

SpeedShop User's Guide

Document Number 007-3311-003

© Copyright 1998 Silicon Graphics, Inc.— All Rights Reserved

The contents of this document may not be copied or duplicated in any form, in whole or in part, without the prior written permission of Silicon Graphics, Inc.

RESTRICTED RIGHTS LEGEND

Use, duplication, or disclosure of the technical data contained in this document by the Government is subject to restrictions as set forth in subdivision (c) (1) (ii) of the Rights in Technical Data and Computer Software clause at DFARS 52.227-7013 and/or in similar or successor clauses in the FAR, or in the DOD or NASA FAR Supplement. Unpublished rights reserved under the Copyright Laws of the United States. Contractor/manufacturer is Silicon Graphics, Inc., 2011 N. Shoreline Blvd., Mountain View, CA 94039-7311.

Silicon Graphics is a registered trademark of Silicon Graphics, Inc. ProDev and IRIX are trademarks of Silicon Graphics, Inc. Ada is a registered trademark of the Ada Joint Program Office, U.S. Government. OSF/Motif is a trademark of the Open Software Foundation. POSIX is a registered trademark of the Institute of Electrical and Electronic Engineers, Inc. Purify is a registered trademark of Pure Software, Inc. R10000 and R4000 are trademarks of MIPS Technologies, Inc. UNIX is a registered trademark of X/Open Company. Ltd. X Window System is a trademark of the Massachusetts Institute of Technology.

Contents

	<i>Page</i>
About This Guide	ix
About This Revision	ix
Introduction to Performance Analysis [1]	1
Sources of Performance Problems	1
CPU-Bound Processes	2
I/O-Bound Processes	2
Memory-Bound Processes	2
Bugs	2
Fixing Performance Problems	3
SpeedShop Tools	3
Main Commands	3
Additional Commands	4
Experiment Types	5
SpeedShop Libraries	6
API	7
Supported Programming Models and Languages	7
Using SpeedShop Tools for Performance Analysis	8
Using ssusage to Evaluate Machine Resource Use	8
Using ssrun and prof to Gather and Analyze Performance Data	9
Collecting Data for Part of a Program	11
Tutorial for C Users [2]	13
Tutorial Overview	13
Contents of the generic Program	13
Output from the generic Program	14

	<i>Page</i>
Tutorial Setup	15
Analyzing Performance Data	15
usertime Experiment	16
Performing a usertime Experiment	16
Generating a Report	17
Analyzing the Report	18
pcsamp Experiment	19
Generating a Report	20
Analyzing the Report	22
Hardware Counter Experiment	22
Performing a Hardware Counter Experiment	22
Generating a Report	23
Analyzing the Report	23
ideal Experiment	25
Performing an ideal Experiment	25
Generating a Report	26
Analyzing the Report	28
fpe Trace	29
Performing an fpe Trace	29
Generating a Report	29
Analyzing the Report	30
Tutorial for Fortran Users [3]	33
Tutorial Overview	33
Tutorial Setup	34
Analyzing Performance Data	35
usertime Experiment	35
Performing a usertime Experiment	35

	<i>Page</i>
Generating a Report	36
Analyzing the Report	38
pcsamp Experiment	39
Performing a pcsamp Experiment	39
Generating a Report	39
Analyzing the Report	40
Hardware Counter Experiment	41
Performing a hardware counter Experiment	41
Generating a Report	41
Analyzing the Report	42
ideal Experiment	43
Performing an ideal Experiment	43
Generating a Report	44
Analyzing the Report	47
Experiment Types [4]	49
Selecting an Experiment	49
usertime Experiment	50
pcsamp Experiment	51
ideal Experiment	51
How SpeedShop Prepares Files	51
How SpeedShop Calculates CPU Time	52
Inclusive Basic Block Counting	52
Using pcsamp and ideal Together	53
Hardware Counter Experiments	54
Two Tools for Hardware Counter Experiments	54
SpeedShop Hardware Counter Experiments	54
[<code>f</code>]gi_hwc	55

	<i>Page</i>
[f]cy_hwc	55
[f]ic_hwc	55
[f]isc_hwc	55
[f]dc_hwc	56
[f]dsc_hwc	56
[f]tlb_hwc	56
[f]gfp_hwc	56
prof_hwc	57
Hardware Counter Numbers	57
fpe Trace	59
Collecting Data on Machine Resource Usage [5]	61
ssusage Syntax	61
ssusage Results	61
Setting Up and Running Experiments: ssrun [6]	63
Building Your Executable	63
Special Information for MP Fortran Programs	64
Setting Up Output Directories and Files	65
Using Runtime Environment Variables	66
User Environment Variables	66
Process Tracking Environment Variables	68
Expert-Mode Environment Variables	68
Running Experiments	70
ssrun Syntax	70
ssrun Examples	71
Example Using the pcsampx Experiment	71
Example Using the -v Option	73
Using ssrun With a Debugger	73

	<i>Page</i>
Running Experiments on MPI Programs	74
Running Experiments on Programs Using Pthreads	75
Using Calipers	75
Setting Calipers With <code>ssrt_caliper_point</code>	76
Setting Calipers With Signals	77
Setting Calipers With a Debugger	78
Effects of <code>ssrun</code>	78
Effects of <code>ssrun -ideal</code>	78
Analyzing Experiment Results: prof [7]	81
Using prof to Generate Performance Reports	81
prof Syntax	81
prof Options	82
prof Output	85
Using prof With <code>ssrun</code>	85
<code>usertime</code> Experiment Reports	86
<code>pcsamp</code> Experiment Reports	87
Hardware Counter Experiment Reports	88
<code>ideal</code> Experiment Reports	89
<code>fpe</code> Trace Reports	92
Using prof Options	93
Using the <code>-dis</code> Option	93
Using the <code>-S</code> Option	98
Using the <code>-calipers</code> Option	101
Using the <code>-gprof</code> Option	101
Generating Reports for Different Machine Types	106
Generating Reports for Multiprocessed Executables	106
Generating Compiler Feedback Files	107

	<i>Page</i>
Interpreting Reports	107
Using SpeedShop in Expert Mode: pixie [8]	109
Using pixie	109
pixie Syntax	109
pixie Options	110
pixie Output	111
Obtaining Basic Block Counts	112
Examples of Basic Block Counting	115
Example Using prof with no Options	115
Example Using prof -heavy	117
Example Using prof -quit	118
Obtaining Inclusive Basic Block Counts	119
Example of prof -gprof	119
Miscellaneous Commands [9]	121
Using the thrash Command	121
thrash Syntax	121
Effects of thrash	122
Using the squeeze Command	122
squeeze Syntax	122
Effects of squeeze	123
Calculating the Working Set of a Program	123
Dumping Performance Data Files	125
ssdump Syntax	125
Experiment File Format	126
Dumping Compiler Feedback Files	131
fbdump Syntax	131

	<i>Page</i>
Glossary	133
Index	135
Figures	
Figure 1. How Basic Block Counting Works	114
Tables	
Table 1. SpeedShop Main Commands	4
Table 2. SpeedShop Additional Commands	4
Table 3. SpeedShop Libraries	6
Table 4. Choosing an Experiment Type	10
Table 5. Letter Codes in Process Experiment ID Numbers	10
Table 6. Summary of Experiments	49
Table 7. Basic Block Counts and PC Profile Counts Compared	53
Table 8. Hardware Counter Numbers	57
Table 9. Letter Codes in Experiment ID Numbers	65
Table 10. General Environment Variables	66
Table 11. Process Tracking Environment Variables	68
Table 12. Expert-Mode Environment Variables	69
Table 13. Flags for ssrun	71
Table 14. Setting Caliper Points	76
Table 15. Options for prof	82
Table 16. Letter Codes in Experiment ID Numbers	107
Table 17. Options for pixie	110
Table 18. Options for fbdump	131

About This Guide

The *SpeedShop User's Guide* describes and illustrates methods for measuring program performance using SpeedShop commands such as `ssrun` and `prof`. It contains tutorials that generate performance statistics for C and Fortran programs.

This manual is a user's guide for the SpeedShop performance tools, Release 1.3. It contains the following chapters:

- Chapter 1, page 1 provides a general introduction to performance analysis concepts and techniques, plus an overview of the SpeedShop tools.
- Chapter 2, page 13 provides a tutorial on how to collect performance data and generate reports for a C program.
- Chapter 3, page 33 provides a tutorial on how to collect performance data and generate reports for Fortran programs running on single-processor machines.
- Chapter 4, page 49 describes the types of experiments that can be performed using SpeedShop tools.
- Chapter 5, page 61 describes how to use the `ssusage` command to collect information about a program's machine resource usage.
- Chapter 6, page 63 explains in detail how to set up and run experiments using `ssrun`, and explains how to use caliper points to generate reports for part of a program.
- Chapter 7, page 81 explains how to generate reports from performance data using `prof`.
- Chapter 8, page 109 explains how to use `pixie` and `prof` directly, without invoking `ssrun`.
- Chapter 9, page 121 explains how to use the `thrash` and `squeeze` commands to determine the memory usage, or working set, of your application. It also covers commands to print performance data files.

About This Revision

This revision of the manual was prepared in the spring of 1998. It includes bug fixes and updates for the SpeedShop 1.3 release.

Introduction to Performance Analysis [1]

This chapter provides a brief introduction to performance analysis techniques for Silicon Graphics systems and describes how to use them to solve performance problems. It includes the following sections:

- Sources of Performance Problems. It provides a general overview of potential performance problems. See Section 1.1, page 1.
- Fixing Performance Problems. This section discusses how you can use SpeedShop to determine what the problems are. See Chapter 2, page 13.
- SpeedShop Tools. This section lists SpeedShop commands, experiment types, and libraries. See Section 1.3, page 3.
- Using SpeedShop Tools for Performance Analysis. This section steps you through the steps to take when using SpeedShop. See Section 1.4, page 8.

1.1 Sources of Performance Problems

To tune a program's performance, you need to determine its consumption of machine resources. At any point (or phase) in a process, there is one limiting resource controlling the speed of execution. Processes can be slowed down by any of the following:

- CPU speed and availability
- I/O processing
- Memory size and availability

Performance problems may span the entire run of a process, or they may occur in just a small portion of the program. For example, a function that performs a lot of I/O processing might be called regularly as the program runs, or a particularly CPU-intensive calculation might occur in just one portion of the program. When there are performance problems in a small portion of the program, collect data for just that part of the program.

Because programs exhibit different behavior during different phases of operation, you need to identify the limiting resource for each phase. A program can be I/O-bound while it reads in data, CPU-bound while it performs computation, and I/O-bound again in its final stage while it writes out data. Once you've identified the limiting resource in a phase, you can perform an

in-depth analysis to find the problem. After you have solved that problem, you can check for other problems within the same or other phases—performance analysis is an iterative process.

1.1.1 CPU-Bound Processes

A *CPU-bound* process spends its time in the CPU and is limited by CPU speed and availability. To improve performance on CPU-bound processes, streamline your code using one or more of the following techniques:

- Modifying algorithms
- Reordering code to avoid interlocks
- Removing nonessential steps
- Blocking to keep data in cache and registers
- Using alternative algorithms

1.1.2 I/O-Bound Processes

An *I/O-bound* process has to wait for I/O to complete and may be limited by disk access speeds or memory caching. To improve the performance of I/O-bound processes, try one of the following techniques:

- Improving overlap of I/O with computation
- Optimizing data usage to minimize disk access
- Using data compression

1.1.3 Memory-Bound Processes

A *memory-bound* program continuously swaps out pages of memory. Page thrashing is often due to accessing virtual memory on a haphazard rather than strategic basis. One telltale indication of page-thrashing with paging to a local disk is noise during disk accesses. To fix a memory-bound process, try to improve the memory reference patterns or, if possible, decrease the memory used by the program.

1.1.4 Bugs

Certain bugs can cause performance problems. Examples include:

- The program is unnecessarily reading from the same file twice in different parts.
- Floating point exceptions are slowing down the program.
- Old code has not been completely removed.
- The program is leaking memory (making `malloc()` calls without the corresponding calls to `free()`).

1.2 Fixing Performance Problems

The SpeedShop performance tools described in this manual can help you to identify specific performance problems described later in this chapter. However, the techniques described in this manual are only a part of performance tuning. Other areas that you can tune, but that are outside the scope of this document, include graphics, I/O, the kernel, system parameters, memory, and real-time system calls.

Although it may be possible to obtain short-term speed increases by relying on unsupported or undocumented quirks of the compiler, it is a bad idea to do so. Any such “features” may break in future compiler releases. The best way to produce efficient code that can be trusted to remain efficient is to follow good programming practices. In particular, choose good algorithms and leave the details to the compiler.

1.3 SpeedShop Tools

The SpeedShop tools allow you to run experiments and generate reports to track down the sources of performance problems. SpeedShop consists of a set of commands that can be run in a shell, an API, and a number of libraries to support the commands.

This section provides an overview of the tools by first discussing the main commands, then providing more detail on additional commands, experiment types, libraries, and supported programs and languages.

1.3.1 Main Commands

SpeedShop provides the commands listed in Table 1.

Table 1. SpeedShop Main Commands

Command	Description
<code>ssusage</code>	Collects information about your program's use of machine resources. Output from <code>ssusage</code> can be used to determine where most resources are being spent.
<code>ssrun</code>	Allows you to run experiments on a program to collect performance data. It establishes the environment to capture performance data for an executable, creates a process from the executable (or from an instrumented version of the executable) and runs it. Input to <code>ssrun</code> consists of an experiment type, control flags, the name of the target, and the arguments to be used in executing the target.
<code>prof</code>	Analyzes the performance data you have recorded using <code>ssrun</code> and provides formatted reports. <code>prof</code> detects the type of experiment you have run, and generates a report specific to the experiment type.

1.3.2 Additional Commands

SpeedShop provides the additional commands shown in Table 2.

Table 2. SpeedShop Additional Commands

Command	Description
<code>pixie</code>	Instruments an executable to enable basic block counting experiments to be performed. If you use <code>ssrun</code> , you will not normally need to call this program directly.
<code>fbdump</code>	Prints out the formatted contents of compiler feedback files generated by <code>prof</code> .
<code>squeeze</code>	Allocates a region of virtual memory and locks the virtual memory down into real memory, making it unavailable to other processes.

Command	Description
thrash	Allows you to allocate a block of memory, then access the allocated memory to explore paging behavior.
ssdump	Prints out formatted performance data that was collected while running <code>ssrun</code> . This program is included for SpeedShop debugging purposes. You don't normally need to use it.

1.3.3 Experiment Types

You can conduct the following types of experiments using the `ssrun` command:

- Statistical PC sampling with `pcsamp` experiments.

Data is measured by periodically sampling the Program Counter (PC) of the target executable when it is executing in the CPU. The PC shows the address of the currently executing instruction in the program. The data that is obtained from the samples is translated to a time that can be displayed at the function, source line, and machine instruction levels. The actual *CPU time* is calculated by multiplying the number of times a specific address is found in the PC by the amount of time between samples. (For a definition of CPU time, wall-clock time, and process virtual time, see the Glossary.)

- Statistical hardware counter sampling with `hwc` experiments.

Hardware counter experiments are available on R10000™ systems that have built-in hardware counters. Data is measured by collecting information each time the specified hardware counter overflows. You can specify the hardware counter and the overflow interval you want to use.

- Statistical call stack profiling with `usertime`.

Data is measured by periodically sampling the call stack. The program's call stack data is used to attribute exclusive user time to the function at the bottom of each call stack (that is, the function being executed at the time of the sample), and to attribute inclusive user time to all the functions above the one currently being executed. Like `pcsamp` experiments, `usertime` experiments return *CPU time*.

- Basic block counting with `ideal`.

Data is measured by counting basic blocks and calculating an ideal CPU time for each function. This involves instrumenting the program to divide

the code into basic blocks, which are sets of instructions with a single entry point, a single exit point, and no branches into or out of the set. Instrumentation also permits a count of all dynamic (function-pointer) calls to be recorded. You can compare this ideal time with the times returned by other experiments to measure the performance of your code against its potential. See Section 4.4.4, page 53).

- Floating-point exception trace with `fpe`.

A floating-point exception trace collects each floating point exception, including the exception type and the call stack, at the time of the exception. `prof` generates a report showing inclusive and exclusive floating-point exception counts.

1.3.4 SpeedShop Libraries

Versions of the SpeedShop libraries `libss.so` and `libssrt.so` are available to support applications built using shared libraries (DSOs) only and the old 32-bit, new 32-bit or 64-bit application binary interfaces (ABIs).

Table 3 provides information about the different SpeedShop libraries.

Table 3. SpeedShop Libraries

Library	Description
<code>libss.so</code>	A shared library (DSO) that supports <code>libssrt.so</code> . <code>libss.so</code> data normally appears in experiment results generated with <code>prof</code> .
<code>libssrt.so</code>	A shared library (DSO) that is linked in to the program you specify when you run an experiment. All the performance data collection with the SpeedShop system is done within the target process(es), by exercising various pieces of functionality using <code>libssrt</code> . Data from <code>libssrt.so</code> does not normally appear in performance data reports generated with <code>prof</code> .
<code>libfpe_ss.so</code>	Supplements the standard <code>libfpe.so</code> for the purposes of collecting floating point exception data. See the <code>fpe_ss</code> reference page for more information.

Library	Description
<code>libmalloc_ss.so</code>	Inserts versions of <code>malloc</code> routines from <code>libc.so.1</code> that allow tracing all calls to <code>malloc</code> , <code>free</code> , <code>realloc</code> , <code>memalign</code> , and <code>valloc</code> . See the <code>malloc_ss</code> reference page for more information.
<code>libpixrt.so</code>	A shared library (DSO) used by pixified programs.

1.3.5 API

The SpeedShop API is primarily available to allow you to use `ssrt_caliper_point` to set caliper points in your source code. See Section 6.7, page 75, for information on using caliper points. For information on other API functions, see the `ssapi` reference page.

1.3.6 Supported Programming Models and Languages

The SpeedShop tools support programs with the following characteristics:

- Shared libraries (DSOs.)
- Non-stripped executables.
- Executables containing `fork`, `sproc`, `system`, or `exec` commands.
- Executables using supported techniques for opening, closing, and/or delay-loading DSOs.
- C, C++, Fortran (Fortran-77, Fortran-90, and High-Performance Fortran), or Ada[®] 95 source code.
- Power Fortran and Power C source code. `prof` understands the syntax and semantics of the multi-processing runtime and displays the data accordingly.
- `pthread`s, supported with data on a per-program basis.
- Message Passing Interface (MPI) or other message-passing paradigms. Currently supported by providing data on the behavior of each process. The behavior of the MPI library itself is monitored just like any other user-level code.

1.4 Using SpeedShop Tools for Performance Analysis

Performance tuning typically consists of

- Examining machine resource usage
- Breaking down the process into phases
- Identifying the resource bottleneck within each phase
- Correcting the cause of the bottleneck

Generally, you run the first experiment to break your program down into phases and run subsequent experiments to examine each phase individually. After you have solved a problem in a phase, you should re-examine machine resource usage to see if there is further opportunity for performance improvement.

The general steps for a performance analysis cycle are:

1. Build the application.
2. Run experiments on the application to collect performance data.
3. Examine the performance data.
4. Generate an improved version of the program.
5. Repeat as needed.

To accomplish this using SpeedShop tools:

- Use `ssusage` to capture information on your program's use of machine resources.
- Use `ssrun` to capture different types of performance data over either your entire program or parts of the program. `ssrun` can be used in conjunction with `dbx` or `WorkShop` debuggers.
- Use `prof` to analyze the data and generate reports.

1.4.1 Using `ssusage` to Evaluate Machine Resource Use

To determine overall resource usage by your program, run the program with `ssusage`. The results of this command allow you to identify high user CPU time, high system CPU time, high I/O time, and a high degree of paging.

```
ssusage prog_name prog_args
```

From the `ssusage` output, you can decide which experiments to run to collect data for further study. For more information on `ssusage`, see Chapter 5, page 61, or see the `ssusage` reference page.

1.4.2 Using `ssrun` and `prof` to Gather and Analyze Performance Data

This section describes the steps involved in a performance analysis cycle when using the main interface to the SpeedShop tools: the `ssrun` command.

You can also call the commands individually. For example, if you are planning to perform basic block counting experiments that involve instrumenting the executable, you can either do this by calling `ssrun` with the appropriate experiment type, or you can set up your environment to call `pixie` directly to instrument your executable. Information on setting up your environment and running `pixie` directly can be found in Chapter 8, page 109.

To perform a performance analysis, follow these general steps:

1. Build the executable.

You can usually build the executable as you would normally. See Section 6.1, page 63, for information on how to build the executable.

2. Specify caliper points if you want to collect data for only a portion of your program. See Section 1.4.3, page 11, for more information.
3. To collect performance data, call `ssrun` with the parameters below. Use the information in Table 4 to determine which experiments to run:

```
ssrun flags exp_type prog_name prog_args
```

flags One or more valid flags. For a complete list of flags, see the `ssrun` reference page.

exp_type Experiment name.

prog_name Executable name.

prog_args Arguments to the executable

Table 4. Choosing an Experiment Type

Performance Problem	Experiment(s) to Run
High user CPU time	<code>usertime</code> <code>pcsamp</code> (four variants) <code>*_hwc</code> experiments <code>ideal</code>
High system CPU time	If floating point exceptions are suspected: <code>fpe</code>
High I/O time	<code>ideal</code> , then examine counts of I/O routines
High paging (<code>majf</code>)	<code>ideal</code> , then <code>prof -feedback</code> and <code>cord</code> to rearrange procedures. If inefficient heap usage is suspected, use WorkShop's Performance Analyzer to gather information.

For each process of the executable, the experiment data is stored in a file with a name of the format `prog_name.exp_type.id`.

The experiment ID, consists of one or two letters (designating the process type) and the process ID number. See Table 5 for letter codes and descriptions.

Table 5. Letter Codes in Process Experiment ID Numbers

Letter Codes	Description
<code>m</code>	Master process created by <code>ssrun</code>
<code>p</code>	Process created by a call to <code>sproc()</code>
<code>f</code>	Process created by a call to <code>fork()</code>
<code>s</code>	Process created by a call to <code>system()</code>
<code>e</code>	Process created by a call to <code>exec()</code>
<code>fe</code>	Process created by a call to <code>fork()</code> and <code>exec()</code>

For more information on the `ssrun` command, see Chapter 6, page 63, or view the `ssrun` reference page.

- To generate a report of the experiment, call `prof` with the following parameters:

`prof flags data_file`

flags One or more valid flags. For a complete list of flags, see the `prof` reference page.

data_file The name of the file in which the experiment data was recorded.

For more information on using `prof`, see Chapter 7, page 81, or see the `prof` reference page.

1.4.3 Collecting Data for Part of a Program

If you have a performance problem in only one part of your program, consider collecting performance data for just that part. You can do this by setting caliper points around the problem area when running an experiment, then using the `prof -calipers` option to generate a report for the problem area.

You can set caliper points using one of the following:

- the SpeedShop API
- the caliper signal environment
- a debugger such as the ProDev WorkShop debugger

For more information on using calipers, see Section 6.7, page 75.

Tutorial for C Users [2]

This chapter presents a tutorial for using the SpeedShop tools to gather and analyze performance data in a C program, and covers these topics:

- Section 2.1, page 13, introduces the sample program and explains the different scenarios in which it will be used.
- Section 2.2, page 15, steps you through the necessary setup for running the experiment.
- Section 2.3, page 15, steps you through five different experiments, discussing first how to do the experiments, then how to interpret the results.

Note: Because of inherent differences between systems and because of concurrent processes that may be running on your system, your experiment will produce different results from the one in this tutorial. However, the basic form of the results should be the same.

2.1 Tutorial Overview

This tutorial uses a sample program called `generic`. There are two versions of the program:

<code>generic</code> directory	Contains files for the n32-bit ABI
<code>generico32</code> directory	Contains files for the (old) 32-bit ABI

When you work with the tutorial, choose the version of `generic` most appropriate for your system. A good guideline is to choose the version that corresponds to the way you expect to develop your programs.

This tutorial was written and tested using the version of `generic` in the `generic` directory.

2.1.1 Contents of the `generic` Program

`generic` was designed as a test and demonstration application. It contains code to run scenarios that each test a different area of SpeedShop. The version of `generic` used in this tutorial performs scenarios that:

- Build a linked list of structures
- Use a lot of usertime

- Scan a directory and run the `stat` command on each file
- Perform file I/O
- Generate a number of floating point exceptions
- Link and call a routine in a DSO

2.1.2 Output from the generic Program

Output from the program looks like the following:

```
0:00:00.000 ===== (27173)          Begin script Fri  06 Feb 1998
15:03:31.
    begin script 'll.u.cvt.d.i.f.dso'
0:00:00.002 ===== (27173)          start of linklist Fri  06 Feb 1998
15:03:31.
    linklist completed.
0:00:00.003 ===== (27173)          start of usertime Fri  06 Feb 1998
15:03:31.
    usertime completed.
0:00:25.572 ===== (27173)          start of cvttrap Fri  06 Feb 1998
15:03:57.
    cvttrap completed, y = 2.60188e+14, z = 2.60188e+14.
0:00:25.806 ===== (27173)          start of dirstat Fri  06 Feb 1998
15:03:57.
    dirstat of /usr/include completed, 304 files.
0:00:26.618 ===== (27173) start of iofile -- stdio Fri  06 Feb 1998
15:03:58.
    stdio iofile on /unix completed, 7307988 chars.
0:00:26.864 ===== (27173)          start of fpetraps Fri  06 Feb 1998
15:03:58.
    fpetraps completed.
0:00:26.865 ===== (27173)          start of libdso Fri  06 Feb 1998
15:03:58.
dlslave_init executed
dlslave_routine executed
    slaveusertime completed, x = 5000000.000000.
    libdso: dynamic routine returned 13
    end of script 'll.u.cvt.d.i.f.dso'
0:00:27.972 ===== (27173)          End script Fri  06 Feb 1998
15:03:59.
```

2.2 Tutorial Setup

Copy the program to a directory where you have write permission and compile it so that you can use it in the tutorial.

1. Change to the `/usr/demos/SpeedShop` directory.
2. Copy the appropriate `generic` directory and its contents to a directory where you have write permission:

```
cp -r generic your_dir
```

3. Change to the directory you just created:

```
cd your_dir/generic
```

4. Compile the program by entering:

```
make all
```

This provides an executable for the experiment.

5. Set the library path so that the program can find shared libraries in the `generic` directory:

```
setenv LD_LIBRARY_PATH your_dir/generic_version
```

2.3 Analyzing Performance Data

This section explains how to run the following experiments on the `generic` program and generate and interpret the results:

- `usertime`. As a first cut at optimization, this may be the most useful experiment. It breaks a program down into its functions and returns the *CPU time* used in each. See Section 2.3.1, page 16.
- `pcsamp`. This experiment returns the CPU time used by each function. See Section 2.3.2, page 19.
- `dsc_hwc`. This experiment basically counts the number of times a required data item was not in secondary data cache. If the data item is not in secondary data cache, it must be fetched from memory, which requires more time. See Section 2.3.3, page 22.
- `ideal`. This experiment calculates the best time achievable. See Section 2.3.4, page 25.

