

Oracle® Solaris Studio 12.4: Performance Analyzer Tutorials

ORACLE®

Part No: E37087
December 2014

Copyright © 2014, Oracle and/or its affiliates. All rights reserved.

This software and related documentation are provided under a license agreement containing restrictions on use and disclosure and are protected by intellectual property laws. Except as expressly permitted in your license agreement or allowed by law, you may not use, copy, reproduce, translate, broadcast, modify, license, transmit, distribute, exhibit, perform, publish, or display any part, in any form, or by any means. Reverse engineering, disassembly, or decompilation of this software, unless required by law for interoperability, is prohibited.

The information contained herein is subject to change without notice and is not warranted to be error-free. If you find any errors, please report them to us in writing.

If this is software or related documentation that is delivered to the U.S. Government or anyone licensing it on behalf of the U.S. Government, the following notice is applicable:

U.S. GOVERNMENT END USERS. Oracle programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, delivered to U.S. Government end users are "commercial computer software" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, use, duplication, disclosure, modification, and adaptation of the programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, shall be subject to license terms and license restrictions applicable to the programs. No other rights are granted to the U.S. Government.

This software or hardware is developed for general use in a variety of information management applications. It is not developed or intended for use in any inherently dangerous applications, including applications that may create a risk of personal injury. If you use this software or hardware in dangerous applications, then you shall be responsible to take all appropriate fail-safe, backup, redundancy, and other measures to ensure its safe use. Oracle Corporation and its affiliates disclaim any liability for any damages caused by use of this software or hardware in dangerous applications.

Oracle and Java are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners.

Intel and Intel Xeon are trademarks or registered trademarks of Intel Corporation. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. AMD, Opteron, the AMD logo, and the AMD Opteron logo are trademarks or registered trademarks of Advanced Micro Devices. UNIX is a registered trademark of The Open Group.

This software or hardware and documentation may provide access to or information on content, products, and services from third parties. Oracle Corporation and its affiliates are not responsible for and expressly disclaim all warranties of any kind with respect to third-party content, products, and services. Oracle Corporation and its affiliates will not be responsible for any loss, costs, or damages incurred due to your access to or use of third-party content, products, or services.

Copyright © 2014, Oracle et/ou ses affiliés. Tous droits réservés.

Ce logiciel et la documentation qui l'accompagne sont protégés par les lois sur la propriété intellectuelle. Ils sont concédés sous licence et soumis à des restrictions d'utilisation et de divulgation. Sauf disposition de votre contrat de licence ou de la loi, vous ne pouvez pas copier, reproduire, traduire, diffuser, modifier, breveter, transmettre, distribuer, exposer, exécuter, publier ou afficher le logiciel, même partiellement, sous quelque forme et par quelque procédé que ce soit. Par ailleurs, il est interdit de procéder à toute ingénierie inverse du logiciel, de le désassembler ou de le décompiler, excepté à des fins d'interopérabilité avec des logiciels tiers ou tel que prescrit par la loi.

Les informations fournies dans ce document sont susceptibles de modification sans préavis. Par ailleurs, Oracle Corporation ne garantit pas qu'elles soient exemptes d'erreurs et vous invite, le cas échéant, à lui en faire part par écrit.

Si ce logiciel, ou la documentation qui l'accompagne, est concédé sous licence au Gouvernement des Etats-Unis, ou à toute entité qui délivre la licence de ce logiciel ou l'utilise pour le compte du Gouvernement des Etats-Unis, la notice suivante s'applique:

U.S. GOVERNMENT END USERS. Oracle programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, delivered to U.S. Government end users are "commercial computer software" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, use, duplication, disclosure, modification, and adaptation of the programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, shall be subject to license terms and license restrictions applicable to the programs. No other rights are granted to the U.S. Government.

Ce logiciel ou matériel a été développé pour un usage général dans le cadre d'applications de gestion des informations. Ce logiciel ou matériel n'est pas conçu ni n'est destiné à être utilisé dans des applications à risque, notamment dans des applications pouvant causer des dommages corporels. Si vous utilisez ce logiciel ou matériel dans le cadre d'applications dangereuses, il est de votre responsabilité de prendre toutes les mesures de secours, de sauvegarde, de redondance et autres mesures nécessaires à son utilisation dans des conditions optimales de sécurité. Oracle Corporation et ses affiliés déclinent toute responsabilité quant aux dommages causés par l'utilisation de ce logiciel ou matériel pour ce type d'applications.

Oracle et Java sont des marques déposées d'Oracle Corporation et/ou de ses affiliés. Tout autre nom mentionné peut correspondre à des marques appartenant à d'autres propriétaires qu'Oracle.

Intel et Intel Xeon sont des marques ou des marques déposées d'Intel Corporation. Toutes les marques SPARC sont utilisées sous licence et sont des marques ou des marques déposées de SPARC International, Inc. AMD, Opteron, le logo AMD et le logo AMD Opteron sont des marques ou des marques déposées d'Advanced Micro Devices. UNIX est une marque déposée d'The Open Group.

Ce logiciel ou matériel et la documentation qui l'accompagne peuvent fournir des informations ou des liens donnant accès à des contenus, des produits et des services émanant de tiers. Oracle Corporation et ses affiliés déclinent toute responsabilité ou garantie expresse quant aux contenus, produits ou services émanant de tiers. En aucun cas, Oracle Corporation et ses affiliés ne sauraient être tenus pour responsables des pertes subies, des coûts occasionnés ou des dommages causés par l'accès à des contenus, produits ou services tiers, ou à leur utilisation.

Contents

Using This Documentation	5
Introduction to the Performance Analyzer Tutorials	7
About the Performance Analyzer Tutorials	7
Getting the Sample Code for the Tutorials	8
Setting Up Your Environment for the Tutorials	9
Introduction to C Profiling	11
About the C Profiling Tutorial	11
Setting Up the lowfruit Sample Code	12
Using Performance Analyzer to Collect Data	12
Using the Performance Analyzer to Examine the lowfruit Data	17
Using the Remote Performance Analyzer	24
Introduction to Java Profiling	27
About the Java Profiling Tutorial	27
Setting Up the jlowfruit Sample Code	28
Using Performance Analyzer to Collect Data from jlowfruit	28
Using Performance Analyzer to Examine the jlowfruit Data	32
Using the Remote Performance Analyzer	41
Java and Mixed Java-C++ Profiling	43
About the Java-C++ Profiling Tutorial	43
Setting Up the jsynprog Sample Code	44
Collecting the Data From jsynprog	45
Examining the jsynprog Data	46
Examining Mixed Java and C++ Code	48
Understanding the JVM Behavior	52
Understanding the Java Garbage Collector Behavior	56

Understanding the Java HotSpot Compiler Behavior	61
Hardware Counter Profiling on a Multithreaded Program	63
About the Hardware Counter Profiling Tutorial	63
Setting Up the mttest Sample Code	64
Collecting Data From mttest for Hardware Counter Profiling Tutorial	65
Examining the Hardware Counter Profiling Experiment for mttest	65
Exploring Clock-Profiling Data	67
Understanding Hardware Counter Instruction Profiling Metrics	69
Understanding Hardware Counter CPU Cycles Profiling Metrics	71
Understanding Cache Contention and Cache Profiling Metrics	73
Detecting False Sharing	77
Synchronization Tracing on a Multithreaded Program	81
About the Synchronization Tracing Tutorial	81
About the mttest Program	82
About Synchronization Tracing	82
Setting Up the mttest Sample Code	82
Collecting Data from mttest for Synchronization Tracing Tutorial	83
Examining the Synchronization Tracing Experiment for mttest	84
Understanding Synchronization Tracing	85
Comparing Two Experiments with Synchronization Tracing	89

Using This Documentation

- **Overview** – This manual provides step-by-step instructions for using the Oracle Solaris Studio 12.4 Performance Analyzer on sample programs.
- **Audience** – Application developers, developer, architect, support engineer
- **Required knowledge** – Programming experience, Program/Software development testing, Aptitude to build and compile software products

Product Documentation Library

The product documentation library is located at http://docs.oracle.com/cd/E37069_01.

System requirements and known problems are included in the “[Oracle Solaris Studio 12.4: Release Notes](#)”.

Access to Oracle Support

Oracle customers have access to electronic support through My Oracle Support. For information, visit <http://www.oracle.com/pls/topic/lookup?ctx=acc&id=info> or visit <http://www.oracle.com/pls/topic/lookup?ctx=acc&id=trs> if you are hearing impaired.

Feedback

Provide feedback about this documentation at <http://www.oracle.com/goto/docfeedback>.

Introduction to the Performance Analyzer Tutorials

Performance Analyzer is the Oracle Solaris Studio tool for examining performance of your Java, C, C++, and Fortran applications. You can use it to understand how well your application is performing and find problem areas. This document provides tutorials that show how to use Performance Analyzer on sample programs using step-by-step instructions.

About the Performance Analyzer Tutorials

This document features several tutorials that show how you can use Performance Analyzer to profile various types of programs. Each tutorial provides steps for using Performance Analyzer with the source files including screen shots at most steps in the tutorial.

The source code for all the tutorials is included in a single distribution. See [“Getting the Sample Code for the Tutorials” on page 8](#) for information about obtaining the sample source code.

The tutorials include the following:

- [“Introduction to C Profiling”](#)

This introductory tutorial uses a target code named `lowfruit`, written in C. The `lowfruit` program is very simple and includes code for two programming tasks which are each implemented in an efficient way and an inefficient way. The tutorial shows how to collect a performance experiment on the C target program and how to use the various data views in Performance Analyzer. You examine the two implementations of each task, and see how Performance Analyzer shows which task is efficient and which is not.
- [“Introduction to Java Profiling”](#)

This introductory tutorial uses a target code named `jlowfruit`, written in Java. Similar to the code used in the C profiling tutorial, the `jlowfruit` program is very simple and includes code for two programming tasks which are each implemented in an efficient way and an inefficient way. The tutorial shows how to collect a performance experiment on the Java target and how to use the various data views in Performance Analyzer. You examine the two implementations of each task, and see how Performance Analyzer shows which task is efficient and which is not.
- [“Java and Mixed Java-C++ Profiling”](#)

This tutorial is based on a Java code named `jsynprog` that performs a number of programming operations one after another. Some operations do arithmetic, one triggers garbage collection, and several use a dynamically loaded C++ shared object, and call from Java to native code and back again. In this tutorial you see how the various operations are implemented, and how Performance Analyzer shows you the performance data about the program.

- [“Hardware Counter Profiling on a Multithreaded Program”](#)

This tutorial is based on a multithreaded program named `mttest` that runs a number of tasks, spawning threads for each one, and uses different synchronization techniques for each task. In this tutorial, you see the performance differences between the computations in the tasks, and use hardware counter profiling to examine and understand an unexpected performance difference between two functions.

- [“Synchronization Tracing on a Multithreaded Program”](#)

This tutorial is also based on the multithreaded program named `mttest` that runs a number of tasks, spawning threads for each one, and uses different synchronization techniques for each task. In this tutorial, you examine the performance differences between the synchronization techniques.

Getting the Sample Code for the Tutorials

The programs used in the Performance Analyzer tutorials are included in a distribution that includes code used for all the Oracle Solaris Studio tools. Use the following instructions to obtain the sample code if you have not previously downloaded it.

1. Go to the Oracle Solaris Studio 12.4 Sample Applications page at <http://www.oracle.com/technetwork/server-storage/solarisstudio/downloads/solaris-studio-12-4-samples-2333090.html>.
2. Read the license from the link on the page and accept by selecting Accept.
3. Download the zip file by clicking its link and unzip using instructions on the download page.

After you download and unpack the sample files, you can find the samples in the `SolarisStudioSampleApplications/PerformanceAnalyzer` directory.

Note that the directory includes some additional samples that are not described in this document: `cachetest`, `ksynprog`, `omptest`, and `synprog`. Each sample directory includes a `Makefile` and a `README` file with instructions that you can use for some additional demonstrations of Performance Analyzer.

Setting Up Your Environment for the Tutorials

Before you try the tutorials, make sure that you have the Oracle Solaris Studio bin directory on your path and have an appropriate Java version in your path as described in [Chapter 5, “After Installing Oracle Solaris Studio 12.4,”](#) in [“Oracle Solaris Studio 12.4: Installation Guide”](#).

The `make` or `gmake` command must also be on your path so you can build the programs.

Introduction to C Profiling

This chapter covers the following topics.

- [“About the C Profiling Tutorial”](#) on page 11
- [“Setting Up the lowfruit Sample Code”](#) on page 12
- [“Using Performance Analyzer to Collect Data”](#) on page 12
- [“Using the Performance Analyzer to Examine the lowfruit Data”](#) on page 17
- [“Using the Remote Performance Analyzer”](#) on page 24

About the C Profiling Tutorial

This tutorial shows the simplest example of profiling with Oracle Solaris Studio Performance Analyzer and demonstrates how to use Performance Analyzer to collect and examine a performance experiment. You use the Overview, Functions view, Source view, and Timeline in this tutorial.

The program `lowfruit` is a simple program that executes two different tasks, one for initializing in a loop and one for inserting numbers into an ordered list. Each task is performed twice, in an inefficient way and in a more efficient way.

Tip - The [“Introduction to Java Profiling”](#) tutorial uses an equivalent Java program and shows similar activities with Performance Analyzer.

The data you see in the experiment that you record will be different from that shown here. The experiment used for the screen-shots in the tutorial was recorded on a SPARC T5 system running Oracle Solaris 11.2. The data from an x86 system running Oracle Solaris or Linux will be different. Furthermore, data collection is statistical in nature and varies from experiment to experiment, even when run on the same system and OS.

The Performance Analyzer window configuration that you see might not precisely match the screen shots. Performance Analyzer enables you to drag separator bars between components of the window, collapse components, and resize the window. Performance Analyzer records its configuration and uses the same configuration the next time it runs. Many configuration changes were made in the course of capturing the screen shots shown in the tutorial.

This tutorial is run locally on a system where Oracle Solaris Studio is installed. You can also run remotely as described in [“Using the Remote Performance Analyzer” on page 24](#).

Setting Up the lowfruit Sample Code

Before You Begin

See the following for information about obtaining the code and setting up your environment.

- [“Getting the Sample Code for the Tutorials” on page 8](#)
 - [“Setting Up Your Environment for the Tutorials” on page 9](#)
1. Copy the contents of the lowfruit directory to your own private working area with the following command:

```
% cp -r SolarisStudioSampleApplications/PerformanceAnalyzer/lowfruit mydirectory
```

where *mydirectory* is the working directory you are using.

2. Change to that working directory.

```
% cd mydirectory/lowfruit
```

3. Build the target executable.

```
% make clobber (needed only if you ran make in the directory before, but safe in any case)
```

```
% make
```

After you run make the directory contains the target program to be used in the tutorial, an executable named lowfruit.

The next section shows how to use Performance Analyzer to collect data from the lowfruit program and create an experiment.

Tip - If you prefer, you can edit the Makefile to do the following: use the GNU compilers rather than the default of the Studio compilers; build in 32-bits rather than the default of 64-bits; and add different compiler flags.

Using Performance Analyzer to Collect Data

This section describes how to use the Profile Application feature of Performance Analyzer to collect data in an experiment.

Tip - If you prefer not to follow these steps to see how to profile applications, you can record an experiment with a make target included in the Makefile for `lowfruit`:

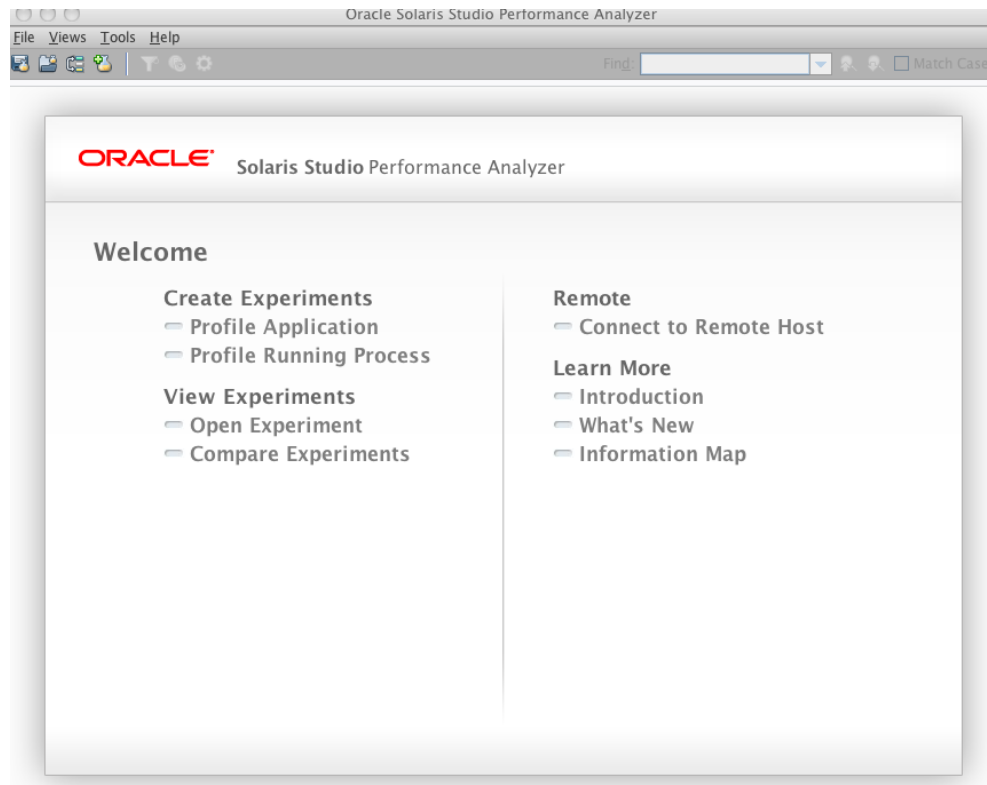
```
make collect
```

The `collect` target launches a `collect` command and records an experiment just like the one that you create using Performance Analyzer in this section. You could then skip to [“Using the Performance Analyzer to Examine the `lowfruit` Data” on page 17](#).

1. While still in the `lowfruit` directory start Performance Analyzer:

```
% analyzer
```

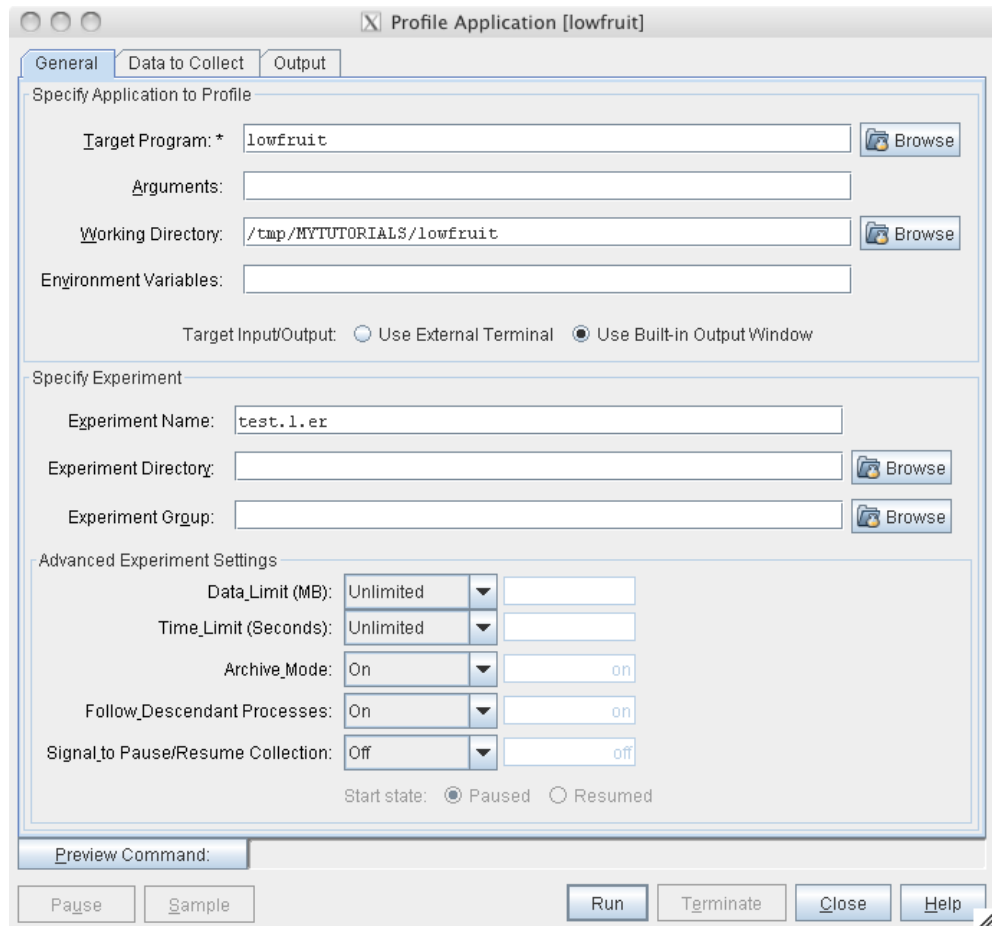
Performance Analyzer starts and displays the Welcome page.



If this is the first time you have used Performance Analyzer, no recent experiments are shown below the Open Experiment item. If you have used it before, you see a list of

the experiments you recently opened from the system where you are currently running Performance Analyzer.

2. Click the Profile Application link under Create Experiments in the Welcome page. The Profile Application dialog box opens with the General tab selected.
3. In the Target Program field, type the program name `lowfruit`.



Tip - You could start Performance Analyzer and open this dialog box directly with the program name already entered by specifying the target name when starting Performance Analyzer with the command `analyzer lowfruit`. This method only works when running Performance Analyzer locally.

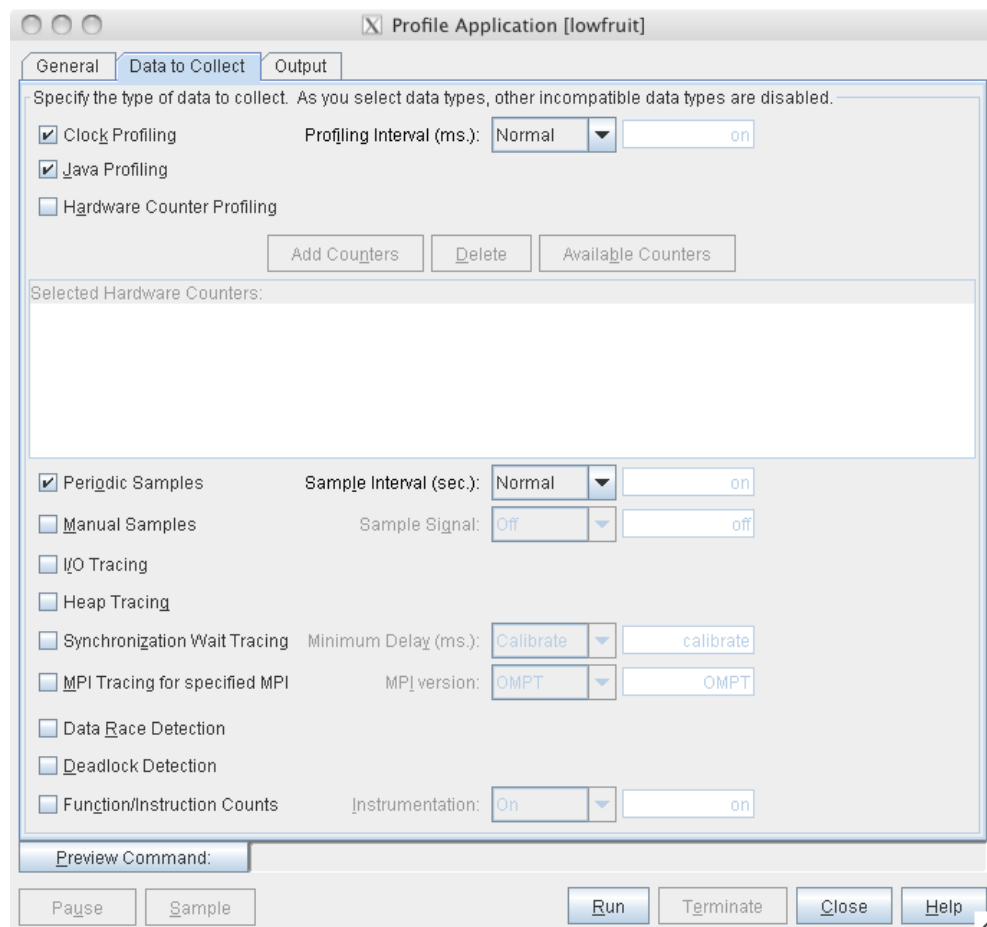
4. For the Target Input/Output option, select Use Built-in Output Window.

Target Input/Output option specifies the window to which the target program `stdout` and `stderr` will be redirected. The default value is Use External Terminal, but in this tutorial the Target Input/Output option was changed to Use Built-in Output Window to keep all the activity in the Performance Analyzer window. With this option the `stdout` and `stderr` is shown in the Output tab in the Profile Application dialog box.

If you are running remotely, the Target Input/Output option is absent because only the built-in output window is supported.

5. For the Experiment Name option, the default experiment name is `test.1.er` but you can change it to a different name as long as the name ends in `.er`, and is not already in use.
6. Click the Data to Collect tab.

The Data to Collect enables you to select the type of data to collect, and shows the defaults already selected.



