

Oracle® Solaris Studio 12.4: Thread Analyzer User's Guide

ORACLE®

Part No: E37080
December 2014

Copyright © 2007, 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 or 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 © 2007, 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	7
1 What is Thread Analyzer and What Does It Do?	9
Getting Started With Thread Analyzer	9
What is a Data Race?	9
What is a Deadlock?	10
Thread Analyzer Usage Model	10
Usage Model for Detecting Data Races	11
Usage Model for Detecting Deadlocks	12
Usage Model for Detecting Data Races and Deadlocks	13
Thread Analyzer Interface	13
2 Data Race Tutorial	15
Data Race Tutorial Source Files	15
Getting the Data Race Tutorial Source Files	15
Source Code for <code>prime_omp.c</code>	16
Source Code for <code>prime_pthr.c</code>	17
How to Use Thread Analyzer to Find Data Races	19
Instrument the Code	20
Create a Data-Race-Detection Experiment	21
Examine the Data-Race-Detection Experiment	21
Understanding the Experiment Results	23
Data Races in <code>prime_omp.c</code>	23
Data Races in <code>prime_pthr.c</code>	26
Call Stack Traces of Data Races	29
Diagnosing the Cause of a Data Race	30
Check Whether or Not the Data Race is a False Positive	30
Check Whether or Not the Data Race is Benign	31
Fix the Bug, Not the Data Race	31
False Positives	34

User-Defined Synchronizations	34
Memory That is Recycled by Different Threads	35
Benign Data Races	36
A Program for Finding Primes	36
A Program that Verifies Array-Value Types	37
A Program Using Double-Checked Locking	38
3 Deadlock Tutorial	41
About Deadlocks	41
Getting the Deadlock Tutorial Source Files	42
Source Code Listing for <code>din_philo.c</code>	42
The Dining Philosophers Scenario	44
How the Philosophers Can Deadlock	45
Introducing a Sleep Time for Philosopher 1	46
How to Use Thread Analyzer to Find Deadlocks	48
Compile the Source Code	49
Create a Deadlock-Detection Experiment	49
Examine the Deadlock-Detection Experiment	50
Understanding the Deadlock Experiment Results	51
Examining Runs That Deadlock	51
Examining Runs That Complete Despite Deadlock Potential	55
Fixing the Deadlocks and Understanding False Positives	58
Regulating the Philosophers With Tokens	59
An Alternative System of Tokens	63
A APIs Recognized by Thread Analyzer	69
Thread Analyzer User APIs	69
Other Recognized APIs	71
POSIX Thread APIs	71
Oracle Solaris Thread APIs	72
Memory Allocation APIs	72
Memory Operations APIs	73
String Operations APIs	73
Realtime Library APIs	74
Atomic Operations (<code>atomic_ops</code>) APIs	74
OpenMP APIs	74
B Tips for Using Thread Analyzer	75

Compiling the Application	75
Instrumenting the Application for Data Race Detection	75
Running the Application With collect	76
Reporting of Data Races	76

Using This Documentation

- **Overview** – Provides an introduction to the Thread Analyzer tool along with two detailed tutorials. One tutorial focuses on data race detection and the other focuses on deadlock detection. The manual also includes an appendix of APIs recognized by Thread Analyzer and an appendix of useful tips.
- **Audience** – Application developers, system developers, architects, support engineers
- **Required knowledge** – Programming experience, 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>.

What is Thread Analyzer and What Does It Do?

Thread Analyzer is an Oracle Solaris Studio tool that you can use to analyze the execution of a multithreaded program. Thread Analyzer can detect multithreaded programming errors such as data races and deadlocks in code that is written using the POSIX thread API, the Oracle Solaris thread API, OpenMP directives, or a mix of these.

This chapter discusses the following topics:

- [“Getting Started With Thread Analyzer” on page 9](#)
- [“What is a Data Race?” on page 9](#)
- [“What is a Deadlock?” on page 10](#)
- [“Thread Analyzer Usage Model” on page 10](#)
- [“Thread Analyzer Interface” on page 13](#)

Getting Started With Thread Analyzer

Thread Analyzer can show data races and deadlocks in experiments that you can create specifically for examining these types of errors, as explained in this document.

Thread Analyzer is a specialized view of Performance Analyzer that is designed for examining thread analysis experiments. See [“Thread Analyzer Interface” on page 13](#) for more information.

What is a Data Race?

Thread Analyzer detects data races that occur during the execution of a multithreaded process. A data race occurs when all of the following are true:

- Two or more threads in a *single process* access the same memory location concurrently
- At least one of the accesses is for writing
- The threads are not using any mutual exclusive locks to control their accesses to that memory

When these three conditions hold, the order of accesses is non-deterministic, and the computation might give different results from run to run depending on that order. Some data races might be benign (for example, when the memory access is used for a busy-wait), but many data races are bugs in the program.

Thread Analyzer works on a multithreaded program written using the POSIX thread API, Oracle Solaris thread API, OpenMP, or a mix of these.

What is a Deadlock?

Deadlock describes a condition in which two or more threads are blocked forever because they are waiting for each other. There are many causes of deadlocks. Thread Analyzer detects deadlocks that are caused by the inappropriate use of mutual exclusion locks. This type of deadlock is commonly encountered in multithreaded applications.

A process with two or more threads can deadlock when all of the following conditions are true:

- Threads that are already holding locks request new locks
- The requests for new locks are made concurrently
- Two or more threads form a circular chain in which each thread waits for a lock which is held by the next thread in the chain

Here is a simple example of a deadlock condition:

- Thread 1 holds lock A and requests lock B
- Thread 2 holds lock B and requests lock A

A deadlock can be of two types: A *potential* deadlock or an *actual* deadlock. A potential deadlock does not necessarily occur in a given run, but can occur in any execution of the program depending on the scheduling of threads and the timing of lock requests by the threads. An actual deadlock is one that occurs during the execution of the program. An actual deadlock causes the threads involved to hang, but might cause the whole process to hang.

Thread Analyzer Usage Model

The following steps show the process by which you can troubleshoot your multithreaded program with Thread Analyzer.

1. Instrument the program, if doing data race detection.
2. Create a data-race-detection or deadlock-detection experiment.
3. Examine the experiment result and establish whether the multithreaded programming issues revealed by Thread Analyzer are legitimate bugs or benign phenomena.

4. Fix the legitimate bugs and create additional experiments (step 2 above) with varied factors such as different input data, a different number of threads, varied loop schedules or even different hardware. This repetition helps locate non-deterministic problems.

Steps 1 through 3 above are described in the following sections.

Usage Model for Detecting Data Races

You must perform three steps to detect data races:

1. Instrument the code to enable data race detection
2. Create an experiment on the instrumented code
3. Examine the experiment for data races

Instrument the Code for Data Race Detection

To enable data race detection in an application, the code must first be instrumented to monitor memory accesses at runtime, meaning calls to the runtime support library `libtha.so` must be inserted in the code to monitor memory accesses at runtime and determine whether there are any data races.

You can instrument your code at the application source-level during compilation, or at the application binary-level by running an additional tool on the binary.

Source-level instrumentation is done by the compiler when you use a special option. You can also specify the optimization level and other compiler options to use. Source-level instrumentation can result in faster runtime since the compiler can do some analysis and instrument fewer memory accesses.

Binary-level instrumentation is useful when the source code is not available. You might also use binary instrumentation if you have the source code, but cannot compile shared libraries that are used by the application. Binary instrumentation using the `discover` tool instruments the binary as well as all shared libraries as they are opened.

Source-level Instrumentation

To instrument at the source level, compile the source code with the special compiler option:

```
-xinstrument=datarace
```

With this compiler option, the code generated by the compiler will be instrumented for data race detection.

The `-g` compiler option should also be used when building application binaries. This option causes extra data to be generated which enables Thread Analyzer to display source code and line number information when reporting data races.

Binary-level Instrumentation

To instrument at the binary level, you must use the `discover` tool. If the binary is named `a.out`, you can create an instrumented binary `a.out.i` by executing:

```
discover -i datarace -o a.out.i a.out
```

The `discover` tool automatically instruments all shared libraries as they are opened, whether they are statically linked in the program or opened dynamically by `dlopen()`. By default, instrumented copies of libraries are cached in the directory `$HOME/SUNW_Bit_Cache`.

Some useful `discover` command line options are shown below. See the `discover(1)` man page for details.

<code>-o file</code>	Output the instrumented binary to the specified file name
<code>-N lib</code>	Do not instrument the specified library
<code>-T</code>	Do not instrument any libraries
<code>-D dir</code>	Change the cache directory to <i>dir</i>

Create an Experiment on the Instrumented Application

To create a data-race-detection experiment, use the `collect` command with the `-r race` flag to run the application and collect experiment data during the execution of the process. When you use the `-r race` option, the collected data includes pairs of data accesses that constitute a race.

Examine the Experiment for Data Races

You can examine the data-race-detection experiment with the `tha` command, which starts the Thread Analyzer graphical user interface. You can also use the `er_print` command-line interface.

Usage Model for Detecting Deadlocks

Two steps are involved in detecting deadlocks:

1. Create a deadlock-detection experiment.
2. Examine the experiment for deadlocks.

Create an Experiment for Detecting Deadlocks

To create a deadlock-detection experiment, use the `collect` command with the `-r deadlock` flag to run the application and collect experiment data during the execution of the process. When you use the `-r deadlock` option, the collected data includes lock holds and lock requests that form a circular chain.

Examine the Experiment for Deadlocks

You can examine the deadlock-detection experiment with the `tha` command, which starts the Thread Analyzer graphical user interface. You can also use the `er_print` command-line interface.

Usage Model for Detecting Data Races and Deadlocks

If you want to detect data races and deadlocks at the same time, follow the three steps in [“Usage Model for Detecting Data Races” on page 11](#) for detecting data races, but use the `collect` command with the `-r race,deadlock` flag to run the application. The experiment will contain both race-detection and deadlock-detection data.

Thread Analyzer Interface

You can start Thread Analyzer by using the `tha` command.

The Thread Analyzer interface is the Performance Analyzer (`analyzer`) interface that is streamlined for multithreaded program analysis. Instead of the usual Performance Analyzer views, you see the Races, Deadlocks, and Dual Source views. If you use Performance Analyzer to look at multithreaded program experiments you see the traditional Performance Analyzer views such as Functions, Callers-Callees, Disassembly, along with the views for data races and deadlocks.

Data Race Tutorial

The following is a detailed tutorial on how to detect and fix data races with Thread Analyzer. The tutorial is divided into the following sections:

- “Data Race Tutorial Source Files” on page 15
- “How to Use Thread Analyzer to Find Data Races” on page 19
- “Understanding the Experiment Results” on page 23
- “Diagnosing the Cause of a Data Race” on page 30
- “False Positives” on page 34
- “Benign Data Races” on page 36

Data Race Tutorial Source Files

This tutorial relies on two programs, both of which contain data races:

- The first program finds prime numbers. It is written with C and is parallelized with OpenMP directives. The source file is called `prime_omp.c`.
- The second program also finds prime numbers and is also written with C. However, it is parallelized with POSIX threads instead of OpenMP directives. The source file is called `prime_pthr.c`.

Getting the Data Race Tutorial Source Files

You can download the source files used in this tutorial from the [Download area \(http://www.oracle.com/technetwork/server-storage/solarisstudio/downloads/index.html\)](http://www.oracle.com/technetwork/server-storage/solarisstudio/downloads/index.html) of the Oracle Solaris Studio developer portal.

After you download and unpack the sample files, you can find the samples in the `SolarisStudioSampleApplications/ThreadAnalyzer` directory. The samples are located in the `prime_omp` and `prime_pthr` subdirectories. Each sample directory includes a `Makefile` and a `DEMO` file of instructions, but this tutorial does not follow those instructions or use the `Makefile`. Instead, you are instructed to execute commands individually.

To follow the tutorial, you can copy the `prime_omp.c` and `prime_pthr.c` files from the samples directories to a different directory, or you can create your own files and copy the code from the following code listings.

Source Code for `prime_omp.c`

This section shows the source code for `prime_omp.c` as follows:

```
1  /*
2  * Copyright (c) 2006, 2010, Oracle and/or its affiliates. All Rights Reserved.
3  * @(#)prime_omp.c 1.3 (Oracle) 10/03/26
4  */
5
6  #include <stdio.h>
7  #include <math.h>
8  #include <omp.h>
9
10 #define THREADS 4
11 #define N 10000
12
13 int primes[N];
14 int pflag[N];
15
16 int is_prime(int v)
17 {
18     int i;
19     int bound = floor(sqrt(v)) + 1;
20
21     for (i = 2; i < bound; i++) {
22         /* no need to check against known composites */
23         if (!pflag[i])
24             continue;
25         if (v % i == 0) {
26             pflag[v] = 0;
27             return 0;
28         }
29     }
30     return (v > 1);
31 }
32
33 int main(int argn, char **argv)
34 {
35     int i;
36     int total = 0;
37
38     #ifdef _OPENMP
39         omp_set_dynamic(0);
40         omp_set_num_threads(THREADS);
41     #endif
42
43     for (i = 0; i < N; i++) {
```



```
44     pflag[i] = 1;
45 }
46
47 #pragma omp parallel for
48 for (i = 2; i < N; i++) {
49     if ( is_prime(i) ) {
50         primes[total] = i;
51         total++;
52     }
53 }
54
55 printf("Number of prime numbers between 2 and %d: %d\n",
56        N, total);
57
58 return 0;
59 }
```

Source Code for prime_pthr.c

This section shows source code for prime_pthr.c as follows:

```
1  /*
2  * Copyright (c) 2006, 2010, Oracle and/or its affiliates. All Rights Reserved.
3  * @(#)prime_pthr.c 1.4 (Oracle) 10/03/26
4  */
5
6  #include <stdio.h>
7  #include <math.h>
8  #include <pthread.h>
9
10 #define THREADS 4
11 #define N 10000
12
13 int primes[N];
14 int pflag[N];
15 int total = 0;
16
17 int is_prime(int v)
18 {
19     int i;
20     int bound = floor(sqrt(v)) + 1;
21
22     for (i = 2; i < bound; i++) {
23         /* no need to check against known composites */
24         if (!pflag[i])
25             continue;
26         if (v % i == 0) {
27             pflag[v] = 0;
28             return 0;
29         }
30     }
31     return (v > 1);
```

```
32 }
33
34 void *work(void *arg)
35 {
36     int start;
37     int end;
38     int i;
39
40     start = (N/THREADS) * *(int *)arg;
41     end = start + N/THREADS;
42     for (i = start; i < end; i++) {
43         if ( is_prime(i) ) {
44             primes[total] = i;
45             total++;
46         }
47     }
48     return NULL;
49 }
50
51 int main(int argn, char **argv)
52 {
53     int i;
54     pthread_t tids[THREADS-1];
55
56     for (i = 0; i < N; i++) {
57         pflag[i] = 1;
58     }
59
60     for (i = 0; i < THREADS-1; i++) {
61         pthread_create(&tids[i], NULL, work, (void *)&i);
62     }
63
64     i = THREADS-1;
65     work((void *)&i);
66
67     for (i = 0; i < THREADS-1; i++) {
68         pthread_join(tids[i], NULL);
69     }
70
71     printf("Number of prime numbers between 2 and %d: %d\n",
72           N, total);
73
74     return 0;
75 }
```

Effect of Data Races in `prime_omp.c` and `prime_pthr.c`

When there is a race condition in the code, the order of memory accesses is non-deterministic so the computation gives different results from run to run. The correct answer in the `prime_omp` and `prime_pthr` programs is 1229.