- `fpe`. This experiment counts the number of floating-point exceptions in each function. See Section 2.3.5, page 29.

You can follow the tutorial from start to finish, or you can choose the experiment(s) you want to perform.

2.3.1 usertime Experiment

This section explains how to perform a `usertime` experiment. The `usertime` experiment allows you to gather data on the amount of CPU time spent in each function in your program. For more information on `usertime`, see Chapter 4, page 49. For definitions of CPU time, wall-clock time, and process-virtual time, see the Glossary.

2.3.1.1 Performing a usertime Experiment

From the command line, enter

```
ssrun -usertime generic
```

This starts the experiment. Output from `generic` and from `ssrun` is printed to `stdout`, as shown in the following example. A data file is also generated. The name consists of the process name (`generic`), the experiment type, `usertime`, and the experiment ID. In this example, the file name is `generic.usertime.m10981`.

```
0:00:00.000 ===== (10981)          Begin script Mon  02 Feb 1998
11:05:02.
    begin script 'll.u.cvt.d.i.f.dso'
0:00:00.002 ===== (10981)          start of linklist Mon  02 Feb 1998
11:05:02.
    linklist completed.
0:00:00.003 ===== (10981)          start of usertime Mon  02 Feb 1998
11:05:02.
    usertime completed.
0:00:22.948 ===== (10981)          start of cvttrap Mon  02 Feb 1998
11:05:25.
    cvttrap completed, y = 2.60188e+14, z = 2.60188e+14.
0:00:23.156 ===== (10981)          start of dirstat Mon  02 Feb 1998
11:05:25.
    dirstat of /usr/include completed, 304 files.
0:00:23.937 ===== (10981) start of iofile -- stdio Mon  02 Feb 1998
11:05:26.
    stdio iofile on /unix completed, 7307988 chars.
```

```

0:00:24.777 ===== (10981)          start of fpetraps Mon  02 Feb 1998
11:05:27.
      fpetraps completed.
0:00:24.777 ===== (10981)          start of libdso Mon  02 Feb 1998
11:05:27.
dlslave_init executed
dlslave_routine executed
      slaveusertime completed, x = 5000000.000000.
      libdso: dynamic routine returned 13
      end of script 'll.u.cvt.d.i.f.dso'
0:00:25.866 ===== (10981)          End script Mon  02 Feb 1998
11:05:28.

```

2.3.1.2 Generating a Report

To generate a report on the data collected, enter at the command line:

```
prof your_output_file_name > usertime.results
```

prof prints results to *stdout*. Because of line width restrictions, the DSO, file name, and line number information at the end of each line is wrapped to the next line.

```

-----
SpeedShop profile listing generated Mon Feb  2 11:07:15 1998
  prof generic.usertime.m10981
      generic (n32): Target program
      usertime: Experiment name
      ut:cu: Marching orders
R4400 / R4000: CPU / FPU
      1: Number of CPUs
      175: Clock frequency (MHz.)

Experiment notes--
  From file generic.usertime.m10981:
  Caliper point 0 at target begin, PID 10981
      /home/hickory28/dep/speedshop/progs.etc/linpack.demos/c/generic
  Caliper point 1 at exit(0)

-----
Summary of statistical callstack sampling data (usertime)--
      809: Total Samples
      0: Samples with incomplete traceback
  24.270: Accumulated Time (secs.)
      30.0: Sample interval (msecs.)

-----

```

Function list, in descending order by exclusive time

```
-----
[index]  excl.secs  excl.%   cum.%   incl.secs  incl.%   samples  procedure(dso:file,line)

   [4]    22.770  93.8%   93.8%    22.770  93.8%    759  anneal
(generic: generic.c, 1573)
   [6]     1.020   4.2%   98.0%     1.020   4.2%     34  slaveusertime
(dlslave.so: dlslave.c, 22)
   [9]     0.210   0.9%   98.9%     0.210   0.9%      7  cvttrap
(generic: generic.c, 317)
  [12]     0.120   0.5%   99.4%     0.120   0.5%      4  __read
(libc.so.1: read.s, 20)
  [14]     0.090   0.4%   99.8%     0.090   0.4%      3  _xstat
(libc.so.1: xstat.s, 12)
  [10]     0.030   0.1%   99.9%     0.180   0.7%      6  iofile
(generic: generic.c, 464)
  [11]     0.030   0.1%  100.0%     0.150   0.6%      5  fread
(libc.so.1: fread.c, 34)
   [1]     0.000   0.0%  100.0%    24.270 100.0%    809  __start
(generic: crt1text.s, 101)
   [2]     0.000   0.0%  100.0%    24.270 100.0%    809  main
(generic: generic.c, 101)
   [3]     0.000   0.0%  100.0%    24.270 100.0%    809  Scriptstring
(generic: generic.c, 184)
   [5]     0.000   0.0%  100.0%    22.770  93.8%    759  usertime
(generic: generic.c, 1377)
  [15]     0.000   0.0%  100.0%     0.090   0.4%      3  dirstat
(generic: generic.c, 348)
  [16]     0.000   0.0%  100.0%     0.090   0.4%      3  _stat
(libc.so.1: stat.c, 31)
  [13]     0.000   0.0%  100.0%     0.120   0.5%      4  __read
(libc.so.1: readSCI.c, 27)
   [7]     0.000   0.0%  100.0%     1.020   4.2%     34  libdso
(generic: generic.c, 619)
   [8]     0.000   0.0%  100.0%     1.020   4.2%     34  dlslave_routine
(dlslave.so: dlslave.c, 7)
```

2.3.1.3 Analyzing the Report

The report shows information for each function. The remaining columns are described below:

- The `index` column, which numbers the functions in the order in which they were encountered in the program, provides an index number for reference.
- The `excl.secs` column shows how much time, in seconds, was spent in the function itself (exclusive time). For example, less than one hundredth of a second was spent in `__start()`, but 0.03 of a second was spent in `fread`.
- The `excl.%` shows the percentage of a program's total time that was spent in the function. The `anneal` function consumed 93.8% of the program's time.
- The `cum.%` column shows the percentage of the complete program time that has been spent when the function completes.
- The `incl.secs` column shows how much time, in seconds, was spent in the function and descendents of the function. For example, 0.21 seconds were spent in `cvtttrap` and the functions that were called through it.
- The `incl.%` column shows the cumulative percentage of inclusive time spent in each function and its descendents. For example, 93.8% of the time was spent in `anneal` and all the functions that were called through it.
- The `samples` column provides the number of samples of the function and all of its descendants.
- The `procedure (dso:file,line)` columns list the function name, its DSO name, its file name, and its line number. For example, the top line reports statistics for the function `anneal`, the DSO `generic`, in the file `generic.c`, at line 1573.

2.3.2 pcsamp Experiment

This section explains how to perform a `pcsamp` experiment. The `pcsamp` experiment allows you to gather information on actual CPU time for each function in your program. For more information on `pcsamp`, see Chapter 4, page 49. For definitions of CPU time, wall-clock time, and process-virtual time, see the Glossary.

From the command line, enter

```
ssrun -fpcsamp generic
```

This starts the experiment. The `f` option is used with `pcsamp` for this program because the program runs quickly and does not gather much data using the default `pcsamp` experiment. Output from `generic` and from `ssrun` is printed to `stdout` as shown in the example below.

A data file is also generated. The name consists of the process name (*generic*), the experiment type (*fpcsamp*), and the experiment ID. In this example, the filename is *generic.fpcsamp.m11140*.

```
0:00:00.000 ===== (11140)          Begin script Mon  02 Feb 1998
10:58:41.
      begin script 'll.u.cvt.d.i.f.dso'
0:00:00.003 ===== (11140)          start of linklist Mon  02 Feb 1998
10:58:41.
      linklist completed.
0:00:00.004 ===== (11140)          start of usertime Mon  02 Feb 1998
10:58:41.
      usertime completed.
0:00:22.437 ===== (11140)          start of cvttrap Mon  02 Feb 1998
10:59:03.
      cvttrap completed, y = 2.60188e+14, z = 2.60188e+14.
0:00:22.638 ===== (11140)          start of dirstat Mon  02 Feb 1998
10:59:03.
      dirstat of /usr/include completed, 304 files.
0:00:23.407 ===== (11140) start of iofile -- stdio Mon  02 Feb 1998
10:59:04.
      stdio iofile on /unix completed, 7307988 chars.
0:00:23.750 ===== (11140)          start of fpetraps Mon  02 Feb 1998
10:59:04.
      fpetraps completed.
0:00:23.751 ===== (11140)          start of libdso Mon  02 Feb 1998
10:59:04.
dlslave_init executed
dlslave_routine executed
      slaveusertime completed, x = 5000000.000000.
      libdso: dynamic routine returned 13
      end of script 'll.u.cvt.d.i.f.dso'
0:00:24.778 ===== (11140)          End script Mon  02 Feb 1998
10:59:05.
```

2.3.2.1 Generating a Report

To generate a report on the data collected, and to redirect the output to a file, enter the following at the command line:

```
prof your_output_file_name > pcsamp.results
```

Output similar to the following is generated:

 SpeedShop profile listing generated Mon Feb 2 11:01:36 1998

prof generic.fpcsamp.ml1140
 generic (n32): Target program
 fpcsamp: Experiment name
 pc,2,1000,0:cu: Marching orders
 R4400 / R4000: CPU / FPU
 1: Number of CPUs
 175: Clock frequency (MHz.)

Experiment notes--

 From file generic.fpcsamp.ml1140:
 Caliper point 0 at target begin, PID 11140
 /home/hickory28/dep/speedshop/progs.etc/linpack.demos/c/generic
 Caliper point 1 at exit(0)

 Summary of statistical PC sampling data (fpcsamp)--

 23828: Total samples
 23.828: Accumulated time (secs.)
 1.0: Time per sample (msecs.)
 2: Sample bin width (bytes)

 Function list, in descending order by time

[index]	secs	%	cum.%	samples	function (dso: file, line)
[1]	22.279	93.5%	93.5%	22279	anneal (generic: generic.c, 1573)
[2]	0.975	4.1%	97.6%	975	slaveusertime (dlslave.so: dlslave.c, 22)
[3]	0.201	0.8%	98.4%	201	__read (libc.so.1: read.s, 20)
[4]	0.198	0.8%	99.3%	198	cvtttrap (generic: generic.c, 317)
[5]	0.121	0.5%	99.8%	121	__xstat (libc.so.1: xstat.s, 12)
[6]	0.010	0.0%	99.8%	10	__open (libc.so.1: open.s, 23)
[7]	0.010	0.0%	99.9%	10	__write (libc.so.1: write.s, 20)
[8]	0.010	0.0%	99.9%	10	__sigfillset (libc.so.1: sigfillset.c, 11)
[9]	0.010	0.0%	99.9%	10	__ecvt_r (libc.so.1: ecvt.c, 70)
[10]	0.003	0.0%	100.0%	3	fread (libc.so.1: fread.c, 34)
[11]	0.003	0.0%	100.0%	3	dirstat (generic: generic.c, 348)
[12]	0.002	0.0%	100.0%	2	__doprnt (libc.so.1: doprnt.c, 285)
[13]	0.001	0.0%	100.0%	1	memcpy (libc.so.1: bcopy.s, 329)
[14]	0.001	0.0%	100.0%	1	__readdir (libc.so.1: readdir.c, 135)
[15]	0.001	0.0%	100.0%	1	__read (libc.so.1: readSCI.c, 27)
[16]	0.001	0.0%	100.0%	1	__sinf (libm.so: fsin.c, 93)
	0.002	0.0%	100.0%	2	**OTHER** (includes excluded DSOs, rld, etc.)

```
23.828 100.0% 100.0%      23828  TOTAL
```

2.3.2.2 Analyzing the Report

The report has the following columns:

- The `secs` column shows the amount of CPU time that was spent in the function.
- The `%` column shows the percentage of the total program time that was spent in the function.
- The `cum.%` column shows the percentage of the complete program time that has been spent when the function completes.
- The `samples` column shows how many samples were taken when the process was executing in the function.
- The `function (dso: file, line)` columns list the function, its DSO name, its file name, and its line number.

2.3.3 Hardware Counter Experiment

Note: This experiment can be performed only on systems that have built-in hardware counters (the R10000 class of machines).

This section takes you through the steps to perform a hardware counter experiment. There are a number of hardware counter experiments, but this tutorial describes the steps involved in performing the `dsc_hwc` experiment. This experiment allows you to capture information about secondary data cache misses. For more information on hardware counter experiments, see Chapter 4, page 49.

2.3.3.1 Performing a Hardware Counter Experiment

From the command line, enter

```
ssrun -dsc_hwc generic
```

This starts the experiment. Output from `generic` and from `ssrun` is printed to `stdout`. A data file is also generated. The name consists of the process name (`generic`), the experiment type, `dsc_hwc`, and the experiment ID. In this example, the file name is `generic.dsc_hwc.m294398`.

2.3.3.2 Generating a Report

To generate a report on the data collected and redirect the output to a file, enter the following at the command line:

```
prof your_output_file_name > dsc_hwc.results
```

The report should look similar to the following listing:

```
-----
SpeedShop profile listing generated Mon Feb  2 11:11:44 1998
prof generic.dsc_hwc.m294398
    generic (n32): Target program
        dsc_hwc: Experiment name
        hwc,26,131:cu: Marching orders
R10000 / R10010: CPU / FPU
        16: Number of CPUs
        195: Clock frequency (MHz.)

Experiment notes--
    From file generic.dsc_hwc.m294398:
    Caliper point 0 at target begin, PID 294398
        /home/hickory28/dep/speedshop/progs.etc/linpack.demos/c/generic
    Caliper point 1 at exit(0)

-----
Summary of R10K perf. counter overflow PC sampling data (dsc_hwc)--
        6: Total samples
    Sec cache D misses (26): Counter name (number)
        131: Counter overflow value
        786: Total counts

-----
Function list, in descending order by counts
-----
[index]      counts      %   cum.%   samples  function (dso: file, line)

    [1]         131   16.7%   16.7%     1  init2da (generic: generic.c, 1430)
    [2]         131   16.7%   33.3%     1  genLog (generic: generic.c, 1686)
    [3]         131   16.7%   50.0%     1  _write (libc.so.1: writeSCI.c, 27)
           393   50.0%  100.0%     3  **OTHER** (includes excluded DSOs, rld, etc.)

           786 100.0% 100.0%     6  TOTAL
```

2.3.3.3 Analyzing the Report

The information immediately above the function list displays the following:

- The `Total samples` is the number of times the program counter was sampled. It is sampled once for each *overflow*, or each time the hardware counter exceeds the specified value.
- The `Counter name (number)` indicates the hardware counter used in the experiment. In this case, hardware counter 26 counts the number of times a value required in a calculation was not available in secondary cache. For a complete list of the hardware counters and their numbers, see Table 8, page 57.
- The `Counter overflow value` is the number at which the hardware counter overflows, or exceeds its preset value. In this case, the value is 131, which is the default. You can change the overflow value by setting the `_SPEEDSHOP_HWC_COUNTER_OVERFLOW` environment variable to a value larger than 0.
- The `Total counts` is the total number of times a value was not in secondary cache when needed. This value is determined by multiplying the total number of samples by the overflow value; extra counts that do not cause an overflow are not recorded.

The function list has the following columns:

- The `counts` column shows the number of times a data item was not in secondary cache when needed for a calculation during the execution of the function. As with `Total counts` (described earlier), a function's `counts` value is determined by multiplying its `samples` value (described later) by the overflow value.
- The `%` column shows the percentage of the program's overflows that occurred in the function.
- The `cum.%` shows the percentage of the program's overflows that occurred by the time the function finished execution. A function might have a low number in its `%` column but a high value in its `cum.%` column if it executed late in the program.
- The `samples` column shows the number of times the program counter was sampled during execution of the function. A sample is taken for each overflow of the hardware counter.
- The `function (dso: file, line)` columns show the name, the DSO, the file name, and line number of the function.

2.3.4 ideal Experiment

This section takes you through the steps to perform an `ideal` experiment. The times returned represent an idealized computation. This experiment ignores potential floating-point interlocks and memory latency time (cache misses and memory bus contention). The CPU times returned will always be lower than for an actual run. For more information on `ideal`, see Chapter 4, page 49.

2.3.4.1 Performing an ideal Experiment

From the command line, enter

```
ssrun -ideal generic
```

This starts the experiment. First the executable and libraries are instrumented using `pixie`. This entails making copies of the libraries and executables, giving the copies an extension that depends on the ABI, and inserting information into the copies. The extension is `.pixie` for the executable, `.pix32` for all 32 libraries, `.pixn32` for all n32 libraries, and `.pix64` for all 64 libraries.

Output from `generic` and from `ssrun` is printed to `stdout`. A data file is also generated. The name consists of the process name (`generic`), the experiment type, `ideal`, and the experiment ID. In this example, the filename is `generic.ideal.m10966`, and the output to `stdout` looks like the following:

```
Beginning libraries
  /usr/lib32/libssrt.so
  /usr/lib32/libss.so
  /usr/lib32/libm.so
  /usr/lib32/libc.so.1
Ending libraries, beginning "generic"
0:00:00.001 ===== (10966)          Begin script Mon  02 Feb 1998
11:28:03.
      begin script `ll.u.cvt.d.i.f.dso'
0:00:00.048 ===== (10966)          start of linklist Mon  02 Feb 1998
11:28:03.
      linklist completed.
0:00:00.072 ===== (10966)          start of usertime Mon  02 Feb 1998
11:28:03.
      usertime completed.
0:00:25.057 ===== (10966)          start of cvttrap Mon  02 Feb 1998
11:28:28.
      cvttrap completed, y = 2.60188e+14, z = 2.60188e+14.
0:00:25.377 ===== (10966)          start of dirstat Mon  02 Feb 1998
```

```
11:28:28.
    dirstat of /usr/include completed, 304 files.
0:00:26.232 ===== (10966) start of iofile -- stdio Mon  02 Feb 1998
11:28:29.
    stdio iofile on /unix completed, 7307988 chars.
0:00:27.716 ===== (10966)          start of fpetraps Mon  02 Feb 1998
11:28:31.
    fpetraps completed.
0:00:27.717 ===== (10966)          start of libdso Mon  02 Feb 1998
11:28:31.
Beginning libraries
Ending libraries, beginning "./dlslave.so"
dlslave_init executed
dlslave_routine executed
    slaveusertime completed, x = 5000000.000000.
    libdso: dynamic routine returned 13
    end of script 'll.u.cvt.d.i.f.dso'
0:00:30.021 ===== (10966)          End script Mon  02 Feb 1998
11:28:33.
```

The output section that starts with Beginning libraries and ends with Ending libraries, beginning 'generic' tells you that ssrun is instrumenting first the libraries listed in the executable, then the generic executable itself. The section that starts Beginning libraries and ends with Ending libraries, beginning 'dlslave.so' is added when the DSO dlslave.so is dlopen'd.

2.3.4.2 Generating a Report

To generate a report on the data collected, enter the following at the command line:

```
prof your_output_file_name > ideal.results
```

This command redirects output to a file called `ideal.results`. The file contains results that look similar to the following partial listing. Because of line length restrictions, the DSO, file name, and line number have been wrapped to the next line.

```
-----
SpeedShop profile listing generated Mon Feb  2 13:23:25 1998
  prof generic.ideal.m10966
    generic (n32): Target program
      ideal: Experiment name
```

```

          it:cu: Marching orders
R4400 / R4000: CPU / FPU
          1: Number of CPUs
          175: Clock frequency (MHz.)
Experiment notes--
  From file generic.ideal.m10966:
  Caliper point 0 at target begin, PID 10966
    /home/hickory28/dep/speedshop/progs.etc/linpack.demos/c/generic.pixie
  Caliper point 1 at exit(0)
-----
Summary of ideal time data (ideal)--
      2062563179: Total number of instructions executed
      3929944273: Total computed cycles
       22.457: Total computed execution time (secs.)
       1.905: Average cycles / instruction
-----
Function list, in descending order by exclusive ideal time
-----
[index]  excl.secs  excl.%   cum.%     cycles  instructions  calls  function
(dso: file, line)

   [1]    21.453   95.5%   95.5%   3754320037   1971220024    1  anneal
(generic: generic.c, 1573)
   [2]     0.829    3.7%   99.2%   145001152    75000732    1  slaveusertime
(dlslave.so: dlslave.c, 22)
   [3]     0.171    0.8%  100.0%   30000081    16000054    1  cvttrap
(generic: generic.c, 317)
   [4]     0.001    0.0%  100.0%    101504      58124    1  init2da
(generic: generic.c, 1430)
   [5]     0.001    0.0%  100.0%     91200   384001600    _drand48
(libc.so.1: drand48.c, 116)
   [6]     0.001    0.0%  100.0%     89072    55011 447  fread
(libc.so.1: fread.c, 34)
   [7]     0.000    0.0%  100.0%     74854    47366  53  _doprnt
(libc.so.1: doprnt.c, 285)
   [8]     0.000    0.0%  100.0%     64035    29479 628  __sinf
(libm.so: fsin.c, 93)
   [9]     0.000    0.0%  100.0%     32355      7182   9  offtime
(libc.so.1: time_comm.c, 180)
  [10]     0.000    0.0%  100.0%     17112    11916 305  _readdir
(libc.so.1: readdir.c, 135)
  [11]     0.000    0.0%  100.0%     16168    10334   1  iofile
(generic: generic.c, 464)

```

```
[12]      0.000    0.0%   100.0%      15232      12544 448  _read
(libc.so.1: readSCI.c, 27)
[13]      0.000    0.0%   100.0%      14530      8498 326  memcpy
(libc.so.1: bcopy.s, 329)
[14]      0.000    0.0%   100.0%      10735      6446  1  dirstat
(generic: generic.c, 348)
[15]      0.000    0.0%   100.0%      6535      2831 106  strlen
(libc.so.1: strlen.s, 58)
[16]      0.000    0.0%   100.0%      6364      4242 304  _xstat
(libc.so.1: xstat.s, 12)
[17]      0.000    0.0%   100.0%      6363      3636 303  _cerror
(libc.so.1: cerror.s, 30)
.
.
.
[129]     0.000    0.0%   100.0%      5          3  1  get_exit_status
(libss.so: sswrap_assembly.s, 6)
[130]     0.000    0.0%   100.0%      4          2  1  __readenv_sigfpe
(libc.so.1: stubfpestart.c, 3)
[131]     0.000    0.0%   100.0%      4          2  1  crtninit.s
(generic: crtninit.s, 3)
[132]     0.000    0.0%   100.0%      1          1  1  __istart
(generic: crtltinit.s, 14)
```

2.3.4.3 Analyzing the Report

The columns in the report provide the following information

- The `excl.secs` column shows the minimum number of seconds that might be spent in the function under ideal conditions. For example, 21.453 seconds is optimal for the `anneal` function. The `pcsamp` experiment actually timed this function at 22.279 seconds.
- The `excl.%` represents how much of the program's total time was spent in the function.
- The `cum.%` column shows the cumulative percentage of time spent so far in the program after the function executes.
- The `cycles` column reports the number of machine cycles used by the function. For example, 3754320037 cycles were spent in the `anneal` function.

- The `instructions` column shows the number of instructions executed by a function. For example, the `anneal` function executed 1971220024 instructions.
- The `calls` column reports the number of calls to the function. For example, there was just one call to the `anneal` function.
- The `procedure (dso: file, line)` column lists the procedure, its DSO name, its file name, and the line number. For example, the first line reports statistics for the function `anneal` in the file `generic.c` in the `generic` executable on line 1573.

2.3.5 fpe Trace

This section takes you through the steps to perform a floating-point exception (fpe) trace, which identifies functions in which floating-point exceptions have occurred. For more information on the fpe trace, see Chapter 4, page 49.

2.3.5.1 Performing an fpe Trace

From the command line, enter

```
ssrun -fpe generic
```

This starts the experiment. Output from `generic` and from `ssrun` is printed to `stdout`. A data file with a name generated by concatenating the process name, `generic`, the experiment type, `fpe`, and the experiment ID is also generated. In this example, the filename is `generic.fpe.m12213`.

2.3.5.2 Generating a Report

To generate a report on the data collected, enter the following at the command line:

```
prof your_output_file_name > fpe.results
```

The report should look similar to the following partial listing:

```
-----
SpeedShop profile listing generated Mon Feb  2 13:26:33 1998
prof generic.fpe.m12213
    generic (n32): Target program
                  fpe: Experiment name
                  fpe:cu: Marching orders
```

```
R4400 / R4000: CPU / FPU
          1: Number of CPUs
          175: Clock frequency (MHz.)
Experiment notes--
  From file generic.fpe.m12213:
  Caliper point 0 at target begin, PID 12213
    /home/hickory28/dep/speedshop/progs.etc/linpack.demos/c/generic
  Caliper point 1 at exit(0)
-----
Summary of FPE callstack tracing data (fpe)--
          4: Total FPEs
          0: Samples with incomplete traceback
-----
Function list, in descending order by exclusive FPEs
-----
[index  excl.FPEs excl.  cum.% incl.FPEs incl.%  function  (dso:file)

[1]      4 100.0% 100.0%      4 100.0% fpetraps (generic: generic.c, 405)
[2]      0  0.0% 100.0%      4 100.0% __start (generic: crt1text.s, 101)
[3]      0  0.0% 100.0%      4 100.0% main (generic: generic.c, 101)
[4]      0  0.0% 100.0%      4 100.0% Scriptstring (generic: generic.c, 184)
```

2.3.5.3 Analyzing the Report

The report shows information for each function. The function names are shown in the right-hand column of the report. The remaining columns are described below:

- The `excl.FPEs` column shows how many floating point exceptions were found in the function. For example, four floating point exceptions were found in `fpetraps`.
- The `excl.%` column shows the percentage of the total number of floating-point exceptions that were found in the function.
- The `cum.%` column shows the percentage of floating-point exceptions in the program that have been encountered so far in the list. The list is sorted by floating-point exceptions, with the most in the top line and the least in the bottom line. Because all of the exceptions are in the first function listed in this example, all entries in this column are 100%.
- The `incl.FPEs` columns shows how many floating-point exceptions were found so far in the program.

- The `incl.%` column provides information on the percentage of the program's total number of floating-point exceptions found so far. Because the `fpetraps` function, which contains 100% of the program's floating-point exceptions, finishes first, all of the subsequent functions also show 100% under this column.

Tutorial for Fortran Users [3]

This chapter presents a tutorial for using the SpeedShop tools to gather and analyze performance data in a Fortran program, and covers these topics:

- Section 3.1, page 33, introduces the sample program and explains the different scenarios in which it will be used.
- Section 3.2, page 34, steps you through the necessary setup for running the experiment.
- Section 3.3, page 35, steps you through five different experiments, discussing first how to do the experiments, then how to interpret the results.

Note: Because of inherent differences between systems and also due to concurrent processes that may be running on your system, your experiment will produce different results from the one in this tutorial. However, the basic form of the results should be the same.

3.1 Tutorial Overview

This tutorial is based on a standard benchmark program called `linpackup`. There are two versions of the program: the `linpack` directory contains files for the n32-bit ABI, and the `linpacko32` directory contains files for the o32-bit ABI. Each `linpack` directory contains versions of the program for a single processor (`linpackup`) and for multiple processors (`linpackd`). When you work with the tutorial, choose the version of the program that is most appropriate for your system. A good guideline is to choose whichever version corresponds to the way you expect to develop your programs.