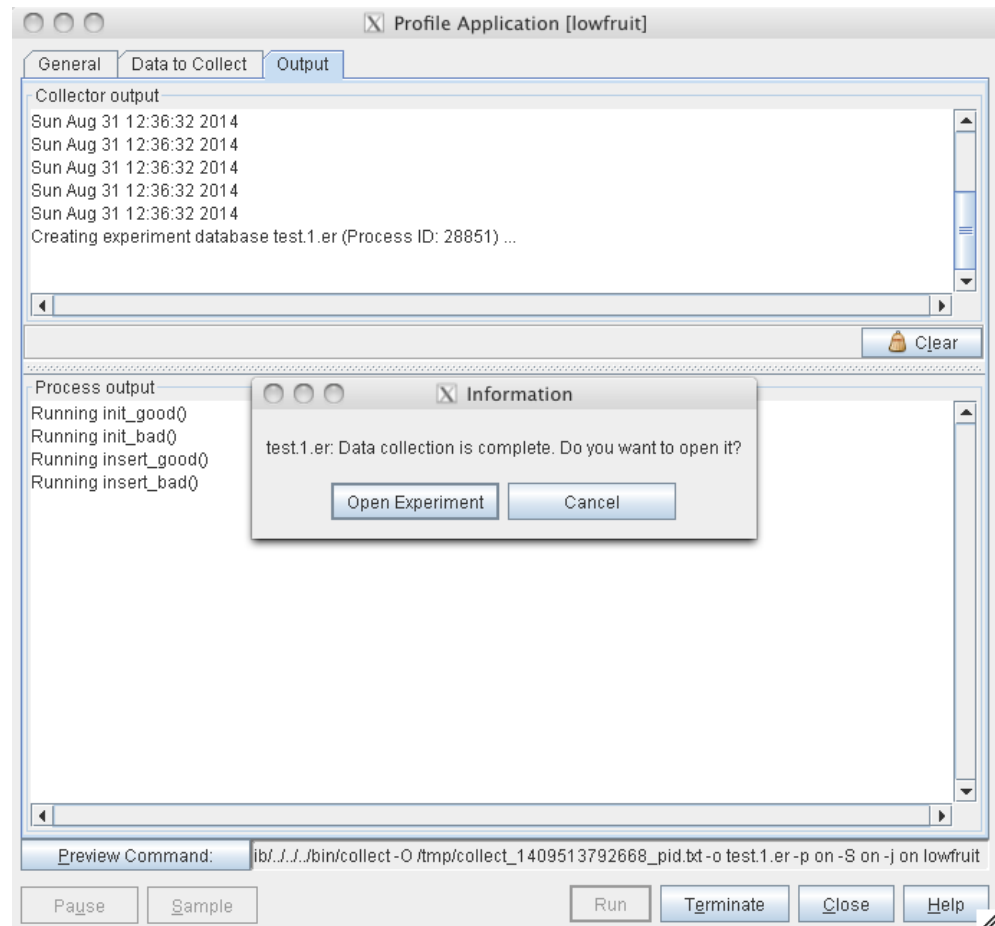
Java profiling is enabled by default as you can see in the screen shot, but it is ignored for a non-Java target such as `lowfruit`.

You can optionally click the Preview Command button and see the `collect` command that will be run when you start profiling.

7. Click the Run button.

The Profile Application dialog box displays the Output tab and shows the program output as it runs in the Process Output panel.

After the program completes, a dialog box asks if you want to open the experiment just recorded.



8. Click Open Experiment in the dialog box.

The experiment opens. The next section shows how to examine the data.

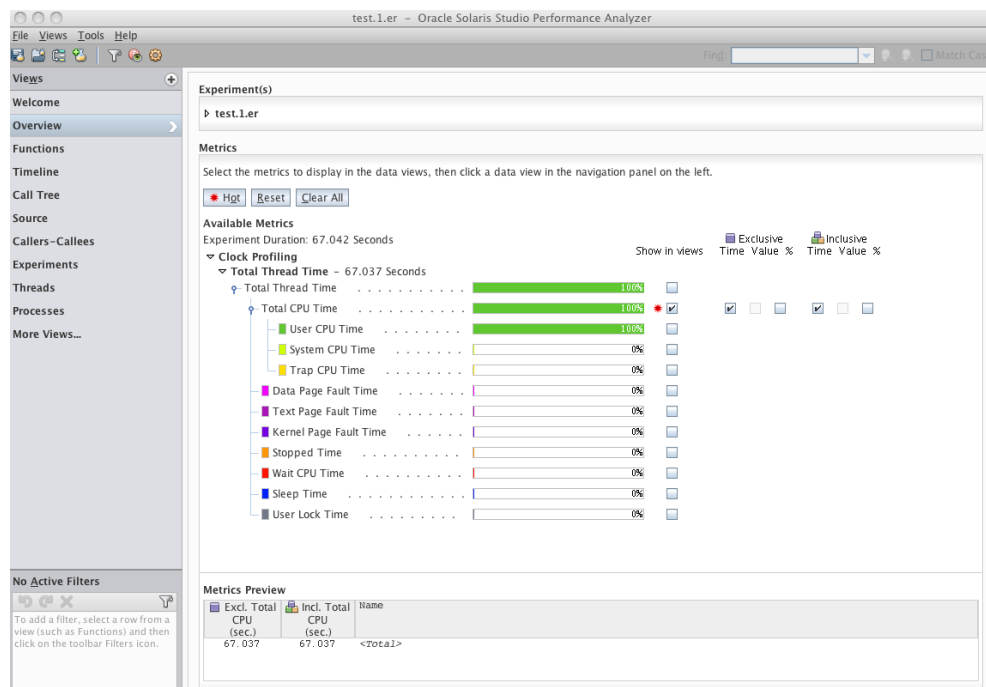
Using the Performance Analyzer to Examine the lowfruit Data

This section shows how to explore the data in the experiment created from the lowfruit sample code.

1. If the experiment you created in the previous section is not already open, you can start Performance Analyzer from the lowfruit directory and load the experiment as follows:

```
% analyzer test.1.er
```

When the experiment opens, Performance Analyzer shows the Overview screen.



In this experiment the Overview shows essentially 100% User CPU time. The program is single-threaded and that one thread is CPU-bound. The experiment was recorded on a Oracle Solaris system, and the Overview shows twelve metrics recorded but only Total CPU Time is enabled by default.

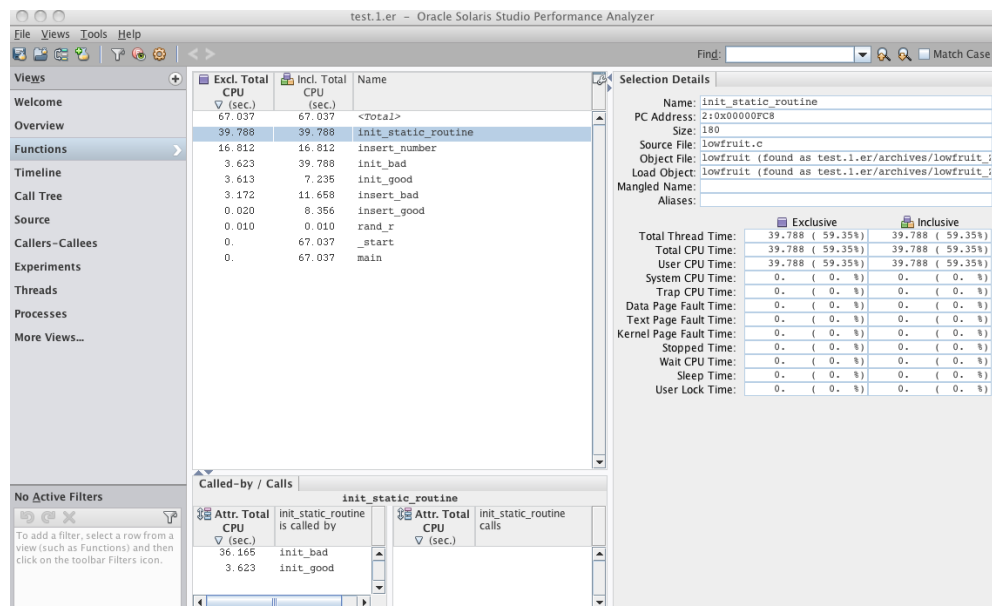
The metrics with colored indicators are the times spent in the ten microstates defined by Oracle Solaris. These metrics include User CPU Time, System CPU Time, and Trap CPU

Time which together are equal to Total CPU Time, as well as various wait times. Total Thread Time is the sum over all of the microstates.

On a Linux machine, only Total CPU Time is recorded because Linux does not support microstate accounting.

By default, both Inclusive and Exclusive Total CPU Time are selected. *Inclusive* for any metric refers to the metric value in that function or method, including metrics accumulated in all the functions or methods that it calls. *Exclusive* refers only to the metric accumulated within that function or method.

2. Click on the Functions view in the Views navigation bar on the left side, or select it using Views > Functions from the menu bar.



The Functions view shows the list of functions in the application, with performance metrics for each function. The list is initially sorted by the Exclusive Total CPU Time spent in each function. The list includes all functions from the target application and any shared objects the program uses. The top-most function, the most expensive one, is selected by default.

The Selection Details window on the right shows all the recorded metrics for the selected function.

The Called-by/Calls panel below the functions list provides more information about the selected function and is split into two lists. The Called-by list shows the callers of the selected function and the metric values show the attribution of the total metric for the function to its callers. The Calls list shows the callees of the selected function and shows how the Inclusive metric of its callees contributed to the total metric of the selected

function. If you double-click a function in either list in the Called-by/Calls panel, the function becomes the selected function in the main Functions view.

3. Experiment with selecting the various functions to see how the windows in the Functions view update with the changing selection.

The Selection Details window shows you that most of the functions come from the lowfruit executable as indicated in the Load Object field.

You can also experiment with clicking on the column headers to change the sort from Exclusive Total CPU Time to Inclusive Total CPU Time, or by Name.

4. In the Functions view compare the two versions of the initialization task, `init_bad()` and `init_good()`.

You can see that the two functions have roughly the same Exclusive Total CPU Time but very different Inclusive times. The `init_bad()` function is slower due to time it spends in a callee. Both functions call the same callee, but they spend very different amounts of time in that routine. You can see why by examining the source of the two routines.

5. Select the function `init_good()` and then click the Source view or choose Views > Source from the menu bar.
6. Adjust the window to allow more space for the code: Collapse the Called-by/Calls panel by clicking the down arrow in the upper margin, and collapse the Selection Details panel by clicking the right-arrow in the side margin.

You should scroll up a little to see the source for both `init_bad()` and `init_good()`. The Source view should look similar to the following screen shot.

Incl. Total CPU (sec.)	Source File: lowfruit.c
0.	46. {
	47. <Function: init_bad>
	48. int i,j,k;
	49. volatile double x;
0.	50. for(i = 0; i < imax; i++) {
36.165	51. init_static_routine(); /* inside loop */
0.	52. for(j= 0; j < 50; j++) {
0.	53. x = 0.0;
0.690	54. for(k=0; k<1000000; k++) {
2.932	55. x = x + 1.0;
	56. }
	57. }
0.	58. }
	59. }
	60. }
	61. void
	62. init_good(int imax)
0.	63. {
	<Function: init_good>
	64. int i,j,k;
	65. volatile double x;
	66. }
3.623	67. init_static_routine(); /* outside loop */
0.	68. for(i = 0; i < imax; i++) {
0.	69. for(j= 0; j < 50; j++) {
0.	70. x = 0.0;
0.801	71. for(k=0; k<1000000; k++) {
2.812	72. x = x + 1.0;
	73. }
	74. }
	75. }
0.	76. }

Notice that the call to `init_static_routine()` is outside of the loop in `init_good()`, while `init_bad()` has the call to `init_static_routine()` inside the loop. The bad version takes about ten times longer (corresponding to the loop count) than in the good version.

This example is not as silly as it might appear. It is based on a real code that produces a table with an icon for each table row. While it is easy to see that the initialization should not be inside the loop in this example, in the real code the initialization was embedded in a library routine and was not obvious.

The toolkit that was used to implement that code had two library calls (APIs) available. The first API added an icon to a table row, and second API added a vector of icons to the entire table. While it is easier to code using the first API, each time an icon was added, the toolkit recomputed the height of all rows in order to set the correct value for the whole table. When the code used the alternative API to add all icons at once, the recomputation of height was done only once.

- Now go back to the Functions view and look at the two versions of the insert task, `insert_bad()` and `insert_good()`.

Note that the Exclusive Total CPU time is significant for `insert_bad()`, but negligible for `insert_good()`. The difference between Inclusive and Exclusive time for each version, representing the time in the function `insert_number()` called to insert each entry into the list, is the same. You can see why by examining the source.

- Select `insert_bad()` and switch to the Source view:

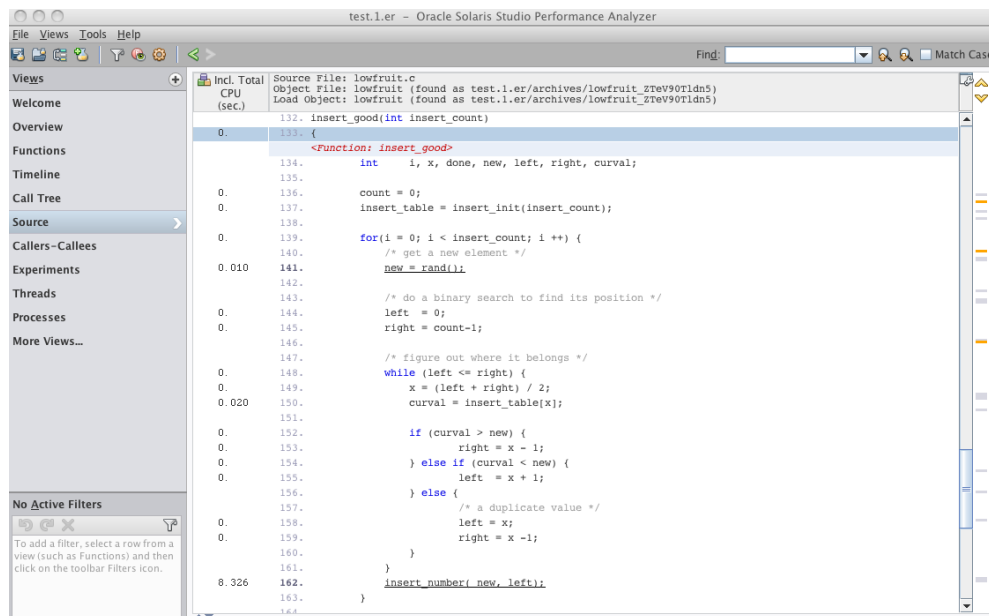
```

test.1.er - Oracle Solaris Studio Performance Analyzer
File Views Tools Help
Find:
Views
Welcome
Overview
Functions
Timeline
Call Tree
Source
Callers-Callees
Experiments
Threads
Processes
More Views...
No Active Filters
To add a filter, select a row from a view (such as Functions) and then click on the toolbar Filters icon.

Incl. Total CPU (sec.) Source File: lowfruit.c
Object File: lowfruit (found as test.1.er/archives/lowfruit_2TeV9071dn5)
Load Object: lowfruit (found as test.1.er/archives/lowfruit_2TeV9071dn5)
97.
98. void
99. insert_bad(int insert_count)
0. 100. {
    <Function: insert_bad>
    101.     int i, done, new;
    102.     count = 0;
    103.     insert_table = insert_init(insert_count);
    104.
    105.     for(i = 0; i < insert_count; i++) {
    106.         /* get a new element */
    107.         new = rand();
    108.         done = 0;
    1.011 109.         for (int j = 0; j < count; j++) {
    2.152 110.             if(new < insert_table[j]) {
    8.486 111.                 insert_number(new, j);
    112.                 done = 1;
    113.                 break;
    114.             }
    115.         }
    0.010 116.         if (done == 0) {
    117.             insert_number(new, count);
    118.         }
    119.     }
    120. }
  
```

Notice that the time, excluding the call to `insert_number()`, is spent in a loop looking with a linear search for the right place to insert the new number.

- Now scroll down to look at `insert_good()`.



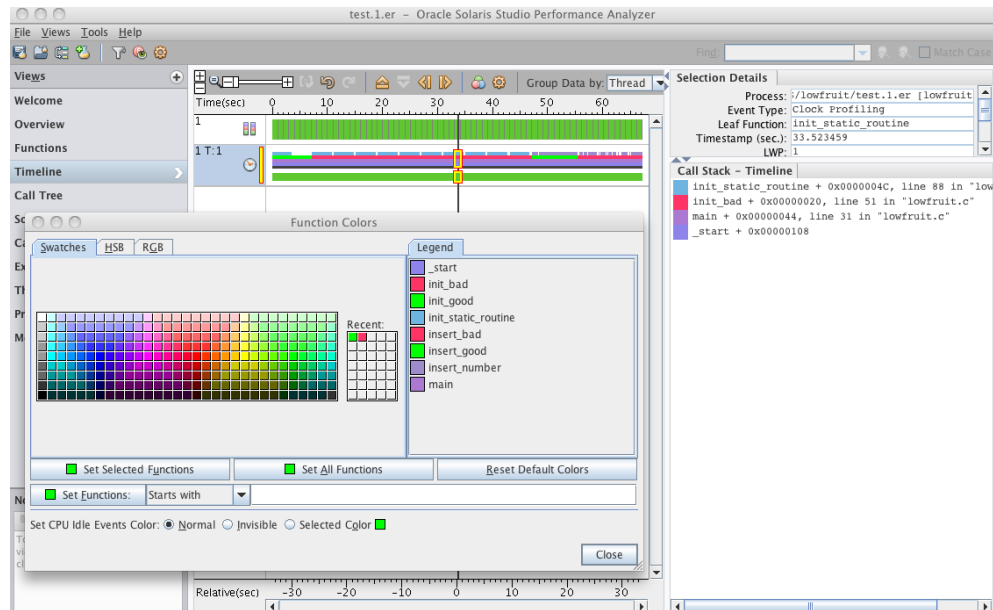
Note that the code is more complicated because it is doing a binary search to find the right place to insert, but the total time spent, excluding the call to `insert_number()`, is much less than in `insert_bad()`. This example illustrates that binary search can be more efficient than linear search.

You can also see the differences in the routines graphically in the Timeline view.

10. Click on the Timeline view or choose Views > Timeline from the menu bar.

The profiling data is recorded as a series of events, one for every tick of the profiling clock for every thread. The Timeline view shows each individual event with the callstack recorded in that event. The callstack is shown as a list of the frames in the callstack, with the leaf PC (the instruction next to execute at the instant of the event) at the top, and the call site calling it next, and so forth. For the main thread of the program, the top of the callstack is always `_start`.

11. In the Timeline tool bar, click the Call Stack Function Colors icon for coloring functions or choose Tools > Function Colors from the menu bar and see the dialog box as shown below.



The function colors were changed to distinguish the good and bad versions of the functions more clearly for the screen shot. The `init_bad()` and `insert_bad()` functions are both now red and the `init_good()` and `insert_good()` are both bright green.

12. To make your Timeline view look similar, do the following in the Function Colors dialog box:

- Scroll down the list of java methods in the Legend to find the `init_bad()` method.
- Select the `init_bad()` method, click on a red color square in Swatches, and click Set Selected Functions button.
- Select the `insert_bad()` method, click on a red color square in Swatches, and click Set Selected Functions button.
- Select the `init_good()` method, click on a green color square in Swatches, and click Set Selected Functions button.
- Select the `insert_good()` method, click on a green color square in Swatches, and click Set Selected Functions button.

13. Look at the top bar of the Timeline.

The top bar of the Timeline is the CPU Utilization Samples bar as you can see in the tool tip if you move your mouse cursor over the first column. Each segment of the CPU Utilization Samples bar represents a one-second interval showing the resource usage of the target during that second of execution.

In this example, all the segments are green because all the intervals were spent accumulating User CPU Time. The Selection Details window shows the mapping of colors to microstate although it is not visible in the screenshot.

14. Look at the second bar of the Timeline.

The second bar is the Clock Profiling Call Stacks bar, labeled "1 T:1" which means Process 1 and Thread 1, the only thread in the example. The Clock Profiling Call Stacks bar shows two bars of data for events occurring during program execution. The upper bar shows color-coded representations of the callstack and the lower bar shows the state of the thread at each event. The state in this example was always User CPU Time so it appears to be a solid green line.

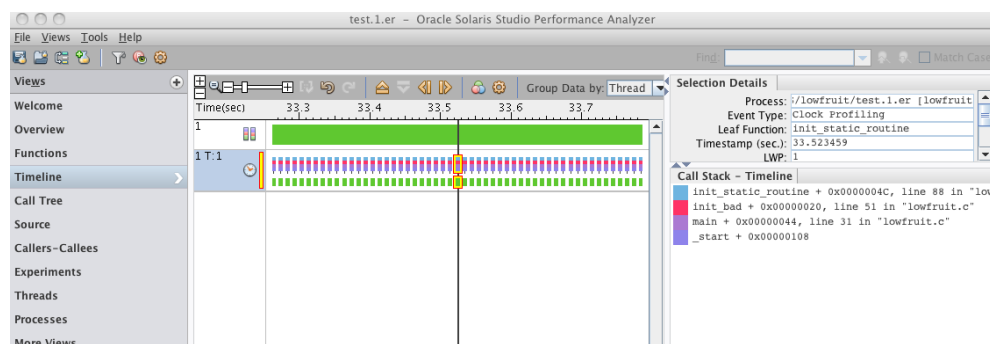
If you click anywhere within that Clock Profiling Call Stacks bar you select the nearest event and the details for that event are shown in the Selection Details window. From the pattern of the call stacks, you can see that the time in the `init_good()` and `insert_good()` routines shown in bright green in the screenshot is considerably shorter than the corresponding time in the `init_bad()` and `insert_bad()` routines shown in red.

15. Select events in the regions corresponding to the good and bad routines in the timeline and look at the call stacks in the Call Stack - Timeline window below the Selection Details window.

You can select any frame in the Call Stack window, and then select the Source view on the Views navigation bar, and go to the source for that source line. You can also double-click a frame in a call stack to go to the Source view or right-click the frame in the call stack and select from a popup menu.

16. Zoom in on the events by using the slider at the top of the Timeline, or using the + key, or by double-clicking with the mouse.

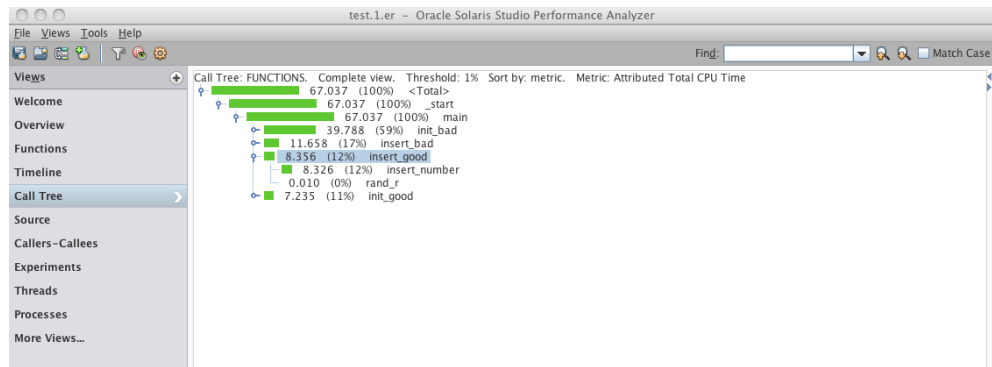
If you zoom in enough you can see that the data shown is not continuous but consists of discrete events, one for each profile tick, which is about 10 ms in this example.



Press the F1 key to see the Help for more information about the Timeline view.

17. Click on the Call Tree view or choose Views > Call Tree to see the structure of your program.

The Call Tree view shows a dynamic call graph of the program, annotated with performance information.



Performance Analyzer has many additional views of the data, such as the Caller-Callees view which enables you to navigate through the program structure, and the Experiments view which shows you details of the recorded experiment. For this simple example the Threads and Processes views are not very interesting.

By clicking on the + button on the Views list you can add other views to the navigation bar. If you are an assembly-language programmer, you might want to look at the Disassembly. Try exploring the other views.

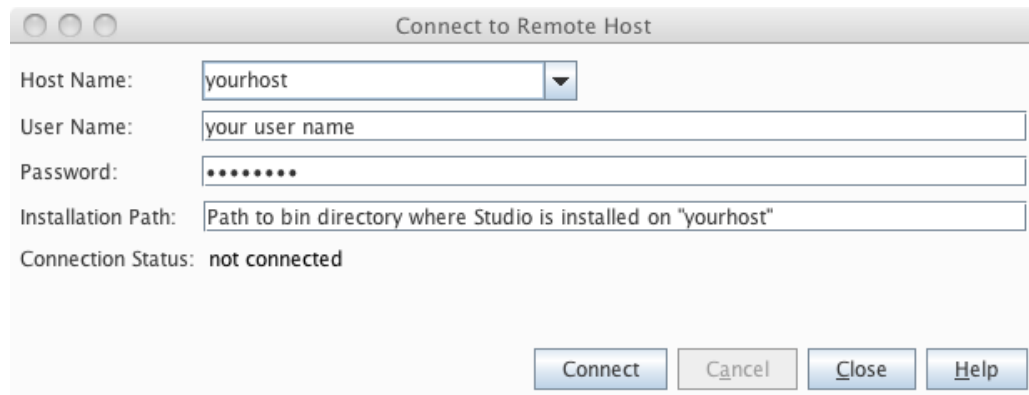
Performance Analyzer also has a very powerful filtering capability. You can filter by time, thread, function, source line, instruction, call stack-fragment, and any combination of them. The use of filtering is outside the scope of this tutorial, since the sample code is so simple that filtering is not needed.

Using the Remote Performance Analyzer

You can use the Remote Performance Analyzer either from a supported system, or from systems where Oracle Solaris Studio cannot be installed, such as Mac OS or Windows. See [“Using Performance Analyzer Remotely”](#) in [“Oracle Solaris Studio 12.4: Performance Analyzer”](#) for information about installing and using this special version of Performance Analyzer.

When you invoke Performance Analyzer remotely, you see the same Welcome page, but the options for creating and viewing experiments are disabled and grayed-out.

Click Connect to Remote Host and Performance Analyzer opens a connection dialog:



Host Name: yourhost

User Name: your user name

Password:

Installation Path: Path to bin directory where Studio is installed on "yourhost"

Connection Status: not connected

Connect Cancel Close Help

Type the name of the system to which you want to connect, your user name and password for that system, and the installation path to the Oracle Solaris Studio installation on that system. Click Connect and Performance Analyzer logs in to the remote system using your name and password, and verifies the connection.

From that point on, the Welcome page will look just as it does with the local Performance Analyzer, except the status area at the bottom shows the name of the remote host to which you connected. Proceed from there in step 2 above.

Introduction to Java Profiling

This chapter covers the following topics.

- [“About the Java Profiling Tutorial”](#) on page 27
- [“Setting Up the `lowfruit` Sample Code”](#) on page 28
- [“Using Performance Analyzer to Collect Data from `lowfruit`”](#) on page 28
- [“Using Performance Analyzer to Examine the `lowfruit` Data”](#) on page 32
- [“Using the Remote Performance Analyzer”](#) on page 41

About the Java Profiling Tutorial

This tutorial shows the simplest example of profiling with Oracle Solaris Studio Performance Analyzer and demonstrates how to use Performance Analyzer to collect and examine a performance experiment. You use the Overview, Functions view, Source view, Timeline view, and Call Tree view in this tutorial.

