

# Debugging Common Issues in Multithreaded Applications

A Comparison of Debugging a Multithreaded Application using  
GNU GDB and TotalView

A White Paper by Rogue Wave Software.



Rogue Wave Software  
5500 Flatiron Parkway,  
Suite 200  
Boulder, CO 80301, USA  
[www.roguewave.com](http://www.roguewave.com)

# Debugging Common Issues in Multithreaded Applications

A Comparison of Debugging a Multithreaded Application using GDB and TotalView

**by Rogue Wave Software**

© 2013 by Rogue Wave Software. All Rights Reserved

Printed in the United States of America

## Trademark Information

The Rogue Wave Software name, logo, and TotalView are registered trademarks of Rogue Wave Software, Inc. or its subsidiaries in the US and other countries. ReplayEngine is a trademark of Rogue Wave Software, Inc. or its subsidiaries. All other company, product, or brand names are the property of their respective owners.

**IMPORTANT NOTICE:** The information contained in this document is subject to change without notice. Rogue Wave Software, Inc. makes no warranty of any kind with regards to this material, including, but not limited to, the implied warranties of merchantability and fitness for a particular purpose. Rogue Wave Software, Inc. shall not be liable for errors contained herein or for incidental, consequential, or other indirect damages in connection with the furnishing, performance, or use of this material.

## TABLE OF CONTENTS

Abstract.....	4
Sample Multithreaded Application with Common Issues .....	4
Debugging a Simple, Multithreaded Application.....	5
Debugging a Multithreaded Application with GDB .....	6
Debugging a Multithreaded Application with TotalView .....	8
Investigating a Nondeterministic Issue.....	12
Investigating a Nondeterministic Issue with GDB .....	12
Investigating a Nondeterministic Issue with TotalView .....	15
Thread Control Capabilities of GDB vs. TotalView.....	20
Thread Control Capabilities of GDB .....	21
Thread Control Capabilities of TotalView .....	26
Displaying Local Variables of Different Threads .....	30
Displaying Local Variables of Different Threads Using GDB .....	30
Displaying Local Variables of Different Threads Using TotalView .....	31
Conclusion.....	34
About TotalView .....	34
About Rogue Wave Software.....	34

## Abstract

This paper describes several challenges that are commonly encountered when debugging multithreaded applications in order to compare the open source GNU GDB debugger (hereafter referred to as GDB) with the TotalView parallel debugger from Rogue Wave Software. The challenges highlighted will include:

- debugging a multithreaded application
- investigating nondeterministic issues
- demonstrating thread control capabilities
- displaying local variables

TotalView offers impressive feature capabilities that are either not included or not possible to perform with GDB. This paper explores some of these features and scenarios to demonstrate the differences between the two debugging tools.

## Sample Multithreaded Application with Common Issues

In the examples found in this paper – we will be referring to an application. The actual code is not included in this discussion, but the following description outlines the important details of the application and the bugs we will be tracking down in the following scenarios.

The application is a simple tcp-server that creates a new thread (super-thread) for processing each incoming client request. Each super-thread spawns a set of sub-threads which approximate the value of pi using the trapezoidal rule. The server allows one to define the maximum number of super-threads and the number of sub-threads. Three implemented bugs can be switched on and off for demonstration purposes. This could be done through the use of global variables, or the use of ifdefs in the code and rebuilding the application for each scenario.

The application is implemented in C and consists of these four functions:

```
int main( int argc, char* argv[] ) /* All super-threads are created here */

void *processRequest( void *sfd ) /* The sub-threads will be spawned in
                                   this function */

void *iterate( void *ind ) /* Sub-threads are calculating PI */

void onError( const char *errorMsg ) /* Error handling */
```

Properties referring to each super-thread are stored in an array of structures which is globally visible. Whether a thread is still alive or has already been terminated, these structures will hold the thread id, the session id, and the results calculated by the super-thread's sub-threads.



Introducing such an array (threadInfo tid[MAXTHREADS];) allows one to provide thread functions with a pointer to an index of this array while creating a thread in order to hand over all relevant super-thread related properties.

The following variables, which are used to store the index's integer value, will be mentioned in this document:

```
int *pindex; /* Pointer to an index of type integer, used in function
             main() as argument in pthread_create() */

void *sfd; /* Void pointer to an index of type integer, used in function
           processRequest() */

int index; /* Local integer variable used in function processRequest() */

void *ind; /* Void pointer to an index of type integer, used in function
           iterate() */

int index; /* Local integer variable used in function iterate() */
```

## Debugging a Simple, Multithreaded Application

The following example compares the capabilities of GDB with TotalView when debugging a simple, multithreaded application.

The first bug should be easy to find. The server calculates reasonable results when only one super-thread is running, as seen in Figure 1. However, when additional super-threads are created the results become incorrect, as shown in Figure 2.

```
luedtke@linux-8ude:~/tests/listener> ./runclient.sh localhost 25000 1000000 1
Try to connect server localhost on port no.: 25000

Please enter number of iterations: 1000000
luedtke@linux-8ude:~/tests/listener>
PI = 3.14159265358987349614
This result was presented to you by thread no.: 3077532528
```

Figure 1: Correct results when one super-thread is run



```
Try to connect server localhost on port no.: 25000

Please enter number of iterations: 1000000
PI = 0.97991465250752252114
This result was presented to you by thread no.: 2976418672

Please enter number of iterations: 1000000
luedtke@linux-8ude:~/tests/listener>
PI = 3.14159265358983974536
This result was presented to you by thread no.: 3001596784
```

Figure 2: Incorrect results when two super-threads are running

It seems that there is a conflict between threads that occurs in cases where more than one super-thread is running. The popular (some would say ubiquitous) GDB debugger<sup>1</sup> will be used to demonstrate this issue. We will also show how that same issue can be addressed using TotalView<sup>2</sup>.