This tutorial was written and tested using the single-processor version of the program (`linpackup`) in the `linpack` directory.

The `linpack` program is a standard benchmark designed to measure CPU performance in solving dense linear equations. The program focuses primarily on floating-point performance.

Output from the `linpackup` program looks like the following:

```
.
.
.
      norm. resid      resid      machep      x(1)      x(n)
5.35882395E+00  7.13873405E-13  2.22044605E-16  1.00000000E+00  1.00000000E+00

      times are reported for matrices of order 300
      dgefa      dgesl      total      mflops      unit      ratio
times for array with leading dimension of 301
3.720E+00  4.000E-02  3.760E+00  4.835E+00  4.136E-01  6.714E+01
3.780E+00  3.000E-02  3.810E+00  4.772E+00  4.191E-01  6.804E+01
3.730E+00  4.000E-02  3.770E+00  4.822E+00  4.147E-01  6.732E+01
3.730E+00  4.000E-02  3.770E+00  4.822E+00  4.147E-01  6.732E+01

times for array with leading dimension of 300
3.800E+00  4.000E-02  3.840E+00  4.734E+00  4.224E-01  6.857E+01
3.810E+00  4.000E-02  3.850E+00  4.722E+00  4.235E-01  6.875E+01
3.770E+00  4.000E-02  3.810E+00  4.772E+00  4.191E-01  6.804E+01
3.782E+00  4.000E-02  3.822E+00  4.757E+00  4.205E-01  6.825E+01
```

3.2 Tutorial Setup

Copy the program to a directory where you have write permission and compile it so that you can use it in the tutorial.

1. Change to the `/usr/demos/SpeedShop` directory.
2. Copy the appropriate `linpack` directory and its contents to a directory in which you have write permission:

```
cp -r linpack your_dir
```

3. Change to the directory you just created:

```
cd your_dir/linpack
```

4. Compile the program by entering:

```
make all
```

This provides an executable for the experiment.

3.3 Analyzing Performance Data

This section provides steps on how to run the following experiments on the `linpackup` program and generate and interpret the results:

- The `usertime` experiment. It returns the *CPU time* (see the Glossary for a definition) used by each routine in your program. See Section 3.3.1, page 35.
- The `pcsamp` experiment. It returns CPU time for each routine in your program. See Section 3.3.2, page 39.
- Hardware counter experiment. In a hardware counter experiment, the program counter (PC) is sampled every time a hardware counter exceeds a specified limit. In the experiment performed in this section, the hardware counter keeps track of the number of times a data item required in a calculation was not present in secondary data cache. When a data item is not in cache, it must be retrieved from memory, which is a more time-consuming process. See Section 3.3.3, page 41.
- The `ideal` experiment. This experiment calculates the best time achievable. See Section 3.3.4, page 43.

3.3.1 usertime Experiment

This section takes you through the steps to perform a `usertime` experiment. The `usertime` experiment allows you to gather data on the amount of *CPU time* spent in each routine in your program. For more information on `usertime`, see Chapter 4, page 49. For definitions of CPU time, wall-clock time, and process-virtual time, see the Glossary.

3.3.1.1 Performing a usertime Experiment

From the command line, enter

```
ssrun -v -usertime linpackup
```

This starts the experiment. The `-v` flag tells `ssrun` to print a log to *stderr*.

Output from `linpackup` and from `ssrun` is printed to *stdout* as shown in the example below. A data file is also generated. The name consists of the process name (`linpackup`), the experiment type, `usertime`, and the experiment ID. In this example, the filename is `linpackup.usertime.m12205`.

```
ssrun: target PID 12205
ssrun: setenv _SPEEDSHOP_MARCHING_ORDERS ut:cu
ssrun: setenv _SPEEDSHOP_EXPERIMENT_TYPE usertime
ssrun: setenv _SPEEDSHOP_TARGET_FILE linpackup
ssrun: setenv _RLD_LIST libss.so:libssrt.so:DEFAULT
Please send the results of this run to:
```

Jack J. Dongarra
Mathematics and Computer Science Division
Argonne National Laboratory
Argonne, Illinois 60439

Telephone: 312-972-7246

ARPAnet: DONGARRA@ANL-MCS

```
      norm. resid      resid      machep      x(1)      x(n)
5.35882395E+00  7.13873405E-13  2.22044605E-16  1.00000000E+00  1.00000000E+00
```

```
      times are reported for matrices of order 300
      dgefa      dgesl      total      mflops      unit      ratio
times for array with leading dimension of 301
3.960E+00  4.000E-02  4.000E+00  4.545E+00  4.400E-01  7.143E+01
3.960E+00  4.000E-02  4.000E+00  4.545E+00  4.400E-01  7.143E+01
3.970E+00  4.000E-02  4.010E+00  4.534E+00  4.411E-01  7.161E+01
3.960E+00  4.000E-02  4.000E+00  4.545E+00  4.400E-01  7.143E+01

times for array with leading dimension of 300
3.910E+00  4.000E-02  3.950E+00  4.603E+00  4.345E-01  7.054E+01
3.880E+00  8.000E-02  3.960E+00  4.591E+00  4.356E-01  7.071E+01
3.930E+00  4.000E-02  3.970E+00  4.579E+00  4.367E-01  7.089E+01
3.922E+00  3.800E-02  3.960E+00  4.591E+00  4.356E-01  7.071E+01
```

3.3.1.2 Generating a Report

To generate a report on the data collected, enter the following at the command line:

```
prof your_output_file_name > usertime.results
```


`prof` interprets the type of experiment you have performed and prints results to *stdout*. The report below shows partial `prof` output, and most lines have been wrapped because of line width restrictions:

```
-----
SpeedShop profile listing generated Mon Feb  2 13:37:38 1998
  prof linpackup.usertime.ml2205
    linpackup (n32): Target program
      usertime: Experiment name
      ut:cu: Marching orders
    R4400 / R4000: CPU / FPU
      1: Number of CPUs
      175: Clock frequency (MHz.)

Experiment notes--
  From file linpackup.usertime.ml2205:
  Caliper point 0 at target begin, PID 12205
    /home/hickory28/dep/speedshop/progs.etc/linpack.demos/fortran/linpackup
  Caliper point 1 at exit(0)
-----

Summary of statistical callstack sampling data (usertime)--
  2777: Total Samples
    0: Samples with incomplete traceback
  83.310: Accumulated Time (secs.)
  30.0: Sample interval (msecs.)
-----

Function list, in descending order by exclusive time
-----
[index]  excl.secs  excl.%   cum.%   incl.secs  incl.%   samples  procedure
(dso: file, line)

   [5]    78.090  93.7%   93.7%    78.090  93.7%    2603  daxpy
(linpackup: linpackup.f, 495)
   [6]     2.730   3.3%   97.0%     2.730   3.3%     91  matgen
(linpackup: linpackup.f, 199)
   [4]     1.920   2.3%   99.3%     1.920  95.6%    2656  dgefa
(linpackup: linpackup.f, 221)
   [8]     0.270   0.3%   99.6%     0.270   0.3%     9  dscal
(linpackup: linpackup.f, 670)
   [9]     0.180   0.2%   99.9%     0.180   0.2%     6  idamax
(linpackup: linpackup.f, 700)
  [10]     0.090   0.1%  100.0%     0.090   0.1%     3  dmxpy
(linpackup: linpackup.f, 826)
   [7]     0.030   0.0%  100.0%     0.810   1.0%     27  dgesl
-----
```

```
(linpackup: linpackup.f, 324)
  [1]    0.000  0.0% 100.0%    83.310 100.0%    2777  __start
(linpackup: crtlttext.s, 101)
  [2]    0.000  0.0% 100.0%    83.310 100.0%    2777  main
(libftn.so: main.c, 76)
  [3]    0.000  0.0% 100.0%    83.310 100.0%    2777  linp
(linpackup: linpackup.f, 3)
```

3.3.1.3 Analyzing the Report

The report shows information for each function. The function names are shown in the right-hand column of the report. The remaining columns are described below:

- The `index` column, which numbers the routines in the order in which they were encountered in the program, provides an index number for reference.
- The `excl.secs` column shows how much time, in seconds, was spent in the routine itself (exclusive time). For example, less than one hundredth of a second was spent in `linp`, but 1.92 seconds were spent in `dgefa`.
- The `excl.%` shows the percentage of a program's total time that was spent in the function. The `daxpy` routine consumed 93.7% of the program's time.
- The `cum.%` column shows the percentage of the complete program time that has been spent when the routine completes. For instance, when the `dgefa` routine completes, 99.3% of the program has completed.
- The `incl.secs` column shows how much time, in seconds, was spent in the function and descendants of the function. For example, 0.81 seconds were spent in `dges` and the routines that were called through it.
- The `incl.%` column shows the cumulative percentage of inclusive time spent in each routine and its descendants. For example, 1% of the time was spent in `dges` and all the routines that were called through it.
- The `samples` column provides the number of samples of the routine and all of its descendants.
- The `procedure (dso:file,line)` column lists the routine name, its DSO name, its file name, and its line number. For example, the top line reports statistics for the routine `daxpy`, the DSO name `linpackup`, in the file `linpackup.f`, at line 495.

Note: Many functions shown here have only one or two “hits.” The data for those functions is not statistically significant.

3.3.2 pcsamp Experiment

This section takes you through the steps to perform a `pcsamp` experiment. The `pcsamp` experiment allows you to gather information on actual CPU time for each source code line, machine line, and function in your program. For more information on `pcsamp`, see Chapter 4, page 49. For definitions of CPU time, wall-clock time, and process-virtual time, see the Glossary.

3.3.2.1 Performing a pcsamp Experiment

From the command line, enter

```
ssrun -pcsamp linpackup
```

This starts the experiment.

Output from `linpackup` and from `ssrun` is printed to `stdout` as shown in the example below. A data file is also generated. The name consists of the process name (`linpackup`), the experiment type, `pcsamp`, and the experiment ID. In this example, the file name is `linpackup.pcsamp.m12333`.

```
.
.
.
      norm. resid      resid      machep      x(1)      x(n)
5.35882395E+00  7.13873405E-13  2.22044605E-16  1.00000000E+00  1.00000000E+00
.
.
.
```

3.3.2.2 Generating a Report

To generate a report on the data collected, enter the following at the command line:

```
prof your_output_file_name > pcsamp.results
```

```
-----
SpeedShop profile listing generated Mon Feb  2 13:52:27 1998
prof linpackup.pcsamp.m12333
      linpackup (n32): Target program
                pcsamp: Experiment name
pc,2,10000,0:cu: Marching orders
      R4400 / R4000: CPU / FPU
                1: Number of CPUs
```

```

                                175: Clock frequency (MHz.)
Experiment notes--
    From file linpackup.pcsamp.m12333:
    Caliper point 0 at target begin, PID 12333
    /home/hickory28/dep/speedshop/progs.etc/linpack.demos/fortran/linpackup
    Caliper point 1 at exit(0)
-----
Summary of statistical PC sampling data (pcsamp)--
    8272: Total samples
    82.720: Accumulated time (secs.)
    10.0: Time per sample (msecs.)
    2: Sample bin width (bytes)
-----
Function list, in descending order by time
-----
[index]      secs      %      cum.%    samples  function (dso: file, line)
-----
    [1]      77.440  93.6%  93.6%     7744  daxpy (linpackup: linpackup.f, 495)
    [2]       2.690   3.3%  96.9%      269  matgen (linpackup: linpackup.f, 199)
    [3]       1.940   2.3%  99.2%      194  dgefa (linpackup: linpackup.f, 221)
    [4]       0.370   0.4%  99.7%       37  idamax (linpackup: linpackup.f, 700)
    [5]       0.210   0.3%  99.9%       21  dscal (linpackup: linpackup.f, 670)
    [6]       0.060   0.1% 100.0%        6  dmxpy (linpackup: linpackup.f, 826)
    0.010   0.0% 100.0%        1  **OTHER** (includes excluded DSOs, rld, etc.)

    82.720 100.0% 100.0%     8272  TOTAL

```

3.3.2.3 Analyzing the Report

The report has the following columns:

- The `secs` column shows the amount of CPU time spent in the routine.
- The `(%)` column shows the percentage of the total program time that was spent in the function.
- The `cum.%` column shows the percentage of the complete program time that has been spent when the function completes.
- The `samples` column shows how many samples were taken when the process was executing in the function.
- The `function (dso:file, line)` columns list the routine name, its DSO name, its file name, and its line number. For example, the first line

reports statistics for the routine `daxpy`, in the DSO `linpackup`, in the file `linpackup.f`, at line number 495.

3.3.3 Hardware Counter Experiment

Note: This experiment can be performed only on systems that have built-in hardware counters (the R10000 class of machines).

This section takes you through the steps to perform a hardware counter experiment. There are a number of hardware counter experiments, but this tutorial describes the steps involved in performing the `dsc_hwc` experiment. This experiment allows you to capture information about secondary data cache misses. For more information on hardware counter experiments, see Chapter 4, page 49.

3.3.3.1 Performing a hardware counter Experiment

From the command line, enter

```
ssrun -dsc_hwc linpackup
```

This starts the experiment. Output from `linpackup` and from `ssrun` will be printed to *stdout*. A data file is also generated. The name consists of the process name (`linpackup`), the experiment type, `dsc_hwc`, and the experiment ID. In this example, the filename is `linpackup.dsc_hwc.m438011`.

3.3.3.2 Generating a Report

To generate a report on the data collected, enter the following at the command line:

```
prof your_output_file_name > dsc_hwc.results
```

Output similar to the following is generated:

```
-----
SpeedShop profile listing generated Mon Feb  2 13:56:59 1998
prof linpackup.dsc_hwc.m438011
  linpackup (n32): Target program
    dsc_hwc: Experiment name
  hwc,26,131:cu: Marching orders
R10000 / R10010: CPU / FPU
    16: Number of CPUs
    195: Clock frequency (MHz.)
```

```

Experiment notes--
    From file linpackup.dsc_hwc.m438011:
    Caliper point 0 at target begin, PID 438011
    /home/hickory28/dep/speedshop/progs.etc/linpack.demos/fortran/linpackup
    Caliper point 1 at exit(0)
-----
Summary of R10K perf. counter overflow PC sampling data (dsc_hwc)--
    2929: Total samples
    Sec cache D misses (26): Counter name (number)
        131: Counter overflow value
    383699: Total counts
-----
Function list, in descending order by counts
-----
[index]   counts    %   cum.% samples  function (dso: file, line)
-----
    [1]   309029  80.5%  80.5%   2359  daxpy (linpackup: linpackup.f, 495)
    [2]   46636  12.2%  92.7%    356  dgefa (linpackup: linpackup.f, 221)
    [3]   25938   6.8%  99.5%    198  matgen (linpackup: linpackup.f, 199)
    [4]    1310   0.3%  99.8%     10  idamax (linpackup: linpackup.f, 700)
    [5]     131   0.0%  99.8%      1  _FWF (libfortran.so: wf90.c, 47)
    [6]     131   0.0%  99.9%      1  memset (libc.so.1: bzero.s, 98)
    [6]     524   0.1% 100.0%      4  **OTHER** (includes excluded DSOs, rld, etc.)

    383699 100.0% 100.0%   2929  TOTAL

```

3.3.3.3 Analyzing the Report

The information immediately above the function list displays the following:

- The `Total samples` is the number of times the program counter was sampled. It is sampled once for each *overflow*, or each time the hardware counter exceeds the specified value.
- The `Counter name (number)` indicates the hardware counter used in the experiment. In this case, hardware counter 26 counts the number of times a value required in a calculation was not available in secondary cache. For a complete list of the hardware counters and their numbers, see Table 8, page 57.
- The `Counter overflow value` is the number at which the hardware counter overflows, or exceeds its preset value. In this case, the value is 131, which is the default. You can change the overflow value by setting the

`_SPEEDSHOP_HWC_COUNTER_OVERFLOW` environment variable to a value larger than 0.

- The `Total counts` is the total number of times a value was not in secondary cache when needed. This value is determined by multiplying the total number of samples by the overflow value; extra counts that do not cause an overflow are not recorded.

The function list has the following columns:

- The `counts` column shows the number of times a data item was not in secondary cache when needed for a calculation during the execution of the routine. As with `Total counts` (described earlier), a routine's `counts` value is determined by multiplying its `samples` value (described later) by the overflow value.
- The `%` column shows the percentage of the program's overflows that occurred in the routine.
- The `cum.%` shows the percentage of the program's overflows that occurred by the time the routine finished execution. For example, although the `matgen` routine had only 6.8% of the program's overflows, by the end of its execution, 99.5% of the program's total overflows had happened.
- The `samples` column shows the number of times the program counter was sampled during execution of the routine. A sample is taken for each overflow of the hardware counter.
- The `function (dso: file, line)` columns show the name, the DSO, the file name, and line number of the routine.

3.3.4 ideal Experiment

This section takes you through the steps to perform an `ideal` experiment. The times returned represent an idealized computation. This experiment ignores potential floating-point interlocks and memory latency time (cache misses and memory bus contention). The CPU times returned will always be lower than for an actual run. For more information on collecting ideal-time data, and basic block counting, see Chapter 4, page 49.

3.3.4.1 Performing an ideal Experiment

From the command line, enter

```
ssrun -ideal linpackup
```

This starts the experiment. First the executable and libraries are instrumented using `pixie`. This entails making copies of the libraries and executables, giving them an extension that depends on the ABI, and inserting information into the copies. The extension is `.pixie` for the executable, `.pix32` for all 32 libraries, `.pixn32` for all n32 libraries, and `.pix64` for all 64 libraries.

Output from `linpackup` and from `ssrun` is printed to `stdout` as shown in the example below. A data file is also generated. The name consists of the process name (`linpackup`), the experiment type, `ideal`, and the experiment ID. In this example, the file name is `linpackup.ideal.n11596`.

```
Beginning libraries
./libssrt.so.pixn32 is up to date.
./libss.so.pixn32 is up to date.
./libfortran.so.pixn32 is up to date.
./libffio.so.pixn32 is up to date.
./libftn.so.pixn32 is up to date.
./libm.so.pixn32 is up to date.
./libc.so.1.pixn32 is up to date.
Ending libraries, beginning "linpackup"
.
.
.
```

3.3.4.2 Generating a Report

To generate a report on the data collected, enter the following at the command line:

```
prof your_output_file_name > ideal.results
```

This command redirects output to a file called `ideal.results`. The file should contain results that look something like the following. Most lines have been wrapped because of line length restrictions.

```
-----
SpeedShop profile listing generated Mon Feb  2 14:04:20 1998
prof linpackup.ideal.m11596
  linpackup (n32): Target program
                   ideal: Experiment name
                   it:cu: Marching orders
R4400 / R4000: CPU / FPU
                   1: Number of CPUs
                   175: Clock frequency (MHz.)
```


Experiment notes--

From file linpackup.ideal.m11596:
 Caliper point 0 at target begin, PID 11596
 /home/hickory28/dep/speedshop/progs.etc/linpack.demos/fortran/linpackup.pixie
 Caliper point 1 at exit(0)

 Summary of ideal time data (ideal)--

4911547956: Total number of instructions executed
 9700441338: Total computed cycles
 55.431: Total computed execution time (secs.)
 1.975: Average cycles / instruction

 Function list, in descending order by exclusive ideal time

[index]	excl.secs	excl.%	cum.%	cycles	instructions	calls
function	(dso: file, line)					
[1]	52.073	93.9%	93.9%	9112833799	4637546756	772633
daxpy (linpackup: linpackup.f, 495)						
[2]	1.937	3.5%	97.4%	339051600	163885662	18
matgen (linpackup: linpackup.f, 199)						
[3]	1.020	1.8%	99.3%	178526333	72336088	17
dgefa (linpackup: linpackup.f, 221)						
[4]	0.180	0.3%	99.6%	31463770	17658342	5083
dscal (linpackup: linpackup.f, 670)						
[5]	0.166	0.3%	99.9%	28990712	15670260	5083
idamax (linpackup: linpackup.f, 700)						
[6]	0.045	0.1%	100.0%	7839357	3605134	1
dmxpy (linpackup: linpackup.f, 826)						
[7]	0.009	0.0%	100.0%	1499774	695929	17
dgesl (linpackup: linpackup.f, 324)						
[8]	0.000	0.0%	100.0%	54065	30649	53
_sd2udee (libffio.so: sd2udee.c, 104)						
[9]	0.000	0.0%	100.0%	44650	28904	1
linp (linpackup: linpackup.f, 3)						
[10]	0.000	0.0%	100.0%	37376	26716	31
_wrfmt (libfortran.so: wrfmt.c, 56)						
[11]	0.000	0.0%	100.0%	11448	8427	159
_IEEE_BINARY_SCALE_I4 (libfortran.so: ieee_binary_scale_r_n.c, 41)						
[12]	0.000	0.0%	100.0%	8784	5116	492
nvmatch (libc.so.1: getenv.c, 46)						
[13]	0.000	0.0%	100.0%	8586	4596	6
getenv (libc.so.1: getenv.c, 25)						

[14]	0.000	0.0%	100.0%	7850	4718	4
_get_next_unit (libfortran.so: fortunite.c, 171)						
[15]	0.000	0.0%	100.0%	7370	5003	15
_FWF (libfortran.so: wf90.c, 47)						
[16]	0.000	0.0%	100.0%	6012	3586	21
_pack (libffio.so: _pack.c, 51)						
[17]	0.000	0.0%	100.0%	5293	2273	16
strlen (libc.so.1: strlen.s, 58)						
[18]	0.000	0.0%	100.0%	3992	2888	8
_map_to_dv (libfortran.so: dopexfer.c, 964)						
[19]	0.000	0.0%	100.0%	3757	2289	9
fflush (libc.so.1: flush.c, 377)						
[20]	0.000	0.0%	100.0%	3394	2352	28
_fwch (libfortran.so: fwch.c, 61)						
[21]	0.000	0.0%	100.0%	3248	1680	8
_stride_dv (libfortran.so: dopexfer.c, 495)						
[22]	0.000	0.0%	100.0%	2809	1908	53
_IEEE_EXPONENT_I4_R (libfortran.so: ieee_exponent_n.c, 73)						
[23]	0.000	0.0%	100.0%	2649	1607	18
_unpack (libffio.so: _unpack.c, 54)						
[24]	0.000	0.0%	100.0%	2544	1908	159
__is_nan64 (libfortran.so: inline.h, 343; compiled in ieee_binary_scale_r_n.c)						
[25]	0.000	0.0%	100.0%	2404	1696	15
setup_format (libfortran.so: f90io.h, 451; compiled in wf90.c)						
[26]	0.000	0.0%	100.0%	2052	1296	54
second_ (linpackup: second.c, 8)						
[27]	0.000	0.0%	100.0%	1988	1764	28
_sw_endrec (libfortran.so: wf.c, 906)						
[28]	0.000	0.0%	100.0%	1712	1252	15
_xfer_iolist (libfortran.so: dopexfer.c, 150)						
[29]	0.000	0.0%	100.0%	848	636	53
__is_nan64 (libfortran.so: inline.h, 343; compiled in ieee_exponent_n.c)						
[30]	0.000	0.0%	100.0%	795	530	53
isdigit (libc.so.1: ctypefcns.c, 62)						
[31]	0.000	0.0%	100.0%	736	72	8
_tripcnt (libfortran.so: dopexfer.c, 1172)						
[32]	0.000	0.0%	100.0%	694	425	3
_s2ui (libffio.so: s2uboiz.c, 394)						
.						
.						
.						
[130]	0.000	0.0%	100.0%	1	1	1
__istart (linpackup: crt1tinit.s, 14)						

3.3.4.3 Analyzing the Report

The report has the following columns:

- The `excl.secs` column shows the minimum number of seconds that might be spent in the routine under ideal conditions. For example, 52.073 seconds is optimal for the `daxpy` routine. The `pcsamp` experiment (see Section 3.3.2, page 39) times this routine at 77.44 seconds.
- The `excl.%` represents how much of the program's total time was spent in the routine.
- The `cum.%` column shows the cumulative percentage of time spent so far in the program after the routine executes.
- The `cycles` column reports the number of machine cycles used by the routine. For example, 9112833799 cycles were spent in the `daxpy` routine.
- The `instructions` column shows the number of instructions executed by a routine. For example, the `dgefa` routine executed 72336088 instructions.
- The `calls` column reports the number of calls to the routine. For example, there was just one call to the `dmxpy` routine.
- The `procedure (dso:file, line)` column lists the name, the DSO name, the file name, and the line number for the routine.

Experiment Types [4]

This chapter provides detailed information on each experiment type available within SpeedShop. It contains the following sections:

- Selecting an experiment. See Section 4.1, page 49.
- The `usertime` experiment. See Section 4.2, page 50.
- The `pcsamp` experiment. See Section 4.3, page 51.
- The ideal experiment. See Section 4.4, page 51.
- The hardware counter experiments. See Section 4.5, page 54.
- The `fpe` trace experiment. See Section 4.6, page 59.

For information on how to run the experiments described in this chapter, see Chapter 6, page 63.

4.1 Selecting an Experiment

Table 6 shows the possible experiments you can perform using the SpeedShop tools and the reasons why you might want to choose a specific experiment. The Clues column shows when you might use an experiment. The Data Collected column indicates performance data collected by the experiment. For detailed information on the experiments listed, see the sections listed in Table 6.

Table 6. Summary of Experiments

Experiment	Clues	Data Collected
<code>usertime</code>	Slow program, nothing else known. Not CPU-bound.	Inclusive and exclusive CPU time for each function by sampling the callstack at 30-millisecond intervals.
<code>pcsamp</code>	High user CPU time.	Actual CPU time at the source line, machine instruction and function levels by sampling the program counter at 10-or 1-millisecond intervals.

Experiment	Clues	Data Collected
ideal	CPU-bound.	Ideal CPU time at the function, source line and machine instruction levels using instrumentation for basic block counting.
ideal	High user CPU time.	On R10000 class machines, exclusive counts at the source line, machine instruction, and function levels for overflows of the following counters: clock cycle, graduated instructions, primary instruction-cache misses, secondary instruction-cache misses, primary data-cache misses, secondary data-cache misses, TLB misses, graduated floating-point instructions.
fpe	High system time. Presence of floating point operations.	All floating point exceptions with the exception type and the callstack at the time of the exception.

4.2 usertime Experiment

The `usertime` experiment uses statistical call stack profiling, based on wall clock time, to measure inclusive and exclusive user time spent in each function while your program runs. This experiment uses an interval of 30 milliseconds.

Data is measured by periodically sampling the callstack. The program's callstack data is used to

- attribute exclusive user time to the function at the bottom of each callstack (that is, the function being executed at the time of the sample)
- attribute inclusive user time to all the functions above the one currently being executed

The time spent in a procedure is determined by multiplying the number of times an instruction for that procedure appears in the stack by the average time interval between call stacks. Call stacks are gathered whether the program was running or blocked; hence, the time computed represents the total time, both within and outside the CPU. If the target process was blocked for a long time as a result of an instruction, that instruction will show up as having a high time.