The program `lowfruit` is a simple program that executes two different tasks, one for initializing in a loop and one for inserting numbers into an ordered list. Each task is performed twice, in an inefficient way and in a more efficient way.

Tip - The [“Introduction to C Profiling”](#) tutorial uses an equivalent C program and shows similar activities with Performance Analyzer.

The data you see in the experiment that you record will be different from that shown here. The experiment used for the screen-shots in the tutorial was recorded on a SPARC T5 system running Oracle Solaris 11.2. The data from an x86 system running Oracle Solaris or Linux will be different. Furthermore, data collection is statistical in nature and varies from experiment to experiment, even when run on the same system and OS.

The Performance Analyzer window configuration that you see might not precisely match the screen shots. Performance Analyzer enables you to drag separator bars between components of the window, collapse components, and resize the window. Performance Analyzer records its configuration and uses the same configuration the next time it runs. Many configuration changes were made in the course of capturing the screen shots shown in the tutorial.

This tutorial is run locally on a system where Oracle Solaris Studio is installed. You can also run remotely as described in [“Using the Remote Performance Analyzer” on page 24](#).

Setting Up the jlowfruit Sample Code

Before You Begin

See the following for information about obtaining the code and setting up your environment.

- [“Getting the Sample Code for the Tutorials” on page 8](#)
 - [“Setting Up Your Environment for the Tutorials” on page 9](#)
1. Copy the contents of the jlowfruit directory to your own private working area with the following command:

```
% cp -r SolarisStudioSampleApplications/PerformanceAnalyzer/jlowfruit mydirectory
```

where *mydirectory* is the working directory you are using.

2. Change to that working directory copy.

```
% cd mydirectory/jlowfruit
```

3. Build the target executable.

```
% make clobber (needed only if you ran make in the directory before, but safe in any case)
```

```
% make
```

After you run make the directory contains the target application to be used in the tutorial, a Java class file named `jlowfruit.class`.

The next section shows how to use Performance Analyzer to collect data from the jlowfruit program and create an experiment.

Using Performance Analyzer to Collect Data from jlowfruit

This section describes how to use the Profile Application feature of Performance Analyzer to collect data in an experiment on a Java application.

Tip - If you prefer not to follow these steps to see how to profile applications from Performance Analyzer, you can record an experiment with a `make` target included in the `Makefile` for `jlowfruit`:

```
% make collect
```

The `collect` target launches a `collect` command and records an experiment just like the one that you create using Performance Analyzer in this section. You could then skip to [“Using Performance Analyzer to Examine the `jlowfruit` Data” on page 32.](#)

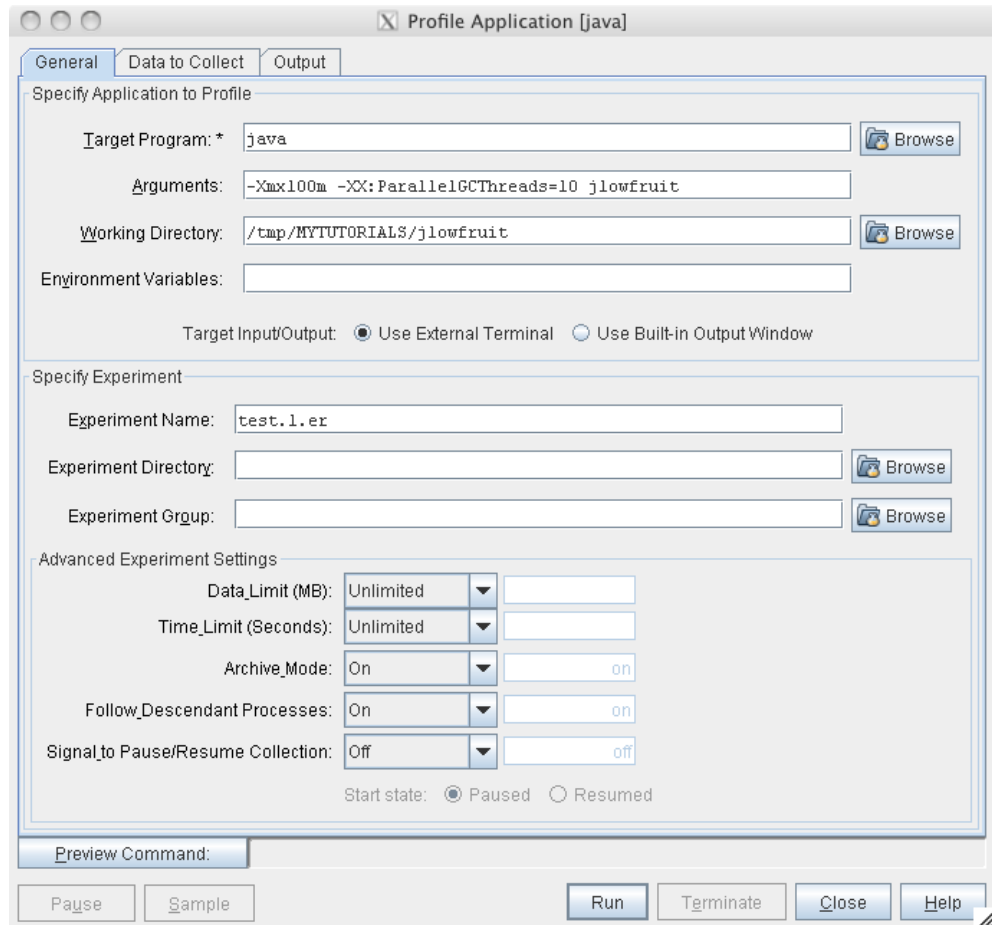
1. While still in the `jlowfruit` directory start Performance Analyzer with the target `java` and its arguments:

```
% analyzer java -Xmx100m -XX:ParallelGCThreads=10 jlowfruit
```

The Profile Application dialog box opens with the General tab selected and several options already filled out using information you provided with the `analyzer` command.

Target Program is set to `java` and Arguments is set to

```
-Xmx100m -XX:ParallelGCThreads=10 jlowfruit
```



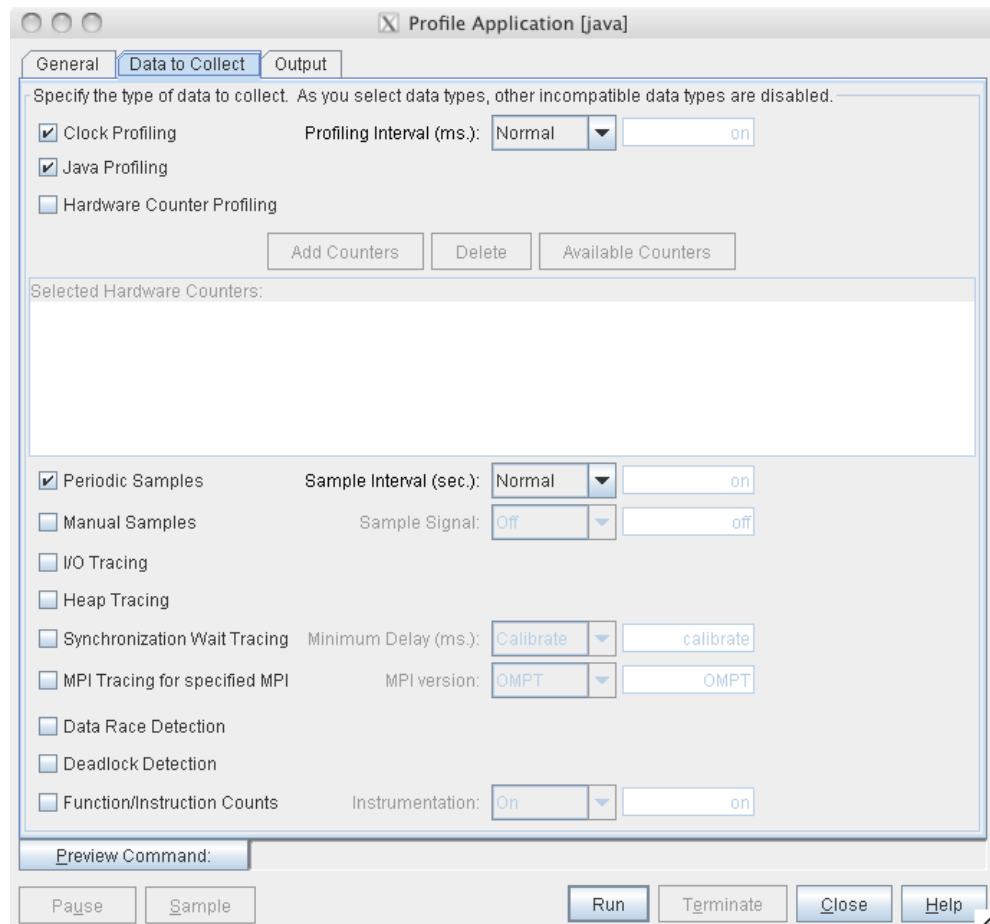
2. For the Target Input/Output option, select Use Built-in Output Window.

Target Input/Output option specifies the window to which the target program `stdout` and `stderr` will be redirected. The default value is Use External Terminal, but in this tutorial you should change the Target Input/Output option to Use Built-in Output Window to keep all the activity in the Performance Analyzer window. With this option the `stdout` and `stderr` is shown in the Output tab in the Profile Application dialog box.

If you are running remotely, the Target Input/Output option is absent because only the built-in output window is supported.

3. For the Experiment Name option, the default experiment name is `test.1.er` but you can change it to a different name as long as the name ends in `.er`, and is not already in use.
4. Click the Data to Collect tab.

The Data to Collect tab enables you to select the type of data to collect, and shows the defaults already selected.



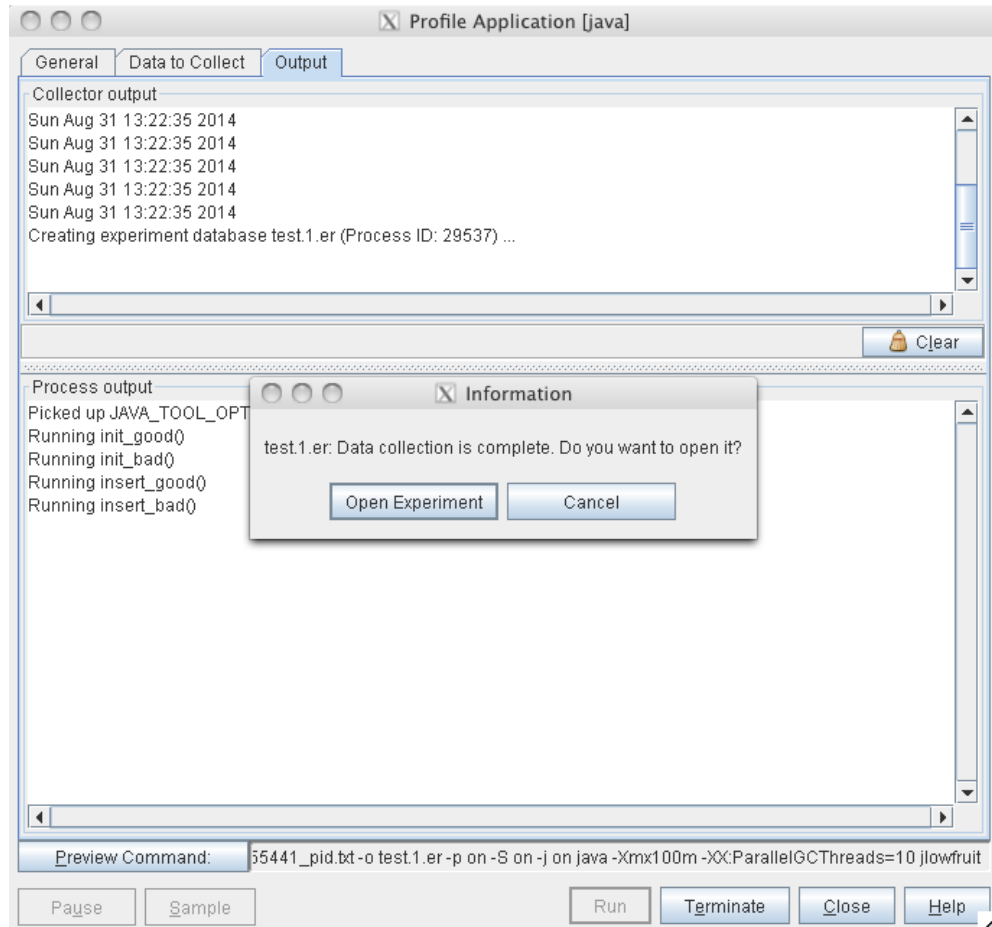
Java profiling is enabled by default as you can see in the screen shot.

You can optionally click the Preview Command button and see the `collect` command that will be run when you start profiling.

5. Click the Run button.

The Profile Application dialog box displays the Output tab and shows the program output in the Process Output panel as the program runs.

After the program completes, a dialog box asks if you want to open the experiment just recorded.



6. Click Open Experiment in the dialog box.

The experiment opens. The next section shows how to examine the data.

Using Performance Analyzer to Examine the jlowfruit Data

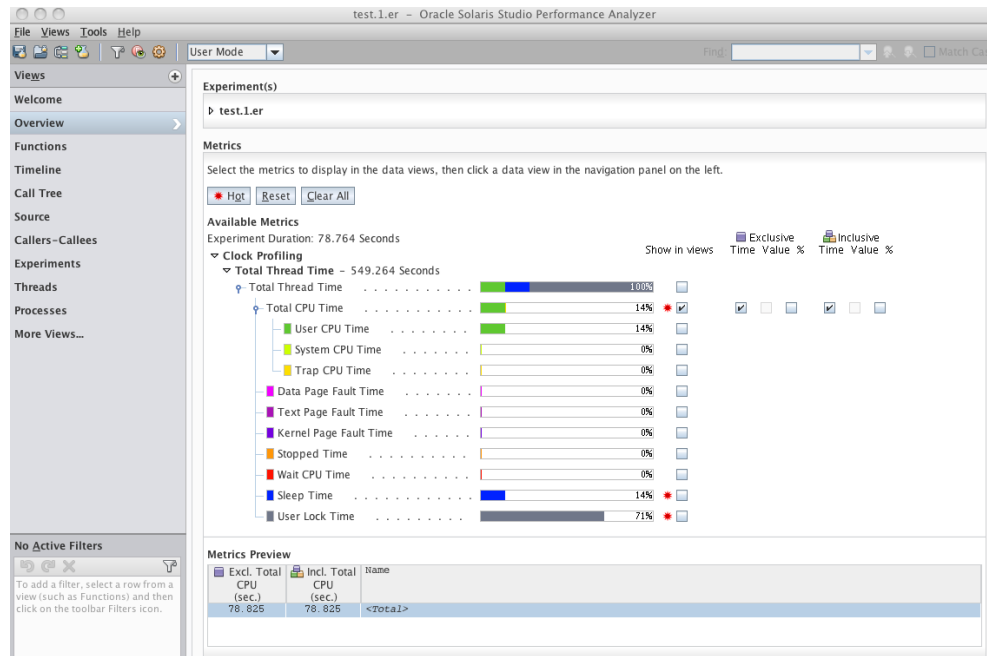
This section shows how to explore the data in the experiment created from the jlowfruit sample code.

1. If the experiment you created in the previous section is not already open, you can start Performance Analyzer from the jlowfruit directory and load the experiment as follows:

```
% analyzer test.1.er
```


When the experiment opens, Performance Analyzer shows the Overview page.

2. Notice the Overview page shows a summary of the metric values and enables you to select metrics.



In this experiment the Overview shows about 14% Total CPU Time which was all User CPU Time, plus about 14% Sleep Time and 71% User Lock Time. The user Java code `jlowfruit` is single-threaded and that one thread is CPU-bound, but all Java programs use multiple threads including a number of system threads. The number of those threads depends on the choice of JVM options, including the Garbage Collector parameters and the size of the machine on which the program was run.

The experiment was recorded on a Oracle Solaris system, and the Overview shows twelve metrics recorded but only Total CPU Time is enabled by default.

The metrics with colored indicators are the times spent in the ten microstates defined by Oracle Solaris. These metrics include User CPU Time, System CPU Time, and Trap CPU Time which together are equal to Total CPU Time, as well as various wait times. Total Thread Time is the sum over all of the microstates.

On a Linux machine, only Total CPU Time is recorded because Linux does not support microstate accounting.

By default, both Inclusive and Exclusive Total CPU Time are selected. *Inclusive* for any metric refers to the metric value in that function or method, including metrics accumulated

in all the functions or methods that it calls. *Exclusive* refers only to the metric accumulated within that function or method.

3. Click the Hot button to select metrics with high values to show them in the data views.

The check boxes next to Sleep Time and User Lock Time are now selected.

The Metrics Preview panel at the bottom is updated to show you how the metrics will be displayed in the data views that present table-formatted data. You will next look to see which threads are responsible for which metrics.

4. Now switch to the Threads view by clicking its name in the Views navigation panel or choosing Views > Threads from the menu bar.

	Excl. Total CPU (sec)	Excl. Sleep (sec)	Excl. User Lock (sec)	Name
	78.825	78.745	391.684	<Total>
	78.715	0.010	0.	Process 1, Thread 2, JThread 3 'main'
	0.040	0.	78.585	Process 1, Thread 20
	0.020	78.735	0.	Process 1, Thread 1
	0.020	0.	78.665	Process 1, Thread 13
	0.020	0.	51.916	Process 1, Thread 18
	0.010	0.	51.916	Process 1, Thread 17
	0.	0.	0.040	Process 1, Thread 15, JThread 1 'Sig'
	0.	0.	78.625	Process 1, Thread 16, JThread 2 'Sig'
	0.	0.	51.936	Process 1, Thread 19

	Exclusive
Total Thread Time:	78.735 (14.33%)
Total CPU Time:	78.715 (99.86%)
User CPU Time:	78.655 (99.91%)
System CPU Time:	0.060 (66.67%)
Trap CPU Time:	0. (0. %)
Data Page Fault Time:	0. (0. %)
Text Page Fault Time:	0. (0. %)
Kernel Page Fault Time:	0. (0. %)
Stopped Time:	0. (0. %)
Wait CPU Time:	0.010 (100.00%)
Sleep Time:	0.010 (0.01%)
User Lock Time:	0. (0. %)

The thread with almost all of the Total CPU Time is Thread 2, which is the only user Java thread in this simple application.

Thread 15 is most likely a user thread even though it is actually created internally by the JVM. It is only active during startup and has very little time accumulated. In your experiment, a second thread similar to thread 15 might be created.

Thread 1 spends its entire time sleeping.

The remaining threads spend their time waiting for a lock, which is how the JVM synchronizes itself internally. Those threads include those used for HotSpot compilation and for Garbage Collection. This tutorial does not explore the behavior of the JVM system, but that is explored in another tutorial, “[Java and Mixed Java-C++ Profiling](#)”.

5. Return to the Overview screen and deselect the check boxes for Sleep Time and User Lock Time.
6. Click on the Functions view in the Views navigation panel, or choose Views > Functions from the menu bar.

The screenshot displays the Oracle Solaris Studio Performance Analyzer interface. The main window is titled "test.1.er - Oracle Solaris Studio Performance Analyzer". The "Views" pane on the left shows the "Functions" view selected. The main area contains a table of functions with columns for "Excl. Total CPU (sec.)", "Incl. Total CPU (sec.)", and "Name". The function "jlowfruit.init_static_routine()" is selected and highlighted in blue. Below the main table, the "Called-by / Calls" panel shows two sub-tables: "jlowfruit.init_static_routine() is called by" and "jlowfruit.init_static_routine() calls". The "is called by" table lists "jlowfruit.init_bad(int)" and "jlowfruit.init_good(int)". The "calls" table is currently empty. On the right side, the "Selection Details" window is open, showing metadata for the selected function, including its name, PC address, size, source file, object file, load object, and mangled name. Below this, a table of performance metrics is shown, with columns for "Exclusive" and "Inclusive" values for various metrics like Total Thread Time, Total CPU Time, User CPU Time, System CPU Time, Trap CPU Time, Data Page Fault Time, Text Page Fault Time, Kernel Page Fault Time, Stopped Time, Wait CPU Time, Sleep Time, and User Lock Time.

	Exclusive	Inclusive
Total Thread Time:	33.774 (6.15%)	33.774 (6.15%)
Total CPU Time:	33.774 (42.85%)	33.774 (42.85%)
User CPU Time:	33.774 (42.90%)	33.774 (42.90%)
System CPU Time:	0. (0. %)	0. (0. %)
Trap CPU Time:	0. (0. %)	0. (0. %)
Data Page Fault Time:	0. (0. %)	0. (0. %)
Text Page Fault Time:	0. (0. %)	0. (0. %)
Kernel Page Fault Time:	0. (0. %)	0. (0. %)
Stopped Time:	0. (0. %)	0. (0. %)
Wait CPU Time:	0. (0. %)	0. (0. %)
Sleep Time:	0. (0. %)	0. (0. %)
User Lock Time:	0. (0. %)	0. (0. %)

The Functions view shows the list of functions in the application, with performance metrics for each function. The list is initially sorted by the Exclusive Total CPU Time spent in each function. There are also a number of functions from the JVM in the Functions view, but they have relatively low metrics. The list includes all functions from the target application and any shared objects the program uses. The top-most function, the most expensive one, is selected by default.

The Selection Details window on the right shows all the recorded metrics for the selected function.

The Called-by/Calls panel below the functions list provides more information about the selected function and is split into two lists. The Called-by list shows the callers of the selected function and the metric values show the attribution of the total metric for the function to its callers. The Calls list shows the callees of the selected function and shows how the Inclusive metric of its callees contributed to the total metric of the selected function. If you double-click a function in either list in the Called-by/Calls panel, the function becomes the selected function in the main Functions view.

- Experiment with selecting the various functions to see how the windows in the Functions view update with the changing selection.

The Selection Details window shows you that most of the functions come from the jlowfruit executable as indicated in the Load Object field.

You can also experiment with clicking on the column headers to change the sort from Exclusive Total CPU Time to Inclusive Total CPU Time, or by Name.

8. In the Functions view compare the two versions of the initialization task, `jlowfruit.init_bad()` and `jlowfruit.init_good()`.

You can see that the two functions have roughly the same Exclusive Total CPU Time but very different Inclusive times. The `jlowfruit.init_bad()` function is slower due to time it spends in a callee. Both functions call the same callee, but they spend very different amounts of time in that routine. You can see why by examining the source of the two routines.

9. Select the function `jlowfruit.init_good()` and then click the Source view or choose Views > Source from the menu bar.
10. Adjust the window to allow more space for the code: Collapse the Called-by/Calls panel by clicking the down arrow in the upper margin, and collapse the Selection Details panel by clicking the right arrow in the side margin.

You should scroll up a little to see the source for both `jlowfruit.init_bad()` and `jlowfruit.init_good()`. The Source view should look similar to the following screen shot.

Function	Incl. Total CPU (sec.)
<code>jlowfruit.init_bad(int)</code>	30.701
<code>jlowfruit.init_good(int)</code>	3.072

```

41.
42.     for(i = 0; i < imax; i++) {
43.         init_static_routine(); /* inside loop */
44.         for(j= 0; j < 50; j++) {
45.             x = 0.0;
46.             for(k=0; k<1000000; k++) {
47.                 x = x + 1.0;
48.             }
49.         }
50.     }
51.     return x;
52. }
53.
54. double
55. init_good(int imax)
56. {
57.     int i,j,k;
58.     double x = 0;
59.
60.     init_static_routine(); /* outside loop */
61.     for(i = 0; i < imax; i++) {
62.         for(j= 0; j < 50; j++) {
63.             x = 0.0;
64.             for(k=0; k<1000000; k++) {
65.                 x = x + 1.0;
66.             }
67.         }
68.     }
69.     return x;
70. }
71.
72. double
73. init_static_routine()
74. {

```

Notice that the call to `jlowfruit.init_static_routine()` is outside of the loop in `jlowfruit.init_good()`, while `jlowfruit.init_bad()` has the call to `jlowfruit.init_static_routine()` inside the loop. The bad version takes about ten times longer (corresponding to the loop count) than in the good version.

This example is not as silly as it might appear. It is based on a real code that produces a table with an icon for each table row. While it is easy to see that the initialization should not be inside the loop in this example, in the real code the initialization was embedded in a library routine and was not obvious.

The toolkit that was used to implement that code had two library calls (APIs) available. The first API added an icon to a table row, and second API added a vector of icons to the entire table. While it is easier to code using the first API, each time an icon was added, the toolkit recomputed the height of all rows in order to set the correct value for the whole table. When the code used the alternative API to add all icons at once, the recomputation of height was done only once.

- Now go back to the Functions view and look at the two versions of the insert task, `jlowfruit.insert_bad()` and `jlowfruit.insert_good()`.

Note that the Exclusive Total CPU time is significant for `jlowfruit.insert_bad()`, but negligible for `jlowfruit.insert_good()`. The difference between Inclusive and Exclusive time for each version, representing the time in the function `jlowfruit.insert_number()` called to insert each entry into the list, is the same. You can see why by examining the source.