### ***Debugging a Multithreaded Application with GDB***

The following steps outline a debug session using GDB:

- 1.) Start the server within GDB: **`gdb ./server`**.
- 2.) Set the breakpoints:  
As shown in Figure 3, it makes sense to place the first breakpoint in `main()`, in front of the function `pthread_create()`, where the super-threads are created.

The GDB command **`list main <return>`** lists the source code with line numbers. The first breakpoint is set at line 266 by typing **`break 266 <return>`**.

<sup>1</sup> <http://www.gnu.org/software/gdb/>

<sup>2</sup> <http://www.roguewave.com/products/totalview.aspx>



```

257         ithreadcount = 0;
258
259 (gdb)
260         tID[itcount].iAllocated = 1;
261         tID[itcount].iPos = itcount;
262         tID[itcount].iSocketfd = iSocketfd2;
263
264
265         /* Create new super-thread for handling the new request */
266         if (pthread_create(&tID[itcount].threadID, NULL, processRequest, (void*) pindex) != 0) {
267             onError("Unable to create thread");
268         }
269     }
270     /* End of infinite loop */
271
272     close(iSocketfd);
273
274     return 0;
275 }
(gdb) break 266
Breakpoint 1 at 0x804903c: file mtserver.c, line 266.
(gdb)

```

**Figure 3: Setting a breakpoint in GDB**

Additional breakpoints are placed at line 132 in the function ‘processRequest(),’ where the sub-threads will be spawned, and before the calculation starts at line 34 in the function ‘iterate().’ These breakpoints will stop the process when they are encountered by the debugger.

While GDB allows you to create thread-specific breakpoints, it does not prevent you from setting breakpoints on lines that are normally considered invalid, such as in front of comments or on blank lines. These incorrectly placed breakpoints will be recorded by the debugger but will not have the intended effect of stopping the program.

- 3.) Run the server by typing the GDB command **run <return>**.
- 4.) Start two instances of the client by opening a separate console and typing **./runclient.sh localhost 25000 100000 2 <return>**.

The server process runs until it hits breakpoint 1. Type the GDB command **info threads <return>** to list the currently running threads. Only the main thread labelled with #1 is listed. Next, check the current value of the variable **pindex**. This variable points to the index into an array where each thread stores its calculated result and is passed as an argument to the ‘pthread\_create()’ function. The GDB command **print \*pindex <return>** de-references the pointer and displays its contents, currently ‘0.’

- 5.) Type the GDB command **c <return>** to continue execution to the next breakpoint.
- 6.) As seen in Figure 4, the execution stops at breakpoint 2, line 132, inside of the function processRequest(). The GDB command **info threads <return>** shows that two threads are now running, indicating that a second thread, a super-thread, was created in response to a request from the client. It is this super-thread (thread 2) that is stopped at breakpoint 2.



```

Try: zypper install -C "debuginfo(build-id)=7eb4e169e926464393ef2e98d99c37f56d5f5858"
No portnumber provided, listening on port no. : 25000
Breakpoint 1, main (argc=1, argv=0xbffff124) at mtserver.c:266
266         if (pthread_create(&tID[itcount], threadID, NULL, processRequest, (void*) pindex) != 0) {
(gdb) info threads
* 1 Thread 0xb7e4f6c0 (LWP 6113): main (argc=1, argv=0xbffff124) at mtserver.c:266
(gdb) print *pindex
$1 = 0
(gdb) c
Continuing.
[New Thread 0xb7e4eb70 (LWP 6122)]
Switching to Thread 0xb7e4eb70 (LWP 6122)

Breakpoint 2, processRequest (sfd=0xbffff064) at mtserver.c:132
132         for (i = 0; i < SUBTHREADS; i++) {
(gdb) print index
$2 = 0
(gdb) print *((int*) sfd)
$3 = 1
(gdb) info threads
* 2 Thread 0xb7e4eb70 (LWP 6122): processRequest (sfd=0xbffff064) at mtserver.c:132
  1 Thread 0xb7e4f6c0 (LWP 6113): 0xffffe422 in __kernel_vsyscall ()
(gdb) █

```

Figure 4: GDB stopped at the second breakpoint (thread 2 in processRequest())

The GDB **print index <return>** command shows the content of the variable *index* is set to 0. In contrast, the value the pointer *sfd* references has changed to 1. In order to print the value referenced by *sfd*, it must be cast to an int pointer and de-referenced: (**print \*((int\*) sfd) <return>**).

This appears to be the cause of the incorrect result. The pointer *sfd* references the value 1. When this value is passed to the newly-created sub-threads it causes them to write their results to *sum[1]* instead of *sum[0]*.

The value for *sfd* was passed to the processRequest() routine from main() via the variable *pindex*. Looking up the string 'pindex' in main() using the GDB command **reverse-search pindex <return>** shows that 'pindex' occurs in lines 252 and 255.

```

251         if (BUG_1) {
252             pindex = &itcount;
253         }
254         else {
255             pindex = &(tID[itcount].ipos);
256         }

```

It is obvious that *pindex* accidentally points to the variable *itcount* instead of pointing to the already stored index *ipos* in the *tID* array of structures.

**Debugging a Multithreaded Application with TotalView**

Next we will take the TotalView Debugger through the same scenario of debugging of a simple, multithreaded application. The TotalView debug session includes the following steps:

- 1.) Start the server within TotalView: **totalview ./server.**





2.) Set the breakpoints:

Just as was done in the GDB session, the first breakpoint is placed in main(), in front of the function pthread\_create(), where the super-threads are created.