User time runs should incur a program execution slowdown of no more than 15%. Data from a `usertime` experiment is statistical in nature and shows some variance from run to run.

Note: For this experiment, o32 executables must explicitly link with `-lexc`.

4.3 pcsamp Experiment

The `pcsamp` experiment uses statistical PC sampling to estimate actual CPU time for each source code line, machine line, and function in your program. The `prof` listing of this experiment shows exclusive PC-sampling time. This experiment is a lightweight, high-speed operation done with kernel support. The actual CPU time is calculated by multiplying the number of times an instruction appears in the PC by the interval specified for the experiment (either 1 or 10 milliseconds.)

To collect the data, the kernel regularly stops the process if it is in the CPU, increments a counter for the current value of the PC, and resumes the process. The default sample interval is 10 milliseconds. If you specify the optional `f` prefix to the experiment, a sample interval of 1 millisecond is used.

By default, the experiment uses 16-bit bins, based on user and system time. If the optional `x` suffix is used, a 32-bit bin size will be used. Using a 32-bit bin provides more accurate information, but requires additional disk space.

- 16-bit bins allow a maximum of 65,000 counts.
- 32-bit bins allow approximately 4,000,000 counts.

PC-sampling time runs should incur a slowdown of execution of the program of no more than 5%. The measurements are statistical in nature, and exhibit variance inversely proportional to the running time.

4.4 ideal Experiment

The `ideal` experiment instruments the executables and any DSOs to permit basic block counting and counting of all dynamic (function-pointer) calls.

4.4.1 How SpeedShop Prepares Files

To permit block counting, SpeedShop

- Divides the code into basic blocks, which are sets of instructions with a single entry point, a single exit point, and no branches into or out of the set.
- Inserts counter code at the beginning of each basic block to increment a counter each time that basic block is executed.

The target executable and all the DSOs it uses are instrumented, including `libc.so.1`, `libexc.so`, `libm.so`, `libss.so`, `libssrt.so`. Instrumented

files with an extension `.pix*`, where `*` depends on the ABI, are written to the current working directory.

After the transformations are complete, the program's symbol table and translation table are updated so that debuggers can map between transformed addresses and the original program's addresses, and reference the measured performance data to the untransformed code.

After instrumentation, `ssrun` executes the instrumented program. Data is generated as long as the process exits normally or receives a fatal signal that the program does not handle.

4.4.2 How SpeedShop Calculates CPU Time

`prof` uses a machine model to convert the block execution counts into an idealized exclusive time at the function, source line, or machine instruction levels. By default, the machine model corresponds to the machine on which the target was run; the user can specify a different machine model for the analysis.

SpeedShop calculates ideal CPU time by using the TDT models. Potential floating-point interlocks are taken into account inside the same basic blocks, but ignored across basic block boundaries. Memory latency time (cache misses and memory bus contention) is ignored. The computed ideal time is therefore always less than the real time that any run would take. See Table 7, page 53, for a comparison of running a `pcsamp` experiment, which generates estimated actual CPU time, and running an `ideal` experiment.

Note that the execution time of an instrumented program is three to six times longer than an uninstrumented one. This timing change may alter the behavior of a program that deals with a graphical user interface (GUI), or depends on events such as `SIGALRM` that are based on an external clock. Also, during analysis, the instrumented executable might appear to be CPU-bound, whereas the original executable was I/O-bound.

Basic block counts are translated to ideal CPU time displayed at the function, source line and machine line levels.

4.4.3 Inclusive Basic Block Counting

The basic block counting explained in the previous section allows you to measure ideal time spent in each procedure, but doesn't propagate the time up to the caller of that procedure. For example, basic block counting may tell you that procedure `sin(x)` took the most time, but significant performance

improvement can only be obtained by optimizing the callers of `sin(x)`. Inclusive basic block counting solves this problem.

Inclusive basic block counting calculates cycles just like regular basic block counting, and then propagates it proportionately to all its callers. The cycles of procedures obtained using regular basic block counting (called exclusive cycles), are divided up among its callers in proportion to the number of times they called this procedure. For example, if `sin(x)` takes 1000 cycles, and its callers, procedures `foo()` and `bar()`, call `sin(x)` 25 and 75 times respectively, 250 cycles are attributed to `foo()` and 750 to `bar()`. By propagating cycles this way, `__start()` ends up with all the cycles counted in the program. For example), the assumption can be very misleading. If `foo()` calls `matmult()` 99 times for 2X2 matrices, while `bar()` calls it once for 100X100 matrices, the inclusive time report will attribute 99% of `matmult()`'s time to `foo()`, but actually almost all the time derives from `bar()`'s one call.

To generate a report that shows inclusive time, specify the `-gprof` flag to `prof`.

4.4.4 Using pcsamp and ideal Together

The ideal experiment can be used together with the `pcsamp` experiment to compare actual and ideal times spent in the CPU. A major discrepancy between `pcsamp` CPU time and `ideal` CPU time indicates one or more of the following situations:

- cache misses and floating point interlocks in a single process application
- secondary cache invalidations in an application with multiple processes that is run on a multiprocessor

A comparison between basic block counts (`ideal` experiment) and PC profile counts (`pcsamp` experiment) is shown in Table 7.

Table 7. Basic Block Counts and PC Profile Counts Compared

Basic Block Counts	PC Profile Counts
Used to compute ideal CPU time	Used to estimate actual CPU time
Data collection by instrumenting	Data collection done with the kernel

Basic Block Counts	PC Profile Counts
Slows program down by factor of three	Has minimal impact on program speed
Generates an exact count	Generates statistical counts

4.5 Hardware Counter Experiments

The experiments described in this section are available for systems that have hardware counters (R10000 class machines). Hardware counters allow you to count various types of events, such as cache misses and counts of issued and graduated instructions.

A hardware counter works as follows: for each event the appropriate hardware counter is incremented on each processor clock cycle when events occur for which there are hardware counters. For example, when a floating point instruction is graduated in a cycle, the graduated floating-point instruction counter is incremented by 1.

4.5.1 Two Tools for Hardware Counter Experiments

There are two tools that allow you to access hardware counter data:

- `perfex` is command-line interface that provides program-level event information. For more information on `perfex`, and on hardware counters, see the `perfex` or `r10k_counters` reference pages.
- SpeedShop allows you to perform the hardware counter experiments described in the next section (Section 4.5.2).

4.5.2 SpeedShop Hardware Counter Experiments

In SpeedShop hardware counter experiments, overflows of a particular hardware counter are recorded. Each hardware counter is configured to count from zero to a number designated as the overflow value. When the counter reaches the overflow value, the system resets it to zero and increments the number of overflows at the present program instruction address. Each experiment provides two possible overflow values; the values are prime numbers, so any profiles that seem the same for both overflow values should be statistically valid.

The hardware counter experiments show where the overflows are being triggered in the program, at the function, source-line, and individual instruction level. When you run `prof` on the data collected during the experiment, the overflow counts are multiplied by the overflow value to compute the total number of events. These numbers are statistical. The generated reports show exclusive hardware counts, that is, information about where the program counter was, not the callstack to get there.

Hardware counter overflow profiling experiments should incur a slowdown of execution of the program of no more than 5%. Count data is kept as 32-bit integers only.

The available hardware experiments are described in the following sections.

4.5.3 [f]gi_hwc

The [f]gi_hwc experiment counts overflows of the graduated instruction counter. The graduated instruction counter is incremented by the number of instructions that were graduated on the previous cycle. The experiment uses statistical PC sampling based on overflows of the counter at an overflow interval of 32771. If the optional `f` prefix is used, the overflow interval is 6553.

4.5.4 [f]cy_hwc

The [f]cy_hwc experiment counts overflows of the cycle counter. The cycle counter is incremented on each clock cycle. The experiment uses statistical PC sampling based on overflows of the counter, at an overflow interval of 16411. If the optional `f` prefix is used, the overflow interval is 3779.

4.5.5 [f]ic_hwc

The [f]ic_hwc experiment counts overflows of the primary instruction-cache miss counter. The primary instruction-cache miss counter is incremented one cycle after an instruction fetch request is entered into the miss handling Table. The experiment uses statistical PC sampling based on the overflow of the counter at an overflow interval of 2053. If the optional `f` prefix is used, the overflow interval is 419.

4.5.6 [f]isc_hwc

The [f]isc_hwc experiment counts overflows of the secondary instruction-cache miss counter. The secondary instruction-cache miss counter is

incremented after the last 16-byte block of a 64-byte primary instruction cache line is written into the instruction cache. The experiment uses statistical PC sampling based on the overflow of the counter at an overflow interval of 131. If the optional `f` prefix is used, the overflow interval is 29.

4.5.7 [f]dc_hwc

The [f]dc_hwc experiment counts overflows of the primary data-cache miss counter. The primary data-cache miss counter is incremented on the cycle after a primary cache data refill is begun. The experiment uses statistical PC sampling based on the overflow of the counter at an overflow interval of 2053. If the optional `f` prefix is used, the overflow interval is 419.

4.5.8 [f]dsc_hwc

The [f]dsc_hwc experiment counts overflows of the secondary data-cache miss counter. The secondary data-cache miss counter is incremented on the cycle after the second 16-byte block of a primary data cache line is written into the data cache. The experiment uses statistical PC sampling, based on the overflow of the counter at an overflow interval of 131. If the optional `f` prefix is used, the overflow interval is 29.

4.5.9 [f]tlb_hwc

The [f]tlb_hwc experiment counts overflows of the TLB (translation lookaside buffer) counter. The TLB counter is incremented on the cycle after the TLB miss handler is invoked. The experiment uses statistical PC sampling based on the overflow of the counter at an overflow interval of 257. If the optional `f` prefix is used, the overflow interval is 53.

4.5.10 [f]gfp_hwc

The [f]gfp_hwc experiment counts overflows of the graduated floating-point instruction counter. The graduated floating-point instruction counter is incremented by the number of floating point instructions which graduated on the previous cycle. The experiment uses statistical PC sampling based on overflows of the counter, at an overflow interval of 32771. If the optional `f` prefix is used, the overflow interval is 6553.

4.5.11 prof_hwc

The `prof_hwc` experiment allows you to set a hardware counter to use in the experiment, and to set a counter overflow interval using the following environment variables:

`_SPEEDSHOP_HWC_COUNTER_NUMBER`

The value of this variable may be any number between 0 and 31. Hardware counters are described in the *MIPS R10000 Microprocessor User's Manual*, Chapter 14, and on the `r10k_counters` reference page. The hardware counter numbers are provided in Table 8, page 57.

`_SPEEDSHOP_HWC_COUNTER_OVERFLOW`

The value of this variable may be any number greater than 0. Some numbers may produce data that is not statistically random, but rather reflects a correlation between the overflow interval and a cyclic behavior in the application. You may want to do two or more runs with different overflow values.

The default counter is the primary instruction-cache miss counter; the default overflow interval is 2053.

The experiment uses statistical PC sampling based on the overflow of the specified counter, at the specified interval. Note that these environment variables cannot be used for other hardware counter experiments. They are examined only when the `prof_hwc` experiment is specified.

4.5.11.1 Hardware Counter Numbers

The possible numeric values for the `_SPEEDSHOP_HWC_COUNTER_NUMBER` variable are shown in Table 8, page 57.

Table 8. Hardware Counter Numbers

0	Cycles
1	Issued instructions
2	Issued loads
3	Issued stores

-
- 4 Issued store conditionals
 - 5 Failed store conditionals
 - 6 Decoded branches
 - 7 Quadwords written back from secondary cache
 - 8 Correctable secondary cache data array ECC errors
 - 9 Primary instruction-cache misses
 - 10 Secondary instruction-cache misses
 - 11 Instruction misprediction from secondary cache way prediction table
 - 12 External interventions
 - 13 External invalidations
 - 14 Virtual coherency conditions (or functional unit completions, depending on hardware version)
 - 15 Graduated instructions
 - 16 Cycles
 - 17 Graduated instructions
 - 18 Graduated loads
 - 19 Graduated stores
 - 20 Graduated store conditionals
 - 21 Graduated floating point instructions
 - 22 Quadwords written back from primary data cache
 - 23 TLB misses
 - 24 Mispredicted branches
 - 25 Primary data-cache misses
 - 26 Secondary data-cache misses
 - 27 Data misprediction from secondary cache way prediction table
 - 28 External intervention hits in secondary cache
 - 29 External invalidation hits in secondary cache

-
-
- 30 Store/prefetch exclusive to clean block in secondary cache
 - 31 Store/prefetch exclusive to shared block in secondary cache
-

4.6 fpe Trace

A floating point exception trace collects each floating point exception with the exception type and the callstack at the time of the exception. Floating-point exception tracing experiments should incur a slowdown of execution of the program of no more than 15%. These measurements are exact, not statistical.

`prof` generates a report that shows inclusive and exclusive floating-point exception counts.

Collecting Data on Machine Resource Usage [5]

This chapter describes how to collect machine resource usage data using the SpeedShop `ssusage` command. Finding out the machine resources that your program uses can help you identify performance bottlenecks and determine which performance experiments you need to run. You can use Table 4, page 10, to identify which experiments to run, based on the results of running `ssusage` on your program.

5.1 `ssusage` Syntax

```
ssusage prog_name [prog_args]
```

prog_name Name of the executable for which you want to collect machine resource usage data.

prog_args Arguments to your executable, if any.

5.2 `ssusage` Results

`ssusage` prints output to *stderr*. For example, the command

```
ssusage generic
```

provides output similar to the following:

```
...
22.03 real, 18.18 user, 0.21 sys, 7 majf, 120 minf, 0 sw, 241 rb, 0
wb, 135 vcx, 648 icx
```

The last two lines of the output is the machine resource usage information that `ssusage` provides. Each field in the report is described below.

<code>real</code>	Elapsed time during the command, in seconds.
<code>user</code>	User CPU time in seconds.
<code>sys</code>	System CPU time in seconds.
<code>majf</code>	Major page faults that cause physical I/O.
<code>minf</code>	Minor page faults that require mapping only.

sw	Process swaps.
rb/wb	Physical blocks read/written. Note that these are attributed to the process that first requests a block, but do not necessarily directly correlate with the process' own I/O operations.
vcx	Voluntary context switches, that is, those caused by the process' own actions.
icx	Involuntary context switches, that is, those caused by the scheduler.

If the program terminates abnormally, a message is printed before the usage line.

Setting Up and Running Experiments: ssrun [6]

This chapter provides information on how to set up and run performance analysis experiments using the `ssrun` command. It consists of the following sections:

- Building Your Executable, see Section 6.1, page 63.
- Setting Up Output Directories and Files, see Section 6.2, page 65.
- Using Runtime Environment Variables, see Section 6.3, page 66.
- Running Experiments, see Section 6.4, page 70.
- Running Experiments on MPI Programs, see Section 6.5, page 74.
- Running Experiments on Programs Using Pthreads, see Section 6.6, page 75.
- Using Calipers, see Section 6.7, page 75.
- Effects of `ssrun`, see Section 6.8, page 78.

6.1 Building Your Executable

The `ssrun` command is designed to be used with normally built executables and default environment settings. However, there are some cases where you need to change the way you build your executable or set certain environment variables.

This section explains when to change the way you build your executable program. For information on setting environment variables, see Section 6.3, page 66.

- If you have used the `ssrt_caliper_point()` function provided in the SpeedShop libraries, you have to explicitly link in the SpeedShop libraries `libss.so` and `libssrt.so`. For more information on setting caliper points, see Section 6.7, page 75.
- If you are planning to build your executable using the `-32` option to the `cc` command, and you want to run the `usertime` experiment, you must add `-lexc` to the link line. For more information on `cc -32`, see the `cc` reference page.

- If you have built a stripped executable, you need to rebuild a non-stripped version to use with SpeedShop. For example, if you are using `ld` to link your C program, do not use the `-s` option because this strips debugging information from the program object and makes the program unusable for performance analysis.
- If you have used compiler optimization level 3, and you are performing experiments that report function-level information, the procedure inlining the optimization performs can result in extremely misleading profiles since the time spent in the inlined procedure will show up in the profile as time spent in the procedure into which it was inlined. It's generally better to use compiler optimization level 2 or less when gathering an execution profile.

6.1.1 Special Information for MP Fortran Programs

If you are compiling MP Fortran programs, you may encounter anomalies in the displayed data:

- For all FORTRAN MP compilations, parallel loops within the program are represented as subroutines with names relating to the source routine in which they are embedded. The naming conventions for these subroutines are different for 32-bit and 64-bit compilations.

For example, in the `linpack` example program, most of the time is spent in the routine `DAXPY`, which can be parallelized.

- In an n32 or 64-bit MP version, the routine has the name “DAXPY,” but most of that work is done in the MP routine named “DAXPY.PREGION1.”
- In a 32-bit version, the DAXPY routine is named “daxpy_,” and the MP routine “_daxpy_519_aaab_.”
- If you perform an `ideal` experiment, the source annotations for 32-bit and 64-bit compilations with the `-g` option differ and are not correct in most cases.
 - In 64-bit source annotations, the exclusive time is correctly shown for each line, but the inclusive time for the first line of the loop (do statement) includes the time spent in the loop body. This same time appears on the lines comprising the loop's body, in effect representing a double-counting.
 - In 32-bit source annotations, the exclusive time is incorrectly shown for the line comprising the loop's body. The line-level data for the loop-body

routine (“`_daxpy_519_aaab_`”) doesn’t refer to proper lines. If the program was compiled with the `-mp_keep` flag, the line-level data should refer to the temporary files that are saved from the compilation, but the temporary files do not contain that information, so no source or disassembly data can be shown. The disassembly data for the main routine does not show the times for the loop-body.

- If the 32-bit program was compiled without the `-mp_keep` flag, the line-level data for the loop-body routine is incorrect. Most lines refer to line 0 of the file, and the rest to other lines at seemingly random places in the file. Consequently, spurious annotations will appear on these other lines. Disassembly correctly shows the instructions and their data, but the line numbers are wrong. This reflects essentially the same double-counting problem as seen in 64-bit compilations, but the extra counts go to other places in the file, rather than to the first line of the loop.

6.2 Setting Up Output Directories and Files

When you run an experiment, performance data files are written to the current working directory by default. They are named using the following convention:

prog_name.exp_name.exp_type.id

The experiment ID, *id*, consists of one or two letters (designating the process type) and the process ID number. See Table 9 for letter codes and descriptions.

Table 9. Letter Codes in Experiment ID Numbers

Letter code	Description
m	Master process created by <code>ssrun</code>
p	Process created by a call to <code>sproc()</code>
f	Process created by a call to <code>fork()</code>
s	Process created by a call to <code>system()</code>
e	Process created by a call to <code>exec()</code>
fe	Process created by a call to <code>fork()</code> and <code>exec()</code>

The following are examples of data file names:

```
statalanal.ideal.m10966  
engines.pcsamp.m14493
```

In a single-process application, `ssrun` generates a single performance data file. In a multi-process application, `ssrun` generates a performance data file for each process.

You can change the default filename or directory for performance data files using environment variables. See `_SPEEDSHOP_OUTPUT_DIRECTORY` and `_SPEEDSHOP_OUTPUT_FILENAME` in Table 10 for more information.

6.3 Using Runtime Environment Variables

This section provides information about available environment variables, grouped by functionality:

- Section 6.3.1, page 66
- Section 6.3.2, page 68
- Section 6.3.3, page 68

6.3.1 User Environment Variables

A number of environment variables are normally used to control the operation of SpeedShop. Table 10 lists these variables.

Table 10. General Environment Variables

Variable	Description
<code>_SPEEDSHOP_VERBOSE</code>	Causes a log of each program's operation to be written to <code>stderr</code> . If this variable is set to an empty string, only major events are logged; if it is set to a non-empty string, more detailed events are logged.
<code>_SPEEDSHOP_SILENT</code>	Suppresses all SpeedShop output, other than fatal error messages.

Variable	Description
<code>_SPEEDSHOP_CALIPER_POINT_SIG <i>sig_num</i></code>	If both <code>_SPEEDSHOP_VERBOSE</code> and <code>_SPEEDSHOP_SILENT</code> are set, <code>_SPEEDSHOP_VERBOSE</code> is ignored. Causes the specified signal number to be used for recording a caliper-point in the experiment.
<code>_SPEEDSHOP_REUSE_FILE_DESCRIPTOR</code>	Opens and closes the file descriptors for the output files every time performance data is to be written.
<code>_SPEEDSHOP_HWC_COUNTER_NUMBER</code>	Specifies the counter to be used for <code>prof_hwc</code> experiments. Counters are numbered between 0 and 31, and are described in the <i>MIPS R10000 Microprocessor's User's Manual</i> , Chapter 14. Counter 0 counters are numbered 0-15, and counter 1 counters are numbered 16-31.
<code>_SPEEDSHOP_HWC_COUNTER_OVERFLOW</code>	Specifies the overflow value for the counter to be used in <code>prof_hwc</code> experiments. The value chosen may be any number greater than 0. Some choices may produce data that is not statistically random, but reflects a correlation between the overflow interval and a cyclic behavior in the application. Users may want to do two or more runs with different overflow values.
<code>_SPEEDSHOP_OUTPUT_NOCOMPRESS</code>	Disables the compression of performance data.
<code>_SPEEDSHOP_OUTPUT_DIRECTORY</code>	Causes the output data files to be placed in the specified directory, rather than the current working directory.
<code>_SPEEDSHOP_OUTPUT_FILENAME</code>	Causes the output file to be saved under the specified name. If <code>_SPEEDSHOP_OUTPUT_DIRECTORY</code> is also specified, it is prepended to the filename you specify.

6.3.2 Process Tracking Environment Variables

A number of environment variables may be used for controlling the treatment of processes spawned from the original target. Table 11, page 68, lists these variables.

Table 11. Process Tracking Environment Variables

Variable	Description
<code>_SPEEDSHOP_TRACE_FORK</code> [<i>True</i> <i>False</i>]	If True, specifies that processes spawned by calls to <code>fork()</code> will be monitored if they don't call <code>exec()</code> . If they do call <code>exec()</code> , and <code>_SPEEDSHOP_TRACE_FORK_TO_EXEC</code> is not set to True, the data covering the time between the <code>fork()</code> and <code>exec()</code> will be discarded. It is true by default. Note: In the current release, data are recorded independent of whether the process calls <code>exec()</code> or not.
<code>_SPEEDSHOP_TRACE_FORK_TO_EXEC</code> [<i>True</i> <i>False</i>]	If True, specifies that a process spawned by calls to <code>fork()</code> will be monitored even if they also call <code>exec()</code> . It is False by default.
<code>_SPEEDSHOP_TRACE_EXEC</code> [<i>True</i> <i>False</i>]	If True, specifies that a process spawned by calls to any of the various flavors of <code>exec()</code> will be monitored. It is true by default.
<code>_SPEEDSHOP_TRACE_SPROC</code> [<i>True</i> <i>False</i>]	If True, specifies that a process spawned by calls to <code>sproc()</code> will be monitored. It is True by default.
<code>_SPEEDSHOP_TRACE_SYSTEM</code> [<i>True</i> <i>False</i>]	If True, specifies that <code>system()</code> calls will be monitored. It is False by default.

6.3.3 Expert-Mode Environment Variables

A number of variables may be used for debugging and finer control of the operation of SpeedShop. Table 12, page 69, lists these variables.

Table 12. Expert-Mode Environment Variables

Variable	Description
<code>_SPEEDSHOP_SAMPLING_MODE</code>	For PC-sampling and hardware-counter profiling. If set to 1, generates data for the base executable only. If not set, or set to a value different from 1, data is generated for the executable and all DSOs it uses.
<code>_SPEEDSHOP_INIT_DEFERRED_SIG <i>sig_num</i></code>	If specified, initialization of the experiment is not performed when the target process starts, but will be delayed until the specified signal is sent to the process. A handler for the given signal is installed when the process starts. It is the user's responsibility to ensure that it is not overridden by the target code.
<code>_SPEEDSHOP_SHUTDOWN_SIG <i>sig_num</i></code>	If specified, termination of the experiment will not be performed when the target process exits, but rather will happen when the specified signal is sent to the process. A handler for the given signal will be installed when the process starts, and it is the user's responsibility to ensure that it is not overridden by the target code.
<code>_SPEEDSHOP_EXPERIMENT_TYPE</code>	Passes the name of the experiment to the runtime. It is normally set by <code>ssrun</code> , but may be overwritten.
<code>_SPEEDSHOP_MARCHING_ORDERS</code>	Passes the marching orders of the experiment to the runtime. It is normally set by <code>ssrun</code> from the experiment type, but may be overwritten.
<code>_SPEEDSHOP_SBRK_BUFFER_LENGTH</code>	Defines the maximum size of the internal malloc arena used. This arena is completely separate from the user's arena, and has a default size of 0x100000.

Variable	Description
<code>_SPEEDSHOP_FILE_BUFFER_LENGTH</code>	Defines the size of the buffer used for writing the experiment files. The default length is 8 KB. The buffer is used only for writing small records to the file; large records are written directly to avoid the buffering overhead.
<code>_SPEEDSHOP_DEBUG_NO_SIG_TRAPS</code>	Disables the normal setting of signal handlers for all fatal and exit signals.
<code>_SPEEDSHOP_DEBUG_NO_STACK_UNWIND</code>	Suppresses the stack unwind as done in usertime experiments, and as is done at caliper-samples for all experiments. The option is used as a workaround for various unwind bugs in <code>libexc</code> .

6.4 Running Experiments

This section describes how to use `ssrun` to perform experiments. For information on using `pixie` directly, see Chapter 8, page 109.

6.4.1 `ssrun` Syntax

```
ssrun flags -exp_type prog_name prog_args
```

<i>flags</i>	Zero or more of the flags described in Table 13, page 71, that control the data collection and the treatment of descendent processes or programs, and how the data is to be externalized.
<i>-exp_type</i>	The experiment type. Experiments are described in detail in Chapter 4, page 49.
<i>prog_name</i>	The name of the program on which you want to run an experiment.
<i>args</i>	Arguments to your program, if any.

`ssrun` generates a performance data file that is named as described in Section 6.2, page 65.

Table 13. Flags for `ssrun`

Name	Result
<code>-hang</code>	Specifies that the process should be left waiting just before executing its first instruction. This allows you to attach the process to a debugger.
<code>-mo <i>marching_orders</i></code>	Allows you to specify marching orders. If this option is used, the environment variable <code>_SSRUNTIME_MARCHING_ORDERS</code> is not examined.
<code>-name <i>target_name</i></code>	Specifies that the target should be run with <code>argv[0]</code> set to <i>target_name</i> .
<code>-purify</code>	Can be used only when the Purify® product is installed. Specifies that <code>purify</code> should be run on the target, and then runs the resulting “purified” executable. Note that <code>-purify</code> and SpeedShop performance experiments cannot be combined.
<code>-v</code>	Prints a log of the operation of <code>ssrun</code> to <code>stderr</code> . The same behavior occurs if the environment variable <code>_SPEEDSHOP_VERBOSE</code> is set to an empty string.
<code>-V</code>	Prints a detailed log of the operation of <code>ssrun</code> to <code>stderr</code> . The same behavior occurs if the environment variable <code>_SPEEDSHOP_VERBOSE</code> is set to a non-zero-length string. This option can be used to see how to set the various environment variables, and how to invoke instrumentation when necessary.