- Select `jlowfruit.insert_bad()` and switch to the Source view:

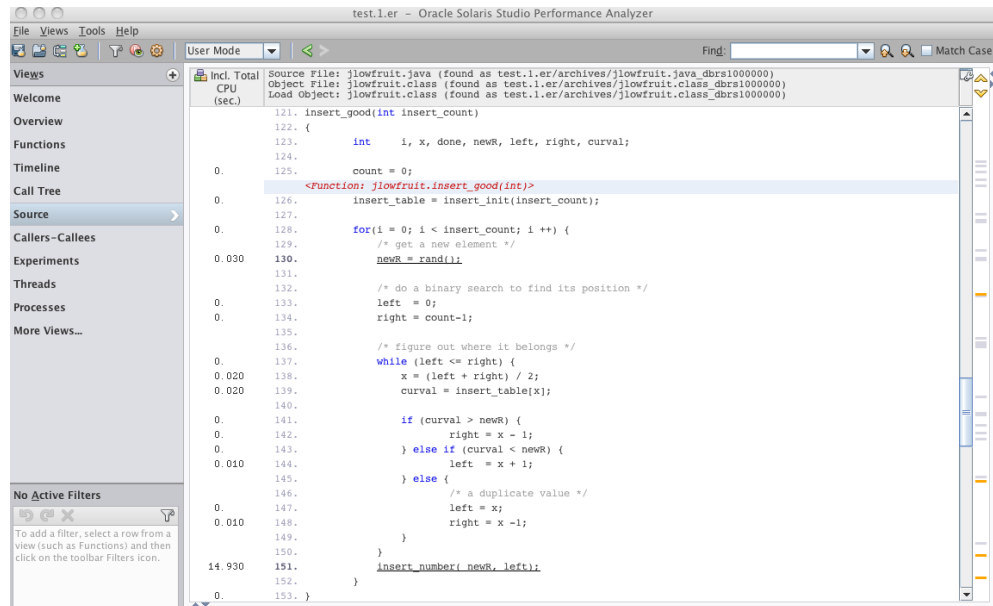
```

test.1.er - Oracle Solaris Studio Performance Analyzer
File Views Tools Help
User Mode
Find:
Match Case
Views
Welcome
Overview
Functions
Timeline
Call Tree
Source
Callers-Callees
Experiments
Threads
Processes
More Views...
No Active Filters
To add a filter, select a row from a view (such as Functions) and then click on the toolbar Filters icon.
Incl. Total CPU (sec.) Source File: jlowfruit.java (found as test.1.er/archives/jlowfruit.java_dbrs1000000)
Object File: jlowfruit.class (found as test.1.er/archives/jlowfruit.class_dbrs1000000)
Load Object: jlowfruit.class (found as test.1.er/archives/jlowfruit.class_dbrs1000000)
93. void
94. insert_bad(int insert_count)
95. {
96.     int i, done, newR;
0. 97.     count = 0;
98.     <Function: jlowfruit.insert_bad(int)>
99.     insert_table = insert_init(insert_count);
0. 100.     for(i = 0; i < insert_count; i++) {
101.         /* get a new element */
0.020 102.         newR = rand(i);
103.         done = 0;
1.651 104.         for (int j = 0; j < count; j++) {
10.087 105.             if(newR < insert_table[j]) {
14.940 106.                 insert_number(newR,j);
107.                 done = 1;
108.                 break;
109.             }
110.         }
0. 111.         if (done == 0) {
0. 112.             insert_number(newR, count);
113.         }
114.     }
115.
116.
117.     free(insert_table);
0. 118. }
119.

```

Notice that the time, excluding the call to `jlowfruit.insert_number()`, is spent in a loop looking with a linear search for the right place to insert the new number.

- Now scroll down to look at `jlowfruit.insert_good()`.



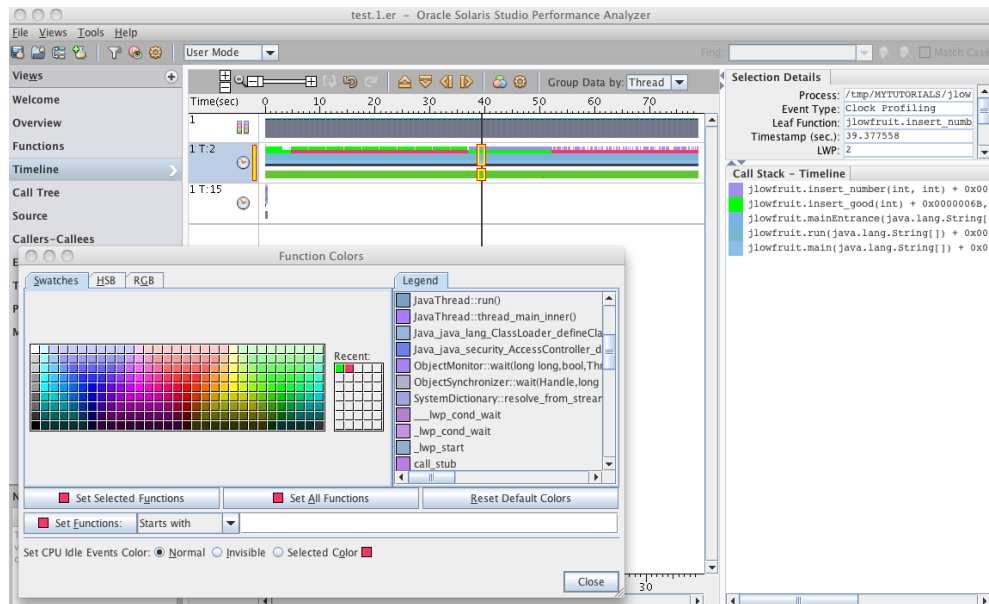
Note that the code is more complicated because it is doing a binary search to find the right place to insert, but the total time spent, excluding the call to `jlowfruit.insert_number()`, is much less than in `jlowfruit.insert_bad()`. This example illustrates that binary search can be more efficient than linear search.

You can also see the differences in the routines graphically in the Timeline view.

14. Click on the Timeline view or choose Views > Timeline from the menu bar.

The profiling data is recorded as a series of events, one for every tick of the profiling clock for every thread. The Timeline view shows each individual event with the call stack recorded in that event. The call stack is shown as a list of the frames in the callstack, with the leaf PC (the instruction next to execute at the instant of the event) at the top, and the call site calling it next, and so forth. For the main thread of the program, the top of the callstack is always `_start`.

15. In the Timeline tool bar, click the Call Stack Function Colors icon for coloring functions or choose Tools > Function Colors from the menu bar and see the dialog box as shown below.



The function colors were changed to distinguish the good and bad versions of the functions more clearly for the screen shot. The `jlowfruit.init_bad()` and `jlowfruit.insert_bad()` functions are both now red and the `jlowfruit.init_good()` and `jlowfruit.insert_good()` are both bright green.

16. To make your Timeline view look similar, do the following in the Function Colors dialog box:
 - Scroll down the list of java methods in the Legend to find the `jlowfruit.init_bad()` method.
 - Select the `jlowfruit.init_bad()` method, click on a red color square in Swatches, and click Set Selected Functions button.
 - Select the `jlowfruit.insert_bad()` method, click on a red color square in Swatches, and click Set Selected Functions button.
 - Select the `jlowfruit.init_good()` method, click on a green color square in Swatches, and click Set Selected Functions button.
 - Select the `jlowfruit.insert_good()` method, click on a green color square in Swatches, and click Set Selected Functions button.
17. Look at the top bar of the Timeline.

The top bar of the Timeline is the CPU Utilization Samples bar as you can see in the tool tip if you move your mouse cursor over the first column. Each segment of the CPU Utilization

Samples bar represents a one-second interval showing the resource usage of the target during that second of execution.

In this example, the segments are mostly gray with some green, reflecting the fact that only a small fraction of the Total Time was spent accumulating User CPU Time. The Selection Details window shows the mapping of colors to microstate although it is not visible in the screenshot.

18. Look at the second bar of the Timeline.

The second bar is the Clock Profiling Call Stacks bar, labeled "1 T:2" which means Process 1 and Thread 2, the main user thread in the example. The Clock Profiling Call Stacks bar shows two bars of data for events occurring during program execution. The upper bar shows color-coded representations of the callstack and the lower bar shows the state of the thread at each event. The state in this example was always User CPU Time so it appears to be a solid green line.

You should see one or two additional bars labeled with different thread numbers but they will only have a few events at the beginning of the run.

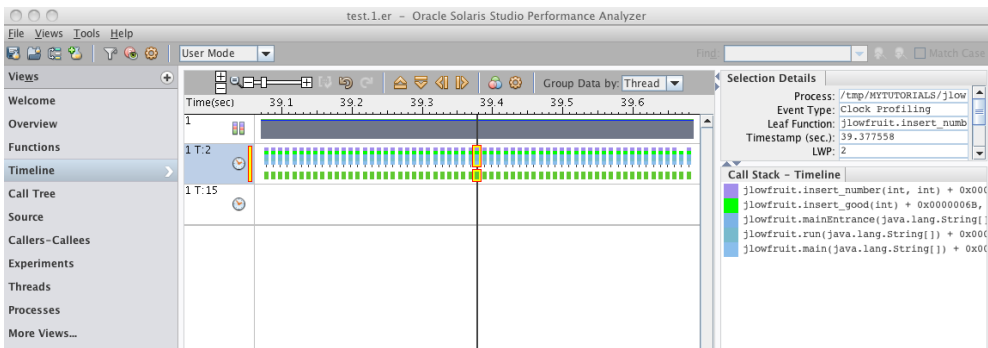
If you click anywhere within that Clock Profiling Call Stacks bar you select the nearest event and the details for that event are shown in the Selection Details window. From the pattern of the call stacks, you can see that the time in the `jlowfruit.init_good()` and `jlowfruit.insert_good()` routines shown in bright green in the screenshot is considerably shorter than the corresponding time in the `jlowfruit.init_bad()` and `jlowfruit.insert_bad()` routines shown in red.

19. Select events in the regions corresponding to the good and bad routines in the timeline and look at the call stacks in the Call Stack - Timeline window below the Selection Details window.

You can select any frame in the Call Stack window, and then select the Source view on the Views navigation bar, and go to the source for that source line. You can also double-click a frame in a call stack to go to the Source view or right-click the frame in the call stack and select from a popup menu.

20. Zoom in on the events by using the slider at the top of the Timeline, or using the + key, or by double-clicking with the mouse.

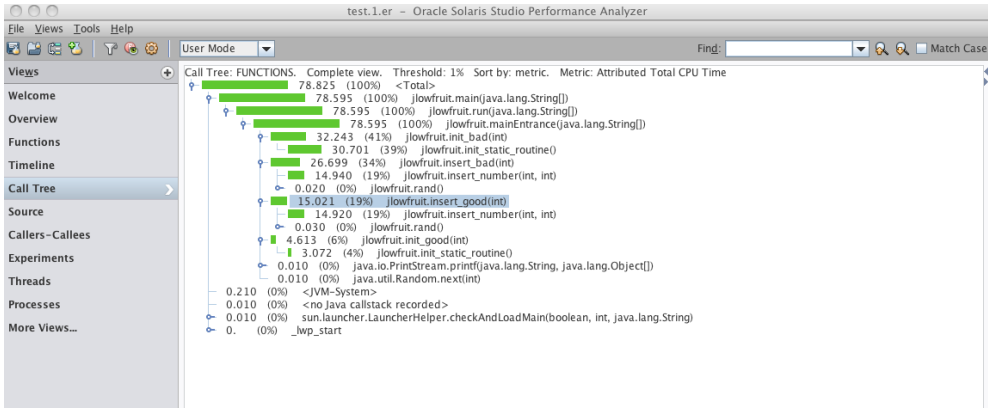
If you zoom in enough you can see that the data shown is not continuous but consists of discrete events, one for each profile tick, which is about 10 ms in this example.



Press the F1 key to see the Help for more information about the Timeline view.

- 21. Click on the Call Tree view or choose Views > Call Tree to see the structure of your program.

The Call Tree view shows a dynamic call graph of the program, annotated with performance information.

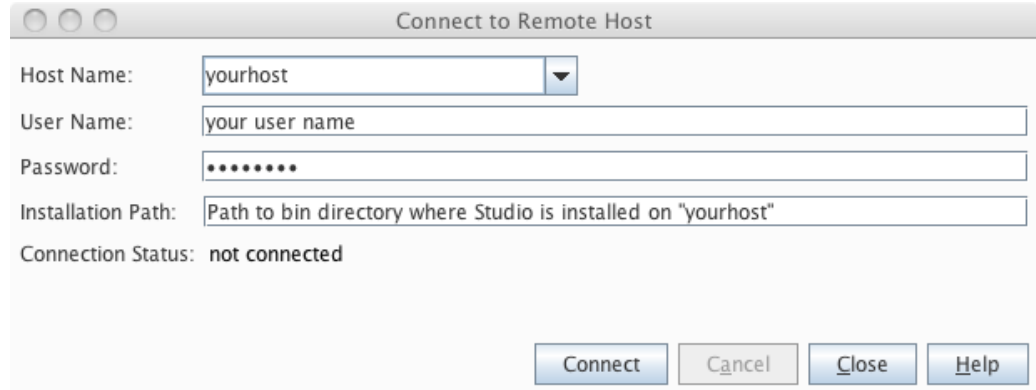


Using the Remote Performance Analyzer

You can use the Remote Performance Analyzer either from a supported system, or from systems where Oracle Solaris Studio cannot be installed, such as Mac OS or Windows. See [“Using Performance Analyzer Remotely”](#) in [“Oracle Solaris Studio 12.4: Performance Analyzer”](#) for information about installing and using this special version of Performance Analyzer.

When you invoke Performance Analyzer remotely, you see the same Welcome page, but the options for creating and viewing experiments are disabled and grayed-out.

Click Connect to Remote Host and Performance Analyzer opens a connection dialog:



Host Name: yourhost

User Name: your user name

Password: *****

Installation Path: Path to bin directory where Studio is installed on "yourhost"

Connection Status: not connected

Connect Cancel Close Help

Type the name of the system to which you want to connect, your user name and password for that system, and the installation path to the Oracle Solaris Studio installation on that system. Click Connect and Performance Analyzer logs in to the remote system using your name and password, and verifies the connection.

From that point on, the Welcome page will look just as it does with the local Performance Analyzer, except the status area at the bottom shows the name of the remote host to which you connected. Proceed from there in step 2 above.

Java and Mixed Java-C++ Profiling

This chapter covers the following topics.

- [“About the Java-C++ Profiling Tutorial” on page 43](#)
- [“Setting Up the jsynprog Sample Code” on page 44](#)
- [“Collecting the Data From jsynprog” on page 45](#)
- [“Examining the jsynprog Data” on page 46](#)
- [“Examining Mixed Java and C++ Code” on page 48](#)
- [“Understanding the JVM Behavior” on page 52](#)
- [“Understanding the Java Garbage Collector Behavior” on page 56](#)
- [“Understanding the Java HotSpot Compiler Behavior” on page 61](#)

About the Java-C++ Profiling Tutorial

This tutorial demonstrates the features of the Oracle Solaris Studio Performance Analyzer for Java profiling. It shows you how to use a sample code to do the following in Performance Analyzer:

- Examine the performance data in various data views including the Overview page, and the Threads, Functions, and Timeline views.
- Look at the Source and Disassembly for both Java code and C++ code.
- Learn the difference between User Mode, Expert Mode, and Machine Mode.
- Drill down into the behavior of the JVM executing the program and see the generated native code for any HotSpot-compiled methods.
- See how the garbage collector can be invoked by user code and how the HotSpot compiler is triggered.

`jsynprog` is a Java program that has a number of subtasks typical of Java programs. The program also loads a C++ shared object and calls various routines from it to show the seamless transition from Java code to native code from a dynamically loaded C++ library, and back again.

`jsynprog.main` is the main method that calls functions from different classes. It uses `gethrtime` and `gethrvtime` through Java Native Interface (JNI) calls to time its own behavior, and writes an accounting file with its own timings, as well as writing messages to `stdout`.

`jsynprog.main` has many methods:

- `Routine.memalloc` does memory allocation, and triggers garbage collection
- `Routine.add_int` does integer addition
- `Routine.add_double` does double (floating point) additions
- `Sub_Routine.add_int` is a derived calls that overrides `Routine.add_int`
- `Routine.has_inner_class` defines an inner class and uses it
- `Routine.recurse` shows direct recursion
- `Routine.recursedeep` does a deep recursion, to show how the tools deal with a truncated stack
- `Routine.bounce` shows indirect recursion, where `bounce` calls `bounce_b` which in turn calls back into `bounce`
- `Routine.array_op` does array operations
- `Routine.vector_op` does vector operations
- `Routine.sys_op` uses methods from the `System` class
- `jsynprog.jni_JavaJavaC`: Java method calls another Java method that calls a C function
- `jsynprog.JavaCJava`: Java method calls a C function which in turn calls a Java method
- `jsynprog.JavaCC`: Java calls a C function that calls another C function

Some of those methods are called from others, so they do not all represent the top-level tasks.

The data you see in the experiment that you record will be different from that shown here. The experiment used for the screen-shots in the tutorial was recorded on a SPARC T5 system running Oracle Solaris 11.2. The data from an x86 system running Oracle Solaris or Linux will be different. Furthermore, data collection is statistical in nature and varies from experiment to experiment, even when run on the same system and OS.

The Performance Analyzer window configuration that you see might not precisely match the screen shots. Performance Analyzer enables you to drag separator bars between components of the window, collapse components, and resize the window. Performance Analyzer records its configuration and uses the same configuration the next time it runs. Many configuration changes were made in the course of capturing the screen shots shown in the tutorial.

Setting Up the jsynprog Sample Code

Before You Begin

See the following for information about obtaining the code and setting up your environment.

- [“Getting the Sample Code for the Tutorials” on page 8](#)

- [“Setting Up Your Environment for the Tutorials” on page 9](#)

You might want to go through the introductory tutorial in [“Introduction to Java Profiling”](#) first to become familiar with Performance Analyzer.

1. Copy the contents of the `jsynprog` directory to your own private working area with the following command:

```
% cp -r SolarisStudioSampleApplications/PerformanceAnalyzer/jsynprog mydirectory
```

where `mydirectory` is the working directory you are using.

2. Change to that working directory copy.

```
% cd mydirectory/jsynprog
```

3. Build the target executable.

```
% make clobber (needed only if you ran make in the directory before, but safe in any case)
```

```
% make
```

After you run `make` the directory contains the target application to be used in the tutorial, a Java class file named `jsynprog.class` and a shared object named `libcloop.so` which contains C++ code that will be dynamically loaded and invoked from the Java program.

Tip - If you prefer, you can edit the `Makefile` to do the following: use the GNU compilers rather than the default of the Studio compilers; build in 32-bits rather than the default of 64-bits; and add different compiler flags.

Collecting the Data From jsynprog

The easiest way to collect the data is to run the following command in the `jsynprog` directory:

```
% make collect
```

The `collect` target of the `Makefile` launches a `collect` command and records an experiment. By default, the experiment is named `test.1.er`.

The `collect` target specifies options `-J "-Xmx100m -XX:ParallelGCThreads=10"` for the JVM and collects clock-profiling data by default.

Alternatively, you can use the Performance Analyzer's Profile Application dialog to record the data. Follow the procedure [“Using Performance Analyzer to Collect Data from jlowfruit” on page 28](#) in the introductory Java tutorial and specify `jsynprog` instead of `jlowfruit` in the Arguments field.

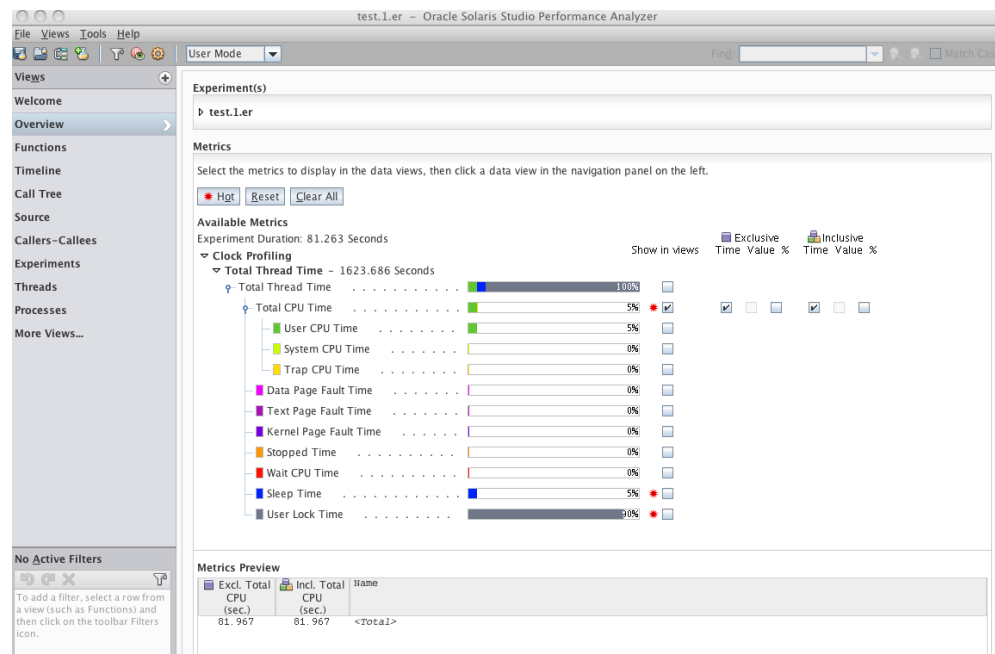
Examining the jsynprog Data

This procedure assumes you have already created an experiment as described in the previous section.

1. Start Performance Analyzer from the jsynprog directory and load the experiment as follows, specifying your experiment name if it is not called test.1.er.

```
% analyzer test.1.er
```


When the experiment opens, Performance Analyzer shows the Overview page.



Notice that the tool bar of Performance Analyzer now has a view mode selector that is initially set to User Mode, showing the user model of the program.

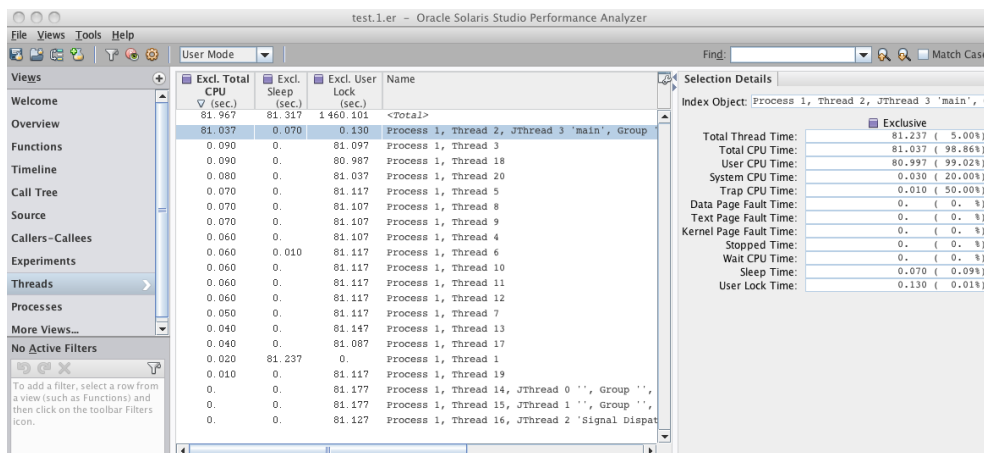
The Overview shows that the experiment ran about 81 seconds but used more than 1600 seconds of total time, implying that on average there were 20 threads in the process.

2. Select the check boxes for the Sleep Time and User Lock Time metrics to add them to the data views.



Notice that the Metrics Preview updates to show you how the data views will look with these metrics added.

3. Select the Threads view in the navigation panel and you will see the data for the threads:



Only Thread 2 accumulated significant Total CPU time. The other threads each had only a few profile events for Total CPU time.

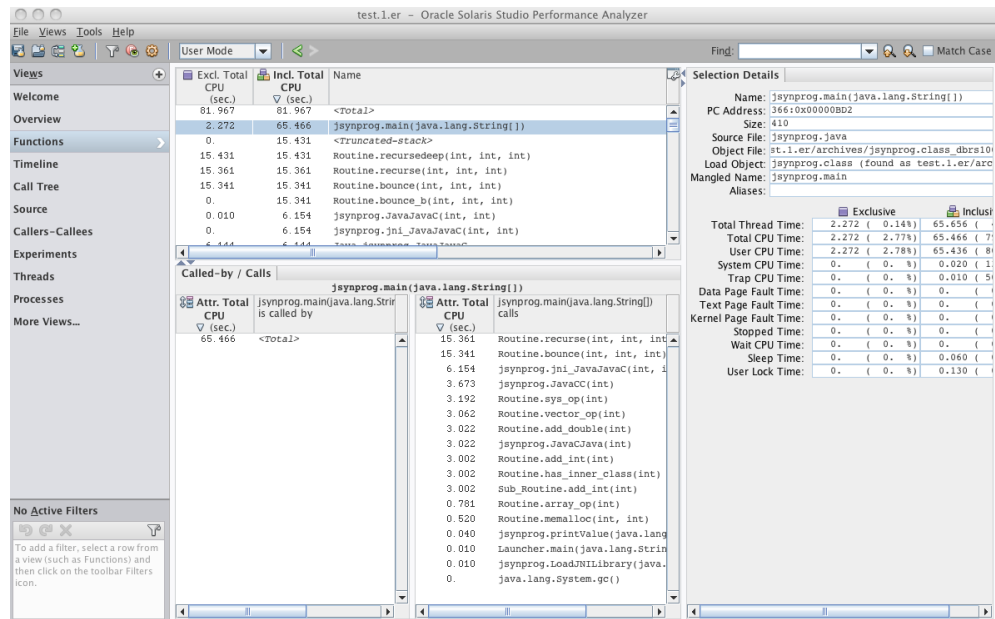
4. Select any thread in the Threads view and see all the information for that thread in the Selection Details window on the right.

You should see that almost all of the threads except Thread 1 and Thread 2 spend all their time in User Lock state. This shows how the JVM synchronizes itself internally. Thread 1 launches the user Java code and then sleeps until it finishes.

5. Go back to the Overview and deselect Sleep Time and User Lock Time.
6. Select the Functions view in the navigation panel, then click on the column headers to sort by Exclusive Total CPU Time, Inclusive Total CPU Time, or Name.

You can sort by descending or ascending order.

Leave the list sorted by Inclusive Total CPU Time in descending order and select the top-most function `jsynprog.main()`. That routine is the initial routine that the JVM calls to start execution.



Notice that the Called-by/Calls panel at the bottom of the Functions view show that the `jsynprog.main()` function is called by `<Total>`, meaning it was at the top of the stack.

The Calls side of the panel shows that `jsynprog.main()` calls a variety of different routines, one for each of the subtasks shown in [“About the Java-C++ Profiling Tutorial” on page 43](#) that are directly called from the main routine. The list also includes a few other routines.

Examining Mixed Java and C++ Code

This section features the Call Tree view and Source view, and shows you how to see the relationships between calls from Java and C++ and back again. It also shows how to add the Disassembly view to the navigation panel.

1. Select each of the functions at the top of the list in the Function view in turn, and examine the detailed information in the Selection Details window.

Note that for some functions the Source File is reported as `jsynprog.java`, while for some others it is reported as `cloop.cc`. That is because the `jsynprog` program has loaded a C++ shared object named `libcloop.so`, which was built from the `cloop.cc` C++ source file. Performance Analyzer reports calls from Java to C++ and vice-versa seamlessly.