You can compile and run the examples so you can see that the execution of `prime_omp` or `prime_pthr` produces incorrect and inconsistent results because of the data races in the code.

In the following example, type the commands at the prompt to compile and run the `prime_omp` program:

```
% cc -xopenmp=noopt -o prime_omp prime_omp.c -lm
%
% ./prime_omp
Number of prime numbers between 2 and 10000: 1229
% ./prime_omp
Number of prime numbers between 2 and 10000: 1228
% ./prime_omp
Number of prime numbers between 2 and 10000: 1180
```

In the following example, type the commands at the prompt to compile and run the `prime_pthr` program:

```
% cc -mt -o prime_pthr prime_pthr.c -lm
%
% ./prime_pthr
Number of prime numbers between 2 and 10000: 1140
% ./prime_pthr
Number of prime numbers between 2 and 10000: 1122
% ./prime_pthr
Number of prime numbers between 2 and 10000: 1141
```

Notice the inconsistency of the results of the three runs of each program. You might need to run the programs more than three times to see inconsistent results.

Next you instrument the code and create experiments so you can find where the data races are occurring.

How to Use Thread Analyzer to Find Data Races

Thread Analyzer follows the same “collect-analyze” model that the Oracle Solaris Studio Performance Analyzer uses.

There are three steps involved in using Thread Analyzer:

1. [“Instrument the Code” on page 20](#)
2. [“Create a Data-Race-Detection Experiment” on page 21](#)
3. [“Examine the Data-Race-Detection Experiment” on page 21](#)

Instrument the Code

In order to enable data race detection in a program, the code must first be instrumented to monitor memory accesses at runtime. The instrumentation can be done on application source code or on the application binary. The tutorial shows how to use both methods of instrumenting the programs.

To Instrument Source Code

To instrument the source code, you must compile the application with the special compiler option `-xinstrument=datarace`. This option instructs the compiler to instrument the generated code for data race detection.

Add the `-xinstrument=datarace` compiler option to the existing set of options you use to compile your program.

Note - Be sure to also specify the `-g` option when you compile your program with `-xinstrument=datarace` to generate additional information to enable Thread Analyzer's full capabilities. Do not specify a high level of optimization when compiling your program for race detection. Compile an OpenMP program with `-xopenmp=noopt`. The information reported, such as line numbers and call stacks, might be incorrect when a high optimization level is used.

You can use the following commands for instrumenting the source code for the tutorial:

```
% cc -xinstrument=datarace -g -xopenmp=noopt -o prime_omp_inst prime_omp.c -lm
% cc -xinstrument=datarace -g -o prime_pthr_inst prime_pthr.c -lm
```

Notice that in the example, the output file is specified with `_inst` at the end so that you can tell that the binary is the instrumented binary. This is not required.

To Instrument Binary Code

To instrument a program's binary code instead of the source code, you need to use the `discover` tool, which is included in Oracle Solaris Studio and is documented in the `discover(1)` man page and [“Oracle Solaris Studio 12.4: Discover and Uncover User’s Guide”](#).

For the tutorial examples, type the following command to compile the code:

```
% cc -xopenmp=noopt -g -o prime_omp prime_omp.c -lm
% cc -g -O2 -o prime_pthr prime_pthr.c -lm
```

Then run `discover` on the `prime_omp` and `prime_pthr` optimized binaries that you created:

```
% discover -i datarace -o prime_omp_disc prime_omp
```

```
% discover -i datarace -o prime_pthr_disc prime_pthr
```

These commands create instrumented binaries, `prime_omp_disc` and `prime_pthr_disc` that you can use with `collect` to create experiments that you can examine with Thread Analyzer.

Create a Data-Race-Detection Experiment

Use the `collect` command with the `-r race` flag to run the program and create a data-race-detection experiment during the execution of the process. For OpenMP programs, make sure that the number of threads used is larger than one. In the tutorial samples, four threads are used.

To create experiments from the binaries that you created by instrumenting the source code:

```
% collect -r race -o prime_omp_inst.er prime_omp_inst
```

```
% collect -r race -o prime_pthr_inst.er prime_pthr_inst
```

To create experiments from the binaries that you created by using the `discover` tool:

```
% collect -r race -o prime_omp_disc.er prime_omp_disc
```

```
% collect -r race -o prime_pthr_disc.er prime_pthr_disc
```

To increase the likelihood of detecting data races, it is recommended that you create several data-race-detection experiments using `collect` with the `-r race` flag. Use a different number of threads and different input data for each experiment.

For example, in `prime_omp.c`, the number of threads is set by the following line:

```
#define THREADS 4
```

The number of threads can be changed by changing 4 in the above to some other integer larger than 1, for example 8.

The following line in `prime_omp.c` limits the program to look for prime numbers between 2 and 3000:

```
#define N 3000
```

You can provide different input data by changing the value of `N` to make the program do more or less work.

Examine the Data-Race-Detection Experiment

You can examine a data-race-detection experiment with Thread Analyzer, Performance Analyzer, or the `er_print` utility. Both Thread Analyzer and Performance Analyzer present a GUI interface; Thread Analyzer presents a simplified set of default views, but is otherwise identical to Performance Analyzer.

Using Thread Analyzer to View the Data Race Experiment

To start Thread Analyzer, type the following command:

```
% tha
```

When you first start Thread Analyzer you see the Welcome screen.

Thread Analyzer has a menu bar, a tool bar, and vertical navigation bar on the left that enables you to select data views.

The following data views are shown by default:

- Overview screen shows a metrics overview of the loaded experiments.
- Races view shows a list of data races detected in the program, and associated call stack traces. This view is selected by default. When you select an item in the Races view, the Race Details window shows detailed information about the data race or call stack trace selected.
- Dual Source view shows the two source locations corresponding to the two accesses of a selected data race. The source line where a data race access occurred is highlighted.
- Experiments view shows the load objects in the experiment, and lists error and warning messages.

You can choose to see other views with the More Views options menu.

Using `er_print` to View the Data Race Experiment

The `er_print` utility presents a command-line interface. You can use the `er_print` utility in an interactive session and specify sub-commands during the session. You can also use command-line options to specify sub-commands non-interactively.

The following sub-commands are useful for examining races with the `er_print` utility:

- `-races`
This reports any data races revealed in the experiment. Specify `races` at the `(er_print)` prompt or `-races` on the `er_print` command line.
- `-rdetail race_id`
This displays detailed information about the data race with the specified `race_id`. Specify `rdetail` at the `(er_print)` prompt or `-rdetail` on the `er_print` command line. If the specified `race_id` is `all`, then detailed information about all data races will be displayed. Otherwise, specify a single race number such as `1` for the first data race.
- `-header`
This displays descriptive information about the experiment, and reports any errors or warnings. Specify `header` at the `(er_print)` prompt or `-header` on the command line.

Refer to the `collect(1)`, `tha(1)`, `analyzer(1)`, and `er_print(1)` man pages for more information.

Understanding the Experiment Results

This section shows how to use both the `er_print` command line and Thread Analyzer to display the following information about each detected data race:

- The unique ID of the data race.
- The virtual address, `Vaddr`, associated with the data race. If there is more than one virtual address, then the label `Multiple Addresses` is displayed in parentheses.
- The memory accesses to the virtual address, `Vaddr` by two different threads. The type of the access (read or write) is shown, as well as the function, offset, and line number in the source code where the access occurred.
- The total number of call stack traces associated with the data race. Each trace refers to the pair of thread call stacks at the time the two data race accesses occurred.

If you are using Thread Analyzer, the two call stacks are displayed in the Race Details window when you select an individual call stack trace in the Races view. If you are using the `er_print` utility, the two call stacks will be displayed by the `rdetail` command.

Data Races in `prime_omp.c`

To examine data races in `prime_omp.c`, you can use one of the experiments you created in [“Create a Data-Race-Detection Experiment” on page 21](#).

To show the data race information in the `prime_omp_instr.er` experiment with `er_print`, type the following command.

```
% er_print prime_omp_instr.er
```

At the (`er_print`) prompt, type `races` to see output similar to the following:

```
(er_print) races

Total Races: 3 Experiment: prime_omp_instr.er

Race #1, Vaddr: 0x219c8
  Access 1: Write, line 25 in "prime_omp.c",
             is_prime
  Access 2: Read,  line 22 in "prime_omp.c",
             is_prime
Total Callstack Traces: 1
```

```
Race #2, Vaddr: (Multiple Addresses)
  Access 1: Write, line 49 in "prime_omp.c",
             main
  Access 2: Write, line 49 in "prime_omp.c",
             main
  Total Callstack Traces: 1
```

```
Race #3, Vaddr: 0xffbfff604
  Access 1: Write, line 50 in "prime_omp.c",
             main
  Access 2: Write, line 50 in "prime_omp.c",
             main
  Total Callstack Traces: 1
```

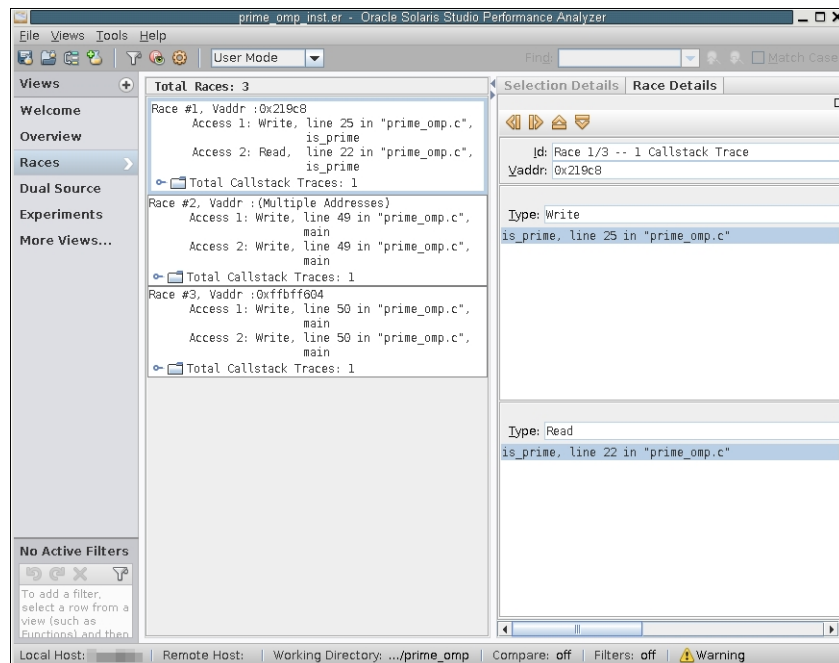
Three data races occurred during this particular run of the program.

To open the `prime_omp_inst.er` experiment in Thread Analyzer, type the following command:

```
% tha prime_omp_inst.er
```

The following screen shot shows the races that were detected in `prime_omp.c` as displayed by Thread Analyzer.

FIGURE 2-1 Data Races Detected in `prime_omp.c`



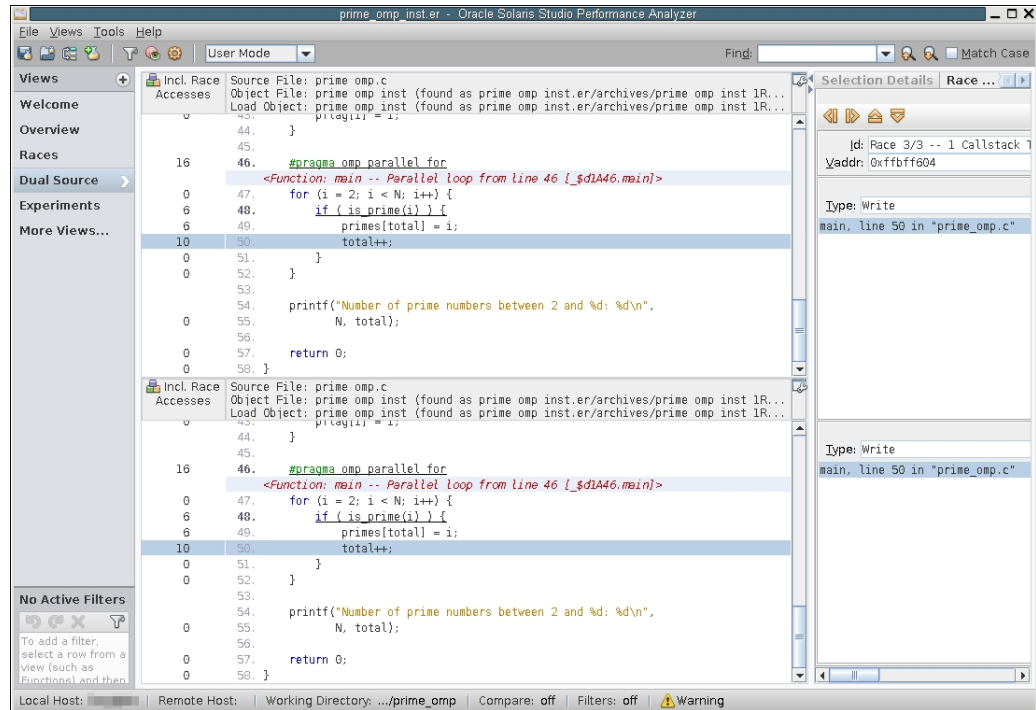
Three data races are shown in `prime_omp.c`:

- Race #1 shows a race between a write in the function `is_prime` on line 25 and a read in the same function on line 22. If you look at the source code you can see that on these lines, the `pf_lag[]` array is being accessed. In Thread Analyzer, you can click the Dual Source view to easily see the source code at both line numbers along with metrics showing the number of race accesses on the affected lines of code.
- Race #2 shows a race between two writes to line 49 of the `main` function. Click the Dual Source view to see that there are multiple attempts to access the value of the `primes []` array in line 49.
- Race #3 shows a race between two writes to line 50 of the `main` function. Click the Dual Source view to see that there are multiple attempts to access the value of the `primes []` array in line 50.