Additional breakpoints are placed at line 132 in the function 'processRequest()' and before the calculation starts at line 34 in the function 'iterate()'.

Breakpoints in TotalView are a specific type of Action Point. All current Action Points are displayed under the 'Action Points' tab at the bottom of the main window. Clicking on an Action Point in the list displays that line in the source code window. To place a breakpoint in TotalView, simply click on the line number for a valid line in the source code window. Lines that are valid locations for breakpoints are marked with rectangles around the line numbers.

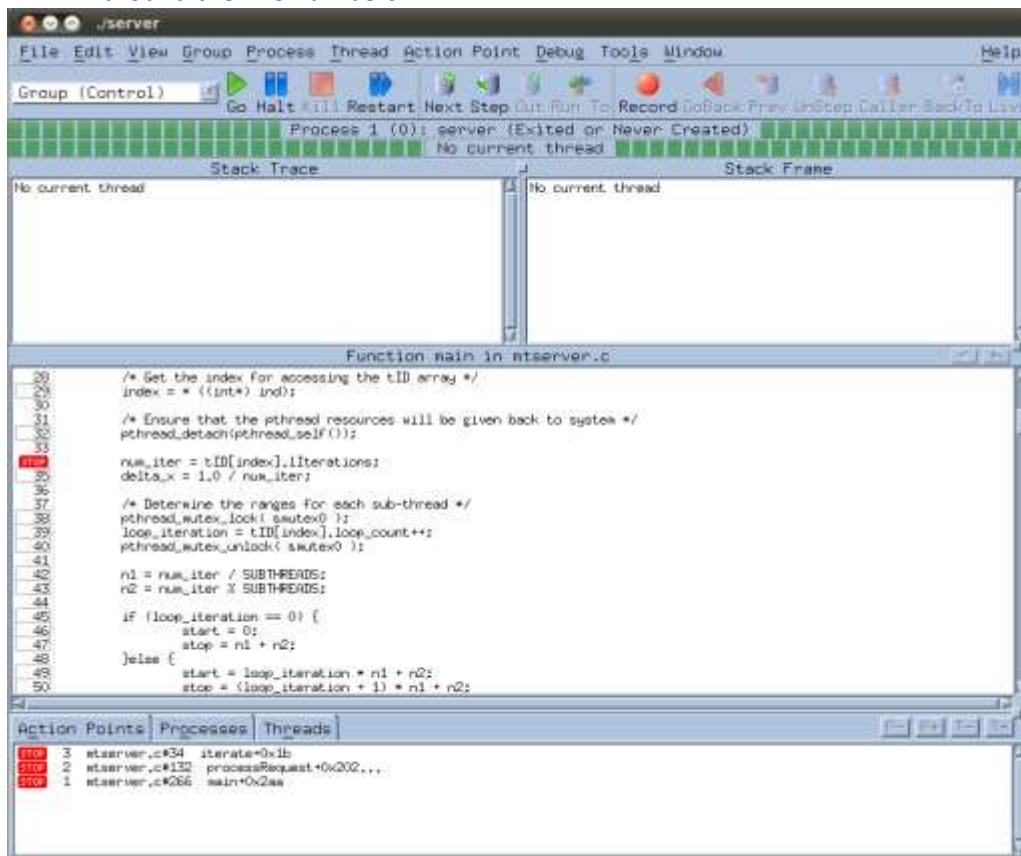


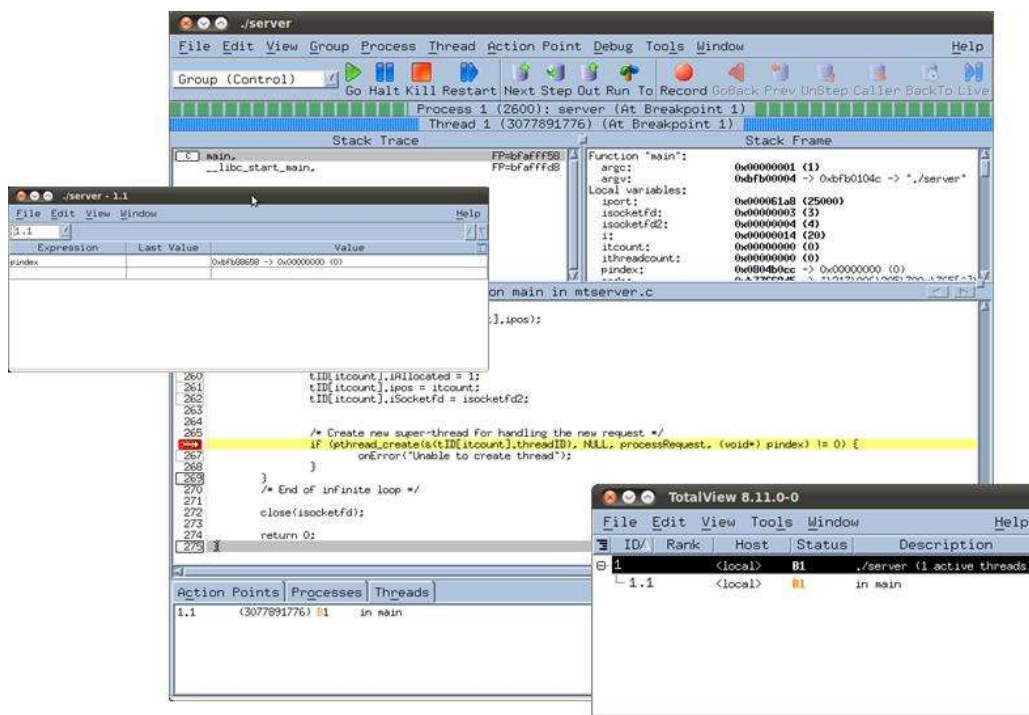
Figure 5: Placing breakpoints in TotalView

- 3.) **TotalView:** Push the 'Go' button in the toolbar to start the server.
- 4.) Start two instances of the client by opening a separate console and typing  
`./runclient.sh localhost 25000 100000 2`  
<return>.