2. Select the Call Tree in the navigation panel.

The Call Tree view shows graphically how these calls between Java and C++ are made.

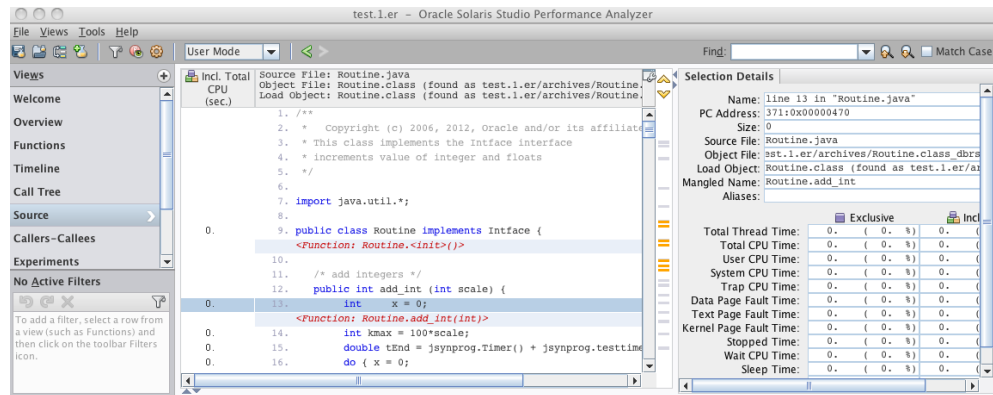
The screenshot shows the Oracle Solaris Studio Performance Analyzer interface. The main window displays the Call Tree view, which is a hierarchical tree of function calls. The tree is sorted by metric, and the total call count is 81,967 (100%). The tree shows a mix of Java and C++ code, with Java code represented by blue icons and C++ code by green icons. The Selection Details panel on the right provides information for the selected function, jsynprog.JavaCJava(int), including its PC Address, Source File, Object File, Load Object, and Mangled Name. Below this information is a table of performance metrics, including Total Thread Time, Total CPU Time, User CPU Time, System CPU Time, Trap CPU Time, Data Page Fault Time, Text Page Fault Time, Kernel Page Fault Time, Stopped Time, Wait CPU Time, Sleep Time, and User Lock Time.

	Exclusive	Inclusi
Total Thread Time:	0. (0. %)	3.022 (
Total CPU Time:	0. (0. %)	3.022 (
User CPU Time:	0. (0. %)	3.022 (
System CPU Time:	0. (0. %)	0. (
Trap CPU Time:	0. (0. %)	0. (
Data Page Fault Time:	0. (0. %)	0. (
Text Page Fault Time:	0. (0. %)	0. (
Kernel Page Fault Time:	0. (0. %)	0. (
Stopped Time:	0. (0. %)	0. (
Wait CPU Time:	0. (0. %)	0. (
Sleep Time:	0. (0. %)	0. (
User Lock Time:	0. (0. %)	0. (

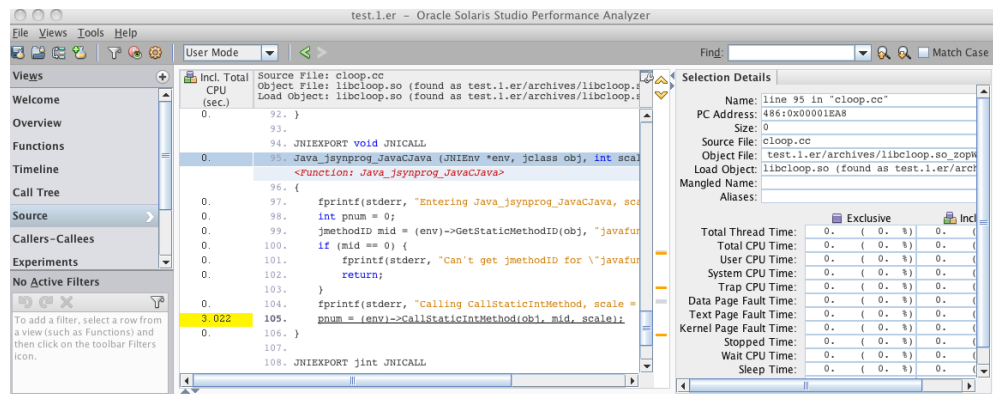
- In the Call Tree view, do the following to see the calls from Java to C++ and back to Java:
 - Expand the lines referring to the various functions with "C" in their name.
 - Select the line for `jsynprog.JavaCC()`. This function comes from the Java code, but it calls into `Java_jsynprog_JavaCC()` which comes from the C++ code.
 - Select the line for `jsynprog.JavaCJava()`. This function also comes from the Java code but calls `Java_jsynprog_JavaCJava()` which is C++ code. That function calls into a C++ method of the `JNIEnv::CallStaticIntMethod()` which calls back into Java to the method `jsynprog.javafunc()`.
- Select a method from either Java or C++ and switch to the Source view to see the source shown in the appropriate language along with performance metrics.

An example of the Source view after selecting a Java method is shown below.

Examining Mixed Java and C++ Code

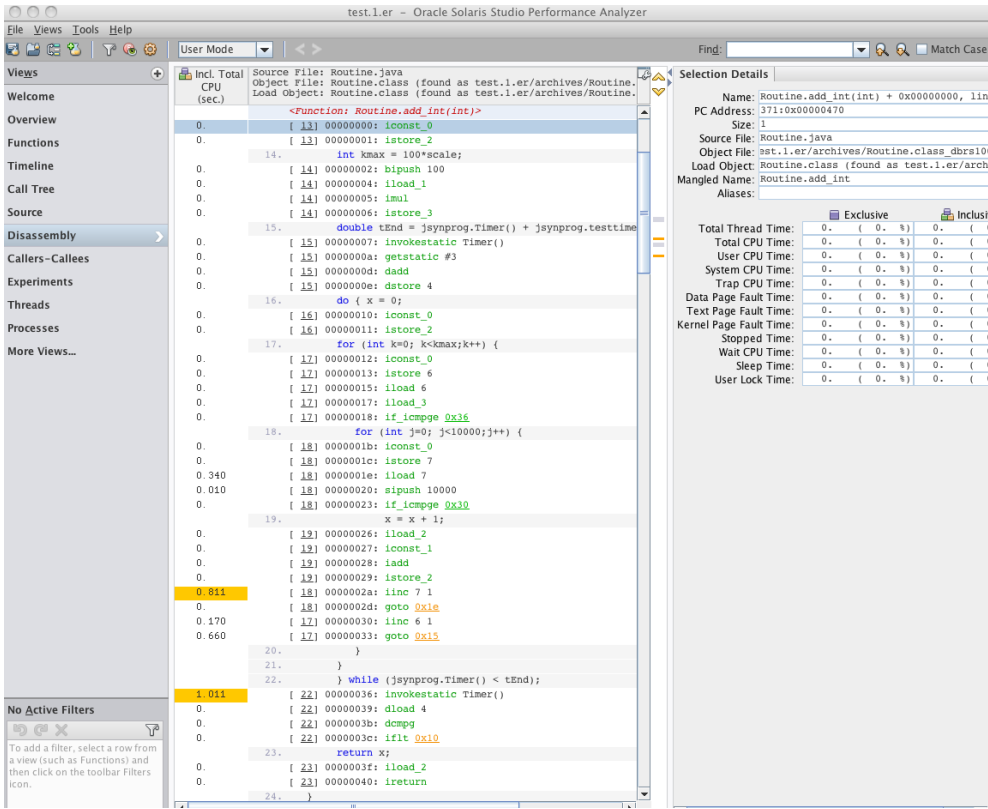


An example of the Source view after selecting a C++ method is shown below.

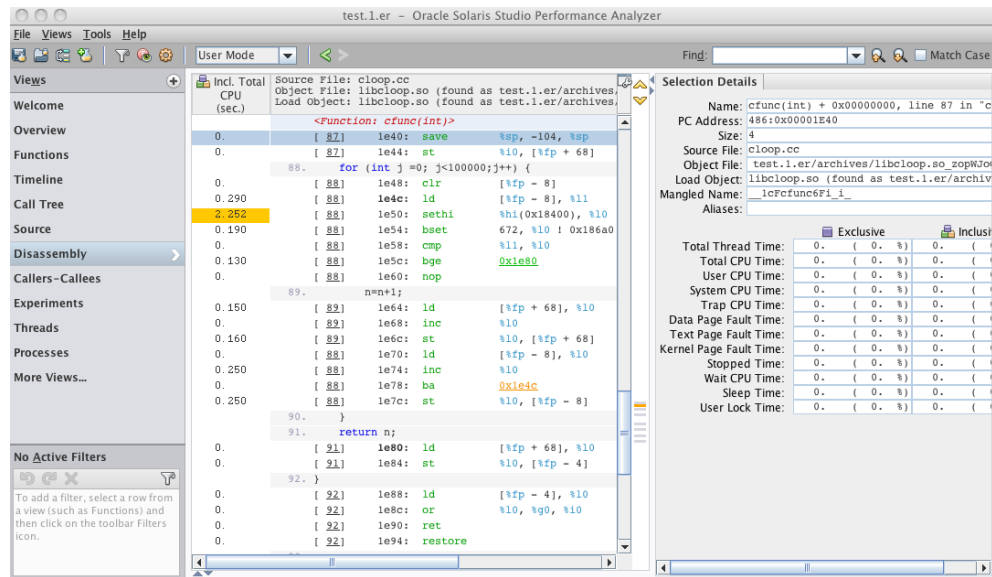


5. At the top of the navigation panel, click the + button next to the Views label and select the checkbox for Disassembly.

The Disassembly view for the function that you last selected is displayed. For a Java function, the Disassembly view shows Java byte code, as shown in the following screenshot.



For a C++ function, the Disassembly view shows native machine code, as shown in the following screenshot.



The next section uses the Disassembly view further.

Understanding the JVM Behavior

This section shows how to examine what is occurring in the JVM by using filters, Expert Mode, and Machine Mode.

1. Select the Functions view and find the routine named <JVM-System>.

You can find it very quickly using the Find tool in the tool bar if you type <JVM and press Enter.

In this experiment, <JVM-System> consumed about four seconds of Total CPU time. Time in the <JVM-System> function represents the workings of the JVM rather than the user code.

2. Right-click on <JVM-System> and select "Add Filter: Include only stacks containing the selected functions".

Notice that the filters panel below the navigation panel previously displayed No Active Filters and now shows 1 Active Filter with the name of the filter that you added. The Functions view refreshes so that only <JVM-System> is remaining.

3. In the Performance Analyzer tool bar, change the view mode selector from User Mode to Expert Mode.

The Functions view refreshes to show many functions that had been represented by <JVM-System> time. The function <JVM-System> itself is no longer visible.

4. Remove the filter by clicking the X in the Active Filters panel.

The Functions view refreshes to show the user functions again, but the functions represented by <JVM-System> are also still visible while the <JVM-System> function is not visible.

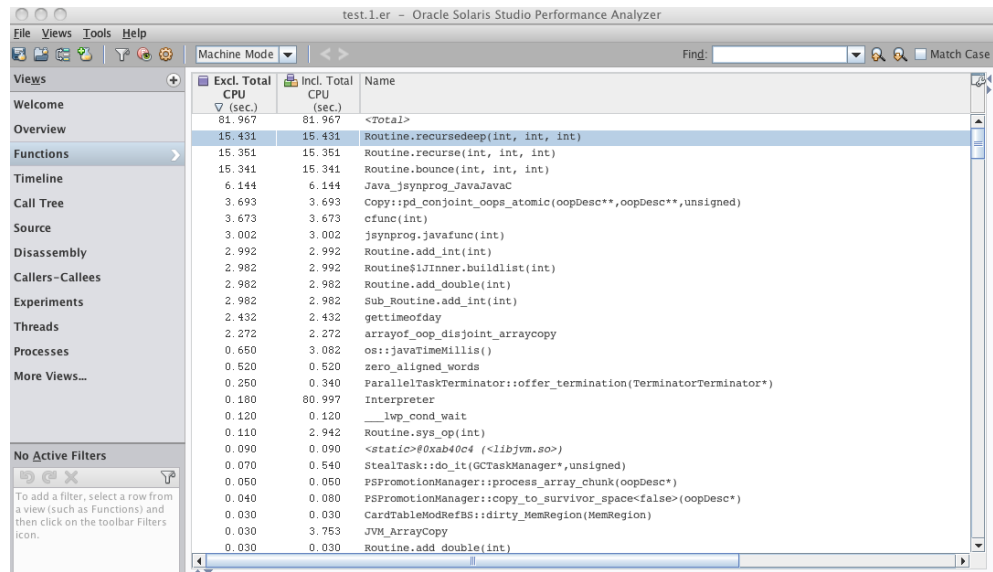
Excl. Total CPU (sec.)	Incl. Total CPU (sec.)	Name
81.967	81.967	<Total>
15.431	15.431	Routine.recursedeep(int, int, int)
15.361	15.361	Routine.recurse(int, int, int)
15.341	15.341	Routine.bounce(int, int, int)
6.144	6.144	Java_jsynprog_JavaJavac
3.763	3.763	java.lang.System.arraycopy(java.lang.Object, int, java.lang.Object, int, int)
3.673	3.673	cfunc(int)
3.192	3.192	Routine.sys_op(int)
3.022	3.022	Routine.add_double(int)
3.022	3.022	jsynprog.javafunc(int)
3.002	3.002	Routine.add_int(int)
2.992	3.002	Routine\$IJInner.buildlist(int)
2.592	2.592	Sub_Routine.addcall(int)
2.272	65.466	jsynprog.main(java.lang.String[])
0.520	0.520	Routine.memalloc(int, int)
0.410	3.002	Sub_Routine.add_int(int)
0.250	0.340	ParallelTaskTerminator::offer_termination(terminator:terminator*)
0.120	0.120	___lwp_cond_wait
0.090	0.090	<static@0xab40c4 (<libjvm.so>)
0.070	0.540	StealTask::do_it(GCTaskManager*, unsigned)
0.050	0.050	FSPromotionManager::process_array_chunk(oopDesc*)
0.040	0.080	FSPromotionManager::copy_to_survivor_space(false)(oopDesc*)
0.040	0.060	Routine.allocate_vector()
0.030	0.030	TaskQueueSetSuper::randomParkAndKiller(int*)
0.020	0.020	Copy::pd_disjoint_words(HeapWord*, HeapWord*, unsigned)
0.020	0.020	Monitor::IUnlock(bool)
0.020	0.070	FSPromotionManager::drain_stacks_depth(bool)

Note that you do not need to perform filtering to expand the <JVM-System>. This procedure includes filtering to more easily show the differences between User Mode and Expert Mode.

To summarize: User Mode shows all the user functions but aggregates all the time spent in the JVM into <JVM-System> while Expert Mode expands that <JVM-System> aggregation.

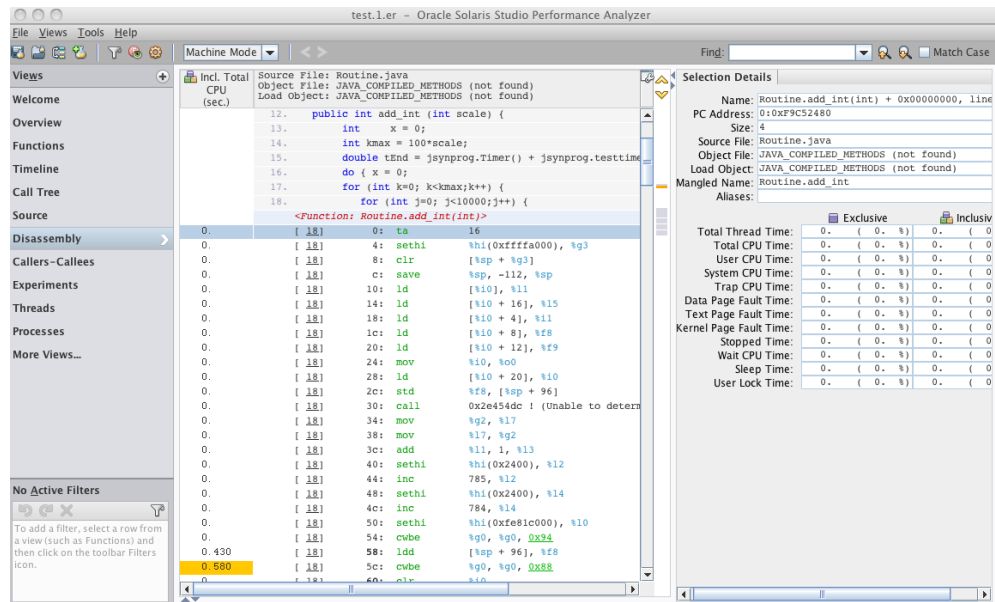
Next you can explore Machine Mode.

5. Select Machine Mode in the view mode list.



In Machine Mode, any user methods that are interpreted are not shown by name in the Functions view. The time spent in interpreted methods is aggregated into the Interpreter entry, which represents that part of the JVM that interpretively executes Java byte code.

However, in Machine Mode the Functions view displays any user methods that were HotSpot-compiled. If you select a compiled method such as `Routine.add_int()`, the Selection Details window shows the method's Java source file as the Source File, but the Object File and Load Object are shown as `JAVA_COMPILED_METHODS`.



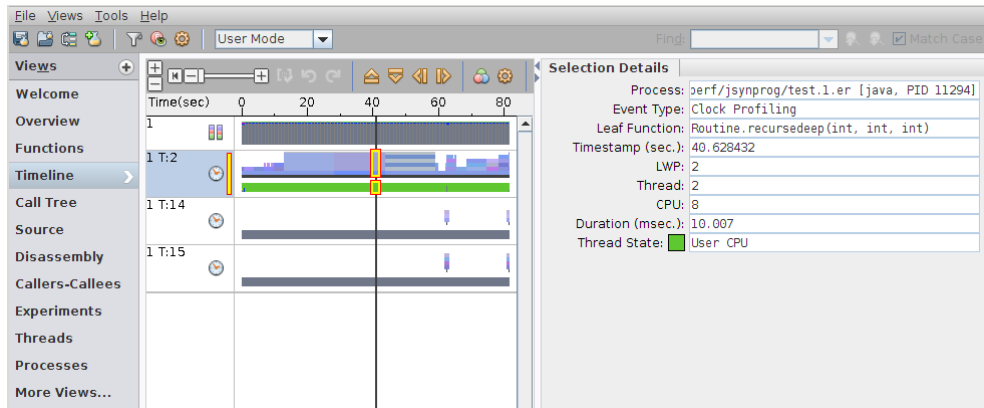
The Total CPU Time shown on most of the visible lines is zero, because most of the work in that function is performed further down in the code.

Continue to the next section.

Understanding the Java Garbage Collector Behavior

This procedure shows you how to use the Timeline view and the affect of the view mode setting on the Timeline, while examining the activities that trigger Java garbage collection.

1. Set the view mode to User Mode and select the Timeline view in the navigation panel to reveal the execution detail of this hybrid Java/native application, `jsynprog`.



You should see the CPU Utilization Samples bar at the top and profile data for three threads. In the screenshot you can see data for Process 1, Threads 2, 14, 15. The numbering and the number of threads you see might depend on the OS, the system, and the version of Java you are using.


Only the first thread, Thread 2 labeled as T:2 in the example, shows its microstate as User CPU. The other two threads spend all their time waiting for a User Lock, part of the JVM synchronization.

2. Set the view mode to Expert Mode.

The Timeline view should now show more threads although the user thread T:2 appears almost unchanged.

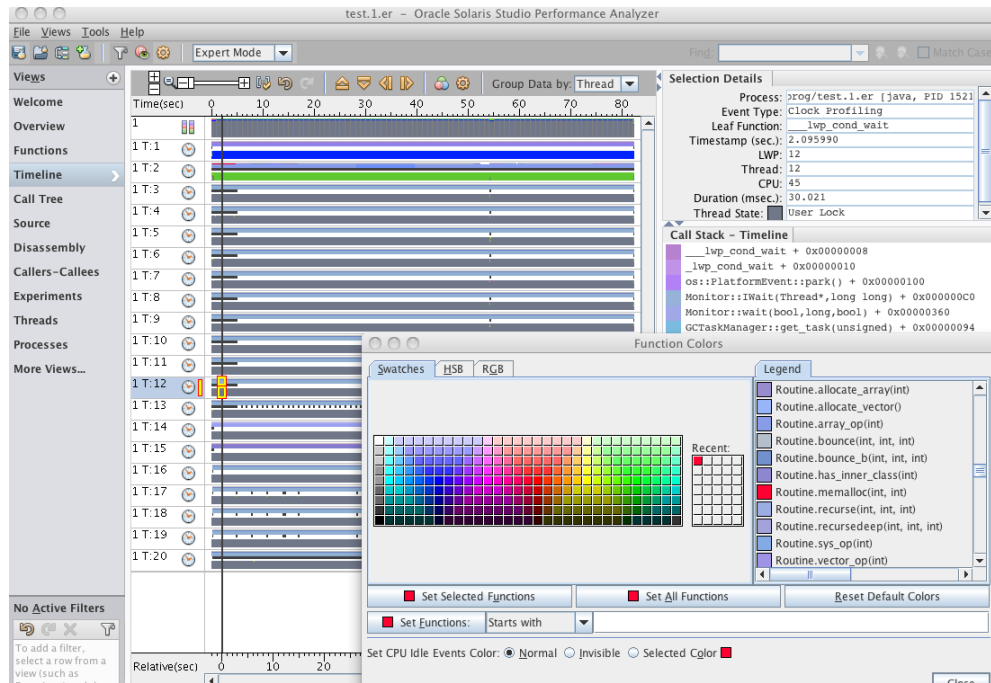
3. Use the vertical zoom control at the top of the timeline to adjust the zoom so that you can see all the threads.

The vertical zoom control is highlighted in red in the following screenshot. Click the minus button to reduce the height of the thread rows until you can see all twenty threads.



4. Click the Call Stack Function Colors button in the Timeline tool bar to set the color of the function `Routine.memalloc()` to red.

In the Function Colors dialog, select the `Routine.memalloc()` function in the Legend, click a red box in Swatches and click Set Selected Functions.

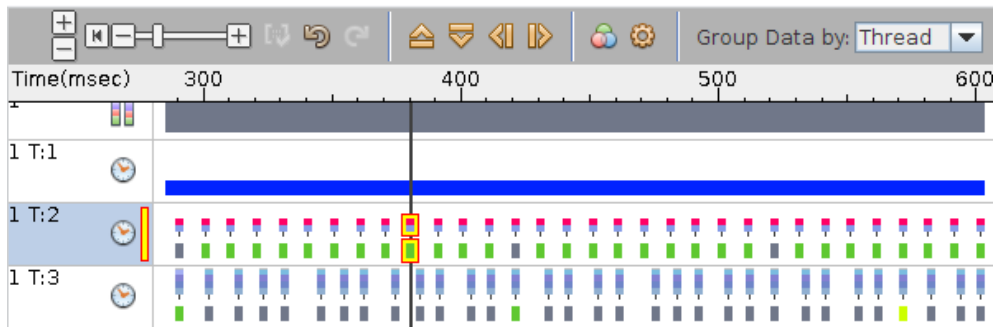


Note that Thread 2 now has a bar of red across the top of its stack. That area represents the portion of time where the `Routine.memalloc()` routine was running.

You might need to zoom out vertically to see more frames of the callstack, and zoom in horizontally to the region of time that is of interest.

5. Use the horizontal slider in the Timeline tool bar to zoom in close enough to see individual events in thread T:2.

You can also zoom by double-clicking or pressing the + key on your keyboard.



Each row of the timeline actually includes three data bars. The top bar is a representation of the callstack for that event. The middle bar shows black tick marks wherever events occur too closely together to show them all. In other words, when you see a tick mark, you know that there are multiple events in that space.

The lower bar is an indicator of the event state. For T:2 the lower bar is green, which indicates User CPU Time was being used. For threads 3 through 12 the lower bar is gray, which indicates User Lock Time.

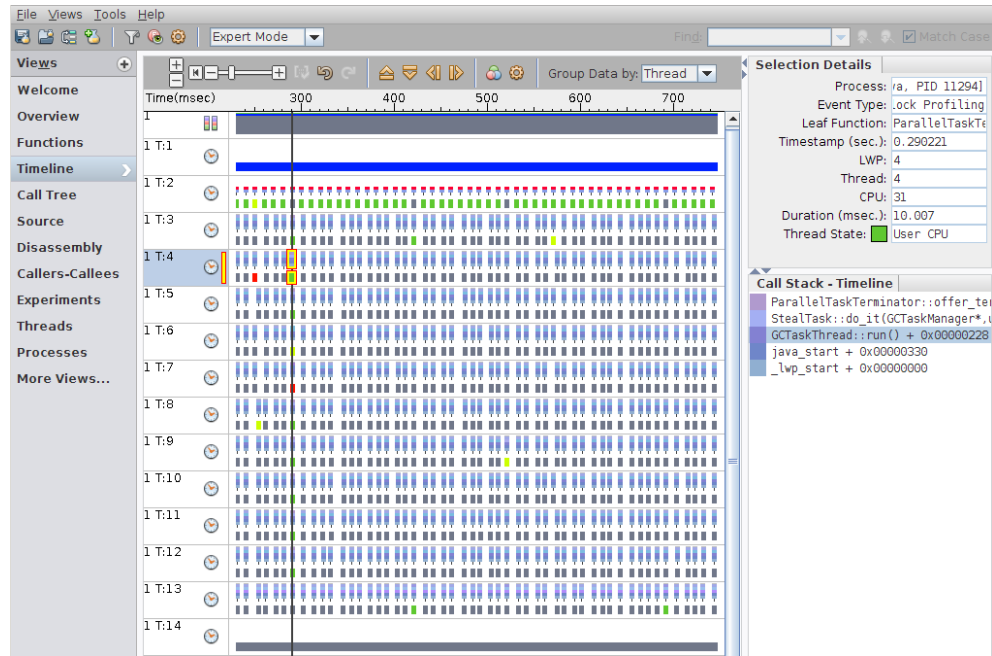
Notice however that all of those threads 3 through 12 have many events clustered together arriving at the same time as the user thread T:2 is in `Routine.memalloc`, the routine shown in red.