6.4.2 `ssrun` Examples

This section provides examples of using `ssrun` with options and experiment types. For additional examples, see Chapter 2, page 13, or Chapter 3, page 33.

6.4.2.1 Example Using the `pcsampx` Experiment

The `pcsampx` experiment collects data to estimate the actual CPU time for each source code line, machine instruction, and function in your program. The optional `x` suffix causes a 32-bit bin size to be used, allowing a larger number of counts to be recorded. For a more detailed description of the `pcsamp` experiment, see the Chapter 4, page 49.

This example performs a `pcsampx` experiment on the `generic` executable:

```
ssrun -pcsampx generic
```

To see the performance data that has been generated, run `prof` on the performance data file, `generic.pcsampx.16064`:

prof generic.pcsampx.m16064

The report is printed to *stdout*. (This layout of this report has been altered slightly to accommodate presentation needs.) For more information on `prof` and the reports generated by `prof`, see Chapter 7, page 81.

 SpeedShop profile listing generated Mon Feb 2 15:08:14 1998

```
prof generic.pcsampx.m12185
    generic (n32): Target program
        pcsampx: Experiment name
    pc,4,10000,0:cu: Marching orders
    R4400 / R4000: CPU / FPU
        1: Number of CPUs
        175: Clock frequency (MHz.)
```

Experiment notes--

```
    From file generic.pcsampx.m12185:
    Caliper point 0 at target begin, PID 12185
        /home/hickory28/dep/speedshop/progs.etc/linpack.demos/c/generic
    Caliper point 1 at exit(0)
```

 Summary of statistical PC sampling data (pcsampx)--

```
    2729: Total samples
    27.290: Accumulated time (secs.)
    10.0: Time per sample (msecs.)
    4: Sample bin width (bytes)
```

 Function list, in descending order by time

[index]	secs	%	cum.%	samples	function (dso: file, line)
[1]	25.470	93.3%	93.3%	2547	anneal (generic: generic.c,
1573)					
[2]	1.100	4.0%	97.4%	110	slaveusertime (dlslave.so: dlslave.c, 22)
[3]	0.310	1.1%	98.5%	31	__read (libc.so.1: read.s, 20)
[4]	0.240	0.9%	99.4%	24	cvtttrap (generic: generic.c, 317)
[5]	0.150	0.5%	99.9%	15	__xstat (libc.so.1: xstat.s,
12)					
[6]	0.010	0.0%	100.0%	1	__write (libc.so.1: write.s, 20)
[7]	0.010	0.0%	100.0%	1	__morecore (libc.so.1: malloc.c, 632)
	27.290	100.0%	100.0%	2729	TOTAL

6.4.2.2 Example Using the `-v` Option

To get information about how a SpeedShop experiment is set up and performed, you can supply the `-v` option to `ssrun`.

This example performs a `pcsampx` experiment on the `generic` executable:

```
ssrun -v -pcsampx generic
```

The `ssrun` command writes the following output to `stderr`. It displays information as the command line is parsed and shows the environment variables that `ssrun` sets.

```
fraser 75% ssrun -v -pcsampx generic
ssrun: setenv _SPEEDSHOP_MARCHING_ORDERS pc,4,10000,0:cu
ssrun: setenv _SPEEDSHOP_EXPERIMENT_TYPE pcsampx
ssrun: setenv _SPEEDSHOP_TARGET_FILE generic
ssrun: setenv _RLD_LIST libss.so:libssrt.so:DEFAULT
...
```

6.4.3 Using `ssrun` With a Debugger

To use the `ssrun` command in conjunction with a debugger such as `dbx` or the ProDev™ WorkShop debugger, you need to call `ssrun` with the `-hang` option and the name of your program.

Follow these steps to run the FPE trace experiment on `generic`, and then run `generic` in a debugger.

1. Call `ssrun` as follows:

```
ssrun -hang -fpe generic
```

`ssrun` parses the command line, sets up the environment for the experiment, calls the target process using `exec`, and hangs the target process on exiting from the call to `exec`.

2. Get the process ID of the call to `ssrun` using a command such as `ps`.
3. Start your debugging session.
4. Attach the process to the debugger.
5. Run the process from the debugger.

You can also invoke `ssrun` from within a debugger. In this case, `ssrun` leaves the target hung on exiting the call to `exec`, and informs the debugger of that fact.

You can also use either `dbx` or the WorkShop debugger to set calipers to record performance data for a part of your program. See Section 6.7, page 75, for more information on setting calipers.

6.5 Running Experiments on MPI Programs

The Message Passing Interface (MPI) is a library specification for message-passing, proposed as a standard by a committee of vendors, implementors, and users. It allows processes to communicate by “mailing” data “messages” to other processes, even those running on distant computers.

If your program uses the MPI, you need to set up SpeedShop experiments a little differently. There are two ways to accomplish this. The first method takes two steps:

1. Set up a shell script that contains the call to `ssrun` and the experiment you want to run.

For example, if you have a program called `testit`, and you want to run the `pcsampx` experiment, a script, named `exp_script`, might look like the following:

```
#!/bin/sh
ssrun -pcsampx testit
```

2. Call `mpirun` with the script name using one of the following:

```
mpirun -np 6 exp_script
mpirun host1 2, host2 2 exp_script
```

The second method is to use one of the following:

```
mpirun -np 6 ssrun -pcsampx testit
mpirun host1 2, host2 2 ssrun -pcsampx testit
```

The master experiment file created on each MPI host might not contain performance data from the application (depending on the MPI version), but rather from a master program that spawns the actual MPI application slaves. You can choose to exclude that file from performance analysis.

When using `ssrun -ideal`, or `ssrun -purify`, you should take care that the code for each separate host executes out of a different physical directory, not out

of the same NFS mounted directory. During process creation, instrumentation is performed, and since different hosts may have different versions of the same named library (`libc.so.1`, for example), conflicts may occur. You may also need to use the `-d` option with `mpirun` to specify the directory on each host.

6.6 Running Experiments on Programs Using Pthreads

Pthreads are the threads defined by the POSIX[®] operating system standard (IEEE1003.1c-1995). This standard contains a set of interfaces and semantics for creating and managing threads within the POSIX operating system definition. The basic Silicon Graphics pthreads implementation consists of a library (one for each o32, n32 and n64 ABI) and a header file.

Applications using pthreads are specifically identified by SpeedShop. Performance data collection is done on a per-program basis, rather than on a per-pthread basis. Under IRIX[™]6.2, 6.3 and 6.4, SpeedShop creates as many experiment files as the number of sprocs used by the pthreads library to create and manage the pthreads. In addition, `cm_usage` data is not supported, and `SIGTERM` is reserved to be used to terminate the application normally. You should analyze all the experiment files together via `prof` to get a valid profile for the code. Under IRIX 6.5, SpeedShop creates only one experiment file. For `usertime` and `fpe` experiments, however, you can specify the `-pthreads` option with `prof` to get per-pthread performance reports.

6.7 Using Calipers

In some cases, you may want to generate performance data reports for only a part of your program. You can do this by setting caliper points to identify the area or areas for which you want to see performance data. When you run `prof`, you can specify a region for which to generate a report by supplying the `-calipers` option and the appropriate caliper numbers. For more information on `prof -calipers`, see Chapter 7, page 81.

Table 14, page 76, shows how you can set caliper points in three different ways.

Table 14. Setting Caliper Points

Use This Approach...	For These Benefits...
Explicitly link with the SpeedShop runtime and call <code>ssrt_caliper_point</code> to record a caliper sample.	Allows you to set a caliper point at a specific location in a file.
Define a signal to be used to record a caliper sample by specifying a signal as a value to the environment variable <code>_SPEEDSHOP_CALIPER_POINT_SIG</code> and then sending the target the given signal.	Useful if you want to be able to set a caliper point as your program is running.
Set a caliper sample trap in <code>dbx</code> or the WorkShop debugger. Setting a trap involves setting a breakpoint and evaluating the expression <code>libss_caliper_point(1)</code> when the process stops.	Useful if you're working with a debugger in conjunction with SpeedShop.

An implicit caliper point is always present at the start of execution of the process. A final caliper-point is recorded when the process calls `_exit`. The implicit caliper point at the beginning of the program is numbered 0, the first caliper point recorded is numbered 1, and any additional caliper points are numbered sequentially.

In addition, caliper points are automatically recorded under the following circumstances to ensure that at least one valid set of data is recorded.

- When a fatal signal is received, such as `SIGQUIT`, `SIGILL`, `SIGTRAP`, `SIGABRT`, `SIGEMT`, `SIGFPE`, `SIGBUS`, `SIGSEGV`, `SIGSYS`, `SIGXCPU` or `SIGXFSZ`. Note that this list does not and cannot include `SIGKILL`.
- When the program calls an `exec` function such as `execve()` or `execvp()`.
- When a program closes a DSO by calling `dlclose()`.
- When an exit signal is received, such as `SIGHUP`, `SIGINT`, `SIGPIPE`, `SIGALRM`, `SIGTERM`, `SIGUSR1`, `SIGUSR2`, `SIGPOLL`, `SIGIO`, `SIGRTMIN` or `SIGRTMAX`.

6.7.1 Setting Calipers With `ssrt_caliper_point`

To set calipers with `ssrt_caliper_point`, follow these steps:

1. Insert calls to `ssrt_caliper_point()` in your source code. Call the function with the argument 1 (True).


```
...
ssrt_caliper_point(1);
...
```

You can insert one or more calls at any point in your code.

2. Link the SpeedShop library `libss.so` into your application.

The library should be placed last on the link line.

3. Run your program with `ssrun` and the desired experiment type.

For example, if you want to run the `ideal` experiment on `generic`:

```
ssrun -ideal generic
```

The caliper points you have set in the source file are recorded in the performance data file that is generated by `ssrun`.

6.7.2 Setting Calipers With Signals

To set calipers with signals, follow these steps:

1. Set the `_SPEEDSHOP_CALIPER_POINT_SIG` variable to the signal number you want to use.

Choose a signal that doesn't terminate the program. The signal should also not be caught by the target program, because this would interfere with its use for triggering a caliper point.

The following signals are good choices because they don't have any semantics already associated with them:

```
SIGUSR1 16      /* user defined signal 1 */
SIGUSR2 17      /* user defined signal 2 */
```

2. Run `ssrun` with your program.
3. Enter a command such as `ps` or `top` to determine the process ID of `ssrun`. This is also the process ID of the program you are working on.
4. Send the signal you used in step 1 to the process using the `kill` command:

```
kill -sig_num pid
```

A caliper point is set at the point in the program where the signal was received by the SpeedShop runtime.

6.7.3 Setting Calipers With a Debugger

From either `dbx` or the `WorkShop` debugger, you can set a caliper point anywhere it is possible to set a breakpoint: function entry or exit, line numbers, execution addresses, watchpoints, pollpoints (timer-based). You can also attach conditions and/or cycle counts.

1. Set a breakpoint in your program at the point at which you want to set a caliper point.
2. When the process stops, evaluate the expression `libss_caliper_point(1)`.

The evaluation of the expression always returns zero, but a side effect of the evaluation is the recording of the appropriate data.

3. Resume execution of the process.

6.8 Effects of `ssrun`

When you call `ssrun`, the system performs the following operations for all experiments:

- Sets various environment variables like `_SPEEDSHOP_MARCHING_ORDERS` and `_SPEEDSHOP_EXPERIMENT_TYPE`.

For more information on these environment variables, see Section 6.3, page 66.

- Inserts the SpeedShop libraries `libss.so` and `libssrt.so` as part of your executable using the environment variable `_RLD_LIST`.
- Invokes the target process by calling `exec()`.
- The SpeedShop runtime library writes the appropriate experiment data to the output file.

6.8.1 Effects of `ssrun -ideal`

When you run an `ideal` experiment, the following additional operations occur:

- `libpixrt.so` is inserted first in the executable's library list.
- `libssrt.so` and `libss.so` are inserted in the executable's library list.

- `ssrun` generates pixified versions of all the libraries that the program uses, as well as the executable.

The generated pixified versions have an extension that depends on the ABI:

- `.pixie` for the executable
- `.pix32` for all 32 libraries
- `.pixn32` for all n32 libraries
- `.pix64` for all 64 libraries

The generated files are written to the current working directory, and include code that allows performance data to be collected for each function and basic block.

For more information on the `ideal` experiment, see Chapter 4, page 49.

Analyzing Experiment Results: prof [7]

This chapter provides information on how to view and analyze experiment results. It consists of the following sections:

- Using `prof` to Generate Performance Reports, see Section 7.1, page 81.
- Using `prof` With `ssrun`, see Section 7.2, page 85.
- Using `prof` Options, see Section 7.3, page 93.
- Generating Reports for Different Machine Types, see Section 7.4, page 106.
- Generating Reports for Multiprocessed Executables, see Section 7.5, page 106.
- Generating Compiler Feedback Files, see Section 7.6, page 107.
- Interpreting Reports, see Section 7.7, page 107.

7.1 Using prof to Generate Performance Reports

Performance data is examined using `prof`, a text-based report generator that prints to `stdout`.

The `prof` command can be used in two modes:

- To generate a report from performance data gathered during experiments recorded by `ssrun`:

```
prof options perf_data_file perf_data_file ...
```

- To generate a report from data files produced by running a program that has been instrumented by `pixie`:

```
prof executable_name [options] [pixie_counts_file]
```

This chapter focuses on the use of `prof` to generate reports from `ssrun` experiments. For information on `prof` for a `pixie` experiment, see Chapter 8, page 109.

7.1.1 prof Syntax

The syntax for `prof` when using it with data files from `ssrun` is:

```
prof options data_file data_file ...
```

<i>options</i>	Zero or more of the options described in Table 15, page 82.
<i>data_file</i>	One or more names of performance data files generated by <code>ssrun</code> .

7.1.2 prof Options

Table 15, page 82, lists `prof` options. For more information, see the `prof(1)` man page.

Table 15. Options for `prof`

Name	Result
<code>-calipers n1 n2</code>	<p>Restricts analysis to a segment of program execution. This option works only for SpeedShop experiments.</p> <p>Causes <code>prof</code> to compute the data between caliper points <i>n1</i> and <i>n2</i>, rather than for the entire experiment.</p> <p>If <i>n1</i> \geq <i>n2</i>, an error is reported.</p> <p>If <i>n1</i> is negative, it is set to the beginning of the experiment.</p> <p>If <i>n2</i> is greater than the maximum number of caliper points recorded, it is set to the maximum.</p> <p>If n1 is omitted, zero (the beginning of the program) is assumed.</p>
<code>-c[lock] n</code>	<p>Lists the number of seconds spent in each routine, based on the CPU clock frequency <i>n</i>, expressed in megahertz. This option is useful when generating reports for <code>ideal</code> experiments, or for basic block counting data obtained with <code>pixie</code>. The default is to use the clock frequency of the machine where the performance data was collected.</p>
<code>-cycle n</code>	<p>Sets the cycle time to <i>n</i> nanoseconds.</p>
<code>-den[sity]</code>	<p>Prints a list of procedures with non-zero instruction cycles sorted by the instruction density, which is the number of cycles per instruction.</p> <p>This option can be used when generating reports for <code>ideal</code> experiments, or for basic block counting data obtained with <code>pixie</code>.</p>

Name	Result																				
-debug: <i>dbg_flags</i>	<p><i>dbg_flags</i> can be combinations of the following:</p> <table> <tbody> <tr> <td>GPROF_FLAG</td> <td>0x00000001</td> </tr> <tr> <td>COUNTS_FLAG</td> <td>0x00000002</td> </tr> <tr> <td>SAMPLE_FLAG</td> <td>0x00000004</td> </tr> <tr> <td>MISS_FLAG</td> <td>0x00000008</td> </tr> <tr> <td>FEEDBACK_FLAG</td> <td>0x00000010</td> </tr> <tr> <td>CORD_FLAG</td> <td>0x00000020</td> </tr> <tr> <td>USERPC_FLAG</td> <td>0x00000040</td> </tr> <tr> <td>MDEBUG_FLAG</td> <td>0x00000080</td> </tr> <tr> <td>BEAD_FLAG</td> <td>0x00000100</td> </tr> <tr> <td>LIBSSRT_FLAG</td> <td>0x00000200</td> </tr> </tbody> </table>	GPROF_FLAG	0x00000001	COUNTS_FLAG	0x00000002	SAMPLE_FLAG	0x00000004	MISS_FLAG	0x00000008	FEEDBACK_FLAG	0x00000010	CORD_FLAG	0x00000020	USERPC_FLAG	0x00000040	MDEBUG_FLAG	0x00000080	BEAD_FLAG	0x00000100	LIBSSRT_FLAG	0x00000200
GPROF_FLAG	0x00000001																				
COUNTS_FLAG	0x00000002																				
SAMPLE_FLAG	0x00000004																				
MISS_FLAG	0x00000008																				
FEEDBACK_FLAG	0x00000010																				
CORD_FLAG	0x00000020																				
USERPC_FLAG	0x00000040																				
MDEBUG_FLAG	0x00000080																				
BEAD_FLAG	0x00000100																				
LIBSSRT_FLAG	0x00000200																				
-dis[assemble]	Disassembles and annotates the analyzed object code with cycle times if you have run an ideal experiment, collected data using <code>pixie</code> , or have run a <code>pcsamp</code> or <code>prof_hwc</code> experiment.																				
-dso [<i>dso_name</i>]	Generates a report only for the named DSO. If you don't specify <i>dso_name</i> , <code>prof</code> prints a list of applicable DSO names. Only the basename of the DSO needs to be specified.																				
-dsolist	List all the DSOs in the program and their start and end text addresses.																				
-e[xclude] <i>procs</i>	Excludes information on the specified procedures. If you specify uppercase <code>-E</code> , <code>prof</code> also omits the specified procedures from the base upon which it calculates percentages.																				
-feedback	<p>Produces files with information that can be used to (a) arrange procedures in the binary in an optimal ordering using <code>cord</code>, and (b) tell the compiler how to optimize compilation of the program using <code>cc -fb filename.cfb</code>. This option can be used when generating reports for ideal experiments, or for basic block counting data obtained with <code>pixie</code>.</p> <p><code>cord</code> feedback files are named <i>program.fb</i> or <i>libso.fb</i>. Compiler feedback files are named <i>progam.cfb</i> or <i>libso.cfb</i>. These are binary files and may be dumped using the <code>fbdump</code> command.</p> <p>Procedures are normally ordered by their measured invocation counts; if <code>-gprof</code> is also specified, procedures are ordered using call graph counts, rather than invocation counts.</p>																				
-gprof	<p>Calculates cycles and propagates basic block counting to a procedure's callers proportionately.</p> <p>This option can be used when generating reports for ideal experiments, or for basic block counting data obtained with <code>pixie</code>. It can also be used for <code>fpe</code> and <code>usertime</code> experiments.</p>																				

Name	Result
-h[eavy]	Lists the most heavily used lines of source code in descending order of use, sorting lines by their frequency of use. This option can be used when generating reports for <code>ideal</code> , <code>pcsamp</code> , or <code>prof_hwc</code> experiments, or for basic block counting data obtained with <code>pixie</code> .
-l[ines]	Lists the most heavily used lines of source code in descending order of use, but lists lines grouped by procedure, sorted by cycles executed per procedure. This option can be used when generating reports for <code>ideal</code> , <code>pcsamp</code> , or <code>prof_hwc</code> experiments, or for basic block counting data obtained with <code>pixie</code> .
-nocounts	Analyzes an executable or a <code>.o</code> file using the <code>pixie</code> machine model, and assuming each instruction is executed once. This analysis cannot match any possible real run of any executable which contains one or more conditional branch instructions.
-nofilenames	Removes <code>a.out</code> , DSO, and source filenames from the listing; useful for scripted analysis of <code>prof</code> output.
-o[nly] <i>procs</i>	Reports information on only the procedures specified. If you specify uppercase <code>-O</code> , <code>prof</code> uses only the procedures, rather than the entire program, as the base upon which it calculates percentages.
-p[rocedures]	Lists the time spent in each procedure.
-pthreads <i>pthrds</i>	Analyzes data only for the specified pthreads (for <code>usertime</code> and <code>fpe</code> experiments on applications that use pthreads (on Irix 6.5 or later)).
-q[uit] <i>n</i>	Condenses output listings by truncating <code>-p[rocedures]</code> , <code>-h[eavy]</code> , <code>-l[ines]</code> , and <code>-gprof</code> listings. You can specify <i>n</i> in three ways: <i>n</i> , an integer, truncates everything after <i>n</i> lines; <i>n</i> %, an integer followed by a percent sign, does not print any procedure or line with less than <i>n</i> in the % column; <i>ncum</i> %, an integer followed by <code>cum%</code> , does not print any procedure or line with more than <i>n</i> in the <code>cum%</code> column. That is, it truncates the listing after the last procedure or line which brings the cumulative total to <i>n</i> %. If <code>-gprof</code> is also specified, it behaves the same as <code>-q n %</code> . For example, <code>-q 15</code> truncates each part of the report after 15 lines of text. <code>-q 15%</code> truncates each part of the report that represents less than 15% of the whole, and <code>-q 15cum%</code> truncates each part of the report that has a cumulative percentage above 15%.

Name	Result
-r10000 -r8000 -r5000 -r4000 -r3000	Overrides the default processor scheduling model that <code>prof</code> uses to generate a report. If this option is not specified, <code>prof</code> uses the scheduling model for the processor on which the experiment is being run.
-S (-source)	Disassembles and annotates the analyzed object code with cycle times, or PC samples, and source code, if you have run an <code>ideal</code> , <code>pcsamp</code> , or <code>prof_hwc</code> experiment, or collected data using <code>pixie</code> .
-z[ero]	Lists the procedures that are never invoked. Use this option when generating reports for <code>ideal</code> experiments, or for basic block counting data obtained with <code>pixie</code> .

7.1.3 prof Output

`prof` generates a performance report that is printed to *stdout*. Warning and fatal errors are printed to *stderr*.

Note: Fortran alternate entry point times are attributed to the `main` function/subroutine, since there is no general way for `prof` to separate the times for the alternate entries.

7.2 Using prof With ssrun

When you call `prof` with one or more SpeedShop performance data files, it collects the data from all the output files and produces a listing. The `prof` command is able to detect which experiment was run and generate an appropriate report. It provides reports for all experiment types.

In cases where `prof` accepts more than one data file as input, it sums up the results. The multiple input data files must be generated from the same executable, using the same experiment type.

`prof` may report times for procedures named with a prefix of `*DF*`, for example `*DF*_hello.init_2`. `DF` stands for “Dummy Function” and indicates cycles spent in parts of text which are not in any function: `init` and `fini` sections, and `MIPS.stubs` sections, for example.

The types of reports that `prof` generates are described in the following sections:

- `usertime` Experiment Reports, see Section 7.2.1, page 86.
- `pcsamp` Experiment Reports, see Section 7.2.2, page 87.

- Hardware Counter Experiment Reports, see Section 7.2.3, page 88.
- `ideal` Experiment Reports, see Section 7.2.4, page 89.
- `fpe` Trace Reports, see Section 7.2.5, page 92.

7.2.1 `usertime` Experiment Reports

For `usertime` experiments, `prof` generates CPU times for individual routines and how those times compare with the rest of the program, as follows:

- The `index` column provides an index number for reference.
- The `excl.secs` column shows how much time, in seconds, was spent in the function itself (exclusive time). For example, less than one hundredth of a second was spent in `__start()`, but 0.03 of a second was spent in `fread`.
- The `excl.%` shows the percentage of a program's total time that was spent in the function.
- The `cum.%` column shows the percentage of the complete program time that has been spent when the function completes.
- The `incl.secs` column shows how much time, in seconds, was spent in the function and descendents of the function.
- The `incl.%` column shows the cumulative percentage of inclusive time spent in each function and its descendents.
- The `samples` column provides the number of samples of the function and all of its descendents.
- The `procedure (dso:file,line)` columns list the function name, its DSO name, its file name, and its line number

This example is a truncated version of the full report. For a complete report see Section 2.3.1.2, page 17.

```
-----  
SpeedShop profile listing generated Mon Feb  2 11:07:15 1998  
prof generic.usertime.m10981  
  generic (n32): Target program  
    usertime: Experiment name  
      ut:cu: Marching orders  
R4400 / R4000: CPU / FPU  
      1: Number of CPUs
```

```

175: Clock frequency (MHz.)
Experiment notes--
  From file generic.usertime.m10981:
  Caliper point 0 at target begin, PID 10981
    /home/hickory28/dep/speedshop/progs.etc/linpack.demos/c/generic
  Caliper point 1 at exit(0)
-----
Summary of statistical callstack sampling data (usertime)--
  809: Total Samples
    0: Samples with incomplete traceback
  24.270: Accumulated Time (secs.)
  30.0: Sample interval (msecs.)
-----
Function list, in descending order by exclusive time
-----
[index]  excl.secs  excl.%   cum.%   incl.secs  incl.%   samples  procedure
(dso: file, line)

   [4]      22.770  93.8%   93.8%    22.770   93.8%     759   anneal
(generic: generic.c, 1573)

```

7.2.2 pcsamp Experiment Reports

For `[f]pcsamp[x]` experiments, `prof` generates a function list annotated with the number of samples taken for the function, and the estimated time spent in the function.

- The `secs` column shows the amount of CPU time that was spent in the function.
- The `%` column shows the percentage of the total program time that was spent in the function.
- The `cum.%` column shows the percentage of the complete program time that has been spent when the function completes.
- The `samples` column shows how many samples were taken when the process was executing in the function.
- The function (`dso:file, line`) columns list the function, its DSO name, its file name, and its line number.

```
-----
SpeedShop profile listing generated Mon Feb  2 11:01:36 1998
```

```

prof generic.fpcsamp.m11140
    generic (n32): Target program
        fpcsamp: Experiment name
pc,2,1000,0:cu: Marching orders
R4400 / R4000: CPU / FPU
    1: Number of CPUs
    175: Clock frequency (MHz.)
Experiment notes--
    From file generic.fpcsamp.m11140:
    Caliper point 0 at target begin, PID 11140
        /home/hickory28/dep/speedshop/progs.etc/linpack.demos/c/generic
    Caliper point 1 at exit(0)
-----
Summary of statistical PC sampling data (fpcsamp)--
    23828: Total samples
    23.828: Accumulated time (secs.)
    1.0: Time per sample (msecs.)
    2: Sample bin width (bytes)
-----
Function list, in descending order by time
-----
[index]      secs      %      cum.%    samples  function (dso: file, line)

    [1]      22.279  93.5%  93.5%     22279  anneal (generic: generic.c,
1573)

```

7.2.3 Hardware Counter Experiment Reports

For the various `hwc` experiments, `prof` generates a function list annotated with the number of overflows generated by the function.

- The `counts` column shows the number of times a data item was not in secondary cache when needed for a calculation during the execution of the function.
- The `%` column shows the percentage of the program's overflows that occurred in the function.
- The `cum.%` shows the percentage of the program's overflows that occurred by the time the function finished execution.
- The `samples` column shows the number of times the program counter was sampled during execution of the function.