When the TotalView process stops at breakpoint 1, the source file containing the breakpoint appears in the source code pane with the breakpoint line highlighted. The main thread is displayed in the Root Window as well as in the tabbed pane. The TotalView thread id is 1.1, the system thread id is displayed in brackets.

To check the value of the variable *pindex*, move the mouse pointer over the variable name in the source code pane. To continuously monitor the value of any variable, add it to a watch list, called an Expression List, by right-clicking on the variable and selecting 'Add to Expression List' in the pop-up menu. As you can see in Figure 6, items in the Expression List display in a separate, thread-specific window.



**Figure6: TotalView stopped at the first breakpoint (in main()) with the variable *pindex* displayed in the Expression List**

- 5.) Push the 'Go' button to continue execution to the next breakpoint.
- 6.) The Root Window shows that the first super-thread is created. The thread is labelled 1.2, and is stopped at breakpoint 2 in the processRequest() function.

As you can see in figure 7, the Expression List shows that *pindex* points to a value of 1.




Figure 7: TotalView Expression List and Root Window after creating the first super-thread

A double-click on thread line 1.1 in the Root Window moves the focus back to the main() function. The source code of main() is displayed in the source code pane.

As you can see in Figure 8, the Edit menu of TotalView's process window contains the option 'Find String' to look up occurrences of the string pindex, which leads to the bug mentioned above.

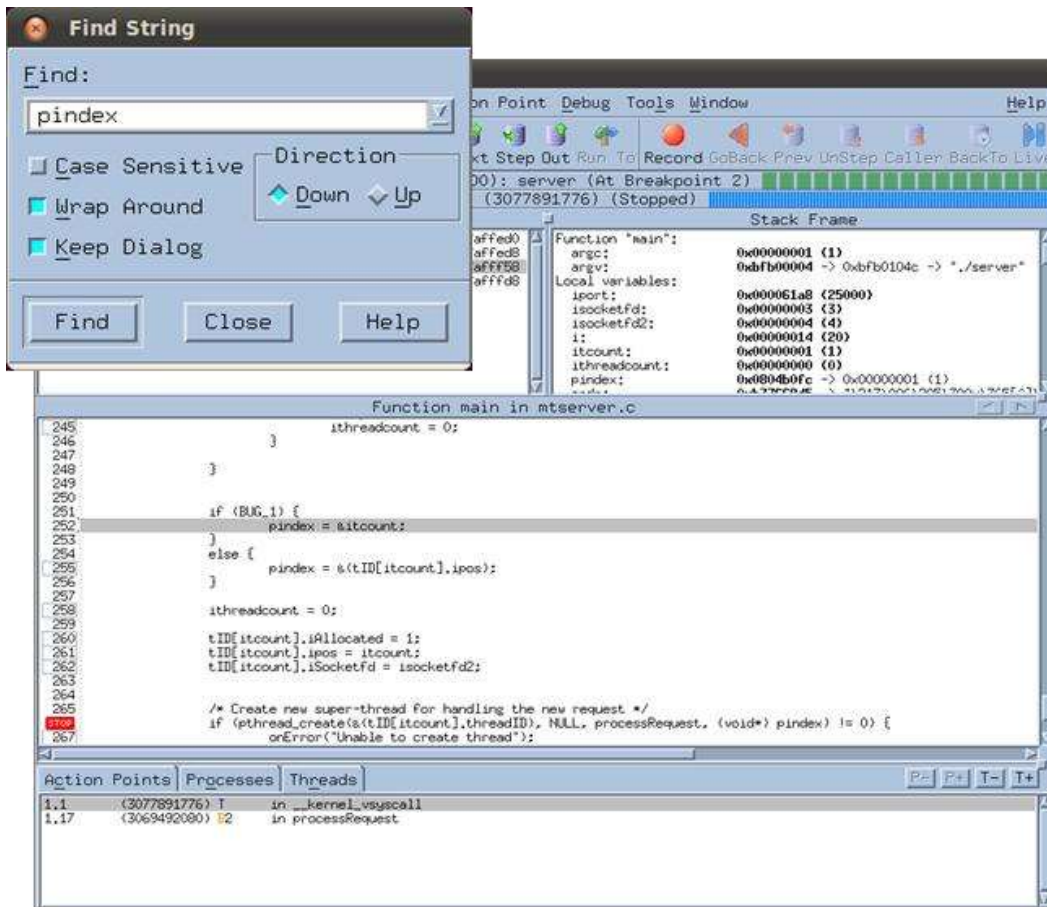


Figure 8: TotalView finding a string in the source code

TotalView's graphical user interface makes it simple to set the breakpoints, look at variable data, and see how the status of the different threads changes.



## Investigating a Nondeterministic Issue

We will now move on to investigating another type of debugging challenge that is commonly encountered when developing multithreaded applications - a nondeterministic issue. The server is rebuilt for the second debugging example with the second implemented bug turned on.