6. Zoom in to the `Routine.memalloc` region and filter to include only that region by doing the following:
 - Click on the T:2 bar close to the beginning of the `Routine.memalloc` region with the red function call on top.
 - Click and drag the mouse to close to the end of that region where the red at the top of the call stack ends.
 - Right-click and select `Zoom > To Selected Time Range`.
 - With the range still selected, right-click and select `Add Filter: Include only events from selected time range`.

After zooming you can see that there are some event states in threads 3-12 that are green to indicate User CPU time, and even a few that are red to indicate Wait CPU Time.

7. Click on any of the events on threads 3-12 and you see in the Call Stack panel that each thread's events include `GCTaskThread::run()` in the stack.

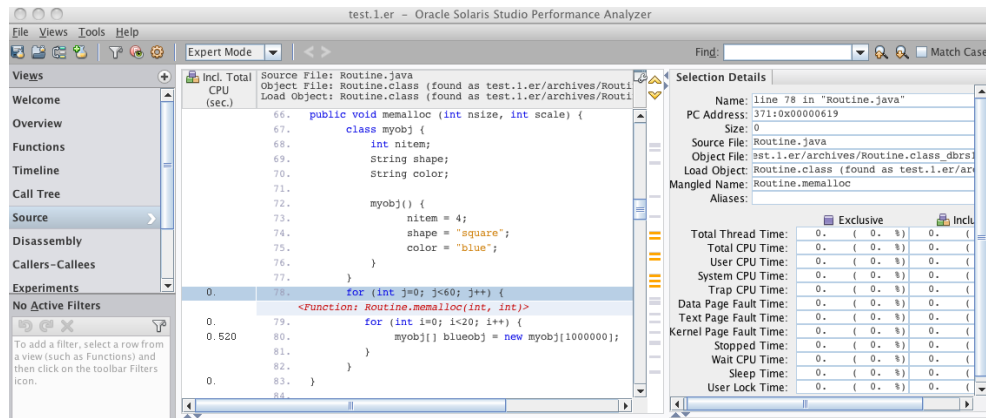
Those threads represent the threads that the JVM uses to run garbage collection. While they do not take a great amount of User CPU Time, the GC threads do run while the user thread is in `Routine.memalloc`, but not otherwise.



8. Go back to the Functions view and click on the Incl. Total CPU column header to sort by inclusive Total CPU Time.

You should see that one of the top functions is the `GCTaskThread::run()` function. This leads you to the conclusion that the user task `Routine.memalloc` is somehow triggering garbage collection.

9. Select the `Routine.memalloc` function and switch to the Source view.



From this fragment of source code it is easy to see why garbage collection is being triggered. The code allocates an array of one million objects and stores the pointers to those objects in the same place with each pass through the loop. This renders the old objects unused, and thus they become garbage.

Continue to the next section.

Understanding the Java HotSpot Compiler Behavior

This procedure continues from the previous section, and shows you how to use the Timeline and Threads views to filter and find the threads responsible for HotSpot compiling.

1. Select the Timeline view and remove the filter by clicking the X in the Active Filters panel, then reset the horizontal zoom to the default by pressing 0 on your keyboard.

You can also click the |< button in front of the horizontal slider in the Timeline tool bar, or right-click in the Timeline and select Reset.

2. Open the Function Colors dialog again, and pick different colors for each of the Routine.* functions.

In the Timeline view, the color changes appear in call stacks of thread 2.


3. Look at all the threads of the Timeline in the period of time where you see the color changes in thread 2.

You should see that there are some threads with patterns of events occurring at just about the same time as the color changes in thread 2. In this example, they are threads 17, 18, and 19.

4. Go to the Threads view and select thread 2 and the threads in your experiment that show activity during the time period where thread 2 shows calls to Routine.* functions.

You might find it easier to first sort by name by clicking the Name column header. Then select the multiple threads by pressing Ctrl as you click the threads.

In this example, threads 2, 17, 18, 19 are selected.

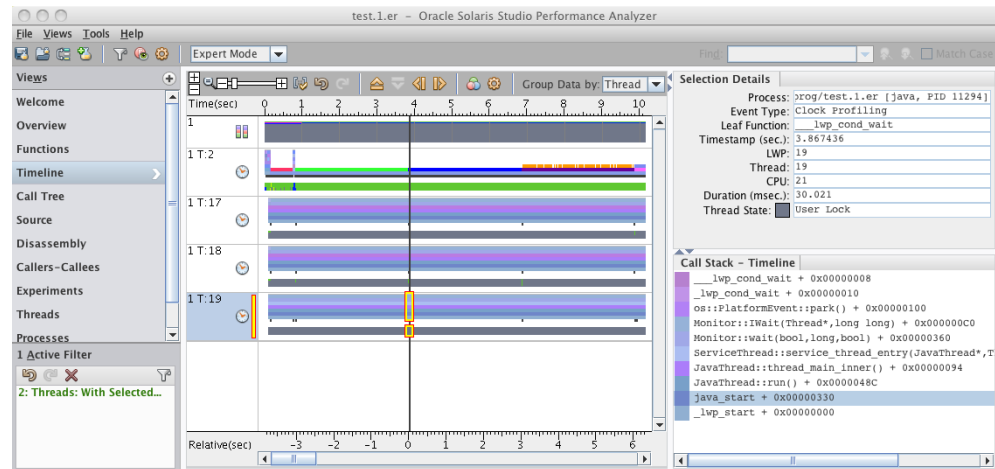
5. Click the filter button  in the toolbar and select Add Filter: Include only events with selected items.

This sets a filter to include only events on those threads. You could also right-click in the Threads view and select the filter.

6. Return to the Timeline View and reset the horizontal zoom to make the pattern easier to see.
7. Click on events in threads 17 and 18.

Note that the Call Stack panel shows `CompileBroker::compiler_thread_loop()`. Those threads are the threads used for the HotSpot compiler.

Thread 19 shows call stacks with `ServiceThread::service_thread_entry()` in them.



The reason the multiple events occur on those threads is that whenever the user code invokes a new method and spends a fair amount of time in it, the HotSpot compiler is triggered to generate machine code for that method. The HotSpot compiler is fast enough that the threads that run it do not consume very much User CPU Time.

The details of exactly how the HotSpot compiler is triggered is beyond the scope of this tutorial.

Hardware Counter Profiling on a Multithreaded Program

This chapter covers the following topics.

- [“About the Hardware Counter Profiling Tutorial” on page 63](#)
- [“Setting Up the `mttest` Sample Code” on page 64](#)
- [“Collecting Data From `mttest` for Hardware Counter Profiling Tutorial” on page 65](#)
- [“Examining the Hardware Counter Profiling Experiment for `mttest`” on page 65](#)
- [“Exploring Clock-Profiling Data” on page 67](#)
- [“Understanding Hardware Counter Instruction Profiling Metrics” on page 69](#)
- [“Understanding Hardware Counter CPU Cycles Profiling Metrics” on page 71](#)
- [“Understanding Cache Contention and Cache Profiling Metrics” on page 73](#)
- [“Detecting False Sharing” on page 77](#)

About the Hardware Counter Profiling Tutorial

This tutorial shows how to use Performance Analyzer on a multithreaded program named `mttest` to collect and understand clock profiling and hardware counter profiling data.

You explore the Overview page and change which metrics are shown, examine the Functions view, Callers-Callees view, and Source and Disassembly views, and apply filters.

You first explore the clock profile data, then the HW-counter profile data with Instructions Executed which is a counter available on all supported systems. Then you explore Instructions Executed and CPU Cycles (available on most, but not all, supported systems) and with D-cache Misses (available on some supported systems).

If run on a system with a precise hardware counter for D-cache Misses (`dcm`), you will also learn how to use the IndexObject and MemoryObject views, and how to detect false sharing of a cache line.

The program `mttest` is a simple program that exercises various synchronization options on dummy data. The program implements a number of different tasks and each task uses a basic algorithm:

- Queue up a number of work blocks, four by default. Each one is an instance of a structure `Workblk`.
- Spawn a number of threads to process the work, also four by default. Each thread is passed its private work block.
- In each task, use a particular synchronization primitive to control access to the work blocks.
- Process the work for the block, after the synchronization.

The data you see in the experiment that you record will be different from that shown here. The experiment used for the screen shots in the tutorial was recorded on a SPARC T5 system running Oracle Solaris 11.2. The data from an x86 system running Oracle Solaris or Linux will be different. Furthermore, data collection is statistical in nature and varies from experiment to experiment, even when run on the same system and OS.

The Performance Analyzer window configuration that you see might not precisely match the screen shots. Performance Analyzer enables you to drag separator bars between components of the window, collapse components, and resize the window. Performance Analyzer records its configuration and uses the same configuration the next time it runs. Many configuration changes were made in the course of capturing the screen shots shown in the tutorial.

Setting Up the `mttest` Sample Code

Before You Begin

See the following for information about obtaining the code and setting up your environment.

- [“Getting the Sample Code for the Tutorials” on page 8](#)
- [“Setting Up Your Environment for the Tutorials” on page 9](#)

You might want to go through the introductory tutorial in [“Introduction to C Profiling”](#) first to become familiar with Performance Analyzer.

1. Copy the contents of the `mttest` directory to your own private working area with the following command:

```
% cp -r SolarisStudioSampleApplications/PerformanceAnalyzer/mttest mydirectory
```

where *mydirectory* is the working directory you are using.

2. Change to that working directory copy.

```
% cd mydirectory/mttest
```

3. Build the target executable.

```
% make clobber (needed only if you ran make in the directory before, but safe in any case)
```

```
% make
```


After you run `make` the directory contains the target application to be used in the tutorial, a C program called `mttest`.

Tip - If you prefer, you can edit the `Makefile` to do the following: use the GNU compilers rather than the default of the Studio compilers; build in 32-bits rather than the default of 64-bits; and add different compiler flags.

Collecting Data From `mttest` for Hardware Counter Profiling Tutorial

The easiest way to collect the data is to run the following command in the `mttest` directory:

```
% make hwcperf
```

The `hwcperf` target of the `Makefile` launches a `collect` command and records an experiment.

The experiment is named `test.1.er` by default and contains clock-profiling data and hardware counter profiling data for three counters: `inst` (instructions), `cycles` (cycles), and `dcm` (data-cache-misses).

If your system does not support a `cycles` counter or a `dcm` counter, the `collect` command will fail. In that case, edit the `Makefile` to move the `#` sign to the appropriate line to enable the `HWC_OPT` variable that specifies only those counters that are supported on your system. The experiment will not have the data from those counters that were omitted.

Tip - You can use the command `collect -h` to determine which counters your system does support. For information about the hardware counters, see [“Hardware Counter Lists”](#) in [“Oracle Solaris Studio 12.4: Performance Analyzer”](#).

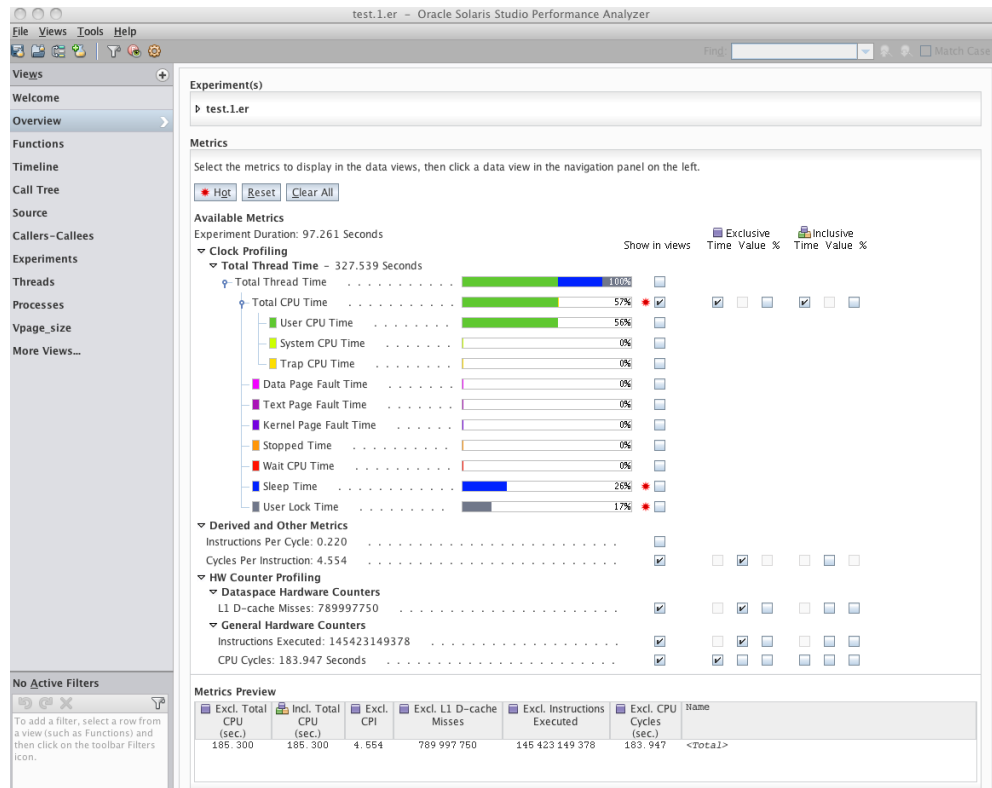
Examining the Hardware Counter Profiling Experiment for `mttest`

This section shows how to explore the data in the experiment you created from the `mttest` sample code in the previous section.

1. Start Performance Analyzer from the `mttest` directory and load the experiment as follows:

```
% analyzer test.1.er
```

When the experiment opens, Performance Analyzer shows the Overview page.



The Clock Profiling metrics are shown first and include colored bars. Most of the thread time is spent in User CPU Time. Some time is spent in Sleep Time or User Lock Time.

The Derived and Other Metrics group is present if you have recorded both cycles and insts counters. The derived metrics represent the ratios of the metrics from those two counters. A high value of Instructions Per Cycle or a low value of Cycles Per Instruction indicates relatively efficient code. Conversely, a low value of Instructions Per Cycle or a high value of Cycles Per Instruction indicates relatively inefficient code.

The HW Counter Profiling group shows two subgroups in this experiment, Dataspace Hardware Counters and General Hardware Counters. The Instructions Executed counter (insts) is listed under General Hardware Counters. If the data you collected included the cycles counter, CPU Cycles is also listed under General Hardware Counters. If the data was collected on a machine with a *precise* dcm counter, L1 D-cache Misses is listed under Dataspace Hardware Counters. If the dcm counter was available but is not a precise counter, L1 D-cache Misses is listed under General Hardware Counters. A precise counter is one whose overflow interrupt is delivered at the execution of the instruction causing the overflow. Non-precise counters are delivered with a variable amount of "skid" past

the instruction causing the overflow. Even if a non-precise counter is memory-related, it cannot be used for dataspace profiling. For more information about dataspace profiling, see [“Dataspace Profiling and Memoryspace Profiling”](#) in [“Oracle Solaris Studio 12.4: Performance Analyzer”](#).

If your system does not support `dcm`, and you edited the `Makefile` to remove the `-h dcm`, you will see the Instructions Executed and CPU Cycles counter. If you edited the `Makefile` to remove both the `-h dcm` and `-h cycles`, you will only see the Instructions Executed counter.

You will explore these metrics and their interpretation in the following sections of the tutorial.

Exploring Clock-Profiling Data

This section explores the clock profiling data using the Overview page and the Functions view with the Called-by/Calls panel.

1. In the Overview page, deselect the check boxes for three HW counter metrics, leaving only the Total CPU Time check boxes selected.
2. Go to the Functions view and click the column heading once for Incl. Total CPU to sort according to inclusive total CPU time.

The function `do_work()` should now be at the top of the list.

Exploring Clock-Profiling Data

The screenshot displays the Oracle Solaris Studio Performance Analyzer interface. The main window shows a list of functions with columns for Excl. Total CPU (sec.), Incl. Total CPU (sec.), and Name. The 'do_work' function is selected, and its details are shown in the 'Selection Details' panel on the right. Below the main list, the 'Called-by / Calls' panel shows the functions that call 'do_work' and the functions it calls.

Excl. Total CPU (sec.)	Incl. Total CPU (sec.)	Name
185.300	185.300	<Total>
0.510	185.300	do_work
0.	173.271	_lwp_start
0.	58.741	cache_trash
58.741	58.741	computeB
0.	44.091	cache_trash_even
0.640	30.011	trylock_global
3.402	17.352	mutex_trylock
0.	14.650	cache_trash_odd
0.	13.950	do_exit_critical
13.950	13.950	take_deferred_direct
0.	12.028	_start
0.	12.028	locktest
0.	12.028	main
12.018	12.018	compute
12.018	12.018	computeB
0.	12.018	nothreads
12.008	12.008	computeA
12.008	12.008	computeC
12.008	12.008	computeE
12.008	12.008	computeG
12.008	12.008	computeI
0.	12.008	cond_global
0.	12.008	lock_global
0.	12.008	lock_local
0.	12.008	lock_none
0.	12.008	sema_global
11.998	11.998	computeH
0.	11.998	cond_timeout_global

Called-by / Calls

Attr. Total CPU (sec.)	do_work is called by	Attr. Total CPU (sec.)	do_work calls
173.271	_lwp_start	58.741	cache_trash
12.028	locktest	30.011	trylock_global
		12.018	nothreads
		12.008	cond_global
		12.008	lock_global
		12.008	lock_local
		12.008	lock_none
		12.008	sema_global
		11.998	cond_timeout_global
		11.978	calladd

Selection Details

Name: do_work
 PC Address: 2:0x00004028
 Size: 352
 Source File: mttest.c
 Object File: mttest (found as test.1.er/archives/mttest_wgvmhd3f0c)
 Load Object: mttest (found as test.1.er/archives/mttest_wgvmhd3f0c)
 Mangled Name:
 Aliases:

	Exclusive	Inclusive
Total Thread Time:	0.510 (0.16%)	242.309 (73.98)
Total CPU Time:	0.510 (0.28%)	185.300 (100.00)
User CPU Time:	0.510 (0.28%)	185.029 (100.00)
System CPU Time:	0. (0. %)	0.120 (100.00)
Trap CPU Time:	0. (0. %)	0.150 (100.00)
Data Page Fault Time:	0. (0. %)	0. (0. %)
Text Page Fault Time:	0. (0. %)	0. (0. %)
Kernel Page Fault Time:	0. (0. %)	0. (0. %)
Stopped Time:	0. (0. %)	0. (0. %)
Wait CPU Time:	0. (0. %)	0. (0. %)
Sleep Time:	0. (0. %)	0. (0. %)
User Lock Time:	0. (0. %)	57.010 (100.00)
L1 D-cache Misses:	0 (0. %)	789997750 (100.00)
Instructions Executed:	0 (0. %)	145423149378 (100.00)
CPU Cycles:	0.009 (0.00%)	183.947 (100.00)
*count:	31600040	662210550587
Instructions Per Cycle:	0. (0. %)	0.220 (100.00)
Cycles Per Instruction:	0. (0. %)	4.554 (100.00)

3. Select the `do_work()` function and look at the Called-by/Calls panel at the bottom of the Functions view.

Note that `do_work()` is called from two places, and it calls ten functions.

The ten functions that `do_work()` calls represent ten different tasks, each with a different synchronization method that the program executed. In some experiments created from `mttest` you might see an eleventh function which uses relatively little time to fetch the work blocks for the other tasks. This function is not shown in the screen shot.

Most often, `do_work()` is called when a thread to process the data is created, and is shown as called from `_lwp_start()`. In one case, `do_work()` calls one single-threaded task called `nothreads()` after being called from `locktest()`.

In the Calls side of the panel, note that except for the first two of the callees, all callees show about the same amount of time (~12 seconds) of Attributed Total CPU.

Understanding Hardware Counter Instruction Profiling Metrics

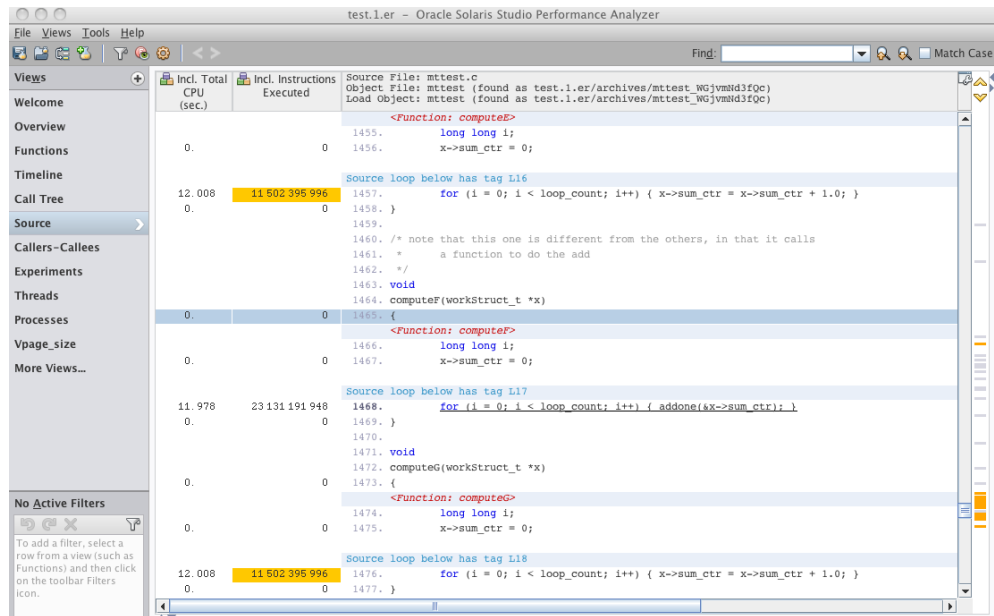
This section shows how to use general hardware counters to see how many instructions are executed for functions.

1. Select the Overview page and enable the HW Counter Profiling metric named Instructions Executed, which is under General Hardware Counters.
2. Return to the Functions view, and click on the Name column header to sort alphabetically.
3. Scroll down to find the functions `compute()`, `computeA()`, `computeB()`, etc.

	Excl. Total CPU (sec.)	Incl. Total CPU (sec.)	Excl. Instructions Executed	Name
	9.156	9.156	13 682 795 237	addone
	0.	58.741	0	cache_trash
	0.	44.091	0	cache_trash_even
	0.	14.650	0	cache_trash_odd
	0.	11.978	0	calladd
	12.018	12.018	11 565 595 974	compute
	12.008	12.008	11 502 395 996	computeA
	58.741	58.741	11 628 795 952	computeB
	12.008	12.008	11 502 395 996	computeC
	12.018	12.018	11 533 995 985	computeD
	12.008	12.008	11 502 395 996	computeE
	2.822	11.978	9 448 396 711	computeF
	12.008	12.008	11 502 395 996	computeG
	11.998	11.998	11 502 395 996	computeH
	12.008	12.008	11 502 395 996	computeI
	0.	12.008	0	cond_global
	0.	0.	0	cond_sleep_queue
	0.	0.	0	cond_timedwait
	0.	11.998	0	cond_timeout_global
	0.	0.	0	cond_wait
	0.	0.	0	cond_wait_common
	0.	0.	0	cond_wait_queue
	0.	13.950	0	do_exit_critical
	0.510	185.300	0	do_work

Note that all of the functions except `computeB()` and `computeF()` have approximately the same amount of Exclusive Total CPU time and of Exclusive Instructions Executed.

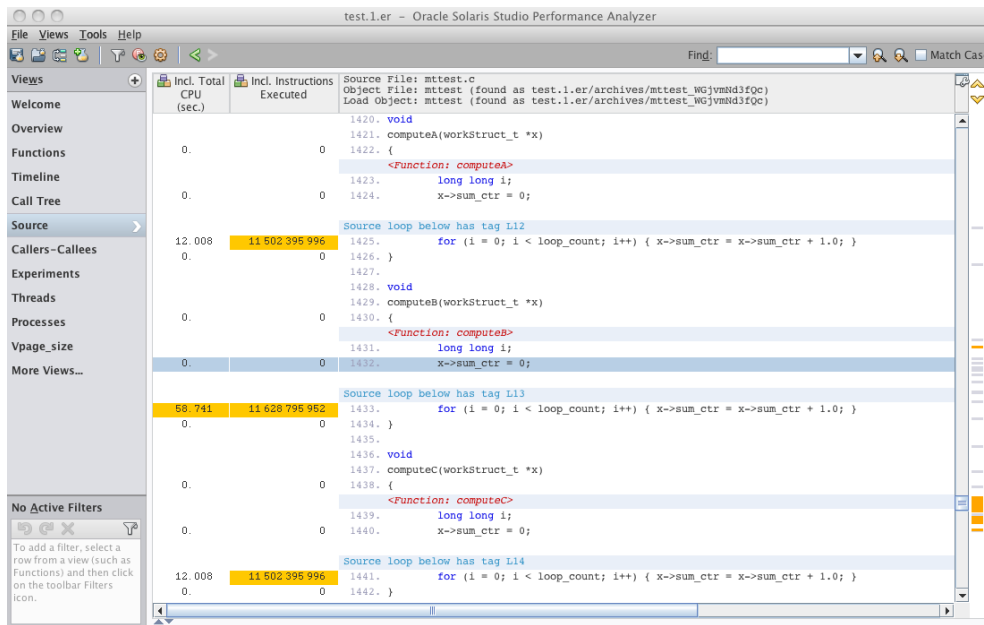
4. Select `computeF()` and switch to the Source view. You can do this in one step by double-clicking `computeF()`.



The computation kernel in computeF() is different because it calls a function addone() to add one, while the other compute*() functions do the addition directly. This explains why its performance is different from the others.

5. Scroll up and down in the Source view to look at all the compute*() functions.

Note that all of the compute*() functions, including computeB(), show approximately the same number of instructions executed. Yet computeB() shows a very different CPU Time cost.



The next section helps show why the Total CPU time is so much higher for computeB().