- The function (dso: file, line) columns show the name, the DSO, the file name, and line number of the function.

```
-----
SpeedShop profile listing generated Mon Feb  2 11:11:44 1998
  prof generic.dsc_hwc.m294398
    generic (n32): Target program
      dsc_hwc: Experiment name
    hwc,26,131:cu: Marching orders
  R10000 / R10010: CPU / FPU
    16: Number of CPUs
    195: Clock frequency (MHz.)

Experiment notes--
  From file generic.dsc_hwc.m294398:
  Caliper point 0 at target begin, PID 294398
    /home/hickory28/dep/speedshop/progs.etc/linpack.demos/c/generic
  Caliper point 1 at exit(0)
-----
```

```
Summary of R10K perf. counter overflow PC sampling data (dsc_hwc)--
    6: Total samples
  Sec cache D misses (26): Counter name (number)
    131: Counter overflow value
    786: Total counts
-----
```

Function list, in descending order by counts

```
-----
[index]      counts      %   cum.%   samples  function (dso: file, line)
-----
  [1]         131  16.7%  16.7%     1  init2da (generic: generic.c, 1430)
  [2]         131  16.7%  33.3%     1  genLog (generic: generic.c, 1686)
  [3]         131  16.7%  50.0%     1  _write (libc.so.1: writeSCI.c, 27)
           393  50.0% 100.0%     3  **OTHER** (includes excluded DSOs, rld, etc.)

           786 100.0% 100.0%     6  TOTAL
-----
```

7.2.4 ideal Experiment Reports

For ideal experiments, prof generates a function list annotated with the number of cycles and instructions attributed to the function, and the estimated time spent in the function.

`prof` does not take into account interactions between basic blocks. Within a single basic block, `prof` computes cycles for one execution and multiplies it with the number of times that basic block is executed.

If any of the object files linked into the application have been stripped of line-number information (with `ld -x` for example), `prof` warns about the affected procedures. The instruction counts for such procedures are shown as a procedure total, not on a per-basic-block basis. Where a line number would normally appear in a report on a function without line numbers, question marks appear instead.

- The `excl.secs` column shows the minimum number of seconds that might be spent in the function under ideal conditions.
- The `excl.%` represents how much of the program's total time was spent in the function.
- The `cum.%` column shows the cumulative percentage of time spent so far in the program after the function executes.
- The `cycles` column reports the number of machine cycles used by the function.
- The `instructions` column shows the number of instructions executed by a function.
- The `calls` column reports the number of calls to the function.
- The `procedure (dso: file, line)` column lists the procedure, its DSO name, its file name, and the line number.

SpeedShop profile listing generated Mon Feb 2 13:23:25 1998

```
prof generic.ideal.m10966
    generic (n32): Target program
        ideal: Experiment name
        it:cu: Marching orders
R4400 / R4000: CPU / FPU
    1: Number of CPUs
    175: Clock frequency (MHz.)
```

Experiment notes--

```
    From file generic.ideal.m10966:
    Caliper point 0 at target begin, PID 10966
        /home/hickory28/dep/speedshop/progs.etc/linpack.demos/c/generic.pixie
    Caliper point 1 at exit(0)
```

Summary of ideal time data (ideal)--

2062563179: Total number of instructions executed
 3929944273: Total computed cycles
 22.457: Total computed execution time (secs.)
 1.905: Average cycles / instruction

 Function list, in descending order by exclusive ideal time

[index]	excl.secs	excl.%	cum.%	cycles	instructions
calls	function	(dso: file, line)			
[1]	21.453	95.5%	95.5%	3754320037	1971220024
1	anneal	(generic: generic.c, 1573)			

If the -gprof flag is added to prof, a list of callers and callees of each function is provided:

 Butterfly function list, in descending order by inclusive ideal time

(callsite) [index]	attrib.%	attrib.time	incl.time	caller	
[index]	incl.%	incl.time	self% attrib.%	self-time attrib.time	procedure [index] incl.time callee
[1]	100.0%	22.456	0.0%	0.000	__start [1]
			100.0%	22.456	22.456 main [2]
			0.0%	0.000	0.000 __readenv_sigfpe [131]
			0.0%	0.000	0.000 __istart [132]
[2]	100.0%	22.456	0.0%	0.000	22.456 __start [1]
	100.0%	22.456	100.0%	22.456	22.456 main [2]
					22.456 Scriptstring [3]
[3]	100.0%	22.456	0.0%	0.000	22.456 main [2]
	100.0%	22.456			Scriptstring
[6]			95.5%	21.454	21.454 usrttime [4]
			3.7%	0.829	0.829 libdso
[9]			0.8%	0.172	0.172 cvttrap

[11]			0.0%	0.001	0.001	iofile
			0.0%	0.000	0.000	dirstat [23]
[12]			0.0%	0.000	0.001	genLog
			0.0%	0.000	0.000	linklist [26]
			0.0%	0.000	0.000	fpetraps [27]
			0.0%	0.000	0.000	fprintf [21]
			0.0%	0.000	0.000	sprintf [17]
[60]			0.0%	0.000	0.000	strcmp

	95.5%	21.454			22.456	Scriptstring [3]
[4]	95.5%	21.454	0.0%	0.000		usrttime [4]
			95.5%	21.454	21.454	anneal
[5]			0.0%	0.000	0.001	genLog
[12]			0.0%	0.000	0.000	fprintf [21]

7.2.5 fpe Trace Reports

The report shows information for each function. The function name is show in the right column of the report. The remaining columns are described below.

- The `excl.FPEs` column shows how many floating point exceptions were found in the function.
- The `excl.%` column shows the percentage of the total number of floating-point exceptions that were found in the function.
- The `cum.%` column shows the percentage of floating-point exceptions in the program that have been encountered so far in the list.
- The `incl.FPEs` columns shows how many floating-point exceptions were found so far in the program.
- The `incl.%` column provides information on the percentage of the program's total number of floating-point exceptions found so far.

 SpeedShop profile listing generated Mon Feb 2 13:26:33 1998
 prof generic.fpe.m12213


```

generic (n32): Target program
      fpe: Experiment name
      fpe:cu: Marching orders
R4400 / R4000: CPU / FPU
      1: Number of CPUs
      175: Clock frequency (MHz.)

Experiment notes--
      From file generic.fpe.ml2213:
      Caliper point 0 at target begin, PID 12213
          /home/hickory28/dep/speedshop/progs.etc/linpack.demos/c/generic
      Caliper point 1 at exit(0)
-----
Summary of FPE callstack tracing data (fpe)--
      4: Total FPEs
      0: Samples with incomplete traceback
-----
Function list, in descending order by exclusive FPEs
-----
[index]  excl.FPEs excl.%   cum.%   incl.FPEs incl.%  function  (dso:file)

      [1]           4 100.0% 100.0%           4 100.0% fpetraps (generic: generic.c, 405)

```

7.3 Using prof Options

This section shows the output from calling `prof` with some of the options available for `prof`.

7.3.1 Using the `-dis` Option

For `pcsamp` and ideal experiments, the `-dis` option to `prof` can be used to obtain machine instruction information. `prof` provides the standard report and then appends the machine instruction information to the end of the report. The example below shows partial output from `prof`, for a `pcsamp` experiment.

```

-----
SpeedShop profile listing generated Tue Feb  3 10:48:59 1998
prof -dis generic.pcsamp.ml4493
      generic (n32): Target program
          pcsamp: Experiment name
      pc,2,10000,0:cu: Marching orders
      R4400 / R4000: CPU / FPU
          1: Number of CPUs

```

```

175: Clock frequency (MHz.)
Experiment notes--
  From file generic.pcsamp.m14493:
Caliper point 0 at target begin, PID 14493
  /home/hickory28/dep/speedshop/c/generic
Caliper point 1 at exit(0)
-----
Summary of statistical PC sampling data (pcsamp)--
      2707: Total samples
    27.070: Accumulated time (secs.)
     10.0: Time per sample (msecs.)
         2: Sample bin width (bytes)
-----
Function list, in descending order by time
-----
[index]      secs      %    cum.%    samples  function (dso: file, line)
  [1]      25.240  93.2%  93.2%     2524  anneal (generic: generic.c, 1573)
  [2]       1.090   4.0%  97.3%      109  slaveusrtime (dlslave.so: dlslave.c, 22)
  [3]       0.390   1.4%  98.7%       39  __read (libc.so.1: read.s, 20)
  [4]       0.230   0.8%  99.6%       23  cvttrap (generic: generic.c, 317)
  [5]       0.090   0.3%  99.9%        9  __xstat (libc.so.1: xstat.s, 12)
  [6]       0.010   0.0%  99.9%        1  __write (libc.so.1: write.s, 20)
  [7]       0.010   0.0% 100.0%        1  _ngetdents (libc.so.1: ngetdents.s, 16)
  [8]       0.010   0.0% 100.0%        1  _doprnt (libc.so.1: doprnt.c, 285)

      27.070 100.0% 100.0%     2707  TOTAL
-----
Disassembly listing, annotated with PC sampling overflows
-----
.
.
.
/home/hickory28/dep/speedshop/progs.etc/linpack.demos/c/generic.c
anneal: <0x10006830-0x10006b3c>    2524 total samples(93.24%)
  [1573] 0x10006830 0x27bdffd0 addiu sp,sp,-48 # 1
  [1573] 0x10006834 0xffbc0020 sd gp,32(sp) # 2
  [1573] 0x10006838 0xffbf0018 sd ra,24(sp) # 3
  [1573] 0x1000683c 0x3c030002 lui v1,0x2 # 4
  [1573] 0x10006840 0x246397e8 addiu v1,v1,-26648 # 5
  [1573] 0x10006844 0x0323e021 addu gp,t9,v1 # 6
  [1575] 0x10006848 0xd7808370 ldc1 $f0,-31888(gp) # 7
  <2 cycle stall for following instruction>

```

```

[1575] 0x1000684c 0xf7a00000 sdc1 $f0,0(sp) # 10
[1577] 0x10006850 0x24010001 li at,1 # 11
[1577] 0x10006854 0x8f82816c lw v0,-32404(gp) # 12
<2 cycle stall for following instruction>
[1577] 0x10006858 0xac410000 sw at,0(v0) # 15
[1578] 0x1000685c 0x8f998148 lw t9,-32440(gp) # 16
[1578] 0x10006860 0x0c00171b jal 0x10005c6c # 17
[1578] 0x10006864 0000000000 nop # 18
<2 cycle stall for following instruction>
[1586] 0x10006868 0xaf000008 sw zero,8(sp) # 21
[1586] 0x1000686c 0x8fa40008 lw a0,8(sp) # 22
<2 cycle stall for following instruction>
[1586] 0x10006870 0x28842710 slti a0,a0,10000 # 25
[1586] 0x10006874 0x108000ac beq a0,zero,0x10006b28 # 26
[1586] 0x10006878 0000000000 nop # 27
<2 cycle stall for following instruction>
[1588] 0x1000687c 0x24070001 li a3,1 # 30
[1588] 0x10006880 0xaf000008 sw a3,12(sp) # 31
[1588] 0x10006884 0x8f868164 lw a2,-32412(gp) # 32
<2 cycle stall for following instruction>
[1588] 0x10006888 0x8cc60000 lw a2,0(a2) # 35
<2 cycle stall for following instruction>
[1588] 0x1000688c 0x24c6ffff addiu a2,a2,-1 # 38
[1588] 0x10006890 0x8fa5000c lw a1,12(sp) # 39
<2 cycle stall for following instruction>
[1588] 0x10006894 0x00a6282a slt a1,a1,a2 # 42
[1588] 0x10006898 0x10a0009c beq a1,zero,0x10006b0c # 43
[1588] 0x1000689c 0000000000 nop # 44
<2 cycle stall for following instruction>
[1589] 0x100068a0 0x240a0001 li t2,1 # 47
^----- 1 samples(0.04%)-----^
[1589] 0x100068a4 0xaf000010 sw t2,16(sp) # 48
^----- 1 samples(0.04%)-----^
[1589] 0x100068a8 0x8f898164 lw t1,-32412(gp) # 49
<2 cycle stall for following instruction>
[1589] 0x100068ac 0x8d290000 lw t1,0(t1) # 52
<2 cycle stall for following instruction>
[1589] 0x100068b0 0x2529ffff addiu t1,t1,-1 # 55
[1589] 0x100068b4 0x8fa80010 lw t0,16(sp) # 56
<2 cycle stall for following instruction>
[1589] 0x100068b8 0x0109402a slt t0,t0,t1 # 59
[1589] 0x100068bc 0x11000089 beq t0,zero,0x10006ae4 # 60
[1589] 0x100068c0 0000000000 nop # 61

```

```
<2 cycle stall for following instruction>
[1590] 0x100068c4 0x8faf000c lw t7,12(sp) # 64
^----- 27 samples(1.00%)-----^
<2 cycle stall for following instruction>
[1590] 0x100068c8 0x25ef0001 addiu t7,t7,1 # 67
^----- 7 samples(0.26%)-----^
[1590] 0x100068cc 0x000f7080 sll t6,t7,2 # 68
^----- 30 samples(1.11%)-----^
[1590] 0x100068d0 0x01cf7021 addu t6,t6,t7 # 69
^----- 8 samples(0.30%)-----^
[1590] 0x100068d4 0x000e70c0 sll t6,t6,3 # 70
^----- 5 samples(0.18%)-----^
[1590] 0x100068d8 0x8faf0010 lw t7,16(sp) # 71
^----- 8 samples(0.30%)-----^
<2 cycle stall for following instruction>
[1590] 0x100068dc 0x01cf7021 addu t6,t6,t7 # 74
^----- 9 samples(0.33%)-----^
[1590] 0x100068e0 0x000e70c0 sll t6,t6,3 # 75
^----- 27 samples(1.00%)-----^
[1590] 0x100068e4 0x8f8f817c lw t7,-32388(gp) # 76
^----- 14 samples(0.52%)-----^
<2 cycle stall for following instruction>
[1590] 0x100068e8 0x01cf7021 addu t6,t6,t7 # 79
^----- 9 samples(0.33%)-----^
[1590] 0x100068ec 0x25ce0008 addiu t6,t6,8 # 80
^----- 28 samples(1.03%)-----^
[1590] 0x100068f0 0xd5c10000 ldc1 $f1,0(t6) # 81
^----- 7 samples(0.26%)-----^
[1590] 0x100068f4 0x8fad000c lw t5,12(sp) # 82
^----- 10 samples(0.37%)-----^
<2 cycle stall for following instruction>
[1590] 0x100068f8 0x25ad0001 addiu t5,t5,1 # 85
^----- 21 samples(0.78%)-----^
[1590] 0x100068fc 0x000d6080 sll t4,t5,2 # 86
^----- 19 samples(0.70%)-----^
[1590] 0x10006900 0x018d6021 addu t4,t4,t5 # 87
^----- 9 samples(0.33%)-----^
[1590] 0x10006904 0x000c60c0 sll t4,t4,3 # 88
^----- 14 samples(0.52%)-----^
[1590] 0x10006908 0x8fad0010 lw t5,16(sp) # 89
^----- 8 samples(0.30%)-----^
<2 cycle stall for following instruction>
[1590] 0x1000690c 0x018d6021 addu t4,t4,t5 # 92
```

```
^----- 8 samples(0.30%)-----^
[1590] 0x10006910 0x000c60c0 sll t4,t4,3 # 93
^----- 30 samples(1.11%)-----^
[1590] 0x10006914 0x8f8d817c lw t5,-32388(gp) # 94
^----- 10 samples(0.37%)-----^
<2 cycle stall for following instruction>
[1590] 0x10006918 0x018d6021 addu t4,t4,t5 # 97
^----- 8 samples(0.30%)-----^
[1590] 0x1000691c 0xd5820000 ldc1 $f2,0(t4) # 98
^----- 28 samples(1.03%)-----^
[1590] 0x10006920 0x8fab000c lw t3,12(sp) # 99
^----- 9 samples(0.33%)-----^
<2 cycle stall for following instruction>
[1590] 0x10006924 0x256b0001 addiu t3,t3,1 # 102
^----- 11 samples(0.41%)-----^
[1590] 0x10006928 0x000b5080 sll t2,t3,2 # 103
^----- 25 samples(0.92%)-----^
[1590] 0x1000692c 0x014b5021 addu t2,t2,t3 # 104
^----- 11 samples(0.41%)-----^
[1590] 0x10006930 0x000a50c0 sll t2,t2,3 # 105
^----- 8 samples(0.30%)-----^
[1590] 0x10006934 0x8fab0010 lw t3,16(sp) # 106
^----- 11 samples(0.41%)-----^
<2 cycle stall for following instruction>
[1590] 0x10006938 0x014b5021 addu t2,t2,t3 # 109
^----- 7 samples(0.26%)-----^
[1590] 0x1000693c 0x000a50c0 sll t2,t2,3 # 110
^----- 26 samples(0.96%)-----^
[1590] 0x10006940 0x8f8b817c lw t3,-32388(gp) # 111
^----- 13 samples(0.48%)-----^
<2 cycle stall for following instruction>
[1590] 0x10006944 0x014b5021 addu t2,t2,t3 # 114
^----- 9 samples(0.33%)-----^
[1590] 0x10006948 0x254afff8 addiu t2,t2,-8 # 115
^----- 26 samples(0.96%)-----^
[1590] 0x1000694c 0xd5430000 ldc1 $f3,0(t2) # 116
^----- 11 samples(0.41%)-----^
[1590] 0x10006950 0x8fa9000c lw t1,12(sp) # 117
^----- 10 samples(0.37%)-----^
<2 cycle stall for following instruction>
[1590] 0x10006954 0x00094080 sll t0,t1,2 # 120
^----- 11 samples(0.41%)-----^
```

.

The listing shows statistics about the procedure `anneal()` in the file `generic.c` and lists the beginning and ending addresses of `anneal()`: `<0x100065b8-0x100068c4>`. The five columns display the following information:

<u>Column...</u>	<u>Displays...</u>
1	Line number of the instruction: [1573].
2	Beginning address of the instruction: 0x10006830.
3	Instruction in hexadecimal: 0x27bdfd0.
4	Assembler form (mnemonic) of the instruction: <code>addiu sp,sp,-48</code> .
5	Cycle in which the instruction executed: # 1.

Other information includes:

- The number of times the immediately preceding branch was executed and taken (`ideal` only).
- The total number of cycles in a basic block and the percentage of the total cycles for that basic block, the number of times the branch terminating that basic block was executed, and the number of cycles for one execution of that basic block (`ideal` only).
- The total number of samples in an instruction (`pcsamp` only).
- Any cycle stalls, that is, cycles that were wasted.

7.3.2 Using the -S Option

For `ideal` experiments, the `-S` option to `prof` can be used to obtain source line information. `prof` provides the standard report and then appends the source line information to the end of the report.

This example shows output from calling `prof` for an `ideal` experiment:

```
-----  
SpeedShop profile listing generated Tue Feb  3 13:49:07 1998  
  prof -S generic.ideal.m15682  
        generic (n32): Target program  
            ideal: Experiment name  
            it:cu: Marching orders
```

```

R4400 / R4000: CPU / FPU
          1: Number of CPUs
          175: Clock frequency (MHz.)

```

Experiment notes--

```

From file generic.ideal.m15682:
Caliper point 0 at target begin, PID 15682
/home/hickory28/dep/speedshop/c/generic.pixie
Caliper point 1 at exit(0)

```

Summary of ideal time data (ideal)--

```

          2062563562: Total number of instructions executed
          3929944935: Total computed cycles
           22.457: Total computed execution time (secs.)
           1.905: Average cycles / instruction

```

.
.
.

disassembly listing

*DF*_generic.MIPS.stubs_1

*DF*_dlslave.text_2@0x5ffe40e0-0x5ffe4ec8: <0x10001ad8-0x10001ec4>

```

 7 total cycles(0.00%) invoked 0 times, average ? cycles/invoation
[1] 0x10001ad8 0x0006000d break 0x6 # 1
^---      0 total cycles(0.00%) executed      0 times, average  1 cycles.---^
[1] 0x10001adc 0x8f998010 lw t9,-32752(gp) # 1
[1] 0x10001ae0 0x03e07825 move t7,ra # 2
<1 cycle stall for following instruction>
[1] 0x10001ae4 0x0320f809 jalr ra,t9 # 4
[1] 0x10001ae8 0x3418003a ori t8,zero,0x3a # 5
<2 cycle stall for following instruction>
^---      7 total cycles(0.00%) executed      1 times, average  7 cycles.---^
[1] 0x10001aec 0000000000 nop # 1
[1] 0x10001af0 0x8f998010 lw t9,-32752(gp) # 2
[1] 0x10001af4 0x03e07825 move t7,ra # 3
<1 cycle stall for following instruction>
[1] 0x10001af8 0x0320f809 jalr ra,t9 # 5
[1] 0x10001afc 0x3418003b ori t8,zero,0x3b # 6
<2 cycle stall for following instruction>
^---      0 total cycles(0.00%) executed      0 times, average  8 cycles.---^
[1] 0x10001b00 0000000000 nop # 1
[1] 0x10001b04 0x8f998010 lw t9,-32752(gp) # 2

```

```

[1] 0x10001b08 0x03e07825 move t7,ra # 3
<1 cycle stall for following instruction>
[1] 0x10001b0c 0x0320f809 jalr ra,t9 # 5
[1] 0x10001b10 0x3418003c ori t8,zero,0x3c # 6
<2 cycle stall for following instruction>
^---      0 total cycles(0.00%) executed      0 times, average  8 cycles.---^
[1] 0x10001b14 0000000000 nop # 1
[1] 0x10001b18 0x8f998010 lw t9,-32752(gp) # 2
[1] 0x10001b1c 0x03e07825 move t7,ra # 3
<1 cycle stall for following instruction>
[1] 0x10001b20 0x0320f809 jalr ra,t9 # 5
[1] 0x10001b24 0x3418003d ori t8,zero,0x3d # 6
<2 cycle stall for following instruction>
^---      0 total cycles(0.00%) executed      0 times, average  8 cycles.---^
[1] 0x10001b28 0000000000 nop # 1
[1] 0x10001b2c 0x8f998010 lw t9,-32752(gp) # 2
[1] 0x10001b30 0x03e07825 move t7,ra # 3
<1 cycle stall for following instruction>
[1] 0x10001b34 0x0320f809 jalr ra,t9 # 5
[1] 0x10001b38 0x3418003e ori t8,zero,0x3e # 6
<2 cycle stall for following instruction>
^---      0 total cycles(0.00%) executed      0 times, average  8 cycles.---^
[1] 0x10001b3c 0000000000 nop # 1
[1] 0x10001b40 0x8f998010 lw t9,-32752(gp) # 2
[1] 0x10001b44 0x03e07825 move t7,ra # 3
<1 cycle stall for following instruction>
[1] 0x10001b48 0x0320f809 jalr ra,t9 # 5
[1] 0x10001b4c 0x3418003f ori t8,zero,0x3f # 6
<2 cycle stall for following instruction>
^---      0 total cycles(0.00%) executed      0 times, average  8 cycles.---^
[1] 0x10001b50 0000000000 nop # 1
[1] 0x10001b54 0x8f998010 lw t9,-32752(gp) # 2
[1] 0x10001b58 0x03e07825 move t7,ra # 3
<1 cycle stall for following instruction>
[1] 0x10001b5c 0x0320f809 jalr ra,t9 # 5
[1] 0x10001b60 0x34180040 ori t8,zero,0x40 # 6
<2 cycle stall for following instruction>
^---      0 total cycles(0.00%) executed      0 times, average  8 cycles.---^
[1] 0x10001b64 0000000000 nop # 1
[1] 0x10001b68 0x8f998010 lw t9,-32752(gp) # 2
[1] 0x10001b6c 0x03e07825 move t7,ra # 3
<1 cycle stall for following instruction>
[1] 0x10001b70 0x0320f809 jalr ra,t9 # 5

```



```

[1] 0x10001b74 0x34180041 ori t8,zero,0x41 # 6
<2 cycle stall for following instruction>
^---      0 total cycles(0.00%) executed      0 times, average  8 cycles.---^
[1] 0x10001b78 0000000000 nop # 1
[1] 0x10001b7c 0x8f998010 lw t9,-32752(gp) # 2
[1] 0x10001b80 0x03e07825 move t7,ra # 3
<1 cycle stall for following instruction>
[1] 0x10001b84 0x0320f809 jalr ra,t9 # 5
[1] 0x10001b88 0x34180042 ori t8,zero,0x42 # 6
<2 cycle stall for following instruction>
.
.
.

```

7.3.3 Using the -calipers Option

When you run `prof` on the output of an experiment in which you have recorded caliper points, you can use the `-calipers` option to specify the area of the program for which you want to generate a performance report. For example, if you set just one caliper point in the middle of your program, `prof` can provide a report from the beginning of the program up to the first caliper point using the following command:

```
prof -calipers 0 1
```

`prof` can also provide a report from the caliper point to the end of the program using the following command:

```
prof -calipers 1 2
```

If you set two caliper points, `prof` can generate a report from the first to the second caliper point:

```
prof -calipers 1 2
```

7.3.4 Using the -gprof Option

For `ideal`, `usertime`, and `fpe` experiments, the `-gprof` option to `prof` can be used to obtain inclusive basic block counting information. `prof` provides the standard report and then appends the inclusive function counts information to the end of the report. The example below shows partial output from `prof`, showing just the inclusive function counts report.

With inclusive cycle counting, `prof` prints a list of functions at the end, which are called but not defined. This list includes functions starting with `_rld` because `rld` is not instrumented. It also includes functions from `libss`; they are instrumented, but their data is normally excluded.

`prof` fails to list cycles of a procedure in the inclusive listing for the following reasons:

- `init` & `fini` sections and MIPS stubs are not part of any procedure.
- Calls to procedures that don't use a "jump and link" are not recognized as procedure calls.
- When global procedures with the same name are executed in different DSOs, only one of them is listed.

These exceptions are listed at the end of the report.