Most requests sent by the client result in correct results, but as you can see in Figure 9, repeated requests produce an intermittent, nondeterministic error in the calculated value for pi.

```
PI = 3.00859149508113032923
This result was presented to you by thread no.: 2909031280

PI = 3.14159265358986239391
This result was presented to you by thread no.: 3001351024

PI = 3.14159265358989125971
This result was presented to you by thread no.: 2841889648

PI = 3.14159265358984329808
This result was presented to you by thread no.: 3034921840

PI = 3.14159265358981576455
```

Figure 9: Nondeterministic error

Very often, nondeterministic issues like this are caused by race conditions. An approach in this case is to trigger the issue by placing a conditional breakpoint. The investigation starts by examining how this would be accomplished by using the GDB debugger.

### *Investigating a Nondeterministic Issue with GDB*

- 1.) After starting the server within GDB, the command **break 151 if approx < 3.14** **<return>** causes the program to stop at line 151 when the result of **approx** is less than 3.14.
- 2.) As you can see in Figure 10, running the server several times eventually triggers the breakpoint when the conditional expression is true.



```

[New Thread 0xae3cb70 (LWP 32567)]
[New Thread 0xade3ab70 (LWP 32569)]
[New Thread 0xae63bb70 (LWP 32568)]
[New Thread 0xad639b70 (LWP 32570)]
[New Thread 0xace38b70 (LWP 32571)]
[New Thread 0xab635b70 (LWP 32572)]
[New Thread 0xaae34b70 (LWP 32573)]
[New Thread 0xa9631b70 (LWP 32576)]
[New Thread 0xaa633b70 (LWP 32574)]
[New Thread 0xa8e30b70 (LWP 32577)]
[New Thread 0xa9e32b70 (LWP 32575)]
[Thread 0xb6e4cb70 (LWP 32550) exited]
[New Thread 0xb6e4cb70 (LWP 32578)]
[New Thread 0xa862fb70 (LWP 32579)]
[New Thread 0xa7e2eb70 (LWP 32580)]
[Thread 0xb5649b70 (LWP 32555) exited]
[New Thread 0xa762db70 (LWP 32581)]
[Switching to Thread 0xb1e42b70 (LWP 32548)]

Breakpoint 1, processRequest (sfd=0x804b2ac) at mtserver.c:151
151          thisThreadID = pthread_self();
(gdb) print approx
$2 = 3.0065293158778639
(gdb)

```

Figure 10: GDB, breakpoint triggered when a conditional expression is met

- 3.) As seen in figure 11, to get an idea about what is going on with the referring sub-threads, use the command **info threads <return>**.

```

New Thread 0xb3e46b70 (LWP 334)]
 50 Thread 0xb3e46b70 (LWP 334)  0xffffe422 in __kernel_vsyscall ()
 49 Thread 0xb3645b70 (LWP 333)  0xffffe422 in __kernel_vsyscall ()
 48 Thread 0xb4e48b70 (LWP 331)  0x08048810 in pthread_mutex_lock@plt
 ()
 47 Thread 0xb1641b70 (LWP 332)  0x08048a55 in iterate (
 ind=0x804b1ec) at mtserver.c:61
 46 Thread 0xad639b70 (LWP 330)  0xb7fb7d03 in pthread_mutex_lock ()
 from /lib/libpthread.so.0
 45 Thread 0xaa633b70 (LWP 329) (Exiting) 0xb7fb4980 in __nptl_death
 event () from /lib/libpthread.so.0
 44 Thread 0xb0e40b70 (LWP 328)  0xffffe422 in __kernel_vsyscall ()
 42 Thread 0xab635b70 (LWP 326)  0x08048a41 in iterate (
 ind=0x804b27c) at mtserver.c:60
 41 Thread 0xb1e42b70 (LWP 325)  0xb7fb7ccf in pthread_mutex_lock ()
 from /lib/libpthread.so.0
 38 Thread 0xabe36b70 (LWP 322)  0xb7fb4970 in __nptl_create_event
 () from /lib/libpthread.so.0
 35 Thread 0xb2e44b70 (LWP 319)  0xb7fb4970 in __nptl_create_event
 () from /lib/libpthread.so.0
 32 Thread 0xa8e30b70 (LWP 316) (Exiting) 0xb7fb4980 in __nptl_death
 event () from /lib/libpthread.so.0
 31 Thread 0xa9631b70 (LWP 315)  processRequest (sfd=0x804b21c)
 at mtserver.c:151

```

Figure 11: Output from the GDB info threads command

A reasonable hypothesis for the incorrect result is the assumption that a sub-thread was not able to finish its work until the result was sent back to the requestor. In this case this sub-thread would be stuck in the `iterate()` function.

The status view on the current threads shows that threads 42 and 47 are still in the `iterate()` function. Do they belong to the current super-thread which has initiated the stop at the breakpoint?

- 4.) The command **backtrace full <return>** shows all local variables of the thread currently in focus. The focus can be switched to another thread by using the command **thread <threadnumber> <return>**.