Race #3 represents a group of data races that occur in different elements of the array `primes[]`. This is indicated by the `Vaddr` specified as `Multiple Addresses`.

The Dual Source view in Thread Analyzer enables you to see the two source locations associated with a data race at the same time. For example, select Race #3 for `prime_pthr.c` in the Races view and then click on the Dual Source view. You will see something similar to the following.

FIGURE 2-2 Source Code of Data Races Detected in `prime_omp.c`



Tip - You might need to drag the mouse on the header of each source panel to see the Race Accesses metrics in the left margin of the Dual Source view.

Data Races in `prime_pthr.c`

To examine data races in `prime_pthr.c`, you can use one of the experiments you created in [“Create a Data-Race-Detection Experiment”](#) on page 21.

To show the data race information in the `prime_pthr_inst.er` experiment with `er_print`, type the following command:

```
% er_print prime_pthr_inst.er
```

At the `(er_print)` prompt, type `races` to see output similar to the following:

```
(er_print) races
```

```
Total Races: 5 Experiment: prime_pthr_inst.er

Race #1, Vaddr: 0x28c28
  Access 1: Write, line 26 in "prime_pthr.c",
            is_prime + 0x0000022C
  Access 2: Write, line 26 in "prime_pthr.c",
            is_prime + 0x0000022C
  Total Callstack Traces: 2

Race #2, Vaddr: (Multiple Addresses)
  Access 1: Read, line 23 in "prime_pthr.c",
            is_prime + 0x00000E4
  Access 2: Write, line 26 in "prime_pthr.c",
            is_prime + 0x0000022C
  Total Callstack Traces: 2

Race #3, Vaddr: 0xffbfff5bc
  Access 1: Write, line 59 in "prime_pthr.c",
            main + 0x000001F4
  Access 2: Read, line 39 in "prime_pthr.c",
            work + 0x0000006C
  Total Callstack Traces: 1

Race #4, Vaddr: 0x216f0
  Access 1: Write, line 44 in "prime_pthr.c",
            work + 0x00000174
  Access 2: Write, line 44 in "prime_pthr.c",
            work + 0x00000174
  Total Callstack Traces: 2

Race #5, Vaddr: (Multiple Addresses)
  Access 1: Write, line 43 in "prime_pthr.c",
            work + 0x0000012C
  Access 2: Write, line 43 in "prime_pthr.c",
            work + 0x0000012C
  Total Callstack Traces: 1
```

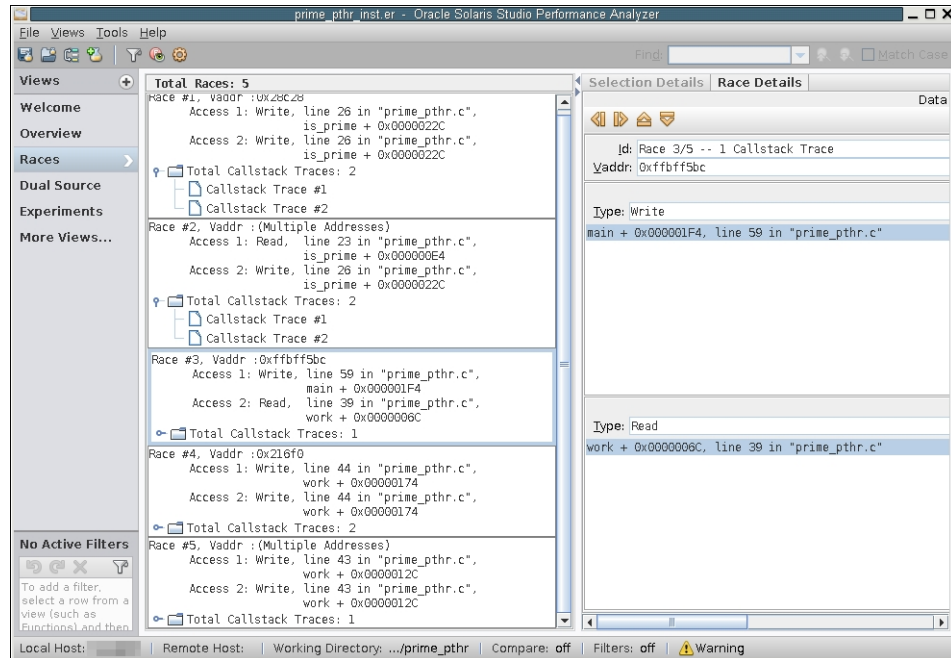
Five data races occurred during this particular run of the program.

To open the `prime_pthr_inst.er` experiment in Thread Analyzer, type the following command:

```
% tha prime_pthr_inst.er
```

The following screen shot shows the races detected in `prime_pthr.c` as displayed by Thread Analyzer. Notice that they are the same as the races shown by `er_print`.

FIGURE 2-3 Data Races Detected in `prime_pthr.c`

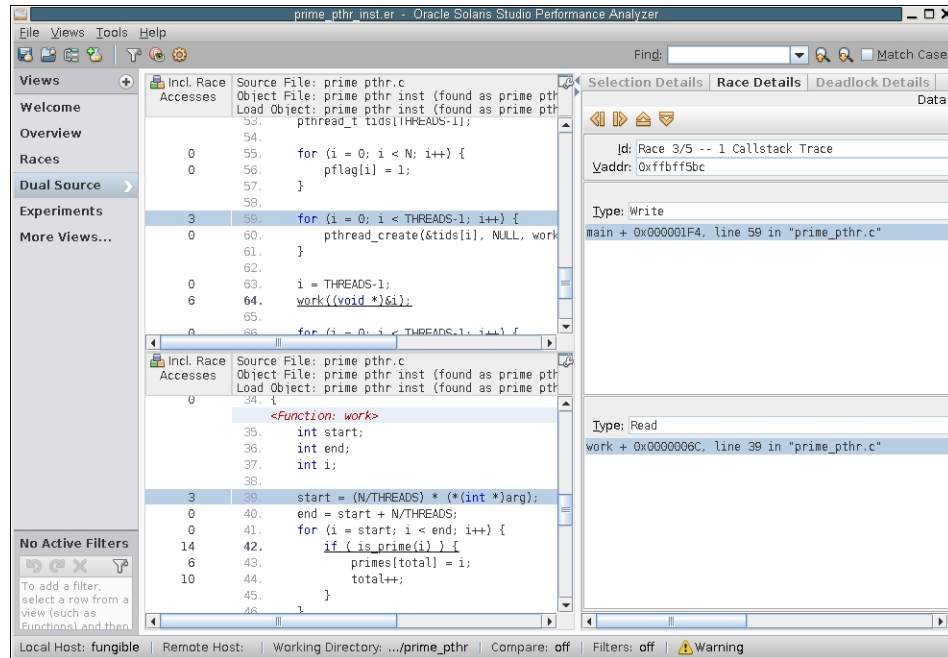


Five data races are shown in `prime_pthr.c`:

- Race #1 is a data race between a write to the `pflag[]` array in function `is_prime` on line 26 and another write to `pflag[]` on the same line.
- Race #2 is a data race between a read on line 23 for `pflag[]` and a write on to `pflag[]` array in function `is_prime` on line 26.
- Race #3 is a data race between a write on line 59 to the memory location named `i` in `main()` and a read on line 39 from the same memory location (named `*arg` in `work()`).
- Race #4 is a data race between a write to `primes[total]` on line 44 and another write to `primes[total]` the same line.
- Race #5 is a data race between a write to `total` on line 46 and another write to `total` on the same line.

If you select Race #3 and then click the Dual Source view, you see the two source locations, similar to the following screen shot.

FIGURE 2-4 Source Code Details of a Data Race

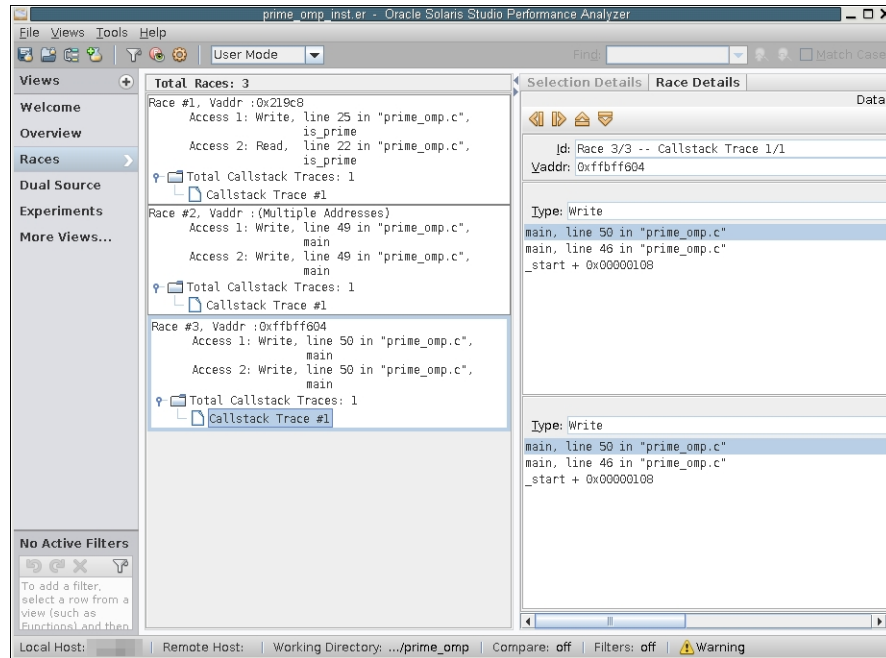


The first access for Race #3 is at line 59 and is shown in the top panel. The second access is at line 40 and is shown in the bottom panel. The Race Accesses metric is highlighted at the left of the source code. This metric gives a count of the number of times a data race access was reported on that line.

Call Stack Traces of Data Races

Each data race listed in the Races view of Thread Analyzer also has one or more associated Call Stack Traces. The call stacks show the execution paths through the code that lead to a data race. When you click a Call Stack Trace, the Race Details window in the right panel shows the function calls that lead to the data race.

FIGURE 2-5 Races View with Call Stack Traces for prime_omp.c



Diagnosing the Cause of a Data Race

This section provides a basic strategy to diagnosing the cause of data races.

Check Whether or Not the Data Race is a False Positive

A false positive data race is a data race that is reported by Thread Analyzer, but has actually not occurred. In other words, it is a "false alarm". Thread Analyzer tries to reduce the number of false positives reported. However, there are cases where the tool is not able to do a precise job and might report false positive data races.

You can ignore a false positive data race because it is not a genuine data race and, therefore, does not affect the behavior of the program.

See [“False Positives” on page 34](#) for some examples of false positive data races. For information on how to remove false positive data races from the report, see [“Thread Analyzer User APIs” on page 69](#).

Check Whether or Not the Data Race is Benign

A benign data race is an intentional data race whose existence does not affect the correctness of the program.

Some multithreaded applications intentionally use code that might cause data races. Since the data races are there by design, no fix is required. In some cases, however, it is quite tricky to get such codes to run correctly. These data races should be reviewed carefully.

See [“False Positives” on page 34](#) for more detailed information about benign races.

Fix the Bug, Not the Data Race

Thread Analyzer can help find data races in the program, but it cannot automatically find bugs in the program nor suggest ways to fix the data races found. A data race might have been introduced by a bug. It is important to find and fix the bug. Merely removing the data race is not the right approach, and could make further debugging even more difficult.

Fixing Bugs in `prime_omp.c`

This section describes how to fix the bug in `prime_omp.c`. See [“Source Code for `prime_omp.c`” on page 16](#) for a complete file listing.

Move lines 50 and 51 into a `critical` section in order to remove the data race on elements of the array `primes[]`.

```

47     #pragma omp parallel for
48     for (i = 2; i < N; i++) {
49         if ( is_prime(i) ) {
50             #pragma omp critical
51                 {
52                     primes[total] = i;
53                     total++;
54                 }
55         }
56     }

```

You could also move lines 50 and 51 into two `critical` sections as follows, but this change fails to correct the program:

```

47     #pragma omp parallel for

```

```
48     for (i = 2; i < N; i++) {
49         if ( is_prime(i) ) {
                    #pragma omp critical
                    {
50             primes[total] = i;
                    }
                    #pragma omp critical
                    {
51             total++;
                    }
52     }
53 }
```

The critical sections around lines 50 and 51 get rid of the data race because the threads are using mutual exclusive locks to control their accesses to the `primes[]` array. However, the program is still incorrect. Two threads might update the same element of `primes[]` using the same value of `total`, and some elements of `primes[]` might not be assigned a value at all.

The second data race, between a read from `pflag[]` from line 23 and a write to `pflag[]` from line 26, is actually a benign race because it does not lead to incorrect results. It is not essential to fix benign data races.

Fixing Bugs in `prime_pthr.c`

This section describes how to fix the bug in `prime_pthr.c`. See [“Source Code for `prime_pthr.c`” on page 17](#) for a complete file listing.

Use a single mutex to remove the data race on `prime[]` at line 44, as well as the data race on `total` at line 45.

The data race between the write to `i` on line 60 and the read from the same memory location (named `*arg`) on line 40, as well as the data race on `pflag[]` on line 27, reveal a problem in the shared access to the variable `i` by different threads. The initial thread in `prime_pthr.c` creates the child threads in a loop in lines 60-62, and dispatches them to work on the function `work()`. The loop index `i` is passed to `work()` by address. Since all threads access the same memory location for `i`, the value of `i` for each thread will not remain unique, but will change as the initial thread increments the loop index. As different threads use the same value of `i`, the data races occur. One way to fix the problem is to pass `i` to `work()` by value, instead of by address.