This example shows output from calling `prof` for a `usertime` experiment:

```
-----
SpeedShop profile listing generated Thu Feb 12 13:52:09 1998
  prof -gprof generic.usertime.m10981
        generic (n32): Target program
            usertime: Experiment name
            ut:cu: Marching orders
R4400 / R4000: CPU / FPU
            1: Number of CPUs
            175: Clock frequency (MHz.)

Experiment notes--
  From file generic.usertime.m10981:
  Caliper point 0 at target begin, PID 10981
    /home/hickory28/dep/speedshop/progs.etc/linpack.demos/c/generic
  Caliper point 1 at exit(0)

-----
Summary of statistical callstack sampling data (usertime)--
            809: Total Samples
            0: Samples with incomplete traceback
        24.270: Accumulated Time (secs.)
            30.0: Sample interval (msecs.)

-----
Function list, in descending order by exclusive time
-----
[index]  excl.secs  excl.%   cum.%   incl.secs  incl.%   samples  procedure
(dso: file, line)
```

```

    [4]      22.770  93.8%  93.8%    22.770  93.8%      759  anneal
(generic: generic.c, 1573)
    [6]       1.020   4.2%  98.0%     1.020   4.2%      34  slaveusrtime
(dlslave.so: dlslave.c, 22)
    [9]       0.210   0.9%  98.9%     0.210   0.9%      7  cvttrap
(generic: generic.c, 317)
   [12]       0.120   0.5%  99.4%     0.120   0.5%      4  _pm_create_special
(libc.so.1: pm.c, 191)
   [14]       0.090   0.4%  99.8%     0.090   0.4%      3  _migr_policy_args_init
(libc.so.1: pm.c, 398)
   [10]       0.030   0.1%  99.9%     0.180   0.7%      6  iofile
(generic: generic.c, 464)
   [11]       0.030   0.1% 100.0%     0.150   0.6%      5  _doscan_f
(libc.so.1: inline_doscan.c, 615)
    [1]       0.000   0.0% 100.0%    24.270 100.0%     809  __start
(generic: crtlttext.s, 101)
    [2]       0.000   0.0% 100.0%    24.270 100.0%     809  main
(generic: generic.c, 101)
    [3]       0.000   0.0% 100.0%    24.270 100.0%     809  Scriptstring
(generic: generic.c, 184)
    [5]       0.000   0.0% 100.0%    22.770  93.8%     759  usrtime
(generic: generic.c, 1377)
   [15]       0.000   0.0% 100.0%     0.090   0.4%      3  dirstat
(generic: generic.c, 348)
   [16]       0.000   0.0% 100.0%     0.090   0.4%      3  _pread
(libc.so.1: preadSCI.c, 33)
   [13]       0.000   0.0% 100.0%     0.120   0.5%      4  _fullocale
(libc.so.1: _locale.c, 77)
    [7]       0.000   0.0% 100.0%     1.020   4.2%      34  libdso
(generic: generic.c, 619)
    [8]       0.000   0.0% 100.0%     1.020   4.2%      34  dlslave_routine
(dlslave.so: dlslave.c, 7)

```

```

-----
Butterfly function list, in descending order by inclusive time
-----

```

```

      attrib.% attrib.time                incl.time  caller
(callsite) [index]
 [index]  incl.%  incl.time  self%  self-time  procedure [index]
                attrib.% attrib.time  incl.time  callee
(callsite) [index]
-----

```

```

    [1] 100.0%    24.270    0.0%    0.000    __start [1]
                                     100.0%    24.270    24.270    main
(@0x10001fac; generic: crt1text.s, 177) [2]
-----
    100.0%    24.270                                24.270    __start
(@0x10001fac; generic: crt1text.s, 177) [1]
    [2] 100.0%    24.270    0.0%    0.000    main [2]
                                     100.0%    24.270    24.270    Scriptstring
(@0x10002040; generic: generic.c, 111) [3]
-----
    100.0%    24.270                                24.270    main
(@0x10002040; generic: generic.c, 111) [2]
    [3] 100.0%    24.270    0.0%    0.000    Scriptstring
[3]
                                     93.8%    22.770    22.770    usrtime
(@0x10002460; generic: generic.c, 214) [5]
                                     4.2%    1.020    1.020    libdso
(@0x10002460; generic: generic.c, 214) [7]
                                     0.9%    0.210    0.210    cvttrap
(@0x10002460; generic: generic.c, 214) [9]
                                     0.7%    0.180    0.180    iofile
(@0x10002460; generic: generic.c, 214) [10]
                                     0.4%    0.090    0.090    dirstat
(@0x10002460; generic: generic.c, 214) [15]
-----
    93.8%    22.770                                22.770    usrtime
(@0x10005c30; generic: generic.c, 1393) [5]
    [4] 93.8%    22.770    93.8%    22.770    anneal [4]
-----
    93.8%    22.770                                24.270    Scriptstring
(@0x10002460; generic: generic.c, 214) [3]
    [5] 93.8%    22.770    0.0%    0.000    usrtime [5]
                                     93.8%    22.770    22.770    anneal
(@0x10005c30; generic: generic.c, 1393) [4]
-----
    4.2%    1.020                                1.020    dlslave_routine
(@0x5ffe0690; dlslave.so: dlslave.c, 9) [8]
    [6] 4.2%    1.020    4.2%    1.020    slaveusrtime
[6]
-----
    4.2%    1.020                                24.270    Scriptstring
(@0x10002460; generic: generic.c, 214) [3]
    [7] 4.2%    1.020    0.0%    0.000    libdso [7]

```

```

                                4.2%      1.020      1.020  dlslave_routine
(@0x100032a0; generic: generic.c, 650) [8]
-----
                                4.2%      1.020      1.020  libdso
(@0x100032a0; generic: generic.c, 650) [7]
      [8]      4.2%      1.020      0.0%      0.000      dlslave_routine [8]
                                4.2%      1.020      1.020  slaveusrtime
(@0x5ffe0690; dlslave.so: dlslave.c, 9) [6]
-----
                                0.9%      0.210      24.270  Scriptstring
(@0x10002460; generic: generic.c, 214) [3]
      [9]      0.9%      0.210      0.9%      0.210      cvttrap [9]
-----
                                0.7%      0.180      24.270  Scriptstring
(@0x10002460; generic: generic.c, 214) [3]
      [10]     0.7%      0.180      0.1%      0.030      iofile [10]
                                0.6%      0.150      0.150  _doscan_f
(@0x10002d48; generic: generic.c, 483) [11]
-----
                                0.6%      0.150      0.180  iofile
(@0x10002d48; generic: generic.c, 483) [10]
      [11]     0.6%      0.150      0.1%      0.030      _doscan_f [11]
                                0.5%      0.120      0.120  _fullocale
(@0x0fadebe4; libc.so.1: inline_doscan.c, 808) [13]
-----
                                0.5%      0.120      0.120  _fullocale
(@0x0fb0b1b8; libc.so.1: _locale.c, 84) [13]
      [12]     0.5%      0.120      0.5%      0.120      _pm_create_special [12]
-----
                                0.5%      0.120      0.150  _doscan_f
(@0x0fadebe4; libc.so.1: inline_doscan.c, 808) [11]
      [13]     0.5%      0.120      0.0%      0.000      _fullocale [13]
                                0.5%      0.120      0.120  _pm_create_special
(@0x0fb0b1b8; libc.so.1: _locale.c, 84) [12]
-----
                                0.4%      0.090      0.090  _pread
(@0x0fb05928; libc.so.1: preadSCI.c, 33) [16]
      [14]     0.4%      0.090      0.4%      0.090      _migr_policy_args_init [14]
-----
                                0.4%      0.090      24.270  Scriptstring
(@0x10002460; generic: generic.c, 214) [3]
      [15]     0.4%      0.090      0.0%      0.000      dirstat [15]
                                0.4%      0.090      0.090  _pread

```

```
(@0x10002a5c; generic: generic.c, 381) [16]
-----
          0.4%      0.090                                0.090  dirstat
(@0x10002a5c; generic: generic.c, 381) [15]
 [16]      0.4%      0.090      0.0%      0.000      _pread [16]
                                0.4%      0.090      0.090  _migr_policy_args_init
(@0xfb05928; libc.so.1: preadSCI.c, 33) [14]
-----
```

7.4 Generating Reports for Different Machine Types

If you need to generate a report for a machine model that is different from the one on which the experiment was performed, you can use several of the `prof` options to specify a machine model.

For example, if you record an `ideal` experiment on an R4000™ processor with a clock frequency of 100 megahertz, but you want to generate a report for an R10000 processor, the `prof` command would be

```
prof -r10000 -clock 196 generic.ideal.m4561
```

7.5 Generating Reports for Multiprocessed Executables

You can gather data from executables that use the `sproc()` and `sprocsp()` system calls, such as those executables generated by POWER Fortran and POWER C. Prepare and run the job using the same method as for uniprocessed executables. For multiprocessed executables, each thread of execution writes its own separate data file. View these data files with `prof`.

The only difference between multiprocessed and regular executables is how the data files are named. The data files are named *prog_name.experiment_type.id*.

The experiment ID, `id`, consists of one or two letters (designating the process type) and the process ID number. See Table 16 for letter codes and descriptions. This naming convention avoids the potential conflict of multiple threads attempting to write simultaneously to the same file.

Table 16. Letter Codes in Experiment ID Numbers

Letter Code	Description
m	Master process created by <code>ssrun</code>
p	Process created by a call to <code>sproc()</code>
f	Process created by a call to <code>fork()</code>
s	Process created by a call to <code>system()</code>
e	Process created by a call to <code>exec()</code>
fe	Process created by a call to <code>fork()</code> and <code>exec()</code>

7.6 Generating Compiler Feedback Files

If you run an `ideal` experiment, run `prof` with the `-feedback` option to generate a feedback file that can be used to arrange procedures more efficiently on the next recompile. You can rearrange procedures using the `-fb` flag to `cc`, or using the `cord` command. For more information, view the `cc` or `cord` reference page.

7.7 Interpreting Reports

If the target process was blocked for a long time as a result of an instruction, that instruction will show up as having a low or zero CPU time. On the other hand, CPU-intensive instructions will show up as having a high CPU time.

One way to sanity-check inclusive cycle counts is to look at the percentage cycles for `__start()`. If the value is anything less than 98-99%, the inclusive report is suspect. Look for other warnings that `prof` didn't take into account certain procedures.

Using SpeedShop in Expert Mode: pixie [8]

This chapter provides information on how to run `pixie` and `prof` without invoking `ssrun`. By calling `pixie` directly, you can generate the following performance data:

- An exact count of the number of times each basic block in your program is executed. A basic block is a sequence of instructions that is entered only at the beginning and exits only at the end.
- Counts for callers of a routine as well as counts for callees. `prof` can provide inclusive basic block counting by propagating regular counts to callers of a routine.

For more information on basic block counting and inclusive basic block counting, see Chapter 7, page 81.

This chapter contains the following sections:

- Using `pixie`, see Section 8.1, page 109.
- Obtaining Basic Block Counts, see Section 8.2, page 112.
- Obtaining Inclusive Basic Block Counts, see Section 8.3, page 119.

8.1 Using `pixie`

Use `pixie` to measure the frequency of code execution. `pixie` reads an executable program, partitions it into basic blocks, and writes (instruments) an equivalent program containing additional code that counts the execution of each basic block.

Note that the execution time of an instrumented program is two to five times longer than that of an uninstrumented one. This timing change may alter the behavior of a program that deals with a graphical user interface (GUI), or depends on events such as `SIGALRM` that are based on an external clock.

8.1.1 `pixie` Syntax

The syntax for `pixie` is

```
pixie prog_name [options]
```

prog_name Name of the input program.
options Zero or more of the keywords listed in Table 17.

8.1.2 pixie Options

Table 17 lists `pixie` options. For a complete list of options, view the `pixie` reference page.

Table 17. Options for `pixie`

Name	Result
<code>-addlibs lib1.so:...libN.so</code>	Adds <code>lib1.so:...libN.so</code> to the library list of the executable. No libraries are added by default.
<code>-copy</code>	Produces a copy of the target with function list (map) and arc list (graph) sections but does not instrument the target.
<code>-counts_file file</code>	Specifies the name to be used for the output <code>.Counts</code> file. By default, <code>.Counts</code> is appended to the original program name.
<code>-dso</code>	Treats executable as an o32 DSO. Performs a search of standard o32 library directories. A <code>.pix32</code> extension is used.
<code>-dso32</code>	Treats executable as an n32 DSO. Performs a search of standard n32 library directories. A <code>.pixn32</code> extension is used.
<code>-dso64</code>	Treats executable as an n64 DSO. Performs a search of standard n64 library directories. A <code>.pix64</code> extension is used.
<code>-directory dir_name</code>	Writes output files to <code>dir_name</code> . Files are written to the current directory by default.
<code>-fcncounts</code>	Produces an instrumented executable that counts function calls and arc calls, but not basic-block or branch counts.

Name	Result
<code>--[no]autopixieo32</code>	Permits or prevents a recursive instrumenting of all dynamic shared libraries used by the input file during run time. <code>pixie</code> keeps the timestamp and checksum from the original executable. Thus, before instrumenting a shared library, <code>pixie</code> checks any pixified files that it finds matching the <code>lib</code> it is to instrument. If the fields match, they are not instrumented. <code>pixie</code> cannot detect shared libraries opened with <code>dlopen()</code> (and hence does not instrument them). All used DSOs need to be instrumented for the pixified executable to work. The default behavior with shared libraries is <code>-noautopixie</code> . The default behavior with an executable is <code>-autopixie</code> .
<code>-[no]longbranch</code>	During instrumentation, some transformations can push a branch offset beyond its legal range and <code>pixie</code> generates warnings about branch offsets being out of range. This option causes <code>pixie</code> to transform these instructions into jumps. The default is <code>-nolongbranch</code> .
<code>-[no]verbose</code>	Prints or suppresses messages summarizing the binary-to-binary translation process. The default is <code>-noverbose</code> .
<code>-pixie_file <i>name</i></code>	Specify the name of the pixified executable.
<code>-suffix <i>.suffix</i></code>	Appends <code>.suffix</code> to the pixified executable and DSOs. The default suffix is <code>.pixie</code> .

8.1.3 pixie Output

The `pixie` command generates a set of files with a `.pixie` extension. These files are essentially copies of your original executable and any DSOs you specified in the call to `pixie` with code inserted to enable the collection of performance data when the `.pixie` version of your program is run.

If you use the `-verbose` flag with `pixie`, it reports the size of the old and new code. The new code size is the size of the code `pixie` will actually execute. It does not count read-only data (including a copy of the original text and another data block the same size as the original text) put into the text section. Calling `size` on the `.pixie` file reports a much larger text size than `pixie -verbose`, because `size` also counts everything in the text segment.

When you run the `.pixie` version of your program, one or more `.Counts` files are generated. The name of an output `.Counts` file is that of the original program with any leading directory names removed and `.Counts` appended. If the program executes calls to `sproc()`, `sprocsp()` or `fork()`, multiple

.Counts files are generated-one for each process in the share group. In this case, each file will have the process ID appended to its name.

8.2 Obtaining Basic Block Counts

Use this procedure to obtain basic block counts. Also refer to Figure 1, page 114, which illustrates how basic block counting works.

1. Compile and link your program. The following example uses the input file `myprog.c`:

```
% cc -o myprog myprog.c
```

The `cc` compiler compiles `myprog.c` into an executable called `myprog`.

2. Run `pixie` to generate the equivalent program containing basic-block-counting code.

```
% pixie myprog
```

`pixie` takes `myprog` and writes an equivalent program, `myprog.pixie`, containing additional code that counts the execution of each basic block. `pixie` also writes an equivalent program for each shared object used by the program (in the form: `libname.so.pix*`), containing additional code that counts the execution of each basic block. For example, if `myprog` uses `libc.so.1`, `pixie` generates `libc.so.1.pix*`. (The value of `*` depends on the ABI).

3. Set the path for your `.pixie` files. `pixie` uses the `rld` search path for libraries (see `rld(1)` for the default paths). If the `.pixie` files are in your local directory, set the path as

```
% setenv LD_LIBRARY_PATH .
```

4. Execute the file(s) generated by `pixie` (`myprog.pixie`) in the same way you executed the original program:

```
% myprog.pixie
```

This program generates a list of basic block counts in files named `myprog.Counts`. If the program executes `fork` or `sproc`, a process ID is appended to the end of the file name (for example, `myprog.Counts.345`) for each process.

Note: Your program may not run as you expect when you invoke it with a `.pixie` extension. Some programs, `uncompress` and `vi` for example, treat their arguments differently when the name of the program changes. You may need to rename the `.pixie` version of your program back to its original name.

To generate a valid `.Counts` file, your program must terminate normally or with a call to `exit()`. If it terminates with a signal such as `SIGINT`, the program must use a signal handler and leave the program through `exit()`.

5. Run the profile formatting program `prof` specifying the name of the original program and the `.Counts` file for the program:

```
% prof myprog myprog.Counts
```


`prof` extracts information from `myprog.Counts` and prints it in an easily readable format. If multiple `.Counts` files exist, you can use the wildcard character (*) to specify all of the files.

```
% prof myprog myprog.Counts*
```

You can run the program several times, altering the input data, to create multiple profile data files.

The time computation assumes a “best case” execution; actual execution takes longer. This is because the time includes predicted stalls within a basic block, but not actual stalls that may occur entering a basic block. It also assumes that all instructions and data are in cache, that is, it excludes the delays due to cache misses and memory fetches and stores.

The complete output of the `-pixie` option is often extremely large. Use the `-quit` option with `prof` to restrict the size of the report. Refer to Chapter 7, page 81, for details about `prof` options.



a11550

Figure 1. How Basic Block Counting Works

8.2.1 Examples of Basic Block Counting

The examples in this section illustrate how to use `prof` to obtain basic block counting information from a C program, `generic`.

8.2.1.1 Example Using `prof` with no Options

The partial listing below illustrates the report generated for basic block counts in `generic`. `prof` first provides a standard report of basic block counts, then provides a report reflecting any options provided to `prof`.

```
-----
SpeedShop profile listing generated Tue Feb  3 14:25:43 1998
  prof generic generic.Counts
      generic (n32): Target program
      pixie-counts: Experiment name
      pixie-counts: Marching orders
R4400 / R4000: CPU / FPU
      1: Number of CPUs
      175: Clock frequency (MHz.)
-----
Summary of ideal time data (pixie-counts)--
      2062563311: Total number of instructions executed
      3929944454: Total computed cycles
      22.457: Total computed execution time (secs.)
      1.905: Average cycles / instruction
-----
Function list, in descending order by exclusive ideal time
-----
 [index]  excl.secs  excl.%   cum.%      cycles  instructions  calls
function (dso: file, line)
  [1]      21.453   95.5%   95.5%    3754320037  1971220024    1
anneal (generic: generic.c, 1573)
  [2]       0.829    3.7%   99.2%    145001152   75000732     1
slaveusrtime (dlslave.so: dlslave.c, 22)
  [3]       0.171    0.8%  100.0%    30000081   16000054     1
cvtttrap (generic: generic.c, 317)
  [4]       0.001    0.0%  100.0%    101504     58124        1
init2da (generic: generic.c, 1430)
  [5]       0.001    0.0%  100.0%     91200     38400       1600
_drand48 (libc.so.1: drand48.c, 116)
  [6]       0.001    0.0%  100.0%     89072     55011        447
fread (libc.so.1: fread.c, 34)
```

```
    [7]      0.000    0.0%   100.0%      74859      47364      53
_doprnt (libc.so.1: doprnt.c, 285)
    [8]      0.000    0.0%   100.0%      64035      29479      628
__sinf (libm.so: fsin.c, 93)
    [9]      0.000    0.0%   100.0%      32355       7182       9
offtime (libc.so.1: time_comm.c, 180)
   [10]      0.000    0.0%   100.0%      17112      11916      305
_readdir (libc.so.1: readdir.c, 135)
.
.
.
```

- The `excl.secs` column shows the number of seconds spent in each procedure. For example, 21.453 seconds were spent in the `anneal` function. The time represents an idealized computation based on modeling the machine. It ignores potential floating-point interlocks and memory latency time (cache misses and memory bus contention).
- The `excl.%` column lists the percentage of the program's total time spent in each function. The `anneal` function takes 95.5% of the total time.
- The `cum%` column shows the cumulative percentage of calls. For example, 99.2% of the total program time was spent in the top two functions in the listing: `anneal` and `slaveusertime`.
- The `cycles` column reports the number and percentage of machine cycles used for the procedure. For example, 3754320037 cycles were spent in the `anneal` function.
- The `instructions` column shows the number of instructions executed by a function. For example, the `anneal` function executed 1971220024 instructions.
- The `calls` column reports the number of calls to each function. For example, there was just one call to the `anneal` function.
- The `procedure (dso: file, line)` columns list the function name, its DSO name, its file name, and its line number. For example, the first line reports statistics for the function `anneal`, in the file `generic`, the DSO `generic.c`, and the line number 1573.

8.2.1.2 Example Using prof -heavy

The following partial listing shows the source code lines responsible for the largest portion of execution time produced with the -heavy option.

```
% prof -heavy generic generic.Counts
```

The partial listing below shows basic block counts sorted in descending order of cycles used. The fields in the report are described in Chapter 7, page 81.

```
-----
SpeedShop profile listing generated Tue Feb  3 15:02:11 1998
```

```
  prof -heavy generic generic.Counts
    generic (n32): Target program
      pixie-counts: Experiment name
      pixie-counts: Marching orders
R4400 / R4000: CPU / FPU
      1: Number of CPUs
      175: Clock frequency (MHz.)
-----
```

```
Summary of ideal time data (pixie-counts)--
```

```
  2062563311: Total number of instructions executed
  3929944454: Total computed cycles
    22.457: Total computed execution time (secs.)
    1.905: Average cycles / instruction
.
.
.
-----
```

```
Line list, in descending order by time
-----
```

excl.secs	%	cum.%	cycles	invocations	function (dso: file, line)
19.800	88.2%	88.2%	3464962830	14440000	anneal (generic: generic.c, 1590)
1.608	7.2%	95.3%	281457170	14440000	anneal (generic: generic.c, 1589)
0.497	2.2%	97.5%	87000454	5000000	slaveusertime (dlslave.so: dlslave.c, 29)
0.331	1.5%	99.0%	57999996	5000000	slaveusertime (dlslave.so: dlslave.c, 30)
0.048	0.2%	99.2%	8437511	500000	cvtttrap (generic: generic.c, 327)
0.048	0.2%	99.4%	8437511	500000	cvtttrap (generic: generic.c, 334)
0.044	0.2%	99.6%	7770000	380000	anneal (generic: generic.c, 1588)
0.037	0.2%	99.8%	6562500	500000	cvtttrap (generic: generic.c, 328)
0.037	0.2%	100.0%	6562500	500000	cvtttrap (generic: generic.c, 335)
0.001	0.0%	100.0%	130009	10000	anneal (generic: generic.c, 1586)
0.000	0.0%	100.0%	43919	1600	init2da (generic: generic.c, 1443)

8.2.1.3 Example Using prof -quit

You can limit the output of `prof` to collect information on only the most time-consuming parts of the program by specifying the `-quit` option. You can instruct `prof` to quit after a particular number of lines of output, after listing the elements consuming more than a certain percentage of the total, or after the portion of each listing whose cumulative use is a certain amount.

Consider the following sample listing:

```
% prof -quit 4 generic generic.Counts

-----
SpeedShop profile listing generated Wed Feb  4 10:18:58 1998
  prof -quit 4 generic generic.Counts
    generic (n32): Target program
      pixie-counts: Experiment name
      pixie-counts: Marching orders
    R4400 / R4000: CPU / FPU
      1: Number of CPUs
      175: Clock frequency (MHz.)
-----
Summary of ideal time data (pixie-counts)--
    2062563311: Total number of instructions executed
    3929944454: Total computed cycles
      22.457: Total computed execution time (secs.)
      1.905: Average cycles / instruction
-----
Function list, in descending order by exclusive ideal time
-----
 [index]  excl.secs  excl.%   cum.%      cycles instructions   calls  function
(dso: file, line)
   [1]    21.453   95.5%   95.5%   3754320037   1971220024       1  anneal
(generic: generic.c, 1573)
   [2]     0.829    3.7%   99.2%   145001152    75000732       1  slaveusrtime
(dlslave.so: dlslave.c, 22)
   [3]     0.171    0.8%  100.0%    30000081    16000054       1  cvttrap
(generic: generic.c, 317)
   [4]     0.001    0.0%  100.0%     101504      58124       1  init2da
(generic: generic.c, 1430)
```

8.3 Obtaining Inclusive Basic Block Counts

Inclusive basic block counting counts basic blocks and generates a call graph. By propagating regular counts to callers of a routine, `prof` provides inclusive basic block counting. For more information on inclusive basic block counting, see Chapter 4, page 49.

To see inclusive data, run the profile formatting program `prof` specifying the name of the original program, the `-gprof` flag, and the `.Counts` file for the program.

```
% prof -gprof myprog myprog.Counts
```

`prof` extracts information from `myprog.Counts` and prints it in an easily readable format. If multiple `.Counts` files exist, you can use the wildcard character (*) to specify all of the files.

```
% prof -gprof myprog myprog.Counts*
```

8.3.1 Example of `prof -gprof`

This section contains part of a sample output obtained by using the `-gprof` option. The fields in the report are explained in detail in the report, but are not provided in this example. For more information on the `-gprof` option, see Chapter 7, page 81, (The format of the output has been adjusted slightly.)

```
% prof -gprof generic generic.Counts
```

```
-----
SpeedShop profile listing generated Wed Feb  4 10:22:01 1998
  prof -gprof generic generic.Counts
    generic (n32): Target program
      pixie-counts: Experiment name
      pixie-counts: Marching orders
R4400 / R4000: CPU / FPU
      1: Number of CPUs
      175: Clock frequency (MHz.)
-----
Summary of ideal time data (pixie-counts)--
  2062563311: Total number of instructions executed
  3929944454: Total computed cycles
    22.457: Total computed execution time (secs.)
    1.905: Average cycles / instruction
-----
```

```

.
.
-----
Butterfly function list, in descending order by inclusive ideal time
-----
      attrib.% attrib.time          incl.time caller (callsite) [index]
[index]  incl.%  incl.time    self%  self-time      procedure [index]
              attrib.% attrib.time  incl.time callee (callsite) [index]
.
.
.
-----
      100.0%    22.456
[3]  100.0%    22.456    0.0%    0.000          22.456 main [2]
              95.5%    21.454          21.454 Scriptstring [3]
              3.7%    0.829          0.829 usertime [4]
              0.8%    0.172          0.172 libdso [6]
              0.0%    0.001          0.172 cvttrap [9]
              0.0%    0.000          0.001 iofile [11]
              0.0%    0.000          0.000 dirstat [23]
              0.0%    0.000          0.001 genLog [12]
              0.0%    0.000          0.000 linklist [26]
              0.0%    0.000          0.000 fpetraps [27]
              0.0%    0.000          0.000 fprintf [21]
              0.0%    0.000          0.000 sprintf [17]
              0.0%    0.000          0.000 strcmp [60]
-----
      95.5%    21.454
[4]  95.5%    21.454    0.0%    0.000          22.456 Scriptstring [3]
              95.5%    21.454          21.454 usertime [4]
              0.0%    0.000          21.454 anneal [5]
              0.0%    0.000          0.001 genLog [12]
              0.0%    0.000          0.000 fprintf [21]
-----

```

Miscellaneous Commands [9]

This chapter describes SpeedShop commands for exploring memory usage and paging, and for printing data files generated by SpeedShop tools. It contains the following sections:

- Using the `thrash` Command, see Section 9.1, page 121.
- Using the `squeeze` Command, see Section 9.3, page 122.
- Calculating the Working Set of a Program, see Section 9.5, page 123.
- Dumping Performance Data Files, see Section 9.6, page 125.
- Dumping Compiler Feedback Files, see Section 9.8, page 131.

9.1 Using the `thrash` Command

The `thrash` command allows you to explore paging behavior by allocating a region of virtual memory, and either randomly or sequentially accessing that memory to explore the system paging behavior.

9.1.1 `thrash` Syntax

`thrash` [*args*]

args

One or more of the following flags:

- | | |
|------------------------------|---|
| <code>-k <i>N</i></code> | The amount of memory to access in kilobytes, where <i>N</i> is the number of kilobytes. |
| <code>-m <i>N</i></code> | The amount of memory to access in megabytes, where <i>N</i> is the number of megabytes. |
| <code>-n <i>count</i></code> | The number of references to make before exiting. The default is 10,000. |
| <code>-p <i>N</i></code> | The amount of memory to access in pages, where <i>N</i> is the number of pages. |
| <code>-s</code> | Sequential thrashing. The default is random. |

`-w time` The amount of time thrash should sleep after thrashing but before exiting.