As you can see in Figure 12, examination of the index variables in the two sub-threads shows they do not belong to the super-thread.

```

151      thisThreadID = pthread_self();
(gdb) backtrace all
No symbol "all" in current
(gdb) backtrace full
#0  processRequest (sf=0)
    __socketfd = 15
    readNum = 10
    writeNum = 37
    i = 4
    index = 7
    thisThreadID = 0
(gdb) backtrace full
num_iter = 1000000
approx_x = 3.109245
delta_x = 9.999999
endp = 0xa9631354
readBuffer = "100"
writeBuffer = "\0"
pth = (2964520560)
#1  0xb7fb56e5 in start_t
No symbol table info avail
#2  0xb7fb5600 in ?? ()
No symbol table info avail
(gdb)

#0  processRequest (sf=0)
    __socketfd = 15
    readNum = 10
    writeNum = 37
    i = 4
    index = 7
    thisThreadID = 0
(gdb) backtrace full
num_iter = 10000000
l = 1236178
n1 = 2500000
n2 = 0
start = 0
stop = 2500000
loop_iteration = 0
index = 9
approx = 0
x = 0.12361775
f_x = 3.92979462680
delta_x = 9.9999999
I 0xb7fb56e5 in start_thread
symbol table info available
I 0xb7fb5600 in ?? ()
No symbol table info available
(gdb)

#1  Thread 0xb7e4f6c0 (LWP 32737) 0xffff422 in __kernel_vsyscall ()
(gdb) thread 47
[Switching to thread 47 (Thread 0xb1641b70 (LWP 322))]#0  0xb048a55 in
iterate (ind=0x804b27c) at mtsrvr.c:60
n iterate (ind=0x804b27c) at mtsrvr.c:60
    x = delta_x * f1 * 0.5;
(gdb) backtrace full
#0  0xb048a55 in iterate (ind=0x804b27c) at mtsrvr.c:61
num_iter = 10000000
l = 8916877
n1 = 2500000
n2 = 0
start = 7500000
stop = 10000000
loop_iteration = 3
index = 6
approx = 8.4879831658373518e-314
x = 0.89168775
f_x = 2.2281705658433677
delta_x = 9.999999999999995e-08
#1  0xb7fb56e5 in start_thread () from /lib/libpthread.so.0
No symbol table info available.
#2  0xb7fb5600 in ?? () from /lib/libpthread.so.0
No symbol table info available.
(gdb)

```

Figure 12: Comparing local variables of different threads in GDB

The next possible candidates are threads 41, 46, and 48. These threads are in the `pthread_mutex_lock()` function that is called inside of `iterate()`. Switching the focus to these threads using the GDB command **thread <threadnumber> <return>**, moving up the backtrace until inside `iterate()` using **up <return>**, and displaying the local variables using **backtrace full <return>** still does not uncover the thread that caused the error.

The only way to solve this issue seems to be by placing breakpoints at the end of function `iterate()`, running the process until the last thread hits the breakpoint, and trying to determine the referring super-thread. This is very time consuming because of the error's sporadic occurrence and GDB's limited support for solving nondeterministic issues. The issue may not be solved in a reasonable timeframe because the `iterate` function may be called many times before the error manifests. Subsequent thread activity may also change the critical state, obscuring the clues needed to make a clear diagnosis. Now we will examine how this works with TotalView.



## Investigating a Nondeterministic Issue with TotalView

The following steps demonstrate TotalView's investigation of the same issue.

- 1.) Analogous to the technique used with GDB, a conditional breakpoint is placed at line 151. As you can see in Figure 13, the process stops when `approx < 3.14`.

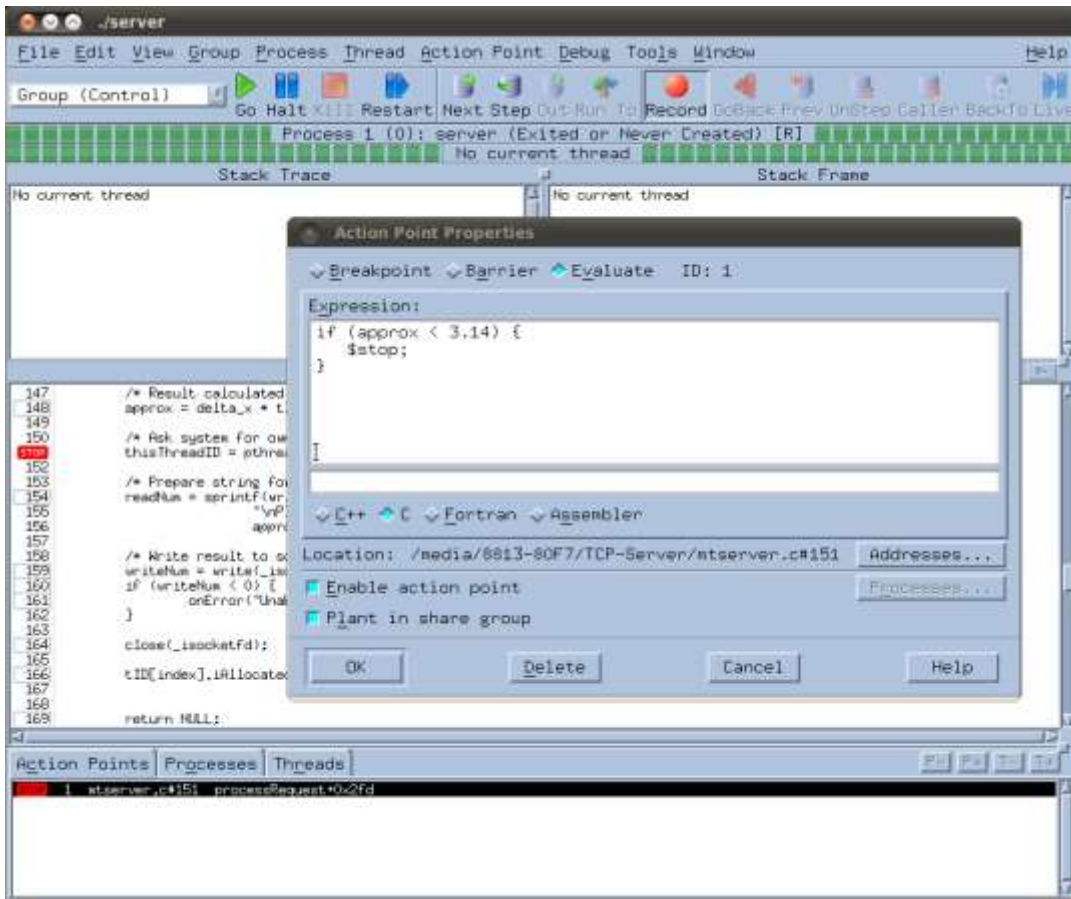


Figure 13: TotalView conditional breakpoint configuration

- 2.) As shown in Figure 14, turn on the TotalView reverse debugging feature called ReplayEngine to record the debugging session. ReplayEngine records the execution history of a program and makes it available for examination. It allows stepping back to any line in the code that was executed while recording was enabled, restoring the state of the entire process to what it was when that line was executed. While very handy in any debugging session, this functionality is indispensable for non-deterministic problems, such as this one.