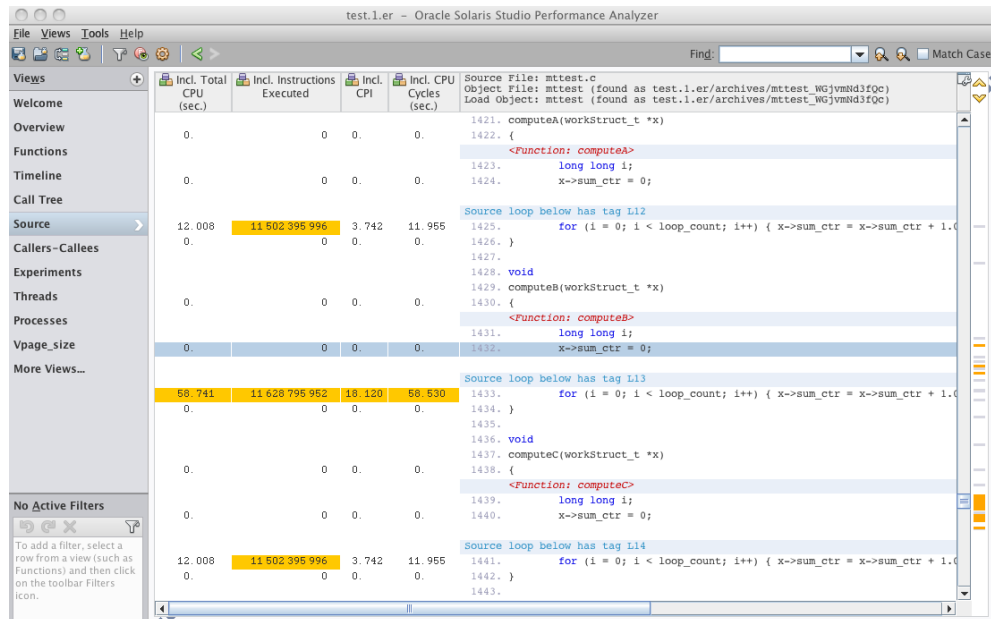
Understanding Hardware Counter CPU Cycles Profiling Metrics

This part of the tutorial requires an experiment with data from the cycles counter. If your system does not support this counter, your experiment cannot be used in this section. Skip to the next section [“Understanding Cache Contention and Cache Profiling Metrics”](#) on page 73.

1. Select the Overview page and enable the derived metric Cycles Per Instruction and the General Hardware Counter metric, CPU Cycles.
You should keep Inclusive Total CPU and Instructions Executed selected.



2. Return to the Source view at computeB().



Note that the Incl. CPU Cycles time and the Incl. Total CPU Time are roughly equivalent in each of the compute*() functions. This indicates that the clock-profiling and CPU Cycles hardware counter profiling are getting similar data.

In the screen shots, the Incl. CPU Cycles and the Incl. Total CPU Time are about 12 seconds for each of the compute*() functions except computeB(). You should also see in your experiment that the Incl. Cycles Per Instruction (CPI) is much higher for computeB() than it is for the other compute*() functions. This indicates that more CPU cycles are needed to execute the same number of instructions, and computeB() is therefore less efficient than the others.

The data you have seen so far shows the difference between that `computeB()` function and the others, but does not show why they might be different. The next part of this tutorial explores why `computeB()` is different.

Understanding Cache Contention and Cache Profiling Metrics

This section and the rest of the tutorial requires an experiment with data from the precise `dcm` hardware counter. If your system does not support the precise `dcm` counter, the remainder of the tutorial is not applicable to the experiment you recorded on the system.

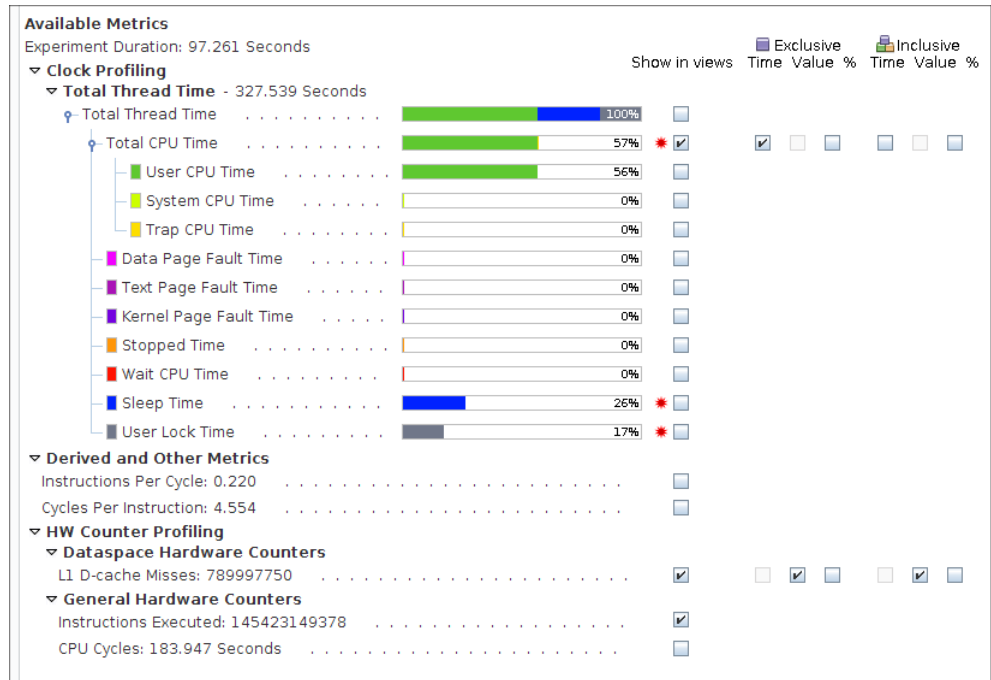
The `dcm` counter is counting cache misses, which are loads and stores that reference a memory address that is not in the cache.

An address might not be in cache for any of the following reasons:

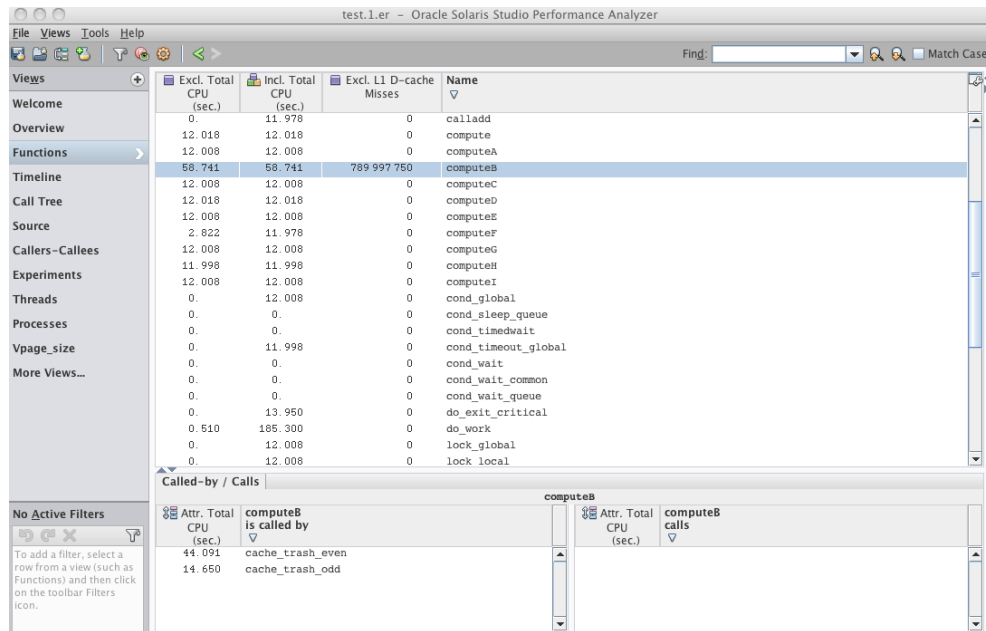
- Because the current instruction is the first reference to that memory location from that CPU. More accurately, it is the first reference to any of the memory locations that share the cache line.
- Because the thread has referenced so many other memory addresses that the current address has been flushed from the cache. This is a capacity miss.
- Because the thread has referenced other memory addresses that map to the same cache line which causes the current address to be flushed. This is a conflict miss.
- Because another thread has written to an address within the cache line which causes the current thread's cache line to be flushed. This is a sharing miss, and could be one of two types of sharing misses:
 - True sharing, where the other thread has written to the same address that the current thread is referencing. Cache misses due to true sharing are unavoidable.
 - *False sharing*, where the other thread has written to a different address from the one that the current thread is referencing. Cache misses due to false sharing occur because the cache hardware operates at a cache-line granularity, not a data-word granularity. False sharing can be avoided by changing the relevant data structures so that the different addresses referenced in each thread are on different cache lines.

This procedure examines a case of false sharing that has an impact on the function `computeB()`.

1. Return to the Overview, and enable the metric for L1 D-cache Misses, and disable the metrics for Cycles Per Instruction and Inclusive Total CPU Time.



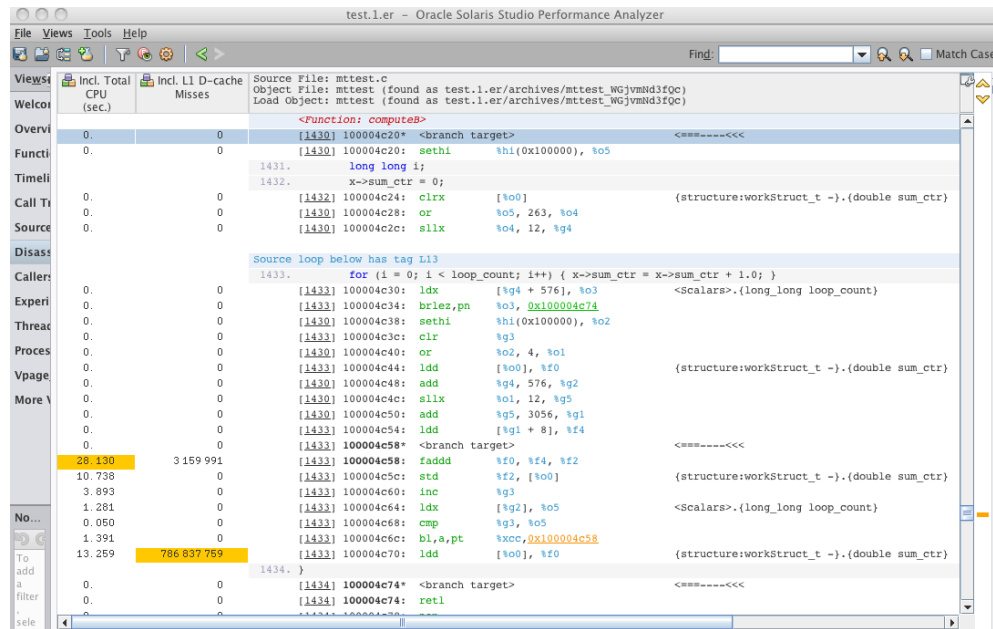
- Switch back to the Functions view and look at the compute*() routines.



Recall that all compute*() functions show approximately the same instruction count, but computeB() shows higher Exclusive Total CPU Time and is the only function with significant counts for Exclusive L1 D-cache Misses.

3. Go back to the Source view and note that in computeB() the cache misses are in the single line loop.
4. In the Views navigation panel, click More Views or the + button and select Disassembly. Scroll the Disassembly view until you see the line with the high number of L1 D-Cache Misses. It will be a load instruction.

Tip - The right margin of views such as Disassembly include shortcuts you can click to jump to the lines with high metrics, or hot lines. Try clicking the Next Hot Line down-arrow at the top of the margin or the Non-Zero Metrics marker to jump quickly to the lines with notable metric values.



On SPARC systems, if you compiled with `-xhwcprof`, loads and stores are annotated with structure information showing that the instruction is referencing a double word, `sum_ctr` in the `workStruct_t` data structure. You also see lines with the same address as the next line, with `<branch target>` as its instruction. Such lines indicate that the next address is the target of a branch, which means the code might have reached an instruction that is indicated as hot without ever executing the instructions above the `<branch target>`.

On x86 systems, the loads and stores are not annotated and `<branch target>` lines are not displayed because the `-xhwcprof` is not supported on x86.

5. Go back and forth between the Functions and Disassembly views, selecting various `compute*()` functions.

Note that for all `compute*()` functions, the instructions with high counts for Instructions Executed reference the same structure field.

You have now seen that `computeB()` takes much longer than the other functions even though it executes the same number of instructions, and is the only function that gets cache misses. The cache misses are responsible for the increased number of cycles to execute the instructions because a load with a cache miss takes many more cycles to complete than a load with a cache hit.

For all the `compute*()` functions except `computeB()`, the double word field `sum_ctr` in the structure `workStruct_t` which is pointed to by the argument from each thread, is contained

within the `Workblk` for that thread. Although the `Workblk` structures are allocated contiguously, they are large enough so that the double words in each structure are too far apart to share a cache line.

For `computeB()`, the `workStruct_t` arguments from the threads are consecutive instances of that structure, which is only one double-word long. As a result the double-words used by the different threads will share a cache line, which causes any store from one thread to invalidate the cache line in the other threads. That is why the cache miss count is so high, and the delay refilling the cache line is why the Total CPU Time and CPU Cycles Metric is so high.

In this example, the data words being stored by the threads do not overlap although they share a cache line. This performance problem is referred to as "false sharing". If the threads were referring to the same data words, that would be true sharing. The data you have looked at so far do not distinguish between false and true sharing.

The difference between false and true sharing is explored in the last section of this tutorial.

Detecting False Sharing

This part of the tutorial is applicable only to systems where the L1 D-Cache Miss `dcm` counter is precise. Such systems include SPARC-T4, SPARC-T5, SPARC-M5 and SPARC-M6, among others. If your experiment was recorded on a system without a precise `dcm` counter, this section does not apply.

This procedure shows how to use Index Object views and Memory Object views along with filtering.

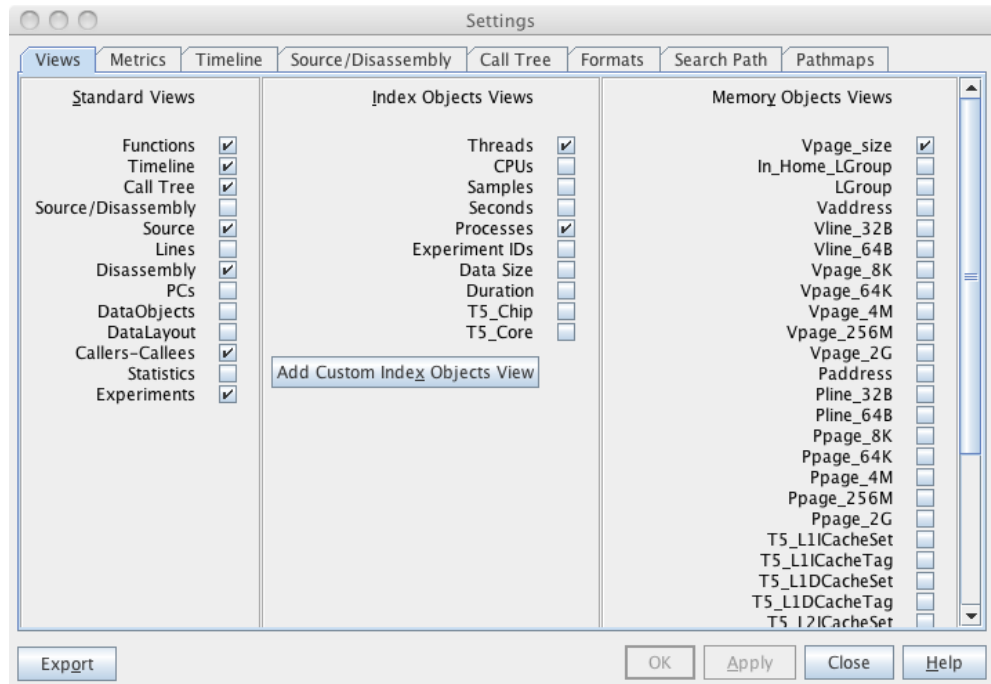
When you create an experiment on a system with precise memory-related counters, a *machine model* is recorded in the experiment. The machine model represents the mappings of addresses to the various components in the memory subsystem of that machine. When you load the experiment in Performance Analyzer or `er_print`, the machine model is automatically loaded.

The experiment used for the screen shots in this tutorial was recorded on a SPARC T5 system and the `t5` machine model for that machine is automatically loaded with the experiment. The machine model adds data views of index objects and memory objects.

1. Go to the Functions view and select `computeB()`, then right-click and select Add Filter: Include only stacks containing the selected functions.

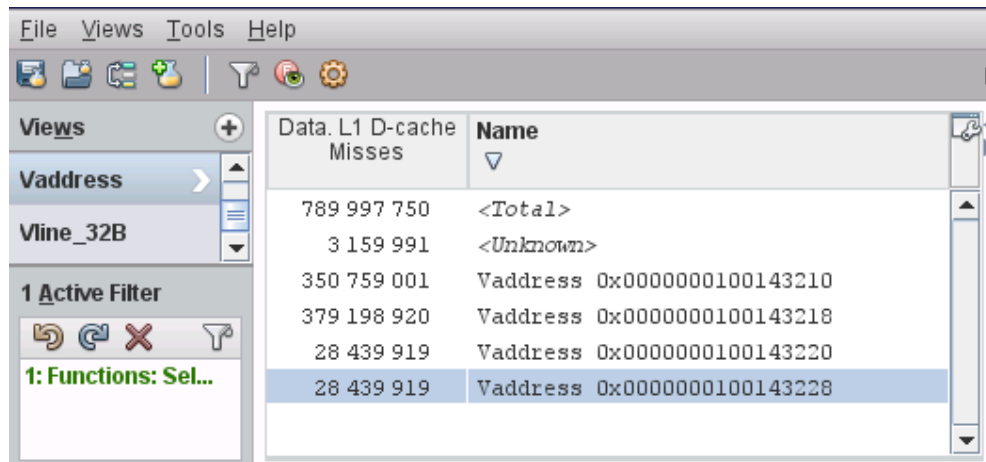
By filtering, you can focus on the performance of the `computeB()` function and the profile events occurring in that function.

2. Click the Settings button in the tool bar or choose Tools > Settings to open the Settings dialog, and select the Views tab in that dialog.



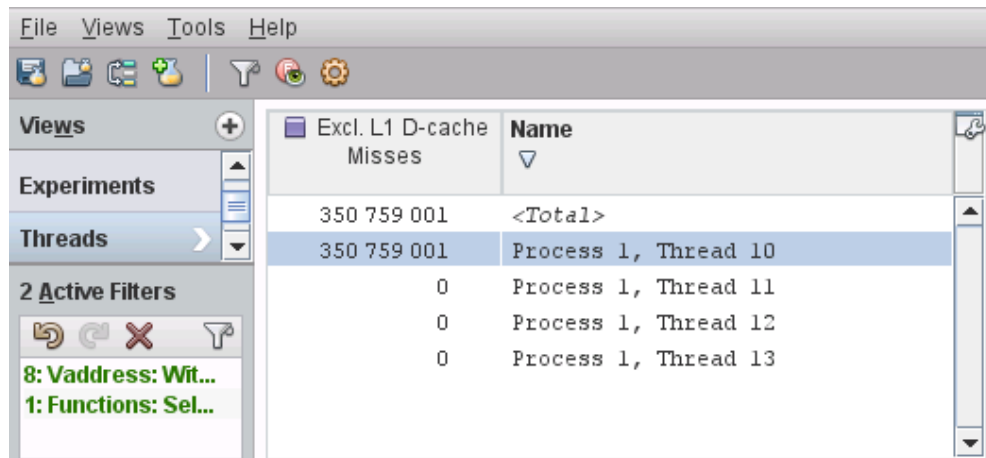
The panel on the right is labeled Memory Objects Views and shows a long list of data views that represent the SPARC T5 machine's memory subsystem structure.

3. Select the check boxes for Vaddress and Vline_32B and click OK.
4. Select the Vaddress view in the Views navigation panel.



In this experiment you can see that there are four different addresses getting the cache misses.

5. Select one of the addresses and then right-click and choose Add Filter: Include only events with the selected item.
6. Select the Threads view.



As you can see in the preceding screen shot, only one thread has cache misses for that address.

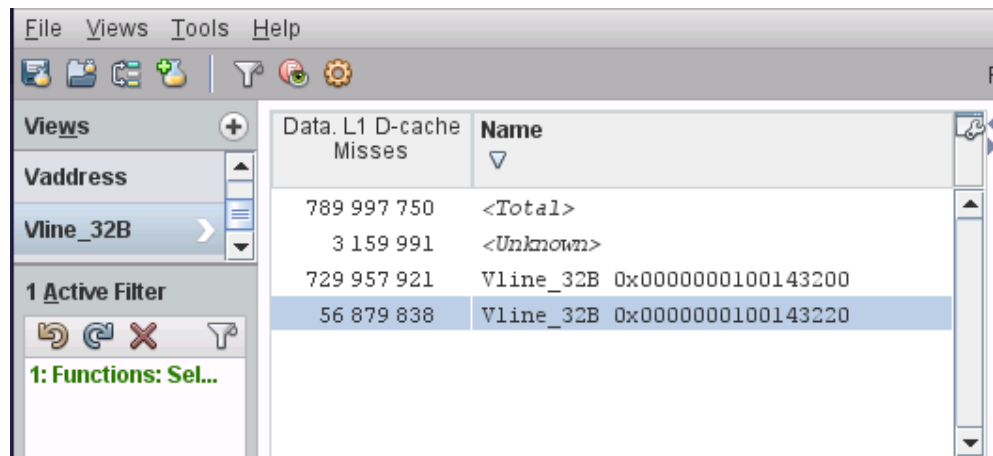
- Remove the address filter by right-clicking in the view and selecting Undo Filter Action from the context menu.

You can alternatively use the Undo Filter Action button in the Active Filters panel to remove the filter.

- Return to the Vaddress view, and select and filter on other addresses and check the associated thread in the Threads view.

By filtering and unfiltering and by switching between the Vaddress and Threads views in this manner, you can confirm that there is a one-to-one relationship between the four threads and the four addresses. That is, the four threads do not share addresses.

- Select the Vline_32B view in the Views navigation panel.



Confirm in the Active Filters panel that there is only the filter active on the function `computerB()`. The filter is shown as `Functions: Selected Functions`. None of the filters on addresses should be active now.

You should see that there are two 32-byte cache lines getting the cache misses of the four threads and their four respective addresses. This confirms that although you saw earlier that the four threads do not share addresses, you see here that they do share cache lines.

False sharing is a very difficult problem to diagnose, and the SPARC T5 chip, along with Oracle Solaris Studio Performance Analyzer enables you to do so.

Synchronization Tracing on a Multithreaded Program

This tutorial includes the following topics.

- [“About the Synchronization Tracing Tutorial” on page 81](#)
- [“Setting Up the mttest Sample Code” on page 82](#)
- [“Collecting Data from mttest for Synchronization Tracing Tutorial” on page 83](#)
- [“Examining the Synchronization Tracing Experiment for mttest” on page 84](#)

About the Synchronization Tracing Tutorial

This tutorial shows how to use Performance Analyzer on a multithreaded program to examine clock profiling and synchronization tracing data.

You use the Overview page to quickly see which performance metrics are highlighted and change which metrics are shown in data views. You use the Functions view, Callers-Callees view, and the Source view to explore the data. The tutorial also shows you how to compare two experiments.

The tutorial helps you understand synchronization tracing data, and explains how to relate it to clock-profiling data.

The data you see in the experiment that you record will be different from that shown here. The experiment used for the screen shots in the tutorial was recorded on a SPARC T5 system running Oracle Solaris 11.2. The data from an x86 system running Oracle Solaris or Linux will be different. Furthermore, data collection is statistical in nature and varies from experiment to experiment, even when run on the same system and OS.

The Performance Analyzer window configuration that you see might not precisely match the screen shots. Performance Analyzer enables you to drag separator bars between components of the window, collapse components, and resize the window. Performance Analyzer records its configuration and uses the same configuration the next time it runs. Many configuration changes were made in the course of capturing the screen shots shown in the tutorial.

About the `mttest` Program

The program `mttest` is a simple program that exercises various synchronization options on dummy data. The program implements a number of different tasks and each task uses the same basic algorithm:

- Queue up a number of work blocks (4, by default).
- Spawn a number of threads to process them (also, 4, by default).
- In each task, use a particular synchronization primitive to control access to the work blocks.
- Process the work for the block, after the synchronization.

Each task uses a different synchronization method. The `mttest` code executes each task in sequence.

About Synchronization Tracing

Synchronization tracing is implemented by interposing on the various library functions for synchronization, such as `mutex_lock()`, `pthread_mutex_lock()`, `sem_wait()`, and so on. Both the `pthread` and Oracle Solaris synchronization calls are traced.

When the target program calls one of these functions, the call is intercepted by the data collector. The current time, the address of the lock, and some other data is captured, and then the interposition routine calls the real library routine. When the real library routine returns, the data collector reads the time again and computes the difference between the end-time and the start-time. If that difference exceeds a user-specified threshold, the event is recorded. If the time does not exceed the threshold, the event is not recorded. In either case, the return value from the real library routine is returned to the caller.

You can set the threshold used to determine whether to record the event by using the `collect` command's `-s` option. If you use Performance Analyzer to collect the experiment, you can specify the threshold as the Minimum Delay for Synchronization Wait Tracing in the Profile Application dialog. You can set the threshold to a number of microseconds or to the keyword `calibrate` or `on`. When you use `calibrate` or `on` the data collector determines the time it takes to acquire an uncontested mutex lock and sets the threshold to five times that value. A specified threshold of `0` or `all` causes all events to be recorded.

In this tutorial, you record synchronization wait tracing in two experiments, with one experiment having a calibrated threshold and one experiment with a zero threshold. Both experiments also include clock profiling.

Setting Up the `mttest` Sample Code

Before You Begin

See the following for information about obtaining the code and setting up your environment.

- [“Getting the Sample Code for the Tutorials” on page 8](#)
- [“Setting Up Your Environment for the Tutorials” on page 9](#)

You might want to go through the introductory tutorial in [“Introduction to C Profiling”](#) first to become familiar with Performance Analyzer.

This tutorial uses the same `mttest` code as the tutorial [“Hardware Counter Profiling on a Multithreaded Program”](#). You should make a separate copy for this tutorial.

1. Copy the contents of the `mttest` directory to your own private working area with the following command:

```
% cp -r SolarisStudioSampleApplications/PerformanceAnalyzer/mttest mydirectory
```

where *mydirectory* is the working directory you are using.

2. Change to that working directory copy.

```
% cd mydirectory/mttest
```

3. Build the target executable.

```
% make clobber (needed only if you ran make in the directory before, but safe in any case)
```

```
% make
```

After you run `make` the directory contains the target application to be used in the tutorial, a C program called `mttest`.

Tip - If you prefer, you can edit the `Makefile` to do the following: use the GNU compilers rather than the default of the Studio compilers; build in 32-bits rather than the default of 64-bits; and add different compiler flags.