The following code is the corrected version of `prime_pthr.c`:

```
1  /*
2   * Copyright (c) 2006, 2010, Oracle and/or its affiliates. All Rights Reserved.
3   * @(#)prime_pthr_fixed.c 1.3 (Oracle) 10/03/26
4   */
5
6  #include <stdio.h>
```



```
7 #include <math.h>
8 #include <pthread.h>
9
10 #define THREADS 4
11 #define N 10000
12
13 int primes[N];
14 int pflag[N];
15 int total = 0;
16 pthread_mutex_t mutex = PTHREAD_MUTEX_INITIALIZER;
17
18 int is_prime(int v)
19 {
20     int i;
21     int bound = floor(sqrt(v)) + 1;
22
23     for (i = 2; i < bound; i++) {
24         /* no need to check against known composites */
25         if (!pflag[i])
26             continue;
27         if (v % i == 0) {
28             pflag[v] = 0;
29             return 0;
30         }
31     }
32     return (v > 1);
33 }
34
35 void *work(void *arg)
36 {
37     int start;
38     int end;
39     int i;
40
41     start = (N/THREADS) * ((int)arg) ;
42     end = start + N/THREADS;
43     for (i = start; i < end; i++) {
44         if ( is_prime(i) ) {
45             pthread_mutex_lock(&mutex);
46             primes[total] = i;
47             total++;
48             pthread_mutex_unlock(&mutex);
49         }
50     }
51     return NULL;
52 }
53
54 int main(int argn, char **argv)
55 {
56     int i;
57     pthread_t tids[THREADS-1];
58
59     for (i = 0; i < N; i++) {
60         pflag[i] = 1;
```

```
61     }
62
63     for (i = 0; i < THREADS-1; i++) {
64         pthread_create(&tids[i], NULL, work, (void *)i);
65     }
66
67     i = THREADS-1;
68     work((void *)i);
69
70     for (i = 0; i < THREADS-1; i++) {
71         pthread_join(tids[i], NULL);
72     }
73
74     printf("Number of prime numbers between 2 and %d: %d\n",
75           N, total);
76
77     return 0;
78 }
```

False Positives

Occasionally, Thread Analyzer might report data races that have not actually occurred in the program. These are called false positives. In most cases, false positives are caused by user-defined synchronizations or by memory that is recycled by different threads. For more information, see [“User-Defined Synchronizations” on page 34](#) and [“Memory That is Recycled by Different Threads” on page 35](#).

User-Defined Synchronizations

Thread Analyzer can recognize most standard synchronization APIs and constructs provided by OpenMP, POSIX threads, and Oracle Solaris threads. However, the tool cannot recognize user-defined synchronizations, and might report false positive data races if your code contains such synchronizations.

Note - In order to avoid reporting this kind of false positive data race, Thread Analyzer provides a set of APIs that can be used to notify the tool when user-defined synchronizations are performed. See [“Thread Analyzer User APIs” on page 69](#) for more information.

To illustrate why you might need to use the APIs, consider the following. Thread Analyzer cannot recognize implementation of locks using CAS instructions, post and wait operations using busy-waits, and so on. Here is a typical example of a class of false positives where the program employs a common way of using POSIX thread condition variables:

```
/* Initially ready_flag is 0 */

/* Thread 1: Producer */
```

```

100  data = ...
101  pthread_mutex_lock (&mutex);
102  ready_flag = 1;
103  pthread_cond_signal (&cond);
104  pthread_mutex_unlock (&mutex);
...
/* Thread 2: Consumer */
200  pthread_mutex_lock (&mutex);
201  while (!ready_flag) {
202      pthread_cond_wait (&cond, &mutex);
203  }
204  pthread_mutex_unlock (&mutex);
205  ... = data;

```

The `pthread_cond_wait()` call is usually made within a loop that tests the predicate to protect against program errors and spurious wake-ups. The test and set of the predicate is often protected by a mutex lock. In the above code, Thread 1 produces the value for the variable `data` at line 100, sets the value of `ready_flag` to one at line 102 to indicate that the data has been produced, and then calls `pthread_cond_signal()` to wake up the consumer thread, Thread 2. Thread 2 tests the predicate (`!ready_flag`) in a loop. When it finds that the flag is set, it consumes the data at line 205.

The write of `ready_flag` at line 102 and read of `ready_flag` at line 201 are protected by the same mutex lock, so there is no data race between the two accesses and the tool recognizes that correctly.

The write of `data` at line 100 and the read of `data` at line 205 are not protected by mutex locks. However, in the program logic, the read at line 205 always happens after the write at line 100 because of the flag variable `ready_flag`. Consequently, there is no data race between these two accesses to `data`. However, the tool reports that there is a data race between the two accesses if the call to `pthread_cond_wait()` (line 202) is actually not called at run time. If line 102 is executed before line 201 is ever executed, then when line 201 is executed, the loop entry test fails and line 202 is skipped. The tool monitors `pthread_cond_signal()` calls and `pthread_cond_wait()` calls and can pair them to derive synchronization. When the `pthread_cond_wait()` at line 202 is not called, the tool does not know that the write at line 100 is always executed before the read at line 205. Therefore, it considers them as executed concurrently and reports a data race between them.

The `libtha(3C)` man page and [“Thread Analyzer User APIs” on page 69](#) explain how to use the APIs to avoid reports of this kind of false positive data race.

Memory That is Recycled by Different Threads

Some memory management routines recycle memory that is freed by one thread for use by another thread. Thread Analyzer is sometimes not able to recognize that the life spans of the same memory location used by different threads do not overlap. When this happens, the tool

might report a false positive data race. The following example illustrates this kind of false positive.

```
/*-----*/           /*-----*/
/* Thread 1 */       /* Thread 2 */
/*-----*/           /*-----*/
ptr1 = mymalloc(sizeof(data_t));
ptr1->data = ...
...
myfree(ptr1);

ptr2 = mymalloc(sizeof(data_t));
ptr2->data = ...
...
myfree(ptr2);
```

Thread 1 and Thread 2 execute concurrently. Each thread allocates a chunk of memory that is used as its private memory. The routine `mymalloc()` might supply the memory freed by a previous call to `myfree()`. If Thread 2 calls `mymalloc()` before Thread 1 calls `myfree()`, then `ptr1` and `ptr2` get different values and there is no data race between the two threads. However, if Thread 2 calls `mymalloc()` after Thread 1 calls `myfree()`, then `ptr1` and `ptr2` might have the same value. There is no data race because Thread 1 no longer accesses that memory. However, if the tool does not know `mymalloc()` is recycling memory, it reports a data race between the write of `ptr1` data and the write of `ptr2` data. This kind of false positive often happens in C++ applications when the C++ runtime library recycles memory for temporary variables. It also often happens in user applications that implement their own memory management routines. Currently, Thread Analyzer is able to recognize memory allocation and free operations performed with the standard `malloc()`, `calloc()`, and `realloc()` interfaces.

Benign Data Races

Some multithreaded applications intentionally enable data races in order to get better performance. A benign data race is an intentional data race whose existence does not affect the correctness of the program. The following examples demonstrate benign data races.

Note - In addition to benign data races, a large class of applications enable data races because they rely on lock-free and wait-free algorithms which are difficult to design correctly. Thread Analyzer can help determine the locations of data races in these applications.

A Program for Finding Primes

The threads in `prime_omp.c` check whether an integer is a prime number by executing the function `is_prime()`.

```

16 int is_prime(int v)
17 {
18     int i;
19     int bound = floor(sqrt(v)) + 1;
20
21     for (i = 2; i < bound; i++) {
22         /* no need to check against known composites */
23         if (!pflag[i])
24             continue;
25         if (v % i == 0) {
26             pflag[v] = 0;
27             return 0;
28         }
29     }
30     return (v > 1);
31 }

```

Thread Analyzer reports that there is a data race between the write to `pflag[]` on line 26 and the read of `pflag[]` on line 23. However, this data race is benign as it does not affect the correctness of the final result. At line 23, a thread checks whether or not `pflag[i]`, for a given value of `i` is equal to zero. If `pflag[i]` is equal to zero, that means that `i` is a known composite number (in other words, `i` is known to be non-prime). Consequently, there is no need to check whether or not `v` is divisible by `i`; you only need to check whether or not `v` is divisible by some prime number. Therefore, if `pflag[i]` is equal to zero, the thread continues to the next value of `i`. If `pflag[i]` is not equal to zero and `v` is divisible by `i`, the thread assigns zero to `pflag[v]` to indicate that `v` is not a prime number.

It does not matter, from a correctness point of view, if multiple threads check the same `pflag[]` element and write to it concurrently. The initial value of a `pflag[]` element is one. When the threads update that element, they assign it the value zero. That is, the threads store zero in the same bit in the same byte of memory for that element. On current architectures, it is safe to assume that those stores are atomic. This means that, when that element is read by a thread, the value read is either one or zero. If a thread checks a given `pflag[]` element (line 23) before it has been assigned the value zero, it then executes lines 25–28. If, in the meantime, another thread assigns zero to that same `pflag[]` element (line 26), the final result is not changed. Essentially, this means that the first thread executed lines 25–28 unnecessarily, but the final result is the same.

A Program that Verifies Array-Value Types

A group of threads call `check_bad_array()` concurrently to check whether any element of array `data_array` is “bad”. Each thread checks a different section of the array. If a thread finds that an element is bad, it sets the value of a global shared variable `is_bad` to true.

```

20 volatile int is_bad = 0;
...

```

```
100 /*
101  * Each thread checks its assigned portion of data_array, and sets
102  * the global flag is_bad to 1 once it finds a bad data element.
103  */
104 void check_bad_array(volatile data_t *data_array, unsigned int thread_id)
105 {
106     int i;
107     for (i=my_start(thread_id); i<my_end(thread_id); i++) {
108         if (is_bad)
109             return;
110         else {
111             if (is_bad_element(data_array[i])) {
112                 is_bad = 1;
113                 return;
114             }
115         }
116     }
117 }
```

There is a data race between the read of `is_bad` on line 108 and the write to `is_bad` on line 112. However, the data race does not affect the correctness of the final result.

The initial value of `is_bad` is zero. When the threads update `is_bad`, they assign it the value one. That is, the threads store one in the same bit in the same byte of memory for `is_bad`. On current architectures, it is safe to assume that those stores are atomic. Therefore, when `is_bad` is read by a thread, the value read will either be zero or one. If a thread checks `is_bad` (line 108) before it has been assigned the value one, then it continues executing the `for` loop. If, in the meantime, another thread has assigned the value one to `is_bad` (line 112), that does not change the final result. It just means that the thread executed the `for` loop longer than necessary.

A Program Using Double-Checked Locking

A singleton ensures that only one object of a certain type exists throughout the program. Double-checked locking is a common, efficient way to initialize a singleton in multithreaded applications. The following code illustrates such an implementation.

```
100 class Singleton {
101     public:
102         static Singleton* instance();
103         ...
104     private:
105         static Singleton* ptr_instance;
106 };
107
108 Singleton* Singleton::ptr_instance = 0;
109 ...
110
111 Singleton* Singleton::instance() {
112     Singleton *tmp;
```

```
202     if (ptr_instance == 0) {
203         Lock();
204         if (ptr_instance == 0) {
205             tmp = new Singleton;
206
207             /* Make sure that all writes used to construct new
208              Singleton have been completed. */
209             memory_barrier();
210
211             /* Update ptr_instance to point to new Singleton. */
212             ptr_instance = tmp;
213
214         }
215         Unlock();
216     }
217     return ptr_instance;
```

The read of `ptr_instance` on line 202 is intentionally not protected by a lock. This makes the check to determine whether or not the `Singleton` has already been instantiated in a multithreaded environment more efficient. Notice that there is a data race on variable `ptr_instance` between the read on line 202 and the write on line 212, but the program works correctly. However, writing a correct program that enables data races requires extra care. For example, in the above double-checked-locking code, the call to `memory_barrier()` at line 209 is used to ensure that `ptr_instance` is not seen to be non-null by the threads until all writes to construct the `Singleton` have been completed.

Deadlock Tutorial

This tutorial explains how to use Thread Analyzer to detect potential deadlocks and actual deadlocks in your multithreaded program.

The tutorial covers the following topics:

- [“About Deadlocks” on page 41](#)
- [“Getting the Deadlock Tutorial Source Files” on page 42](#)
- [“The Dining Philosophers Scenario” on page 44](#)
- [“How to Use Thread Analyzer to Find Deadlocks” on page 48](#)
- [“Understanding the Deadlock Experiment Results” on page 51](#)
- [“Fixing the Deadlocks and Understanding False Positives” on page 58](#)

About Deadlocks

The term *deadlock* describes a condition in which two or more threads are blocked forever because they are waiting for each other. There are many causes of deadlocks such as erroneous program logic and inappropriate use of synchronizations such as locks and barriers. This tutorial focuses on deadlocks that are caused by the inappropriate use of *mutexes*, or mutual exclusion locks. This type of deadlock is commonly encountered in multithreaded applications.

A process with two or more threads can enter deadlock when the following three conditions hold:

- Threads that are already holding locks request new locks
- The requests for new locks are made concurrently
- Two or more threads form a circular chain in which each thread waits for a lock which is held by the next thread in the chain

Here is a simple example of a deadlock condition:

- Thread 1 holds lock A and requests lock B
- Thread 2 holds lock B and requests lock A

A deadlock can be of two types: A *potential* deadlock or an *actual* deadlock and they are distinguished as follows:

- A potential deadlock does not necessarily occur in a given run, but can occur in any execution of the program depending on the scheduling of threads and the timing of lock requests by the threads.
- An actual deadlock is one that occurs during the execution of a program. An actual deadlock causes the threads involved to hang, but might cause the whole process to hang.

Getting the Deadlock Tutorial Source Files

You can download the source files used in this tutorial from the [Download area \(http://www.oracle.com/technetwork/server-storage/solarisstudio/downloads/index.html\)](http://www.oracle.com/technetwork/server-storage/solarisstudio/downloads/index.html) of the Oracle Solaris Studio developer portal.

After you download and unpack the sample files, you can find the samples in the `SolarisStudioSampleApplications/ThreadAnalyzer` directory. The samples are located in the `din_philo` subdirectory. The `din_philo` directory includes a `Makefile` and a `DEMO` file of instructions, but this tutorial does not follow those instructions or use the `Makefile`. Instead, you are instructed to execute commands individually.

To follow the tutorial, you can copy the `din_philo.c` file from the `SolarisStudioSampleApplications/ThreadAnalyzer/din_philo` directory to a different directory, or you can create your own file and copy the code from the following code listing.

The `din_philo.c` sample program which simulates the dining-philosophers problem is a C program that uses POSIX threads. The program can exhibit both potential and actual deadlocks.