Figure 14: Switching on the TotalView reverse debugging feature, ReplayEngine



- 3.) Push the 'Go' button to start the server and send a number of client requests. As you can see in Figure 15, the process by stops because a super-thread met the condition at the conditional breakpoint.

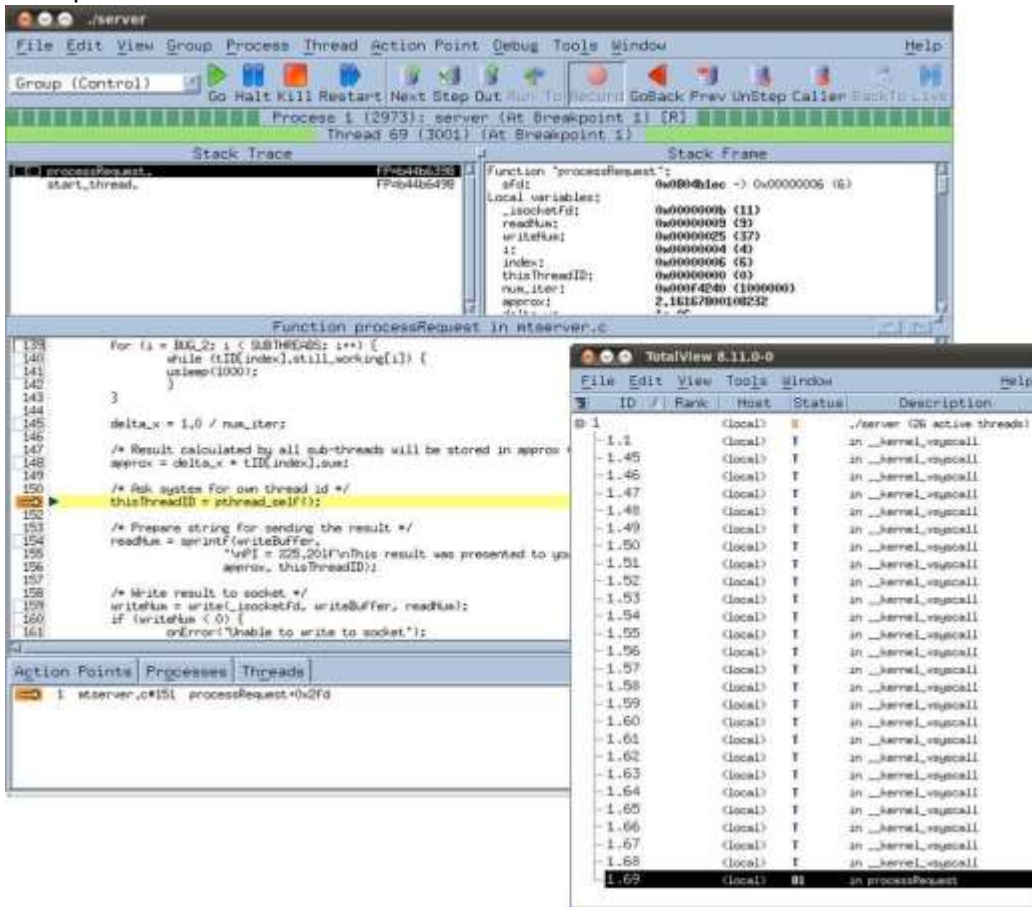


Figure 15: TotalView stops at the conditional breakpoint

- 4.) As seen in Figure 16, use the TotalView ReplayEngine to step back to line 148 and restore the entire process to its state when that line was executed. Examine the status of each thread at the moment the error occurred.