Collecting Data from `mttest` for Synchronization Tracing Tutorial

The easiest way to collect the data is to run the following command in the `mttest` directory:

```
% make syncperf
```

The `syncperf` target of the `Makefile` launches the `collect` command twice and creates two experiments.

The two experiments are named `test.1.er` and `test.2.er` and each contains synchronization tracing data and clock profile data. For the first experiment, `collect` uses a calibrated threshold

for recording events by specifying the `-s on` option. For the second experiment, `collect` sets the threshold to zero to record all events by specifying the `-s all` option. In both experiments, clock-profiling is enabled through the `-p on` option.

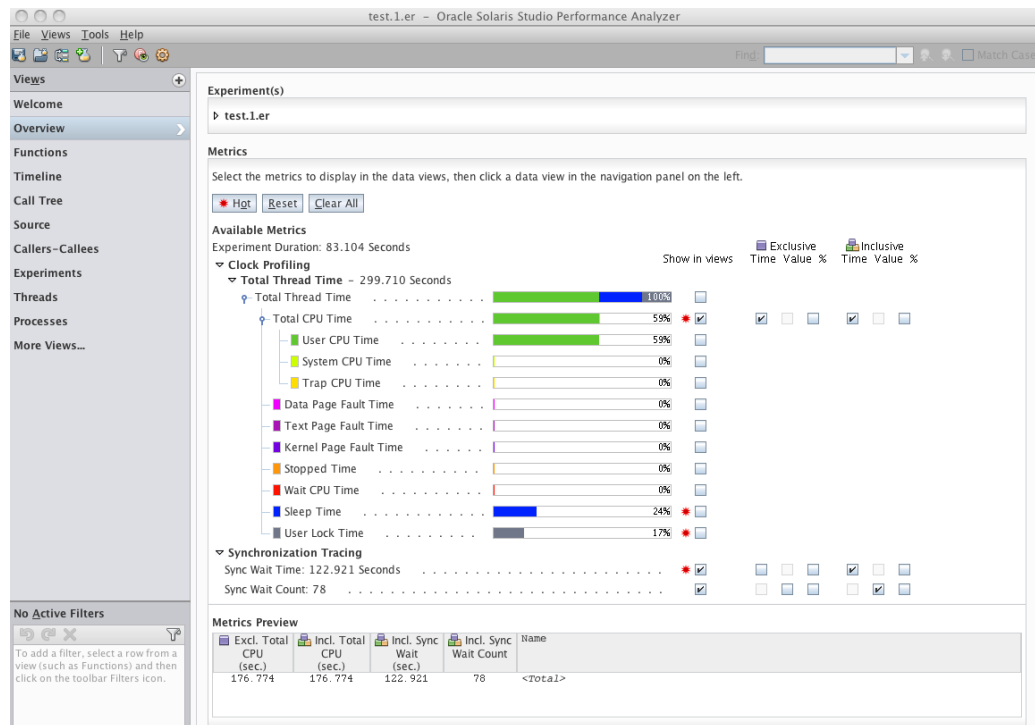
Examining the Synchronization Tracing Experiment for mttest

This section shows how to explore the data in the experiments you created from the `mttest` sample code in the previous section.

Start Performance Analyzer from the `mttest` directory and load the first experiment as follows:

```
% analyzer test.1.er
```

When the experiment opens, Performance Analyzer shows the Overview page.



Clock Profiling metrics are shown first and include colored bars. Most of the thread time is spent in User CPU Time. Some time is spent in Sleep Time or User Lock Time.

Synchronization Tracing metrics are shown in a second group that includes two metrics, Sync Wait Time and Sync Wait Count.

You can explore these metrics and their interpretation in the following sections of the tutorial.

Understanding Synchronization Tracing

This section explores the synchronization tracing data and explains how to relate it to clock-profiling data.

1. Go to the Functions view and sort according to inclusive Total CPU Time by clicking the column header Inclusive Total CPU.
2. Select the `do_work()` function at the top of the list.

The screenshot displays the Oracle Solaris Studio Performance Analyzer interface. The main window shows the Functions view, sorted by Inclusive Total CPU Time. The `do_work` function is selected at the top of the list. The Called-by/Calls panel at the bottom shows that `do_work` is called by `_lwp_start` and `locktest`, and it calls ten other functions.

Excl. Total CPU (sec.)	Incl. Total CPU (sec.)	Incl. Sync Wait (sec.)	Incl. Sync Wait Count	Name
0.610	176.774	50.493	41	<Total>
0.	166.126	50.493	40	do_work
0.	64.565	0.	0	_lwp_start
64.565	64.565	0.	0	cache_trash
0.	31.782	0.	0	compute8
0.	31.782	0.	0	cache_trash_even
0.	31.782	0.	0	cache_trash_odd
0.430	26.569	0.	0	trylock_global
2.402	15.521	0.	0	mutex_trylock
0.	13.119	0.	0	do_exit_critical
13.119	13.119	0.	0	take_deferred_direct
0.	10.647	72.428	38	_start
0.	10.647	0.	0	calladd
2.472	10.647	0.	0	computeF
0.	10.647	72.428	37	locktest
0.	10.647	72.428	38	main
10.637	10.637	0.	0	compute
10.637	10.637	0.	0	computeI
0.	10.637	0.	0	computeE
0.	10.627	2.658	2	sema_global
10.627	10.627	0.	0	nothreads
10.627	10.627	0.	0	computeC
10.627	10.627	0.	0	computeG
0.	10.607	0.	0	fetch_work

Attr. Total CPU (sec.)	do_work is called by	Attr. Total CPU (sec.)	do_work calls
166.126	_lwp_start	64.565	cache_trash
10.647	locktest	26.569	trylock_global
		10.647	calladd
		10.637	nothreads
		10.637	sema_global
		10.627	cond_global
		10.627	cond_timeout_global
		10.627	lock_global
		10.617	lock_none
		10.607	lock_local
		0.	fetch_work

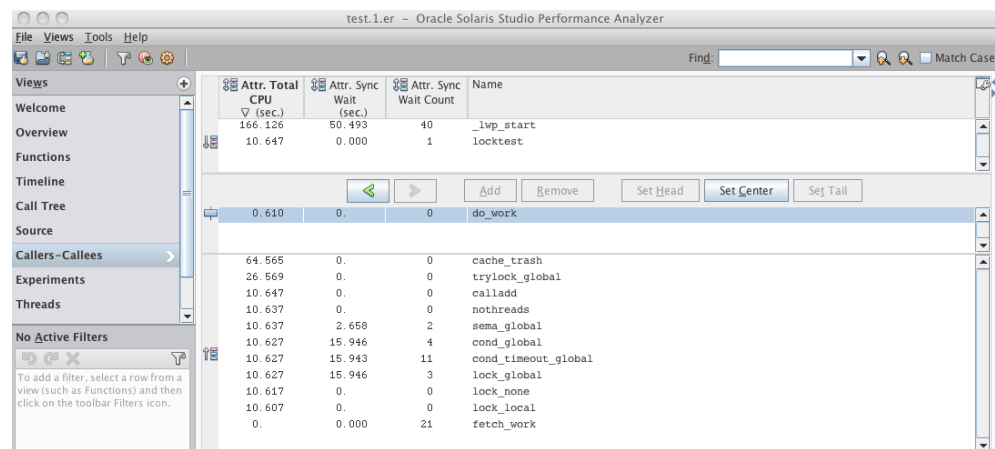
3. Look at the Called-by/Calls panel at the bottom of the Functions view and note that `do_work()` is called from two places, and it calls ten functions.

Most often, `do_work()` is called when a thread to process the data is created, and is shown as called from `_lwp_start()`. In one case, `do_work()` calls one single-threaded task called `nothreads()` after being called from `locktest()`.

The ten functions that `do_work()` calls represent ten different tasks, and each task uses a different synchronization method that the program executed. In some experiments created from `mttest` you might see an eleventh function which uses relatively little time to fetch the work blocks for the other tasks. This function `fetch_work()` is displayed in the Calls panel in the preceding screen shot.

Note that except for the first two of the callees in the Calls panel, all callees show approximately the same amount of time (~10.6 seconds) of Attributed Total CPU.

4. Switch to the Callers-Callees view.



Callers-Callees view shows the same callers and callees as the Called-by/Calls panel, but it also shows the other metrics that were selected in the Overview page, including Attributed Sync Wait Time.

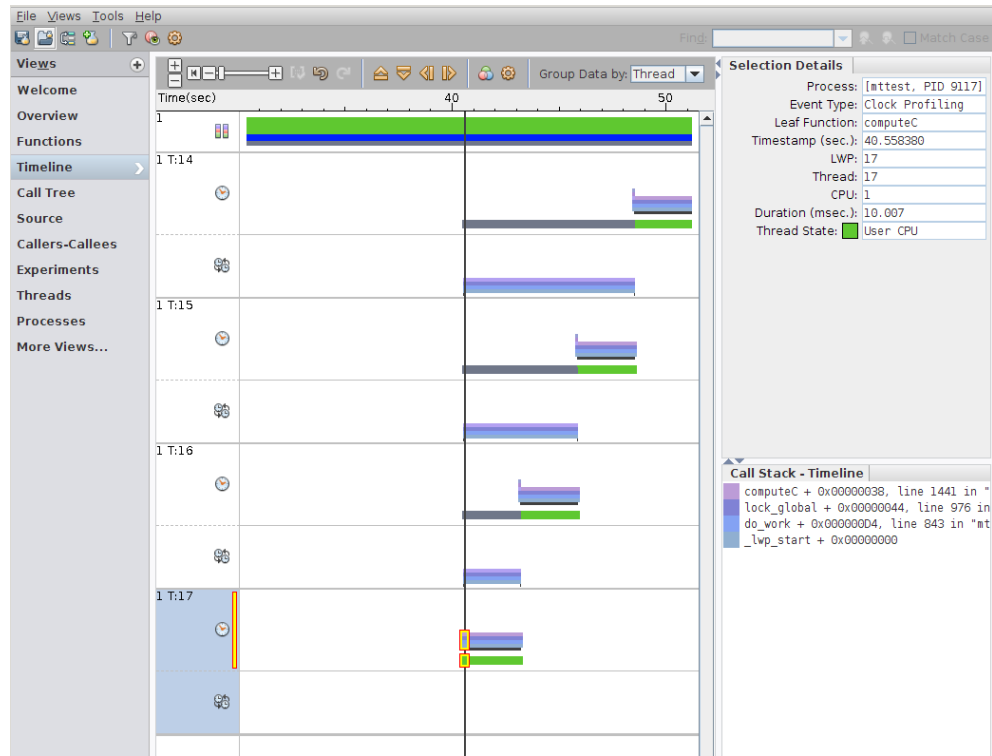
Look for the two functions `lock_global()` and `lock_local()`, and note that they show about the same amount of Attributed Total CPU time, but very different amounts of Attributed Sync Wait Time.


5. Select the `lock_global()` function and switch to Source view.

Incl. CPU (sec.)	Incl. Sync Wait (sec.)	Incl. Sync Wait Count	Source File: /tmp/HYPERUTORIALS/mttest/mttest.c
0.	0.	0	956. /* lock_global: use a global lock to process array's data */
0.	0.	0	957. void
0.	0.	0	958. lock_global(Workblk *array, struct scripttab *k)
0.	0.	0	959. {
0.	0.	0	<Function: lock_global>
0.	0.	0	960. /* acquire the global lock */
0.	0.	0	961.
0.	0.	0	962. #ifdef SOLARIS
0.	0.	0	963. mutex_lock(&global_lock);
0.	0.	0	964. #endif
0.	0.	0	965. #ifdef POSIX
0.	15.946	3	966. pthread_mutex_lock(&global_lock);
0.	0.	0	967. #endif
0.	0.	0	968.
0.	0.	0	969. array->ready = gethrtime();
0.	0.	0	970. array->vready = gethrtime();
0.	0.	0	971.
0.	0.	0	972. array->compute_ready = array->ready;
0.	0.	0	973. array->compute_vready = array->vready;
0.	0.	0	974.
10.627	0.	0	975. /* do some work on the current array */
0.	0.	0	976. k->called_func(array->list[i]);
0.	0.	0	977.
0.	0.	0	978. array->compute_done = gethrtime();
0.	0.	0	979. array->compute_vdone = gethrtime();
0.	0.	0	980.
0.	0.	0	981. /* free the global lock */
0.	0.	0	982. #ifdef SOLARIS
0.	0.	0	983. mutex_unlock(&global_lock);
0.	0.	0	984. #endif
0.	0.	0	985. #ifdef POSIX
0.	0.	0	986. pthread_mutex_unlock(&global_lock);
0.	0.	0	987. #endif
0.	0.	0	988.
0.	0.	0	989. /* make another call to preclude tail-call optimization on the unlock */
0.	0.	0	990. (void) gethrtime();
0.	0.	0	991.
0.	0.	0	992.

Note that all the Sync Wait time is on the line with the call to `pthread_mutex_lock(&global_lock)` which has 0 Total CPU Time. As you might guess from the function names, the four threads executing this task all do their work when they acquire a global lock, which they acquire one by one.

- Go back to the Functions view and select `lock_global()`, then click the Filter icon and select Add Filter: Include only stacks containing the selected functions.
- Select the Timeline view and you should see the four threads.



The first thread to get the lock (T:17 in the screen shot) works for 2.65 seconds, then gives up the lock. You can see that the thread state bar is green for that thread which means all its time was spent in User CPU Time, with none in User Lock Time. Notice also that the second bar for Synchronization Tracing Call Stacks marked with the  show no call stacks for this thread.

The second thread (T:16 in the screen shot) has waited 2.6 seconds in User Lock Time, and then it computes for 2.65 seconds and gives up the lock. The Synchronization Tracing Call Stacks coincide with the User Lock Time.

The third thread (T:15) has waited 5.3 seconds in User Lock Time and it then computes for 2.65 seconds, and gives up the lock.

The last thread (T:14) has waited 7.9 seconds in User Lock Time, and it computes for 2.65 seconds. The total computation was 2.65×4 or ~ 10.6 seconds.

The total synchronization wait was $2.6 + 5.3 + 7.9$ or ~ 15.9 seconds, which you can confirm in the Functions view.

8. Remove the filter by clicking the X in the Active Filters panel.

- Go back to the Functions view, select the function `lock_local()`, and switch to the Source view.


Incl. Total CPU (sec.)	Incl. Sync Wait (sec.)	Incl. Sync Wait Count	Source File: /tmp/MTUTORIALS/mttest/mttest.c
0.	0.	0	1034. /* lock_local: use a local lock to process array's data */
0.	0.	0	1035. void
0.	0.	0	1036. lock_local(Workblk *array, struct scripttab *k)
0.	0.	0	1037. {
			<Function: lock_local>
			1038. /* acquire the local lock */
			1039. #ifdef SOLARIS
			1040. mutex_lock(&(array->lock));
			1041. #endif
0.	0.	0	1042. #ifdef POSIX
			1043. pthread_mutex_lock(&(array->lock));
			1044. #endif
0.	0.	0	1045. array->ready = gettimeofday();
0.	0.	0	1046. array->vready = gethrvtime();
0.	0.	0	1047.
0.	0.	0	1048. array->compute_ready = array->ready;
0.	0.	0	1049. array->compute_vready = array->vready;
			1050.
10.607	0.	0	1051. /* do some work on the current array */
			1052. k->called_func(&array->ls&{0});
0.	0.	0	1053.
0.	0.	0	1054. array->compute_done = gettimeofday();
0.	0.	0	1055. array->compute_vdone = gethrvtime();
			1056.
			1057. /* free the local lock */
			1058. #ifdef SOLARIS
			1059. mutex_unlock(&(array->lock));
			1060. #endif
0.	0.	0	1061. #ifdef POSIX
			1062. pthread_mutex_unlock(&(array->lock));
			1063. #endif
0.	0.	0	1064. /* make another call to preclude tail-call optimization on the unlock */
0.	0.	0	1065. (void) gettimeofday();
			1066. }
			1067.
0.	0.	0	1068. /* cond_global: use a global condition variable to process array's data */
			1069. void
0.	0.	0	1070. cond_global(Workblk *array, struct scripttab *k)
			1071. {

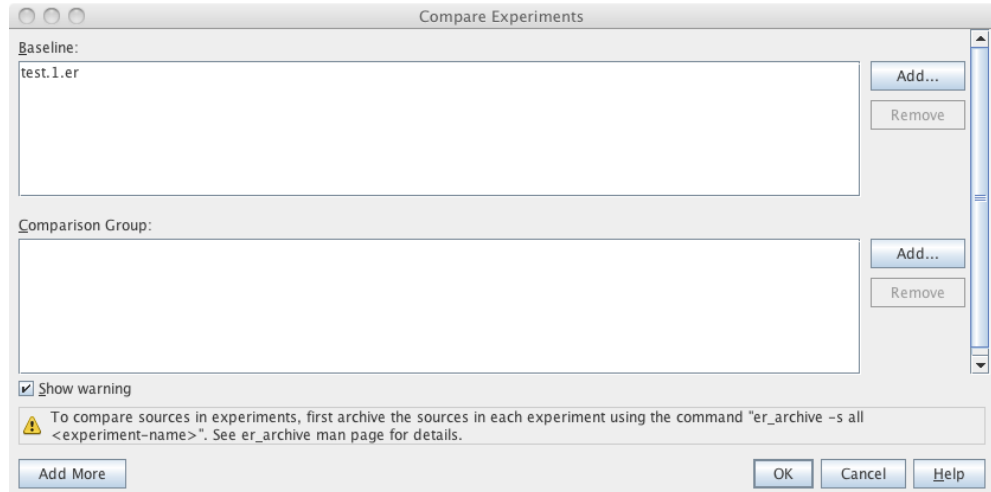
Note that the Sync Wait Time is 0 on the line with the call to `pthread_mutex_lock(&array->lock)`, line 1043 in the screen shot. This is because the lock is local to the workblock, so there is no contention and all four threads compute simultaneously.

The experiment you looked at was recorded with a calibrated threshold. In the next section, you compare to a second experiment which was recorded with zero threshold when you ran the make command.

Comparing Two Experiments with Synchronization Tracing

In this section you compare the two experiments. The `test.1.er` experiment was recorded with a calibrated threshold for recording events, and the `test.2.er` experiment was recorded with zero threshold to include all synchronization events that occurred in the `mttest` program execution.

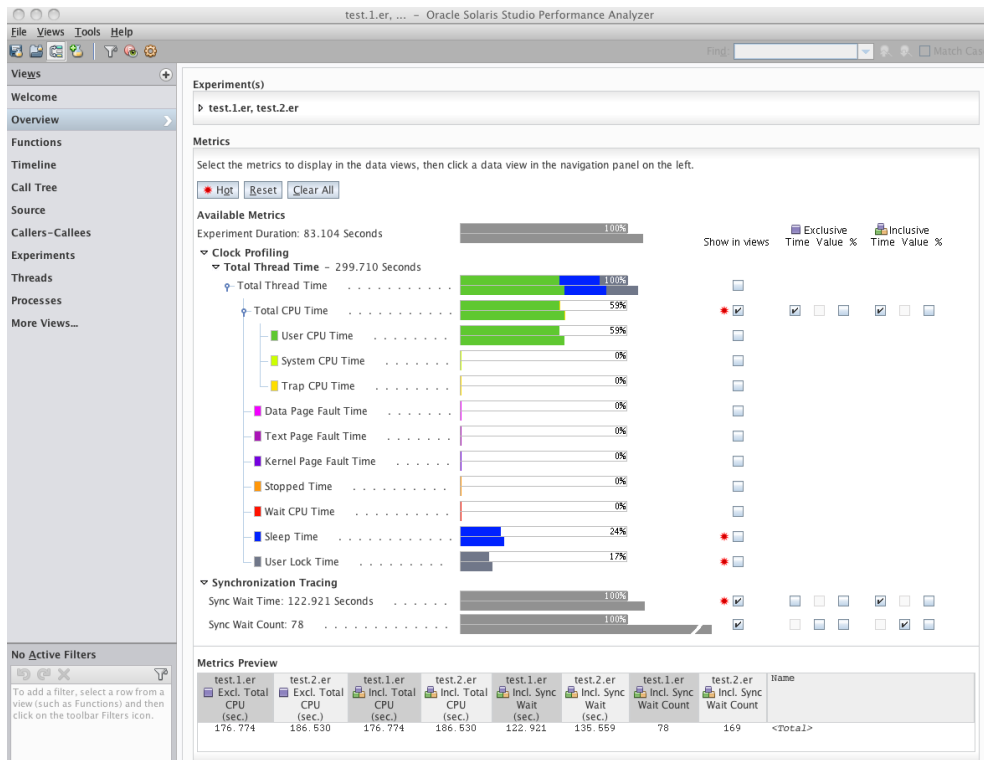
1. Click the Compare Experiments button  on the tool bar or choose File > Compare Experiments.
The Compare Experiments dialog box opens.



The test.1.er experiment that you already have open is listed in the Baseline group. You must create a list of experiments to compare to the baseline experiment in the Comparison Group panel.

In this tutorial, each group contains only one experiment.

2. Click the Add button next to the Comparison Group, and open the test.2.er experiment in the Select Experiment dialog.
3. Click OK in the Compare Experiments dialog to load the second experiment.
The Overview page reopens with the data of both experiments included.



The Clock Profiling metrics display two colored bars for each metric, one bar for each experiment. The data from the test.1.er Baseline experiment is on top.

If you move the mouse cursor over the data bars, popup text shows the data from the Baseline and Comparison groups and difference between them in numbers and percentage.

Note that the Total CPU Time recorded is a little larger in the second experiment, but there are more than twice as many Sync Wait Counts, and about 10% more Sync Wait Time.

4. Switch to the Functions view, and click the column header labeled "test.1.er Incl. Sync Wait Count" to sort the functions by the number of events in the first experiment.

Examining the Synchronization Tracing Experiment for mt test

	test.1.er Excl. Total CPU (sec.)	test.2.er Excl. Total CPU (sec.)	test.1.er Incl. Total CPU (sec.)	test.2.er Incl. Total CPU (sec.)	test.1.er Incl. Sync Wait (sec.)	test.2.er Incl. Sync Wait (sec.)	test.1.er Incl. Sync Wait Count	test.2.er Incl. Sync Wait Count	Name
	176.774	186.630	176.774	186.630	122.921	138.569	78	169	<Total>
	0.610	0.490	176.774	186.520	60.493	56.810	41	131	do_work
	0.	0.	166.126	174.542	60.493	56.810	40	126	_lvp_start
	0.	0.	10.647	11.988	72.428	78.749	38	43	_start
	0.	0.	10.647	11.988	72.428	78.749	38	43	main
	0.	0.	10.647	11.988	72.428	78.749	37	41	locktest
	0.	0.	0.	0.010	72.428	78.749	36	36	pthread_join
	0.	0.	0.	0.010	72.428	78.749	36	36	thread_work
	0.	0.	0.	0.010	15.946	17.939	26	118	pthread_mutex_lock
	0.	0.	0.	0.010	0.000	0.005	21	77	fetch_work
	0.	0.	10.627	11.958	15.943	17.939	11	27	cond_timeout_global
	0.	0.	0.	0.	15.943	17.939	10	11	pthread_cond_timedwait
	0.	0.	10.627	11.958	15.946	17.939	4	15	cond_global
	0.	0.	10.627	11.958	15.946	17.938	3	4	lock_global
	0.	0.	0.	0.	15.946	17.939	3	3	pthread_cond_wait
	0.	0.	0.	0.	2.658	2.989	3	5	sem_wait
	0.	0.	10.637	11.968	2.658	2.989	2	4	sema_global
	0.	0.	0.	0.	0.000	0.000	1	2	resolve_symbols
	0.	0.010	0.	0.010	0.	0.	0	0	__collector_getUserCtx
	0.	0.	0.	0.010	0.	0.	0	0	__collector_write_packet
	0.	0.	0.	0.	0.	0.	0	0	__cond_timedwait
	0.	0.	0.	0.	0.	0.	0	0	__lvp_park
	0.	0.	0.	0.	0.	0.	0	0	__lvp_wait
	0.	0.	0.	0.	0.	0.	0	0	__thrp_join

The function `pthread_mutex_lock()` shows the second largest discrepancy between `test.1.er` and `test.2.er` in the number of events. The largest discrepancy is in `do_work()`, which includes the discrepancies from all the functions it calls, directly or indirectly, including `pthread_mutex_lock()`.

Tip - You can compare the discrepancies even more easily if you change the comparison format. Click the Settings button in the tool bar, select the Formats tab, and choose Deltas for the Comparison Style. After you apply the change, the metrics for `test.2.er` display as the + or - difference from the metrics in `test.1.er`. In the preceding screenshot, the selected `pthread_mutex_lock()` function would show +88 in the `test.2.er` Incl Sync Wait Count column

5. Select Callers-Callees view.

	test.1.er Attr. Total CPU (sec.)	test.2.er Attr. Total CPU (sec.)	test.1.er Attr. Sync Wait (sec.)	test.2.er Attr. Sync Wait (sec.)	test.1.er Attr. Sync Wait Count	test.2.er Attr. Sync Wait Count	Name
	0.	0.010	0.000	0.005	21	77	fetch_work
	0.	0.	15.946	17.938	3	4	lock_global
	0.	0.	0.000	0.000	1	12	cond_global
	0.	0.	0.000	0.000	1	16	cond_timeout_global
	0.	0.	0.	0.	0	4	lock_local
	0.	0.	0.	0.	0	1	resolve_symbols
<input type="button" value="Add"/> <input type="button" value="Remove"/> <input type="button" value="Set Head"/> <input type="button" value="Set Center"/> <input type="button" value="Set Tail"/>							
	0	0	15.946	17.942	26	114	pthread_mutex_lock
	0.	0.010	0.	0.	0	0	__collector_write_packet
	0.	0.	0.	0.	0	0	mutex_lock_queue

Look at two of the callers, `lock_global()` and `lock_local()`.

The `lock_global()` function shows 3 events for Attributed Sync Wait Count in `test.1.er`, but 4 events in `test.2.er`. The reason is that the first thread to acquire the lock in the `test.1.er` was not stalled, so the event was not recorded. In the `test.2.er` experiment the threshold was set to record all events, so even the first thread's lock acquisition was recorded.

Similarly, in the first experiment there were no recorded events for `lock_local()` because there was no contention for the lock. There were 4 events in the second experiment, even though in aggregate they had negligible delays.