Source Code Listing for `din_philo.c`

The source code for `din_philo.c` is shown below:

```
1 /*
2  * Copyright (c) 2006, 2010, Oracle and/or its affiliates. All Rights Reserved.
3  * @(#)din_philo.c 1.4 (Oracle) 10/03/26
4  */
5
6 #include <pthread.h>
7 #include <stdio.h>
8 #include <unistd.h>
9 #include <stdlib.h>
10 #include <errno.h>
11 #include <assert.h>
12
13 #define PHILOS 5
14 #define DELAY 5000
15 #define FOOD 100
16
17 void *philosopher (void *id);
```

```
18 void grab_chopstick (int,
19                     int,
20                     char *);
21 void down_chopsticks (int,
22                      int);
23 int food_on_table ();
24
25 pthread_mutex_t chopstick[PHILOS];
26 pthread_t philo[PHILOS];
27 pthread_mutex_t food_lock;
28 int sleep_seconds = 0;
29
30
31 int
32 main (int argn,
33       char **argv)
34 {
35     int i;
36
37     if (argn == 2)
38         sleep_seconds = atoi (argv[1]);
39
40     pthread_mutex_init (&food_lock, NULL);
41     for (i = 0; i < PHILOS; i++)
42         pthread_mutex_init (&chopstick[i], NULL);
43     for (i = 0; i < PHILOS; i++)
44         pthread_create (&philo[i], NULL, philosopher, (void *)i);
45     for (i = 0; i < PHILOS; i++)
46         pthread_join (philo[i], NULL);
47     return 0;
48 }
49
50 void *
51 philosopher (void *num)
52 {
53     int id;
54     int i, left_chopstick, right_chopstick, f;
55
56     id = (int)num;
57     printf ("Philosopher %d is done thinking and now ready to eat.\n", id);
58     right_chopstick = id;
59     left_chopstick = id + 1;
60
61     /* Wrap around the chopsticks. */
62     if (left_chopstick == PHILOS)
63         left_chopstick = 0;
64
65     while (f = food_on_table ()) {
66
67         /* Thanks to philosophers #1 who would like to take a nap
68          * before picking up the chopsticks, the other philosophers
69          * may be able to eat their dishes and not deadlock.
70          */
71         if (id == 1)
```

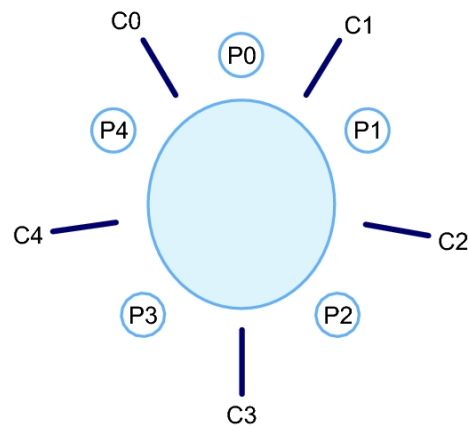
```
72         sleep (sleep_seconds);
73
74         grab_chopstick (id, right_chopstick, "right ");
75         grab_chopstick (id, left_chopstick, "left");
76
77         printf ("Philosopher %d: eating.\n", id);
78         usleep (DELAY * (FOOD - f + 1));
79         down_chopsticks (left_chopstick, right_chopstick);
80     }
81
82     printf ("Philosopher %d is done eating.\n", id);
83     return (NULL);
84 }
85
86 int
87 food_on_table ()
88 {
89     static int food = FOOD;
90     int myfood;
91
92     pthread_mutex_lock (&food_lock);
93     if (food > 0) {
94         food--;
95     }
96     myfood = food;
97     pthread_mutex_unlock (&food_lock);
98     return myfood;
99 }
100
101 void
102 grab_chopstick (int phil,
103                int c,
104                char *hand)
105 {
106     pthread_mutex_lock (&chopstick[c]);
107     printf ("Philosopher %d: got %s chopstick %d\n", phil, hand, c);
108 }
109
110 void
111 down_chopsticks (int c1,
112                 int c2)
113 {
114     pthread_mutex_unlock (&chopstick[c1]);
115     pthread_mutex_unlock (&chopstick[c2]);
116 }
```

The Dining Philosophers Scenario

The dining philosophers scenario is a classic which is structured as follows. Five philosophers, numbered zero to four, are sitting at a round table, thinking. As time passes, different individuals become hungry and decide to eat. There is a platter of noodles on the table but each

philosopher only has one chopstick to use. In order to eat, they must share chopsticks. The chopstick to the right of each philosopher (as they sit facing the table) has the same number as that philosopher.

FIGURE 3-1 Dining Philosophers



P = Philosopher
C = Chopstick

Each philosopher first reaches for his own chopstick which is the one with his number. When he has his assigned chopstick, he reaches for the chopstick assigned to his neighbor. After he has both chopsticks, he can eat. After eating, he returns the chopsticks to their original positions on the table, one on either side. The process is repeated until there are no more noodles.

How the Philosophers Can Deadlock

An actual deadlock occurs when every philosopher is holding his own chopstick and waiting for the one from his neighbor to become available:

- Philosopher 0 is holding chopstick 0, but is waiting for chopstick 1
- Philosopher 1 is holding chopstick 1, but is waiting for chopstick 2
- Philosopher 2 is holding chopstick 2, but is waiting for chopstick 3
- Philosopher 3 is holding chopstick 3, but is waiting for chopstick 4
- Philosopher 4 is holding chopstick 4, but is waiting for chopstick 0

In this situation, nobody can eat and the philosophers are in a deadlock. Run the program a number of times. You will see that the program might hang sometimes, and run to completion at other times. The program might hang as shown in the following sample run:

```
prompt% cc din_philo.c
prompt% a.out
Philosopher 0 is done thinking and now ready to eat.
Philosopher 2 is done thinking and now ready to eat.
Philosopher 2: got right chopstick 2
Philosopher 2: got left chopstick 3
Philosopher 0: got right chopstick 0
Philosopher 0: got left chopstick 1
Philosopher 0: eating.
Philosopher 4 is done thinking and now ready to eat.
Philosopher 4: got right chopstick 4
Philosopher 2: eating.
Philosopher 3 is done thinking and now ready to eat.
Philosopher 1 is done thinking and now ready to eat.
Philosopher 0: got right chopstick 0
Philosopher 3: got right chopstick 3
Philosopher 2: got right chopstick 2
Philosopher 1: got right chopstick 1
(hang)
```

Execution terminated by pressing CTRL-C

Introducing a Sleep Time for Philosopher 1

One way to avoid deadlocks is for Philosopher 1 to wait before reaching for his chopstick. In terms of the code, he can be put to sleep for a specified amount of time (`sleep_seconds`) before reaching for his chopstick. If he sleeps long enough, then the program might finish without any actual deadlock. You can specify the number of seconds he sleeps as an argument to the executable. If you do not specify an argument, the philosopher does not sleep.

The following pseudo-code shows the logic for each philosopher:

```
while (there is still food on the table)
{
    if (sleep argument is specified and I am philosopher #1)
    {
        sleep specified amount of time
    }

    grab right chopstick
    grab left chopstick
    eat some food
    put down left chopstick
    put down right chopstick
}
```



```
Philosopher 0: got left chopstick 1
Philosopher 0: eating.
Philosopher 0: got right chopstick 0
Philosopher 0: got left chopstick 1
Philosopher 0: eating.
Philosopher 0: got right chopstick 0
Philosopher 0: got left chopstick 1
Philosopher 0: eating.
Philosopher 0: got right chopstick 0
Philosopher 0: got left chopstick 1
Philosopher 0: eating.
Philosopher 0: got right chopstick 0
Philosopher 0: got left chopstick 1
Philosopher 0: eating.
Philosopher 0 is done eating.
Philosopher 4: got left chopstick 0
Philosopher 4: eating.
Philosopher 4 is done eating.
Philosopher 3: got left chopstick 4
Philosopher 3: eating.
Philosopher 3 is done eating.
Philosopher 2: got left chopstick 3
Philosopher 2: eating.
Philosopher 2 is done eating.
Philosopher 1: got right chopstick 1
Philosopher 1: got left chopstick 2
Philosopher 1: eating.
Philosopher 1 is done eating.
%
```

Execution terminated normally

Try running the program several times and specifying different sleep arguments. What happens when Philosopher 1 waits only a short time before reaching for his chopstick? How about when he waits longer? Try specifying different sleep arguments to the executable `a.out`. Rerun the program with or without a sleep argument several times. Sometimes the program hangs, while it runs to completion at other times. Whether the program hangs or not depends on the scheduling of threads and the timings of requests for locks by the threads.

How to Use Thread Analyzer to Find Deadlocks

You can use Thread Analyzer to check for potential and actual deadlocks in your program. Thread Analyzer follows the same collect-analyze model that Oracle Solaris Studio Performance Analyzer uses.

There are three steps involved in using Thread Analyzer:

- Compile the source code.
- Create a deadlock-detection experiment.
- Examine the experiment results.

Compile the Source Code

Compile your code and be sure to specify `-g`. Do not specify a high-level of optimization because information such as line numbers and call stacks, might be reported incorrectly at a high optimization level. Compile an OpenMP program with `-g -xopenmp=noopt`, and compile a POSIX threads program with just `-g -mt`.

See `cc(1)`, `CC(1)`, or `f95(1)` man pages for more information about these compiler options.

For this tutorial, compile the code using the following command:

```
% cc -g -o din_philo din_philo.c
```

Create a Deadlock-Detection Experiment

Use the `collect` command with the `-r deadlock` option. This option creates a deadlock-detection experiment during the execution of the program.

For this tutorial, create a deadlock-detection experiment named `din_philo.1.er` using the following command:

```
% collect -r deadlock -o din_philo.1.er din_philo
```

The `collect` command accepts the following options, which are useful when creating a deadlock-detection experiment:

<code>terminate</code>	If an unrecoverable error is detected, terminate the program.
<code>abort</code>	If an unrecoverable error is detected, terminate the program with a core dump.
<code>-continue</code>	If an unrecoverable error is detected, enable the program to continue.

The default behavior is `terminate`.

You can use any of the previous options with the `collect` command to get the behavior you want. For example, to cause the program to terminate with a core dump when an actual deadlock occurs use the following `collect` command.

```
% collect -r deadlock,abort -o din_philo.1.er din_philo
```

To cause the program to hang when an actual deadlock occurs, use the following `collect` command:

```
% collect -r deadlock,continue -o din_philo.1.er din_philo
```

You can increase the likelihood of detecting deadlocks by creating several deadlock-detection experiments. Use a different number of threads and different input data for the various experiments. For example, in the `din_philo.c` code, you could change the values in the following lines:

```
13 #define PHILOS 5
14 #define DELAY 5000
15 #define FOOD 100
```

You could then compile as before and collect another experiment.

See `collect(1)` and `collector(1)` man pages for more information.

Examine the Deadlock-Detection Experiment

You can examine a deadlock-detection experiment with Thread Analyzer, Performance Analyzer, or the `er_print` utility. Both Thread Analyzer and Performance Analyzer present a GUI interface; Thread Analyzer presents a simplified set of default views, but is otherwise identical to Performance Analyzer.

Using Thread Analyzer to View the Deadlock-Detection Experiment

To start Thread Analyzer and open the `din_philo.1.er` experiment, type the following command:

```
% tha din_philo.1.er
```

Thread Analyzer has a menu bar, a tool bar, and vertical navigation bar on the left that enables you to select data views.

The following data views are shown by default when you open an experiment that was collected for deadlock detection:

- Overview screen shows the metrics overview of the loaded experiments.
- Deadlocks view shows a list of potential and actual deadlocks that Thread Analyzer detected in the program. This view is selected by default. The threads involved for each deadlock are shown. These threads form a circular chain where each thread holds a lock and requests another lock that the next thread in the chain holds.

When you select a deadlock, the Deadlock Details window in the right panel shows detailed information about the threads involved.

- Dual Source view shows the source location where the thread held a lock, and the source location where the same thread requested a lock. The source lines where the thread held and requested locks are highlighted. To display this view, select a thread in the circular chain on the Deadlocks view and then click on the Dual Source view.

- Experiments view shows the load objects in the experiment and lists any error and warning messages.

You can choose to see other views with the More Views options menu.

Using `er_print` to View the Deadlock-Detection Experiment

The `er_print` utility presents a command-line interface. You can use the `er_print` utility in an interactive session and specify sub-commands during the session. You can also use command-line options to specify sub-commands non-interactively.

The following sub-commands are useful for examining deadlocks with the `er_print` utility:

- `-deadlocks`
This option reports any potential and actual deadlocks detected in the experiment. Specify deadlocks at the `(er_print)` prompt or `-deadlocks` on the `er_print` command line.
- `-ddetail deadlock-ID`
This option returns detailed information about the deadlock with the specified *deadlock-ID*. Specify `ddetail` at the `(er_print)` prompt or `-ddetail` on the `er_print` command line. If the specified *deadlock-ID* is `all`, then detailed information about all deadlocks is displayed. Otherwise, specify a single deadlock number such as `1` for the first deadlock.
- `-header`
This option displays descriptive information about the experiment and reports any errors or warnings. Specify `header` at the `(er_print)` prompt or `-header` on the command line.

Refer to the `collect(1)`, `tha(1)`, `analyzer(1)`, and `er_print(1)` man pages for more information.

Understanding the Deadlock Experiment Results

This section explains how to use Thread Analyzer to investigate the deadlocks in the dining philosopher program.

Examining Runs That Deadlock

The following listing shows a run of the dining philosophers program that results in an actual deadlock.

```
% cc -g -o din_philo din_philo.c
% collect -r deadlock -o din_philo.1.er din_philo
```

```
Creating experiment database din_philo.1.er ...
Philosopher 1 is done thinking and now ready to eat.
Philosopher 2 is done thinking and now ready to eat.
Philosopher 3 is done thinking and now ready to eat.
Philosopher 0 is done thinking and now ready to eat.
Philosopher 1: got right chopstick 1
Philosopher 3: got right chopstick 3
Philosopher 0: got right chopstick 0
Philosopher 1: got left chopstick 2
Philosopher 3: got left chopstick 4
Philosopher 4 is done thinking and now ready to eat.
Philosopher 1: eating.
Philosopher 3: eating.
Philosopher 3: got right chopstick 3
Philosopher 4: got right chopstick 4
Philosopher 2: got right chopstick 2
Philosopher 0: got left chopstick 1
Philosopher 0: eating.
Philosopher 1: got right chopstick 1
Philosopher 4: got left chopstick 0
Philosopher 4: eating.
Philosopher 0: got right chopstick 0
Philosopher 3: got left chopstick 4
Philosopher 3: eating.
Philosopher 4: got right chopstick 4
Philosopher 2: got left chopstick 3
Philosopher 2: eating.
Philosopher 3: got right chopstick 3
Philosopher 1: got left chopstick 2
Philosopher 1: eating.
Philosopher 2: got right chopstick 2
Philosopher 0: got left chopstick 1
Philosopher 0: eating.
Philosopher 1: got right chopstick 1
Philosopher 4: got left chopstick 0
Philosopher 4: eating.
Philosopher 0: got right chopstick 0
Philosopher 3: got left chopstick 4
Philosopher 3: eating.
...
Philosopher 4: got right chopstick 4
Philosopher 2: got left chopstick 3
Philosopher 2: eating.
Philosopher 2: got right chopstick 2
Philosopher 3: got right chopstick 3
(hang)
```

