg, 70
  - cuda=gmem, 70
  - cuda=list, 70
  - custom event specification, 68
  - dcache option, 68
  - es-function option, 68
  - event-based, 67
  - event option, 68
  - exe option, 68
  - help, 66
  - hwttime option, 68
  - imisses option, 68
  - invoking, 66, 66
  - list-events option, 68
  - OProfile, 67, 67
  - post-function option, 68
  - shutdown option, 68
  - time-based, 66
  - time option, 66
  - version, 66
  - V option, 66
- PGPROF
  - Command-line options, 35
  - commands, 59
  - command summary, 59
  - command syntax, 59
  - Compilation, 29
  - Edit menu, 47
  - environment variables, 39
  - features reference, 45, 45
  - File menu, 46

- focus panel, 45
- Help, 49
- information bar, 46
- menu bar, 45
- menus, 46
- optimization, 43
- overview, 1, 45
- Overview, 1
- scalability comparison, 19
- Sorting Profile Data, 13
- Sort menu, 49
- statistics table, 45
- toolbar, 45
- toolbar, 50
- Using, 7
- View Menu, 48
- PGPROF\_DEPTH, 39
- PGPROF\_EVENTS, 39
- PGPROF\_NAME, 39
- pgprof.out, 4
- PGPROF command line option, 35, 35, 36
  - datafile, 35
  - feedbackonly, 35
  - help, 36
  - help, 35
  - jarg, 36
  - text, 36
  - V, 36
- PGPROF Commands
  - ! (history), 63
  - asm, 60
  - ccff, 61
  - display, 61
  - help, 61
  - history, 61
  - lines, 61
  - load, 61
  - merge, 62
  - print, 62
  - process, 62
  - quit, 62
  - select, 62
  - sort, 62
  - srcdir, 62
  - stat, 62
  - thread, 63
  - times, 63
- Print
  - PGPROF, 47
  - PGPROF command, 62
- process
  - PGPROF command, 62
- Profile
  - calls, 42
  - cost information, 42
  - count executions, 42
  - data statistics, 41
  - hardware counters, 20
  - line information, 42
  - MPI, 41, 42, 42, 42, 42, 42
  - pgcollect, 65
  - time, 42
- Profiler
  - invoke, 5
  - sample launch, 5
- Profiles
  - collect data, 4
  - edit, 47
  - Event-based, 3
  - Instrumentation-based, 2
  - multi-threaded programs, 15
  - navigation, 9
  - open, 46
  - output file, 4
  - parallel programs, 15
  - platforms, 4
  - sort, 13, 49
  - Time-based, 3
  - use hardware event counters, 4
- ProfilesL data, print, 47
- Profiling
  - clock resolution, 43
  - Command-level interface, 59
  - compilation, 29
  - line-level, 30
  - measurements, 51
  - MPICH, 30, 30
  - MPICH-1, 30
  - MPICH-2, 30
  - MVAPICH-1, 30
  - optimization, 43
  - overhead, 42
  - PGPROF, 1, 39
  - routine-level, 30
  - sample-based, 30
  - time-based, 30
  - Virtual Timer, 41
- profiling
  - event-based, 67
  - time-based, 66
- Program execution
  - Using Hardware Event Counters, 4

## Q

- quit
  - PGPROF command, 62

## R

- Region
  - accelerator, 55
  - accelerator time, 55
- Related Publications, xvi
- Repeat
  - search, 47
- Reset
  - Bar Chart colors, 48
- Resource
  - utilization, profiling, 20
- Restore
  - settings, 48
- Revert
  - settings, 48
- Routines
  - execution time, 42
  - inlined messages, 15
  - Mprof, 30

## S

- Save
  - current settings, 48
- Scalability
  - comparison, 19, 47
- Scale
  - PGPROF command line option, 36
- srcdir
  - PGPROF command, 62
- Search

- again, 47
- backward, 47
- clear, 48
- forward, 47
- Select
  - columns, 48
  - PGPROF command, 62
- Set
  - colors in GUI, 48
  - font, PGPROF, 49
- Settings
  - restore default, 48
  - revert to saved, 48
  - save on exit, 48
- Sort
  - menu, PGPROF, 49
  - PGPROF command, 62
  - profile data, 13, 13
- Source
  - line numbering, 52
- source file
  - search path, 35
- Specify
  - custom events pgcollect, 68
- stat
  - PGPROF command, 62
- Statistics
  - data profile, 41
  - select data, 51
  - table, PGPROF, 45
  - table in PGPROF, 51
- System
  - environment, 4
  - variables, 4
- System Requirements, xvi

## T

- Tables
  - Focus Panel, PGPROF, 52
  - statistics in PGPROF, 45
- tabs
  - Histogram, PGPROF, 53
- text
  - PGPROF command line option, 36
- Threads
  - PGPROF command, 63

- Time
  - execution, 42
  - measure, 41
  - Mprof option, 30
- time-based profiling, 66
- times
  - PGPROF command, 63
- Toolbar
  - PGPROF, 45, 50
- Tool tips
  - show, 49
- Troubleshoot
  - PGPROF, 5
- Tune
  - application, 5

## V

- V
  - PGPROF command line option, 36
- Variables
  - system environment, 4
- Versions
  - V option, 36
- View
  - menu, 48
  - performance data, 51
  - previous profile, 12
  - profiles with hardware event counters, 21
  - select profile, 12

## W

- Window
  - new, PGPROF, 46

## NOTICE

ALL NVIDIA DESIGN SPECIFICATIONS, REFERENCE BOARDS, FILES, DRAWINGS, DIAGNOSTICS, LISTS, AND OTHER DOCUMENTS (TOGETHER AND SEPARATELY, "MATERIALS") ARE BEING PROVIDED "AS IS." NVIDIA MAKES NO WARRANTIES, EXPRESSED, IMPLIED, STATUTORY, OR OTHERWISE WITH RESPECT TO THE MATERIALS, AND EXPRESSLY DISCLAIMS ALL IMPLIED WARRANTIES OF NONINFRINGEMENT, MERCHANTABILITY, AND FITNESS FOR A PARTICULAR PURPOSE.

Information furnished is believed to be accurate and reliable. However, NVIDIA Corporation assumes no responsibility for the consequences of use of such information or for any infringement of patents or other rights of third parties that may result from its use. No license is granted by implication of otherwise under any patent rights of NVIDIA Corporation. Specifications mentioned in this publication are subject to change without notice. This publication supersedes and replaces all other information previously supplied. NVIDIA Corporation products are not authorized as critical components in life support devices or systems without express written approval of NVIDIA Corporation.

## TRADEMARKS

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 NVIDIA Corporation in the U.S. and other countries. Other company and product names may be trademarks of the respective companies with which they are associated.

## COPYRIGHT

© 2013 NVIDIA Corporation. All rights reserved.