9.2 Effects of thrash

Once the memory is allocated, thrash prints a message on *stdout* saying how much memory it is using and then proceeds to thrash over it. Here's an example:

```
fraser 82% thrash -m 4

thrashing randomly: 4.00 MB (= 0x00400000 = 4194304 bytes = 1024 pages)

10000 iterations
```

You can use `thrash` in conjunction with `ssusage` and `squeeze` to determine the approximate available working memory on a system, as described in Section 9.5, page 123.

9.3 Using the `squeeze` Command

The `squeeze` command allows you to specify an amount of virtual memory to lock down into real memory, thus making it unavailable to other processes. This command can only be used only by superuser.

9.3.1 `squeeze` Syntax

`squeeze [flag] amount`

flag One of the following flags. If no flag is specified, the default is megabytes.

<code>-k</code>	Kilobytes
<code>-m</code>	Megabytes
<code>-p</code>	Pages
<code>-%</code>	A percentage of the installed memory

amount The amount of memory to be locked.

9.4 Effects of squeeze

`squeeze` performs the following operations:

- Locks down the amount of virtual memory you supply as an argument to the command.
- Prints a message to *stdout* that provide information on how much memory has been locked, and how much working memory is available.
- Sleeps indefinitely, or until interrupted by SIGINT or SIGTERM. At that time, it frees up the memory and exits with an exit message.

Wait until after the exit message is printed before doing any experiments.

Here's an example:

```
fraser 1# squeeze 4
squeeze: leaving 60.00 MB ( = 0x03c01000 = 62918656 ) available memory;
        pinned 4.00 MB ( = 0x00400000 = 4194304 ) at address 0x1000e000;
        from 64.00 MB ( = 0x04001000 = 67112960 ) installed memory.
```

Use Ctrl-C to exit `squeeze`. The following message is printed:

```
squeeze exiting
```

9.5 Calculating the Working Set of a Program

You can use the `thrash`, `squeeze`, and `ssusage` commands together to determine the approximate working set of a program as follows. For all practical purposes, the working set of your program is the size of memory allocated.

The process involves three steps. First you determine the working set of the kernel and other applications:

1. Choose a machine that has a large amount of physical memory (enough to allow your target application to run without any paging other than at start-up).
2. Make sure that the machine is running a minimal number of applications that will remain fairly consistent for the duration of these steps.

3. Run `thrash` with `ssusage` to determine the working set of the kernel and any other applications you have running.

In this example, the `thrash` command uses 4 MB of memory:

```
ssusage thrash -m 4
```

When the `thrash` command completes, `ssusage` prints the resource usage of `thrash`; the value labelled `majf` gives the number of major page faults (i.e. the number of faults that required a physical read.) When you run on a machine with a large amount of physical memory, this value is the number of faults needed to start the program, which is the minimum number for any run. For more information on `ssusage`, see Chapter 5, page 61.

4. As superuser in a separate window, run the `squeeze` command to lock down an amount of memory.
5. Rerun `thrash` with `ssusage`:

```
ssusage thrash -m 4
```
6. Repeat steps 1 and 2, increasing the amount of memory for `squeeze`, until the `majf` number begins to rise.

The amount of working memory available reported by `squeeze` at the point at which page faults begin to rise for `thrash` tells you the combined working set of `thrash` (approximately 4 MB), the kernel and any other applications you have running.

7. Deduct the 4 MB that `thrash` uses from the amount of working memory reported by `squeeze` at the point the page faults began to rise.

This computation helps you find out the approximate working set of the kernel and any other applications that are running on the machine. You'll need this number when you reach the next steps.

8. Determine the working set of the program you're interested in. Make sure the applications that the machine is running remain consistent with the setup from step 2.
9. Run `ssusage` with your program to ensure that the machine has the amount of memory your program needs.

```
ssusage prog_name
```

When your program exits, `ssusage` prints the application's resource usage: the `majf` field gives the number of major page faults. When run on a

machine with a large amount of physical memory, this value is the number of faults needed to start the program, which is the minimum number for any run.

10. Switch to superuser.
11. Run `squeeze` to lock down an amount of memory. The following example locks down 15 megabytes of memory:

```
squeeze 15
```

12. Rerun your program with `ssusage`.
13. Repeat steps 11 and 12 until the majf number begins to rise.
14. Deduct the amount squeezed at the point at which the application begins to page fault from the total amount of physical memory in the system.

This computation determines the combined working set of your program, the kernel and any other applications you have running.

15. Deduct the amount of working memory calculated in step 7 from the total amount of physical memory in the system.

This computation determines the approximate working set of your program.

9.6 Dumping Performance Data Files

All the performance data for a single process is in one file. The file begins with a prologue and continues with a mixture of performance data, sample records, and control records.

The `ssdump` command can be used for printing performance data files. It provides a formatted ASCII dump of one or more performance experiment data files. This command is most likely to be useful in verifying performance data that does not seem accurate when reported through `prof`.

9.6.1 `ssdump` Syntax

```
ssdump [options] data_file [... data_file]
```

options Zero or more of the following print options:

-d	Prints detailed information for each bead (defined in Section 9.7, page 126). For compressed beads, the compressed form will be dumped.
----	---

-D	Prints detailed information for each bead. For compressed beads, the uncompressed form will be dumped.
-h	Prints the hex contents of the body of each bead.
-i <i>index</i>	Prints only one bead at <i>index</i> in the file.
-q	Suppresses printing of those fields that will normally change from run to run such as process IDs and time stamps. This option is useful for QA work, to enable automatic comparisons of recorded experiments.
-s <i>offset</i>	Prints only one bead at <i>offset</i> into the file.

9.7 Experiment File Format

The file is written as a string of “beads,” each of which is a record with

- a 32-bit type
- a 32-bit byte count
- a body whose length is given by the byte-count, rounded up to a double-word boundary

The file prologue consists of these beads:

- file-identifier bead, which acts as a magic number, indicating that the file is a SpeedShop data file
- machine and executable name
- hardware inventory describing the machine
- machine page size
- O/S revision, date, and checksum information about the executable
- target name (the target is the executable after instrumentation)
- arguments with which the target was invoked
- instrumentation performed
- types of performance data that are to be recorded in the remainder of the file

The following example calls `ssdump` on performance data for a `pcsamp` experiment:

```
ssdump generic.pcsamp.m847
```

Below is some partial output from `ssdump`. The format has been adjusted slightly to meet presentation needs.

```
Printing experiment record file ``generic.pcsamp.m847`` (2688 bytes), last written
on Tue 15 Apr 1997 15:27:02
```

```
SpeedShop File Preface 1, offset 0 = 0x00000000 (size 32)
```

```
file type 1 (SSRUN); version 4
```

```
process control flags: 0xd
```

```
  _SPEEDSHOP_TRACE_FORK=True
```

```
  _SPEEDSHOP_TRACE_FORK_TO_EXEC=False
```

```
  _SPEEDSHOP_TRACE_SPROC=True
```

```
  _SPEEDSHOP_TRACE_EXEC=True
```

```
  _SPEEDSHOP_TRACE_SYSTEM=False
```

```
ancestor exp file name:
```

```
created: Tue 15 Apr 1997 15:26:10.719
```

```
Hardware Inventory 2, offset 40 = 0x00000028 (size 280)
```

```
hardware inventory: 17 items
```

```
class 1, type 1, contrlr 100, unit 255, state 12
```

```
class 1, type 3, contrlr 0, unit 0, state 8192
```

```
class 1, type 2, contrlr 0, unit 0, state 8208
```

```
class 4, type 8, contrlr 0, unit 0, state 2
```

```
class 5, type 5, contrlr 0, unit 0, state 1
```

```
class 3, type 3, contrlr 0, unit 0, state 16384
```

```
class 3, type 4, contrlr 0, unit 0, state 16384
```

```
class 3, type 9, contrlr 0, unit 0, state 64
```

```
class 3, type 1, contrlr 0, unit 0, state 67108864
```

```
class 12, type 3, contrlr 0, unit 0, state 16
```

```
class 8, type 7, contrlr 17, unit 0, state 16777472
```

```
class 10, type 3, contrlr 0, unit 0, state 16400
```

```
class 8, type 0, contrlr 0, unit 0, state 1
```

```
class 2, type 1, contrlr 0, unit 13, state 2
```

```
class 2, type 2, contrlr 0, unit 2, state 0
```

```
class 2, type 2, contrlr 0, unit 1, state 0
```

```
class 7, type 14, contrlr 0, unit 0, state 0
```

```
Experiment name 3, offset 328 = 0x00000148 (size 8)
```

```
pcsamp
```

```
Experiment marching orders 4, offset 344 = 0x00000158 (size 16)
```

```

pc,2,10000,0:cu

Capture module symbol          5, offset 368 = 0x00000170 (size 16)
pc,2,10000,0

Capture module symbol          6, offset 392 = 0x00000188 (size 8)
cu

Executable file                7, offset 408 = 0x00000198 (size 8)
generic

Target file                    8, offset 424 = 0x000001a8 (size 8)
generic

Target arguments               9, offset 440 = 0x000001b8 (size 32)
Time: Tue 15 Apr 1997 15:26:10.719, process pid = 847
arguments: ""

Target begin                   10, offset 480 = 0x000001e0 (size 40)
process # -1, pid = 847, event # 0
event type = 0,0
at time = Tue 15 Apr 1997 15:26:10.719

Program Object List           11, offset 528 = 0x00000210 (size 312)
process # -1, pid = 847, event # 0, -- 5 DSOs
Program Object 0, Named`generic'
  Link Time Address: 0x0000000010000000
  Run Time Address: 0x0000000010000000
  Size: 0x00000000000007000 (28672)
  Base Pointer: 0x0000000000000000

Program Object 1, Named`/usr/lib32/libss.so'
  Link Time Address: 0x0000000009e50000
  Run Time Address: 0x0000000009e50000
  Size: 0x00000000000002000 (8192)
  Base Pointer: 0x0000000000000000

Program Object 2, Named`/usr/lib32/libssrt.so'
  Link Time Address: 0x0000000009da0000
  Run Time Address: 0x0000000009da0000
  Size: 0x00000000000008b000 (569344)
  Base Pointer: 0x0000000000000000

Program Object 3, Named`/usr/lib32/libm.so'
  Link Time Address: 0x000000000f840000

```

```
Run Time Address: 0x00000000f840000
      Size: 0x00000000028000 (163840)
      Base Pointer: 0x0000000000000000

Program Object 4, Named`/usr/lib32/libc.so.1'
  Link Time Address: 0x00000000fa00000
  Run Time Address: 0x00000000fa00000
      Size: 0x000000000108000 (1081344)
      Base Pointer: 0x0000000000000000

Target DSO open          12, offset 848 = 0x00000350 (size 56)
  process # -1, pid = 847, event # 0
    at time = Tue 15 Apr 1997 15:27:00.716
  fname = ./dlslave.so

Program Object List      13, offset 912 = 0x00000390 (size 360)
  process # -1, pid = 847, event # 0, -- 6 DSOs
  Program Object 0, Named`generic'
    Link Time Address: 0x0000000010000000
    Run Time Address: 0x0000000010000000
      Size: 0x0000000000007000 (28672)
      Base Pointer: 0x0000000000000000

  Program Object 1, Named`/usr/lib32/libss.so'
    Link Time Address: 0x0000000009e50000
    Run Time Address: 0x0000000009e50000
      Size: 0x0000000000002000 (8192)
      Base Pointer: 0x0000000000000000

  Program Object 2, Named`/usr/lib32/libssrt.so'
    Link Time Address: 0x0000000009da0000
    Run Time Address: 0x0000000009da0000
      Size: 0x0000000000008b000 (569344)
      Base Pointer: 0x0000000000000000

  Program Object 3, Named`/usr/lib32/libm.so'
    Link Time Address: 0x00000000f840000
    Run Time Address: 0x00000000f840000
      Size: 0x0000000000028000 (163840)
      Base Pointer: 0x0000000000000000

  Program Object 4, Named`/usr/lib32/libc.so.1'
    Link Time Address: 0x00000000fa00000
```

Run Time Address: 0x00000000fa00000
Size: 0x000000000108000 (1081344)
Base Pointer: 0x0000000000000000

Program Object 5, Named`./dlslave.so'
Link Time Address: 0x000000005ffe0000
Run Time Address: 0x000000005ffe0000
Size: 0x000000000001000 (4096)
Base Pointer: 0x0000000000000000

Sample event trigger 14, offset 1280 = 0x00000500 (size 40)
process # -1, trap index # -1
at time = Tue 15 Apr 1997 15:27:01.989, #-1

Compressed PC sampling array (16-bit) 15, offset 1328 = 0x00000530 (size 320)
compressed short array, dso index = 0, array size = 7168, 156
compressed

Compressed PC sampling array (16-bit) 16, offset 1656 = 0x00000678 (size 16)
compressed short array, dso index = 1, array size = 2048, 4 compressed

Compressed PC sampling array (16-bit) 17, offset 1680 = 0x00000690 (size 40)
compressed short array, dso index = 2, array size = 142336, 16
compressed

Compressed PC sampling array (16-bit) 18, offset 1728 = 0x000006c0 (size 16)
compressed short array, dso index = 3, array size = 40960, 4 compressed

Compressed PC sampling array (16-bit) 19, offset 1752 = 0x000006d8 (size 64)
compressed short array, dso index = 4, array size = 270336, 28
compressed

Compressed PC sampling array (16-bit) 20, offset 1824 = 0x00000720 (size 48)
compressed short array, dso index = 5, array size = 1024, 20 compressed

PC sampling array (16-bit) 21, offset 1880 = 0x00000758 (size 16)
short array, dso index = -1, array size = 1

Resource usage 22, offset 1904 = 0x00000770 (size 680)

Sample data end marker 23, offset 2592 = 0x00000a20 (size 40)

Target termination 24, offset 2640 = 0x00000a50 (size 40)

```

process # -1, pid = 847, event # 0
event type = 0,0 (normal termination, exit status 0)
at time = Tue 15 Apr 1997 15:27:02.231

```

```
** End-of-File                25, offset 2688 = 0x00000a80 (size 0)
```

```
**** End of experiment record file ``generic.pcsamp.m847``
```

9.8 Dumping Compiler Feedback Files

The `fbdump` command can be used to print out the compiler feedback files generated by running `prof -feedback`. For more information on using compiler feedback files, view the `cord` or `cc` reference pages.

9.8.1 fbdump Syntax

```
fbdump options file_name
```

options Zero or more of the options described in table Table 18.

file_name The feedback filename. This file has a `.fb` extension.

Table 18. Options for `fbdump`

Option	Prints.
<code>-all</code>	Feedback using all options. This is the default.
<code>-ascii</code>	Feedback in the same style as earlier version of the feedback dump program.
<code>-bb</code>	Feedback per basic block table as described in “ <code>cmplrs/fb.h</code> ”. If <code>-verbose</code> is specified, all basic blocks are printed, even those with zero execution counts. If <code>-verbose</code> is not specified, <code>fbdump</code> prints only the basic blocks that have non-zero execution counts.
<code>-call</code>	Feedback call table as described in “ <code>cmplrs/fb.h</code> ”. If <code>-verbose</code> is specified, all the points of call are printed, even if they have not been called. If <code>-verbose</code> is not specified, <code>fbdump</code> prints only the relevant information on the calls.
<code>-header</code>	Feedback file header as described in “ <code>cmplrs/fb.h</code> ”.

Option	Prints.
-proc	Feedback procedure table as described in "cmplrs/fb.h". If <code>-verbose</code> is specified, all procedures will be printed, even if they are not invoked. If <code>-verbose</code> is not specified, <code>fbdump</code> prints only the relevant information on the procedures that have been invoked.
-sections	Feedback file section headers table as described in "cmplrs/fb.h".
-str	Feedback string table.
-verbose	All the information in verbose mode including a table with all zero entries.

Glossary

This glossary defines terms used in this document.

basic block	A set of instructions with a single entry point, a single exit point, and no branches into or out of the set.
call stack	A software stack of functions and routines used by the compiler. The functions and routines are listed in the reverse order, from top to bottom, in which they were called. If function <i>a</i> is immediately below function <i>b</i> in the stack, then <i>a</i> was called by <i>b</i> . The function at the bottom of the stack is the one currently executing.
context switch	The act of saving the state of one process and replacing it with that of another when both processes time-share a single processor.
CPU time	Process virtual time (see the glossary entry) plus time spent when the system is running on behalf of the process, performing such tasks as executing a system call. This is the time returned in <code>usertime</code> and <code>pcsamp</code> experiments. It can be specified in an experiment by using the <code>ut, 30000, 2</code> marching orders.
overflow interval	As used by the hardware counter experiments, it is the number at which a hardware counter exceeds a preset value. See the <code>speedshop(1)</code> man page, <code>dsc_hwc</code> experiment.
PC	Program counter.
process virtual time	Time spent when a program is actually running. This does not include either 1) the time spent when the program is swapped out and waiting for a CPU or 2) the time when the operating system is in control, such as executing a system call for the program. The marching orders <code>ut, 30000, 1</code> return process virtual time.
statistical data	Sampling. The results from this method of data gathering vary from run to run.

wall-clock time

Total time a program takes to execute, including the time it takes waiting for a CPU. This is real time, not computer time. The marching orders `ut, 30000, 0` return wall-clock time.

A

API

- setting calipers, 11

B

basic block counting

- overview ideal experiment
- overview, 5

C

calipers

- automatic, 76
- sample traps calipers
 - using the debugger debugger, 76, 78

calipers -calipers option prof

- calipers option setting calipers, 11

calipers setting calipers, 75

compiler feedback files rearranging procedures

- cord, 107

compiler optimization restrictions, 64

.Counts file pixie

- .Counts file, 112

cy_hwc experiment experiments

- cy_hwc, 55

D

data display anomalies, 64

dc_hwc experiment experiments

- dc_hwc, 56

debugger

- setting calipers, 11
- using ssrun ssrun

- using a debugger, 73

demo program SpeedShop

- C and C++, 13

dsc_hwc experiment experiments

- dsc_hwc, 56

- DSOs shared libraries, 7

E

examples

- c tutorial

- c c, 13

- fortran tutorial, 33

executables

- calculating a working set, 123

experiments

- hardware counter, 88

experiment data

- file format files

- performance data, 126

- filenames experiment data

- controlling output file, 65

experiment data files

- performance data, 10

experiments

- choosing, 10, 49

F

fbdump

- overview, 4

fbdump files

- compiler feedback cord prof

- compiler feedback, 131

floating-point exception trace

- overview fpe trace experiment

- overview, 6
- fork processes
 - forking sproc system exec, 7
- fpe trace experiment
 - tutorial experiments
 - fpe trace floating-point exceptions, 29
- fpe trace experiment experiments
 - fpe trace floating-point exceptions, 59
- fpe trace experiment experiments
 - fpe trace floating-point exception trace tracing floating-point exceptions, 10

G

- gfp_hwc experiment experiments
 - gfp_hwc, 56
- gi_hwc experiment experiments
 - gi_hwc, 55
- gprof
 - example, 119

H

- hardware counter experiment reports, 88
- hardware counter experiments
 - tutorial experiments
 - hardware counter hardware counter overflows, 22, 41
- hardware counter experiments experiments
 - hardware counter hwc experiments
 - hardware counter overflows, 54
 - hardware counters hwc experiments, 10
- hardware counter numbers, 57

I

- ic_hwc experiment experiments
 - ic_hwc, 55
- ideal experiment
 - tutorial experiments

- ideal basic block counting, 25, 43
- ideal experiment experiments
 - ideal basic block counting pixie, 51
- isc_hwc experiment experiments
 - isc_hwc, 55

L

- LD_LIBRARY_PATH environment variable
 - LD_LIBRARY_PATH, 15
- libraries
 - linking in speedshop speedShop libraries linking libss.so libssrt.so, 77
- linpack benchmark SpeedShop demo
 - linpack, 33

M

- machine resource usage, 61
- mp fortran limitations fortran
 - limitations multi-processor executables, 64
- MPI message-passing paradigms, 8
- multi-processor executables, 7
 - profiling profiling
 - multiprocessor executables, 106

P

- pcsamp experiment
 - example, 71
 - tutorial experiments
 - pcsamp PC sampling, 19, 39
- pcsamp experiment experiments
 - pcsamp, 10
 - pcsamp PC sampling program counter sampling fpcsampx experiment, 51
- perfex, 54
- performance analysis
 - phases, 8

- theory performance problems, , 1
- performance problems, 10
 - Bugs, 3
 - cpu cpu-bound processes, 2
 - i/o i/O-bound processes, 2
 - memory memory-bound processes, 2
- pixie, 109
 - autopixie option, 111
 - command syntax pixie
 - command option, 109
 - examples, 112
 - overview, 4
 - restricting output profiling
 - quit option pixie, 113
 - setting search path search path
 - rld rld, 112
 - verbose option, 111
- prof
 - gprof example profiling
 - inclusive basic block counts, 91
 - options, 82
 - output, 85
 - overview, 4
 - syntax, 81
 - using with pixie, 81
 - using with ssrun, 81
- prof_hwc experiment experiments
 - prof_hwc, 57
- profiling
 - calipers option prof
 - calipers example calipers, 101
 - clock option, 82
 - command syntax prof
 - Also see profiling, , 81
 - dis option, 83
 - dis option prof
 - dis example, 93
 - dso option, 83
 - dsolist option, 83
 - heavy option
 - example prof, 117
 - quit option, 118
 - exclude option, 83
 - feedback option, 83
 - fpe trace experiment experiments
 - fpe fpe trace experiment reports, 92
 - gprof option, 83
 - gprof option prof
 - gprof example, 101
 - hardware counter experiments, 88
 - heavy option, 84
 - ideal experiment experiments
 - ideal ideal experiment reports, 89
 - lines option, 84
 - machine scheduler option reports
 - for different machine models, 106
 - only option, 84
 - pcsamp experiment experiments
 - pcsamp pcsamp experiment reports, 87
 - procedure invocation example, 115
 - procedures option, 84
 - processor scheduler option option, 85
 - quit option, 84
 - S option, 85
 - S option prof
 - S example, 98
 - usertime experiment experiments
 - usertime usertime experiment reports, 86
 - zero option, 85
- pthreads, 7

R

- reports
 - interpreting profiles
 - interpreting, 107
 - _RLD_LIST variable environment variables
 - _RLD_LIST, 78

S

- signals
 - setting calipers, 11

SpeedShop
 overview, 3
 speedshop api api ssrt_caliper_point, 7
 SpeedShop libraries
 libfpe_ss.so libraries
 libfpe_ss.so libfpe_ss.so, 6
 libmalloc_ss.so libraries
 libmalloc_ss.so libmalloc_ss.so, 7
 libss.so libraries
 libss.so libss.so, 6
 libssrt.so libraries
 libssrt.so libssrt.so, 6
 SpeedShop libraries libraries
 libss.so libraries
 libssrt.so libss.so libssrt.so, 78
 _SPEEDSHOP_CALIPER_POINT_SIG variable
 environment variables
 _SPEEDSHOP_CALIPER_POINT_SIG, 67
 _SPEEDSHOP_CALIPER_POINT_SIG
 signals
 setting calipers calipers, 76, 77
 _SPEEDSHOP_DEBUG_NO_SIG_TRAPS
 variable environment variables
 _SPEEDSHOP_DEBUG_NO_SIG_TRAPS, 70
 _SPEEDSHOP_DEBUG_NO_STACK_UNWIND
 variable environment variables
 _SPEEDSHOP_DEBUG_NO_STACK_UNWIND, 70
 _SPEEDSHOP_EXPERIMENT_TYPE variable
 environment variables
 _SPEEDSHOP_EXPERIMENT_TYPE, 69
 _SPEEDSHOP_FILE_BUFFER_LENGTH
 variable environment variables
 _SPEEDSHOP_FILE_BUFFER_LENGTH, 70
 _SPEEDSHOP_HWC_COUNTER_NUMBER
 variable environment variables
 _SPEEDSHOP_HWC_COUNTER_NUMBER, 57,
 67
 _SPEEDSHOP_HWC_COUNTER_OVERFLOW
 variable environment variables
 _SPEEDSHOP_HWC_COUNTER_OVERFLOW, 57,
 67
 _SPEEDSHOP_INIT_DEFERRED_SIGNAL
 variable environment variables
 _SPEEDSHOP_INIT_DEFERRED_SIGNAL, 69
 _SPEEDSHOP_MARCHING_ORDERS variable
 _SPEEDSHOP_EXPERIMENT_TYPE
 variable _SPEEDSHOP_TARGET_FILE
 variable environment variables
 _SPEEDSHOP_MARCHING_ORDERS
 environment variables
 _SPEEDSHOP_EXPERIMENT_TYPE
 environment variables, 78
 _SPEEDSHOP_MARCHING_ORDERS variable
 environment variables
 _SPEEDSHOP_MARCHING_ORDERS, 69
 _SPEEDSHOP_OUTPUT_DIRECTORY variable
 environment variables
 _SPEEDSHOP_OUTPUT_DIRECTORY, 67
 _SPEEDSHOP_OUTPUT_FILENAME variable
 environment variables
 _SPEEDSHOP_OUTPUT_FILENAME, 67
 _SPEEDSHOP_OUTPUT_NOCOMPRESS
 variable environment variables
 _SPEEDSHOP_OUTPUT_NOCOMPRESS, 67
 _SPEEDSHOP_REUSE_FILE_DESCRIPTOR
 variable environment variables
 _SPEEDSHOP_REUSE_FILE_DESCRIPTOR, 67
 _SPEEDSHOP_SAMPLING_MODE variable
 environment variables
 _SPEEDSHOP_SAMPLING_MODE, 69
 _SPEEDSHOP_SBRK_BUFFER_LENGTH
 variable environment variables
 _SPEEDSHOP_SBRK_BUFFER_LENGTH, 69
 _SPEEDSHOP_SILENT variable environment
 variables
 _SPEEDSHOP_SILENT, 66
 _SPEEDSHOP_TRACE_EXEC variable
 environment variables
 _SPEEDSHOP_TRACE_EXEC, 68
 _SPEEDSHOP_TRACE_FORK variable
 environment variables
 _SPEEDSHOP_TRACE_FORK, 68
 _SPEEDSHOP_TRACE_FORK_TO_EXEC
 variable environment variables
 _SPEEDSHOP_TRACE_FORK_TO_EXEC, 68

Tell Us About This Manual

As a user of Silicon Graphics products, you can help us to better understand your needs and to improve the quality of our documentation.

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

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

Please send the title and part number of the document with your comments. The part number for this document is 007-3311-003.

Thank you!

Three Ways to Reach Us

- To send your comments by **electronic mail**, use either of these addresses:
 - On the Internet: techpubs@sgi.com
 - For UUCP mail (through any backbone site): *[your_site]!sgi!techpubs*
- To **fax** your comments (or annotated copies of manual pages), use this fax number: 650-932-0801
- To send your comments by **traditional mail**, use this address:

Technical Publications
Silicon Graphics, Inc.
2011 North Shoreline Boulevard, M/S 535
Mountain View, California 94043-1389