Execution terminated by pressing CTRL-C

Type the following commands to examine the experiment with `er_print` utility:

```
% er_print din_philo.1.er
(er_print) deadlocks

Deadlock #1, Potential deadlock
```

```

    Thread #2
    Lock being held:      0x21380, at: grab_chopstick + 0x00000024, line 105 in
"din_philo.c"
    Lock being requested: 0x21398, at: grab_chopstick + 0x00000024, line 105 in
"din_philo.c"
    Thread #3
    Lock being held:      0x21398, at: grab_chopstick + 0x00000024, line 105 in
"din_philo.c"
    Lock being requested: 0x213b0, at: grab_chopstick + 0x00000024, line 105 in
"din_philo.c"
    Thread #4
    Lock being held:      0x213b0, at: grab_chopstick + 0x00000024, line 105 in
"din_philo.c"
    Lock being requested: 0x213c8, at: grab_chopstick + 0x00000024, line 105 in
"din_philo.c"
    Thread #5
    Lock being held:      0x213c8, at: grab_chopstick + 0x00000024, line 105 in
"din_philo.c"
    Lock being requested: 0x213e0, at: grab_chopstick + 0x00000024, line 105 in
"din_philo.c"
    Thread #6
    Lock being held:      0x213e0, at: grab_chopstick + 0x00000024, line 105 in
"din_philo.c"
    Lock being requested: 0x21380, at: grab_chopstick + 0x00000024, line 105 in
"din_philo.c"

```

Deadlock #2, Actual deadlock

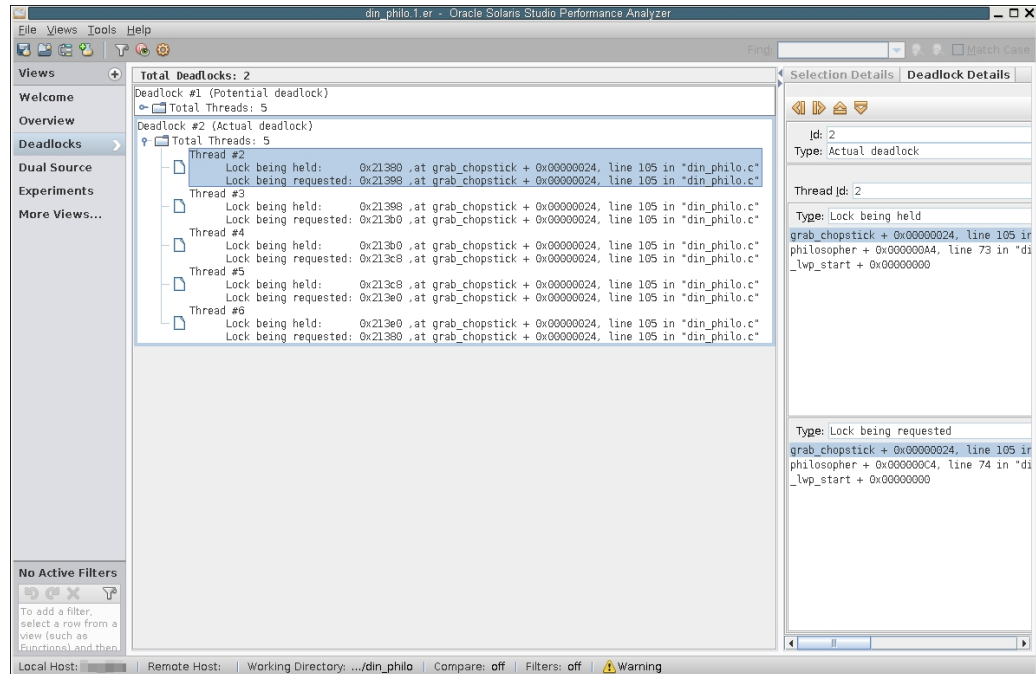
```

    Thread #2
    Lock being held:      0x21380, at: grab_chopstick + 0x00000024, line 105 in
"din_philo.c"
    Lock being requested: 0x21398, at: grab_chopstick + 0x00000024, line 105 in
"din_philo.c"
    Thread #3
    Lock being held:      0x21398, at: grab_chopstick + 0x00000024, line 105 in
"din_philo.c"
    Lock being requested: 0x213b0, at: grab_chopstick + 0x00000024, line 105 in
"din_philo.c"
    Thread #4
    Lock being held:      0x213b0, at: grab_chopstick + 0x00000024, line 105 in
"din_philo.c"
    Lock being requested: 0x213c8, at: grab_chopstick + 0x00000024, line 105 in
"din_philo.c"
    Thread #5
    Lock being held:      0x213c8, at: grab_chopstick + 0x00000024, line 105 in
"din_philo.c"
    Lock being requested: 0x213e0, at: grab_chopstick + 0x00000024, line 105 in
"din_philo.c"
    Thread #6
    Lock being held:      0x213e0, at: grab_chopstick + 0x00000024, line 105 in
"din_philo.c"
    Lock being requested: 0x21380, at: grab_chopstick + 0x00000024, line 105 in
"din_philo.c"

```

The following screen shot shows Thread Analyzer's presentation of the deadlock information.

FIGURE 3-2 Deadlock Detected in `din_philo.c`



Thread Analyzer reports two deadlocks for `din_philo.c`, one potential and the other actual. On closer inspection, you find that the two deadlocks are identical.

The circular chain involved in the deadlock is as follows:

- Thread 2: holds lock at address `0x21380`, requests lock at address `0x21398`
- Thread 3: holds lock at address `0x21398`, requests lock at address `0x213b0`
- Thread 4: holds lock at address `0x213b0`, requests lock at address `0x213c8`
- Thread 5: holds lock at address `0x213c8`, requests lock at address `0x213e0`
- Thread 6: holds lock at address `0x213e0`, requests lock at address `0x21380`

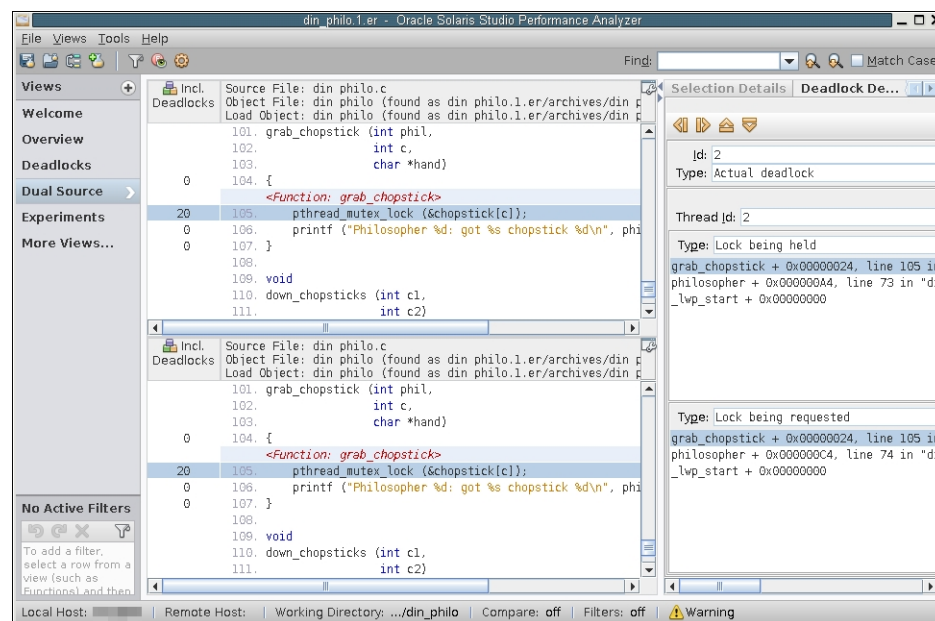
Select the first thread in the chain (Thread #2) and then click on the Dual Source view to see where in the source code Thread #2 acquired the lock at address `0x21380`, and where in the source code it requested the lock at address `0x21398`.

The following screen shot shows the Dual Source view for Thread #2. The top half of the screen shot shows that Thread #2 acquired the lock at address `0x21380` by calling `pthread_mutex_lock()` on line 105. The bottom half of the screen shot shows that the same thread requested the lock at address `0x21398` by calling `pthread_mutex_lock()` on line 105.

Each of the two calls to `pthread_mutex_lock()` used a different lock as the argument. In general, the lock-acquire and lock-request operations might not be on the same source line.

The default metric (Exclusive Deadlocks metric) is shown to the left of each source line in the screen shot. This metric gives a count of the number of times a lock-acquire or lock-request operation, which was involved in a deadlock, was reported on that source line. Only source lines that are part of a deadlock chain would have a value for this metric that is larger than zero.

FIGURE 3-3 Potential Deadlock in `din_philo.c`



Examining Runs That Complete Despite Deadlock Potential

The dining philosophers program can avoid actual deadlock and terminate normally if you supply a large enough sleep argument. Normal termination, however, does not mean the program is safe from deadlocks. It simply means that the locks that were held and requested did not form a deadlock chain during a given run. If the timing changes in other runs, an actual deadlock can occur. The following listing shows a run of the dining philosophers program that terminates normally because of the 40 second sleep time. However, the `er_print` utility and Thread Analyzer report potential deadlocks.

```
% cc -g -o din_philo_pt din_philo.c
% collect -r deadlock -o din_philo_pt.1.er din_philo_pt 40
Creating experiment database tha.2.er ...
Philosopher 0 is done thinking and now ready to eat.
Philosopher 2 is done thinking and now ready to eat.
Philosopher 1 is done thinking and now ready to eat.
Philosopher 3 is done thinking and now ready to eat.
Philosopher 2: got right chopstick 2
Philosopher 3: got right chopstick 3
Philosopher 0: got right chopstick 0
Philosopher 4 is done thinking and now ready to eat.
Philosopher 0: got left chopstick 1
Philosopher 0: eating.
Philosopher 3: got left chopstick 4
Philosopher 3: eating.
Philosopher 0: got left chopstick 1
Philosopher 0: eating.
Philosopher 0: got right chopstick 0
Philosopher 2: got left chopstick 3
Philosopher 2: eating.
...
Philosopher 4: got right chopstick 4
Philosopher 3: got right chopstick 3
Philosopher 2: got right chopstick 2
Philosopher 4: got left chopstick 0
Philosopher 4: eating.
Philosopher 4 is done eating.
Philosopher 3: got left chopstick 4
Philosopher 3: eating.
Philosopher 0: got right chopstick 0
Philosopher 0: got left chopstick 1
Philosopher 0: eating.
Philosopher 3 is done eating.
Philosopher 2: got left chopstick 3
Philosopher 2: eating.
Philosopher 0 is done eating.
Philosopher 2 is done eating.
Philosopher 1: got right chopstick 1
Philosopher 1: got left chopstick 2
Philosopher 1: eating.
Philosopher 1 is done eating.
```

%

Execution terminated normally

Type the following commands shown at the prompts to examine the experiment with `er_print` utility:

```
% er_print din_philo_pt.1.er
(er_print) deadlocks
Deadlock #1, Potential deadlock
      Thread #2
            Lock being held:          0x21388, at: grab_chopstick + 0x00000024, line 105 in
            "din_philo.c"
```



```

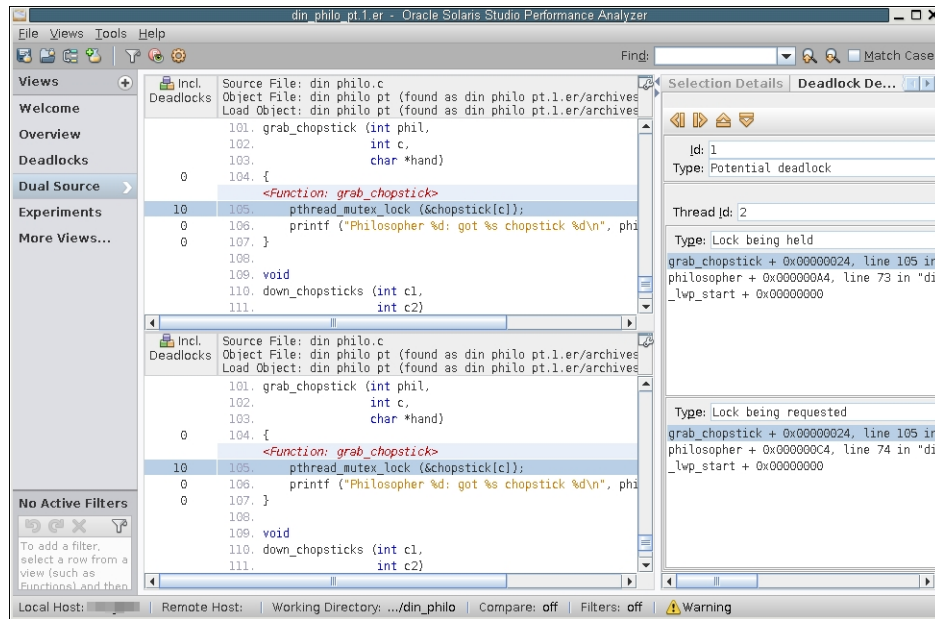
        Lock being requested: 0x213a0, at: grab_chopstick + 0x00000024, line 105 in
"din_philo.c"
    Thread #3
        Lock being held:      0x213a0, at: grab_chopstick + 0x00000024, line 105 in
"din_philo.c"
        Lock being requested: 0x213b8, at: grab_chopstick + 0x00000024, line 105 in
"din_philo.c"
    Thread #4
        Lock being held:      0x213b8, at: grab_chopstick + 0x00000024, line 105 in
"din_philo.c"
        Lock being requested: 0x213d0, at: grab_chopstick + 0x00000024, line 105 in
"din_philo.c"
    Thread #5
        Lock being held:      0x213d0, at: grab_chopstick + 0x00000024, line 105 in
"din_philo.c"
        Lock being requested: 0x213e8, at: grab_chopstick + 0x00000024, line 105 in
"din_philo.c"
    Thread #6
        Lock being held:      0x213e8, at: grab_chopstick + 0x00000024, line 105 in
"din_philo.c"
        Lock being requested: 0x21388, at: grab_chopstick + 0x00000024, line 105 in
"din_philo.c"

```

Deadlocks List Summary: Experiment: din_philo_pt.1.er Total Deadlocks: 1

The following screen shot shows the potential deadlock information in Thread Analyzer.

FIGURE 3-4 Potential deadlock in din_philo.c



Fixing the Deadlocks and Understanding False Positives

One way to remove potential and actual deadlocks is to use a system of tokens so that a philosopher must receive a token before attempting to eat. The number of available tokens must be less than the number of philosophers at the table. After a philosopher receives a token, he can attempt to eat in accordance with the rules of the table. After eating, each philosopher returns the token and repeats the process. The following pseudo-code shows the logic for each philosopher when using the token system.

```
while (there is still food on the table)
{
    get token
    grab right fork
    grab left fork
    eat some food
    put down left fork
    put down right fork
    return token
}
```

The following sections detail two different implementations for the system of tokens.

Regulating the Philosophers With Tokens

The following listing shows the fixed version of the dining philosophers program that uses the token system. This solution incorporates four tokens, one less than the number of diners, so no more than four philosophers can attempt to eat at the same time. This version of the program is called `din_philo_fix1.c`:

Tip - If you downloaded the sample applications, you can copy the `din_philo_fix1.c` file from the `SolarisStudioSampleApplications/ThreadAnalyzer/din_philo` directory.

```

1  /*
2   * Copyright (c) 2006, 2010, Oracle and/or its affiliates. All Rights Reserved.
3   * @(#)din_philo_fix1.c 1.3 (Oracle) 10/03/26
4   */
5
6  #include <pthread.h>
7  #include <stdio.h>
8  #include <unistd.h>
9  #include <stdlib.h>
10 #include <errno.h>
11 #include <assert.h>
12
13 #define PHILOS 5
14 #define DELAY 5000
15 #define FOOD 100
16
17 void *philosopher (void *id);
18 void grab_chopstick (int,
19                     int,
20                     char *);
21 void down_chopsticks (int,
22                      int);
23 int food_on_table ();
24 void get_token ();
25 void return_token ();
26
27 pthread_mutex_t chopstick[PHILOS];
28 pthread_t philo[PHILOS];
29 pthread_mutex_t food_lock;
30 pthread_mutex_t num_can_eat_lock;
31 int sleep_seconds = 0;
32 uint32_t num_can_eat = PHILOS - 1;
33
34
35 int
36 main (int argn,
37       char **argv)
38 {
39     int i;
40
41     pthread_mutex_init (&food_lock, NULL);

```

```
42 pthread_mutex_init (&num_can_eat_lock, NULL);
43 for (i = 0; i < PHILOS; i++)
44     pthread_mutex_init (&chopstick[i], NULL);
45 for (i = 0; i < PHILOS; i++)
46     pthread_create (&philo[i], NULL, philosopher, (void *)i);
47 for (i = 0; i < PHILOS; i++)
48     pthread_join (philo[i], NULL);
49 return 0;
50 }
51
52 void *
53 philosopher (void *num)
54 {
55     int id;
56     int i, left_chopstick, right_chopstick, f;
57
58     id = (int)num;
59     printf ("Philosopher %d is done thinking and now ready to eat.\n", id);
60     right_chopstick = id;
61     left_chopstick = id + 1;
62
63     /* Wrap around the chopsticks. */
64     if (left_chopstick == PHILOS)
65         left_chopstick = 0;
66
67     while (f = food_on_table ()) {
68         get_token ();
69
70         grab_chopstick (id, right_chopstick, "right ");
71         grab_chopstick (id, left_chopstick, "left");
72
73         printf ("Philosopher %d: eating.\n", id);
74         usleep (DELAY * (FOOD - f + 1));
75         down_chopsticks (left_chopstick, right_chopstick);
76
77         return_token ();
78     }
79
80     printf ("Philosopher %d is done eating.\n", id);
81     return (NULL);
82 }
83
84 int
85 food_on_table ()
86 {
87     static int food = FOOD;
88     int myfood;
89
90     pthread_mutex_lock (&food_lock);
91     if (food > 0) {
92         food--;
93     }
94     myfood = food;
95     pthread_mutex_unlock (&food_lock);
```

```
96     return myfood;
97 }
98
99 void
100 grab_chopstick (int phil,
101                int c,
102                char *hand)
103 {
104     pthread_mutex_lock (&chopstick[c]);
105     printf ("Philosopher %d: got %s chopstick %d\n", phil, hand, c);
106 }
107
108
109
110 void
111 down_chopsticks (int c1,
112                 int c2)
113 {
114     pthread_mutex_unlock (&chopstick[c1]);
115     pthread_mutex_unlock (&chopstick[c2]);
116 }
117
118
119 void
120 get_token ()
121 {
122     int successful = 0;
123
124     while (!successful) {
125         pthread_mutex_lock (&num_can_eat_lock);
126         if (num_can_eat > 0) {
127             num_can_eat--;
128             successful = 1;
129         }
130         else {
131             successful = 0;
132         }
133         pthread_mutex_unlock (&num_can_eat_lock);
134     }
135 }
136
137 void
138 return_token ()
139 {
140     pthread_mutex_lock (&num_can_eat_lock);
141     num_can_eat++;
142     pthread_mutex_unlock (&num_can_eat_lock);
143 }
```

Try compiling this fixed version of the dining philosophers program and running it several times. The system of tokens limits the number of diners attempting to use the chopsticks and thus avoids actual and potential deadlocks.

To compile, use the following command:

```
% cc -g -o din_philo_fix1 din_philo_fix1.c
```

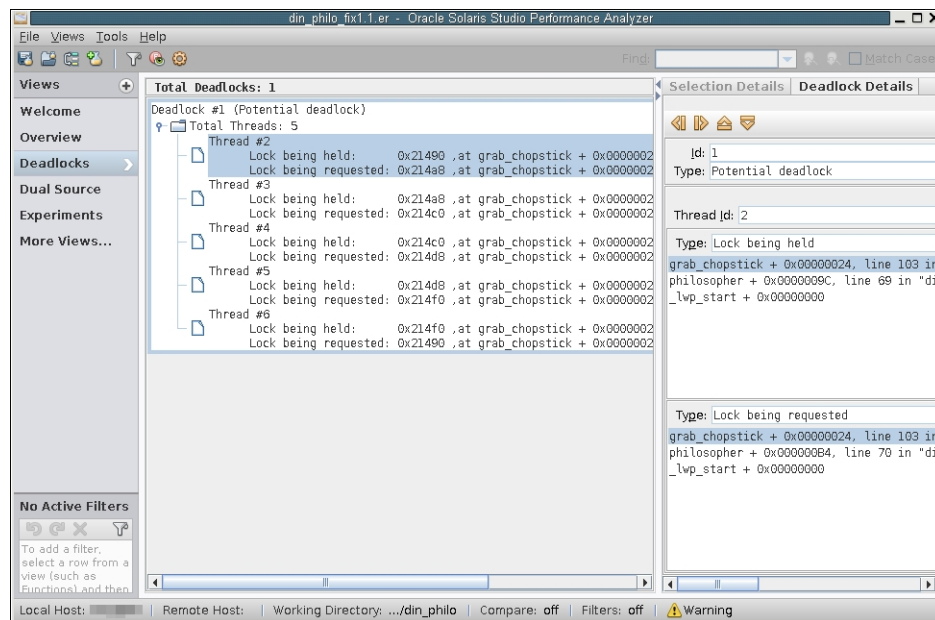
To collect an experiment:

```
% collect -r deadlock -o din_philo_fix1.1.er din_philo_fix1
```

A False Positive Report

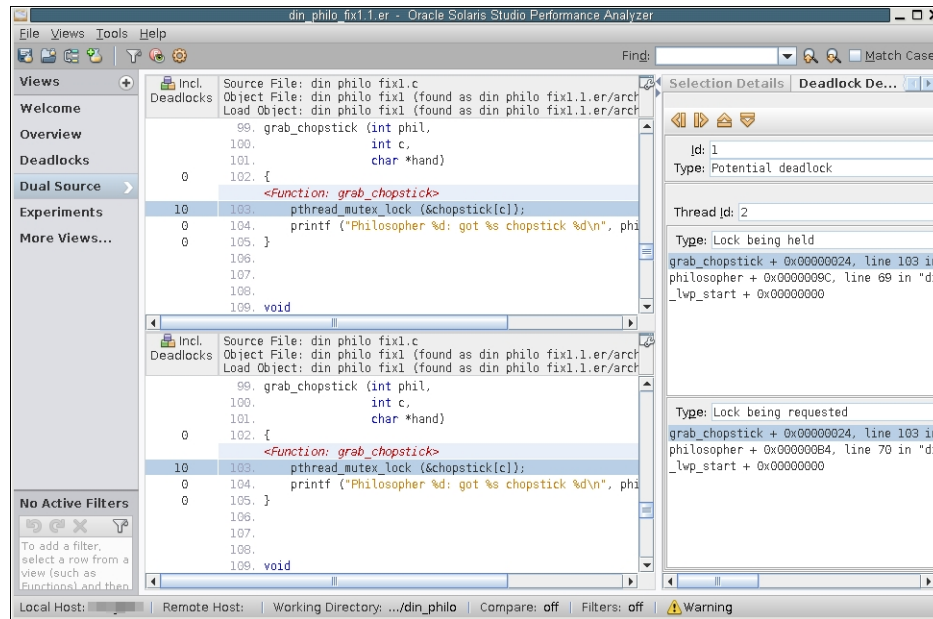
Even when using the system of tokens, Thread Analyzer reports a potential deadlock for this implementation when none exists. This is a false positive. Consider the following screen shot which details the potential deadlock.

FIGURE 3-5 False Positive Report of a Potential Deadlock



Select the first thread in the chain (Thread #2) and then click the Dual Source view to see the source code location in which Thread #2 held the lock at address 0x216a8, and where in the source code it requested the lock at address 0x216c0. The following figure shows the Dual Source view for Thread #2.

FIGURE 3-6 False Positive Potential Deadlock's Source



The `get_token()` function in `din_philo_fix1.c` uses a while loop to synchronize the threads. A thread will not leave the while loop until it successfully gets a token (this occurs when `num_can_eat` is greater than zero). The while loop limits the number of simultaneous diners to four. However, the synchronization implemented by the while loop is not recognized by Thread Analyzer. It assumes that all five philosophers attempt to grab the chopsticks and eat concurrently, so it reports a potential deadlock. The following section details how to limit the number of simultaneous diners by using synchronizations that Thread Analyzer recognizes.

An Alternative System of Tokens

The following listing shows an alternative implementation of the system of tokens. This implementation still uses four tokens, so no more than four diners attempt to eat at the same time. However, this implementation uses the `sem_wait()` and `sem_post()` semaphore routines to limit the number of eating philosophers. This version of the source file is called `din_philo_fix2.c`.

Tip - If you downloaded the sample applications, you can copy the `din_philo_fix2.c` file from the `SolarisStudioSampleApplications/ThreadAnalyzer/din_philo` directory.

The following listing details `din_philo_fix2.c`:

```
1 /*
2  * Copyright (c) 2006, 2010, Oracle and/or its affiliates. All Rights Reserved.
3  * @(#)din_philo_fix2.c 1.3 (Oracle) 10/03/26
4  */
5
6 #include <pthread.h>
7 #include <stdio.h>
8 #include <unistd.h>
9 #include <stdlib.h>
10 #include <errno.h>
11 #include <assert.h>
12 #include <semaphore.h>
13
14 #define PHILOS 5
15 #define DELAY 5000
16 #define FOOD 100
17
18 void *philosopher (void *id);
19 void grab_chopstick (int,
20                     int,
21                     char *);
22 void down_chopsticks (int,
23                      int);
24 int food_on_table ();
25 void get_token ();
26 void return_token ();
27
28 pthread_mutex_t chopstick[PHILOS];
29 pthread_t philo[PHILOS];
30 pthread_mutex_t food_lock;
31 int sleep_seconds = 0;
32 sem_t num_can_eat_sem;
33
34
35 int
36 main (int argn,
37       char **argv)
38 {
39     int i;
40
41     pthread_mutex_init (&food_lock, NULL);
42     sem_init(&num_can_eat_sem, 0, PHILOS - 1);
43     for (i = 0; i < PHILOS; i++)
44         pthread_mutex_init (&chopstick[i], NULL);
45     for (i = 0; i < PHILOS; i++)
46         pthread_create (&philo[i], NULL, philosopher, (void *)i);
47     for (i = 0; i < PHILOS; i++)
48         pthread_join (philo[i], NULL);
49     return 0;
50 }
51
52 void *
```



```
53 philosopher (void *num)
54 {
55     int id;
56     int i, left_chopstick, right_chopstick, f;
57
58     id = (int)num;
59     printf ("Philosopher %d is done thinking and now ready to eat.\n", id);
60     right_chopstick = id;
61     left_chopstick = id + 1;
62
63     /* Wrap around the chopsticks. */
64     if (left_chopstick == PHILLOS)
65         left_chopstick = 0;
66
67     while (f = food_on_table ()) {
68         get_token ();
69
70         grab_chopstick (id, right_chopstick, "right ");
71         grab_chopstick (id, left_chopstick, "left");
72
73         printf ("Philosopher %d: eating.\n", id);
74         usleep (DELAY * (FOOD - f + 1));
75         down_chopsticks (left_chopstick, right_chopstick);
76
77         return_token ();
78     }
79
80     printf ("Philosopher %d is done eating.\n", id);
81     return (NULL);
82 }
83
84 int
85 food_on_table ()
86 {
87     static int food = FOOD;
88     int myfood;
89
90     pthread_mutex_lock (&food_lock);
91     if (food > 0) {
92         food--;
93     }
94     myfood = food;
95     pthread_mutex_unlock (&food_lock);
96     return myfood;
97 }
98
99 void
100 grab_chopstick (int phil,
101                int c,
102                char *hand)
103 {
104     pthread_mutex_lock (&chopstick[c]);
105     printf ("Philosopher %d: got %s chopstick %d\n", phil, hand, c);
106 }
```

```
107
108 void
109 down_chopsticks (int c1,
110                 int c2)
111 {
112     pthread_mutex_unlock (&chopstick[c1]);
113     pthread_mutex_unlock (&chopstick[c2]);
114 }
115
116
117 void
118 get_token ()
119 {
120     sem_wait(&num_can_eat_sem);
121 }
122
123 void
124 return_token ()
125 {
126     sem_post(&num_can_eat_sem);
127 }
```

This new implementation uses the semaphore `num_can_eat_sem` to limit the number of philosophers who can eat at the same time. The semaphore `num_can_eat_sem` is initialized to four, one less than the number of philosophers. Before attempting to eat, a philosopher calls `get_token()` which in turn calls `sem_wait(&num_can_eat_sem)`. The call to `sem_wait()` causes the calling philosopher to wait until the semaphore's value is positive, then changes the semaphore's value by subtracting one from the value. When a philosopher is done eating, he calls `return_token()` which in turn calls `sem_post(&num_can_eat_sem)`. The call to `sem_post()` changes the semaphore's value by adding one. Thread Analyzer recognizes the calls to `sem_wait()` and `sem_post()`, and determines that not all philosophers attempt to eat concurrently.

Note - You must compile `din_philo_fix2.c` with `-lrt` to link with the appropriate semaphore routines.

To compile `din_philo_fix2.c`, use the following command:

```
% cc -g -lrt -o din_philo_fix2 din_philo_fix2.c
```

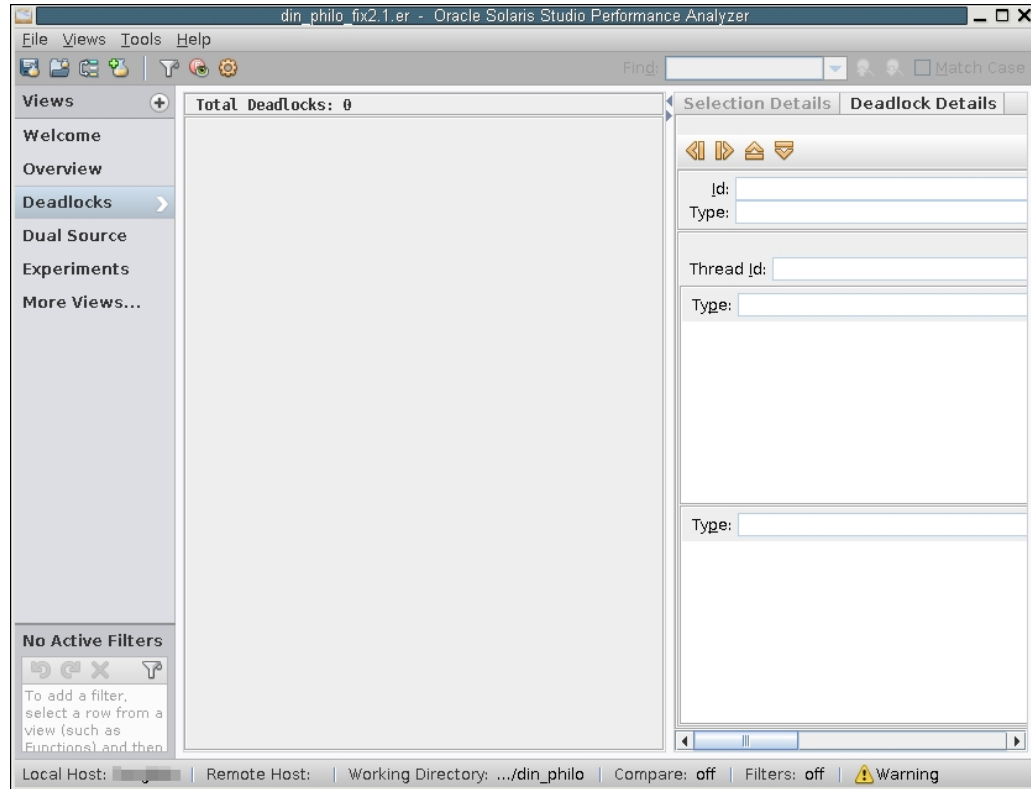
If you run this new implementation of the program `din_philo_fix2` several times, you will find that it terminates normally each time and does not hang.

To create an experiment on this new binary:

```
% collect -r deadlock -o din_philo_fix2.1.er din_philo_fix2
```

You will find that Thread Analyzer does not report any actual or potential deadlocks in the `din_philo_fix2.1.er` experiment, as the following figure shows.

FIGURE 3-7 Deadlocks Not Reported in `din_philo_fix2.c`



See [Appendix A, “APIs Recognized by Thread Analyzer”](#) for a listing of the threading and memory allocation APIs that Thread Analyzer recognizes.

APIs Recognized by Thread Analyzer

Thread Analyzer can recognize most standard synchronization APIs and constructs provided by OpenMP, POSIX threads, and Solaris threads. However, the tool cannot recognize user-defined synchronizations, and might report false positive data races if you employ such synchronizations. For example, the tool cannot recognize spin locking that is implemented through hand-coded assembly-language code.

Thread Analyzer User APIs

If your code includes user-defined synchronizations, insert user APIs supported by Thread Analyzer into the program to identify those synchronizations. This identification enables Thread Analyzer to recognize the synchronizations and reduce the number of false positives. Thread Analyzer user APIs are defined in `libtha.so` and are listed below.

TABLE A-1 Thread Analyzer User APIs

Routine Name	Description
<code>tha_notify_acquire_lock()</code>	This routine can be called immediately before the program tries to acquire a user-defined lock.
<code>tha_notify_lock_acquired()</code>	This routine can be called immediately after a user-defined lock is successfully acquired.
<code>tha_notify_acquire_writelock()</code>	This routine can be called immediately before the program tries to acquire a user-defined read/write lock in write mode.
<code>tha_notify_writelock_acquired()</code>	This routine can be called immediately after a user-defined read/write lock is successfully acquired in write mode.
<code>tha_notify_acquire_readlock()</code>	This routine can be called immediately before the program tries to acquire a user-defined read/write lock in read mode.
<code>tha_notify_readlock_acquired()</code>	This routine can be called immediately after a user-defined read/write lock is successfully acquired in read mode.
<code>tha_notify_release_lock()</code>	This routine can be called immediately before a user-defined lock or read/write lock is to be released.
<code>tha_notify_lock_released()</code>	This routine can be called immediately after a user-defined lock or read/write lock is successfully released.
<code>tha_notify_sync_post_begin()</code>	This routine can be called immediately before a user-defined post synchronization is performed.

Routine Name	Description
<code>tha_notify_sync_post_end()</code>	This routine can be called immediately after a user-defined post synchronization is performed.
<code>tha_notify_sync_wait_begin()</code>	This routine can be called immediately before a user-defined wait synchronization is performed.
<code>tha_notify_sync_wait_end()</code>	This routine can be called immediately after a user-defined wait synchronization is performed.
<code>tha_check_datarace_mem()</code>	This routine instructs Thread Analyzer to monitor or ignore accesses to a specified block of memory when doing data race detection.
<code>tha_check_datarace_thr()</code>	This routine instructs Thread Analyzer to monitor or ignore memory accesses by one or more threads when doing data race detection.

A C/C++ version and a Fortran version of the APIs are provided. Each API call takes a single argument `id`, whose value should uniquely identify the synchronization object.

In the C/C++ version of the APIs, the type of the argument is `uintptr_t`, which is 4 bytes long in 32-bit mode and 8 bytes long in 64-bit mode. You need to add `#include <tha_interface.h>` to your C/C++ source file when calling any of the APIs.

In the Fortran version of the APIs, the type of the argument is integer of kind `tha_sobj_kind` which is 8-bytes long in both 32-bit and 64-bit mode. You need to add `#include "tha_finterface.h"` to your Fortran source file when calling any of the APIs.

To uniquely identify a synchronization object, the argument `id` should have a different value for each different synchronization object. One way to do this is to use the value of the address of the synchronization object as the `id`. The following code example shows how to use the API to avoid a false positive data race.

EXAMPLE A-1 Example Using Thread Analyzer APIs to Avoid False Positive Data Races

```
# include <tha_interface.h>
...
/* Initially, the ready_flag value is zero */
...
/* Thread 1: Producer */
100 data = ...
101 pthread_mutex_lock (&mutex);
    tha_notify_sync_post_begin ((uintptr_t) &ready_flag);
102 ready_flag = 1;
    tha_notify_sync_post_end ((uintptr_t) &ready_flag);

103 pthread_cond_signal (&cond);
104 pthread_mutex_unlock (&mutex);

/* Thread 2: Consumer */
200 pthread_mutex_lock (&mutex);
    tha_notify_sync_wait_begin ((uintptr_t) &ready_flag);
201 while (!ready_flag) {
```

```

202     pthread_cond_wait (&cond, &mutex);
203 }
    tha_notify_sync_wait_end ((uintptr_t) &ready_flag);
204 pthread_mutex_unlock (&mutex);
205 ... = data;

```

For more information on the user APIs, see the `libtha(3)` man page.

Other Recognized APIs

The following sections detail the threading APIs which Thread Analyzer recognizes.

POSIX Thread APIs

See the “Multithreaded Programming Guide” in the Oracle Solaris documentation for more information about these APIs.

```

pthread_detach()
pthread_mutex_init()
pthread_mutex_lock()
pthread_mutex_timedlock()
pthread_mutex_reltimedlock_np()
pthread_mutex_timedlock()
pthread_mutex_trylock()
pthread_mutex_unlock()
pthread_rwlock_rdlock()
pthread_rwlock_tryrdlock()
pthread_rwlock_wrlock()
pthread_rwlock_trywrlock()
pthread_rwlock_unlock()
pthread_create()
pthread_join()
pthread_cond_signal()
pthread_cond_broadcast()
pthread_cond_wait()
pthread_cond_timedwait()
pthread_cond_reltimedwait_np()
pthread_barrier_init()
pthread_barrier_wait()
pthread_spin_lock()
pthread_spin_unlock()
pthread_spin_trylock()

```

```
pthread_rwlock_init()
pthread_rwlock_timedrdlock()
pthread_rwlock_reltimedrdlock_np()
pthread_rwlock_timedwrlock()
pthread_rwlock_reltimedwrlock_np()
sem_post()
sem_wait()
sem_trywait()
sem_timedwait()
sem_reltimedwait_np()
```

Oracle Solaris Thread APIs

See the “Multithreaded Programming Guide” in the Oracle Solaris documentation for more information about these APIs.

```
mutex_init()
mutex_lock()
mutex_trylock()
mutex_unlock()
rw_rdlock()
rw_tryrdlock()
rw_wrlock()
rw_trywrlock()
rw_unlock()
rwlock_init()
thr_create()
thr_join()
cond_signal()
cond_broadcast()
cond_wait()
cond_timedwait()
cond_reltimedwait()
sema_post()
sema_wait()
sema_trywait()
```

Memory Allocation APIs

```
calloc()
```



```
malloc()  
realloc()  
valloc()  
memalign()  
free()
```

See the `malloc(3C)` man page for information about the memory allocation APIs.

Memory Operations APIs

```
memcpy()  
memccpy()  
memmove()  
memchr()  
memcmp()  
memset()
```

See the `memcpy(3C)` man page for information about the memory operations APIs.

String Operations APIs

```
strcat()  
strncat()  
strlcat()  
strcasecmp()  
strncasecmp()  
strchr()  
strrchr()  
strcmp()  
strncmp()  
strcpy()  
strncpy()  
strncpy()  
strlcpy()  
strcspn()  
strspn()  
strdup()  
strlen()  
strpbrk()  
strstr()  
strtok()
```

See the `strcat(3C)` man page for information about the string operations APIs.

Realtime Library APIs

```
sem_post()
sem_wait()
sem_trywait()
sem_timedwait()
```

Atomic Operations (`atomic_ops`) APIs

```
atomic_add()
atomic_and()
atomic_cas()
atomic_dec()
atomic_inc()
atomic_or()
atomic_swap()
```

OpenMP APIs

Thread Analyzer recognizes OpenMP synchronizations, such as barriers, locks, critical regions, atomic regions, and taskwait.

See the [“Oracle Solaris Studio 12.4: OpenMP API User’s Guide”](#) for more information.

Tips for Using Thread Analyzer

This appendix includes some tips for using Thread Analyzer.

Compiling the Application

Tips for compiling an application before collecting an experiment:

- Use the `-g` compiler option when building application binaries. This enables Thread Analyzer to report line number information for data races and deadlocks.
- Compile with an optimization level less than `-xO3` when building application binaries. Compiler transformations might distort line number information and make the results difficult to understand.
- Thread Analyzer interposes on the routines shown in [“Memory Allocation APIs” on page 72](#). Linking to archive versions of memory allocation libraries might result in false positive data races being reported.

Instrumenting the Application for Data Race Detection

Tips for instrumenting an application for data race detection before collecting an experiment:

- The `collect -r race` command issues a warning if the binary is not instrumented for data race detection, as shown here:

```
% collect -r races a.out  
WARNING: Target `a.out' is not instrumented for datarace  
detection; reported datarace data may be misleading
```

- You can determine whether a binary is instrumented for data race detection by using the `nm` command and looking for calls to `tha` routines. If routines whose names begin with `__tha_` are shown, the binary is instrumented. Example output is shown below.

Source-level instrumentation:

```
% cc -xopenmp -g -xinstrument=datarace source.c
```

```
% nm a.out | grep __tha_  
[71] | 135408 | 0 | FUNC | GLOB | 0 | UNDEF | __tha_get_stack_id  
[53] | 135468 | 0 | FUNC | GLOB | 0 | UNDEF | __tha_src_read_w_frame  
[61] | 135444 | 0 | FUNC | GLOB | 0 | UNDEF | __tha_src_write_w_frame
```

Binary-level instrumentation:

```
% cc -xopenmp -g source.c  
% discover -i datarace -o a.out.i a.out  
% nm a.out.i | grep __tha_  
[88] | 0 | 0 | NOTY | GLOB | 0 | UNDEF | __tha_read_w_pc_frame  
[49] | 0 | 0 | NOTY | GLOB | 0 | UNDEF | __tha_write_w_pc_frame
```

Running the Application With collect

Tips for running an instrumented application to detect data races and deadlocks.

- Make sure that the Oracle Solaris system has all the required patches installed. The `collect` command lists any missing required patches. For OpenMP applications, the latest version of `libmtsk.so` is required.
- Instrumentation might cause a significant slowdown in execution time, 50 times or more, and an increase in memory consumption. You can try reducing the execution time by using a smaller data set. You can also try reducing the execution time by increasing the number of threads.
- To detect data races, make sure that the application is using more than one thread. For OpenMP, the number of threads can be specified by setting the environment variable `OMP_NUM_THREADS` to the desired number of threads, and setting the environment variable `OMP_DYNAMIC` to `FALSE`.

Reporting of Data Races

Tips for reporting of data races:

- Thread Analyzer detects data races at runtime. The runtime behavior of an application depends on the input data set used and operating system scheduling. Run the application under `collect` with different numbers of threads and with different input data sets. Also repeat experiments with a single data set to maximize the tool's chance of detecting data races.
- Thread Analyzer detects data races between different threads that are spawned from a single process. It does not detect data races between different processes.
- Thread Analyzer does not report the name of the variable accessed in a data race. However, you can determine the name of the variable by inspecting the source lines where the two

data race accesses occurred, and determining which variables are written to and read from on those source lines.

- In some cases, Thread Analyzer might report data races that did not actually occur in the program. These data races are called false positives. This usually happens when a user-implemented synchronization is used or when memory is recycled between threads. For example, if your code includes hand-coded assembly that implements spin locks, Thread Analyzer will not recognize these synchronization points. Insert calls to Thread Analyzer user APIs in your source code to notify Thread Analyzer about user-defined synchronizations. See [“False Positives” on page 34](#) and [Appendix A, “APIs Recognized by Thread Analyzer”](#) for more information.
- Data races reported using source-level instrumentation and binary-level instrumentation might not be identical. In binary-level instrumentation, shared libraries are instrumented by default as they are opened, whether they are statically linked in the program or opened dynamically by `dlopen()`. In source-level instrumentation, libraries are instrumented only if their sources are compiled with `-xinstrument=datarace`.