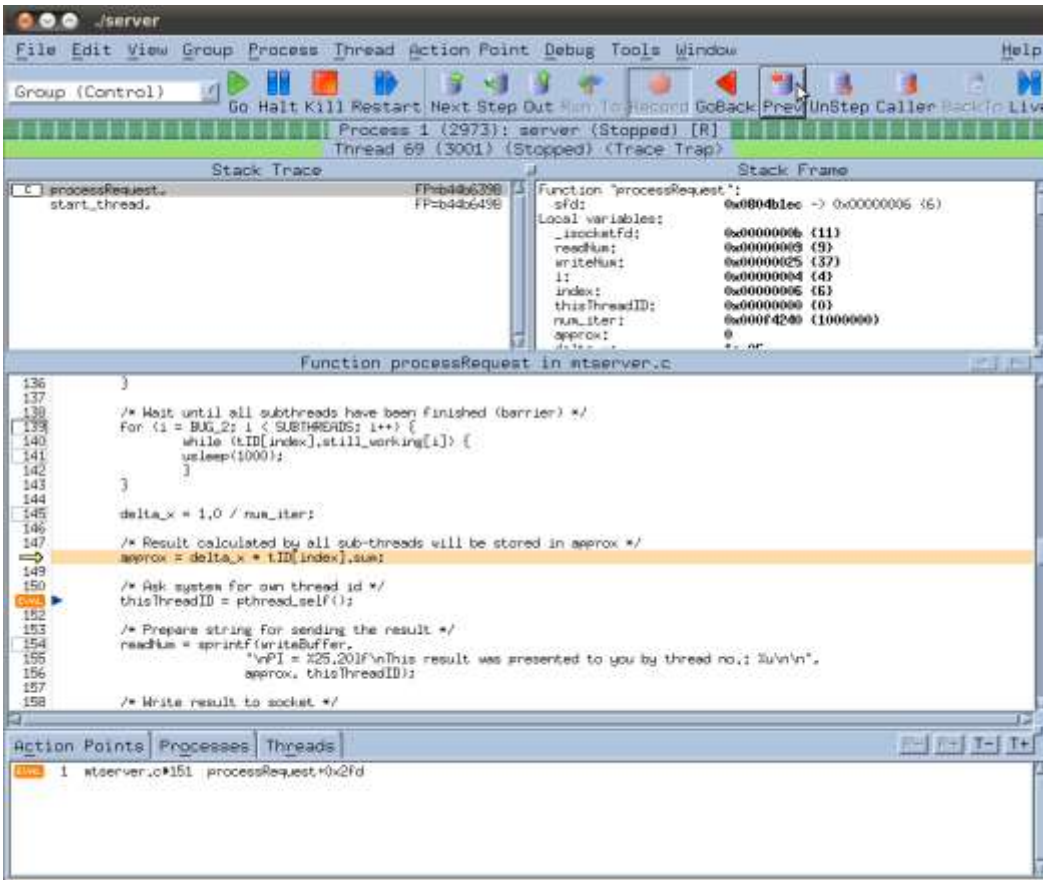


Figure 16: Stepping backwards using the TotalView ReplayEngine

- 5.) Following the same hypothesis as in the GDB session, test to see if one of these threads is still executing in the `iterate()` function. The value of its local variable `index` should be 6 because that is the value of the referring super-thread's `index` variable. Unfortunately, no thread seems to be currently running in `iterate()`, but, as you can see in Figure 17, the Stack Trace pane shows that `__kernel_vsyscall()` was called from within `iterate()` by several threads.

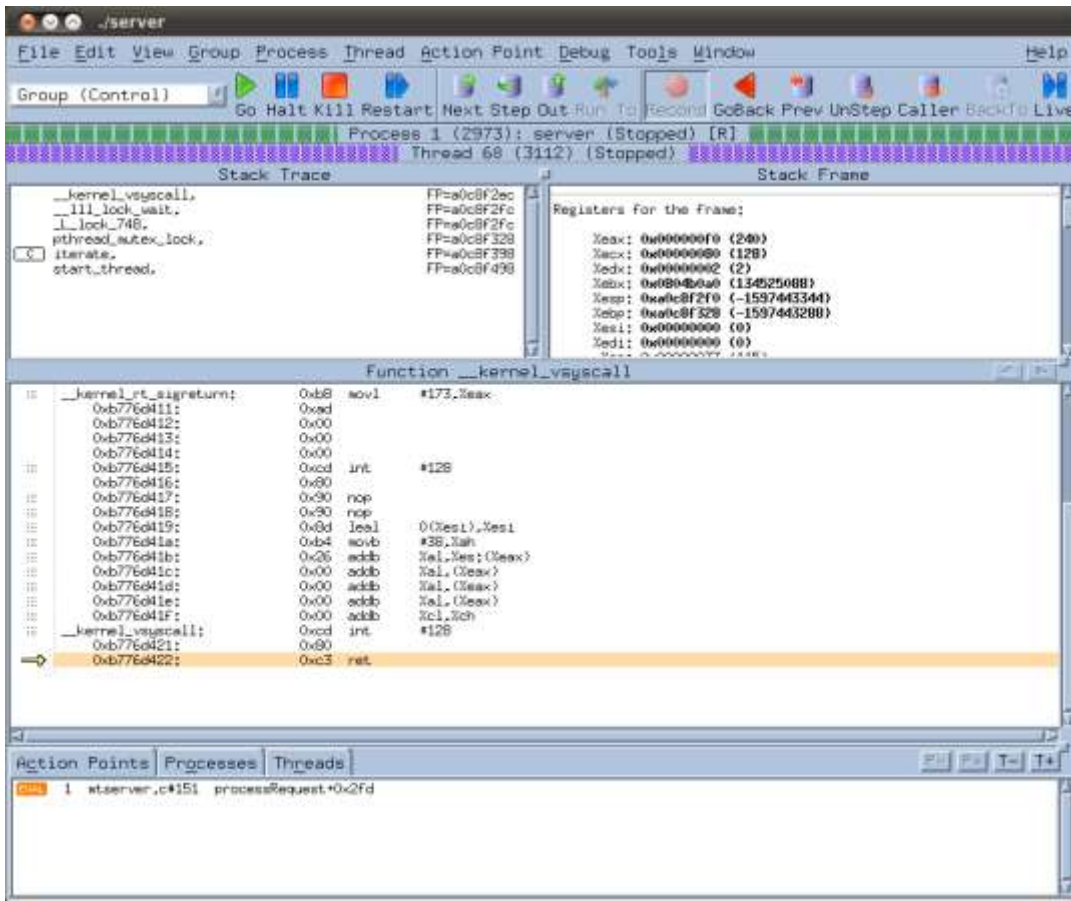


Figure 17: Thread in \_\_kernel\_vsyscall() in TotalView

- 6.) Click on 'iterate' in the Stack Trace pane and check the local variables in the referring Stack Frame pane to see that thread 68, which belongs to the super-thread at the breakpoint, has not yet finished the iterate() function. As you can see in Figure 18, this is the reason for the incorrect result.

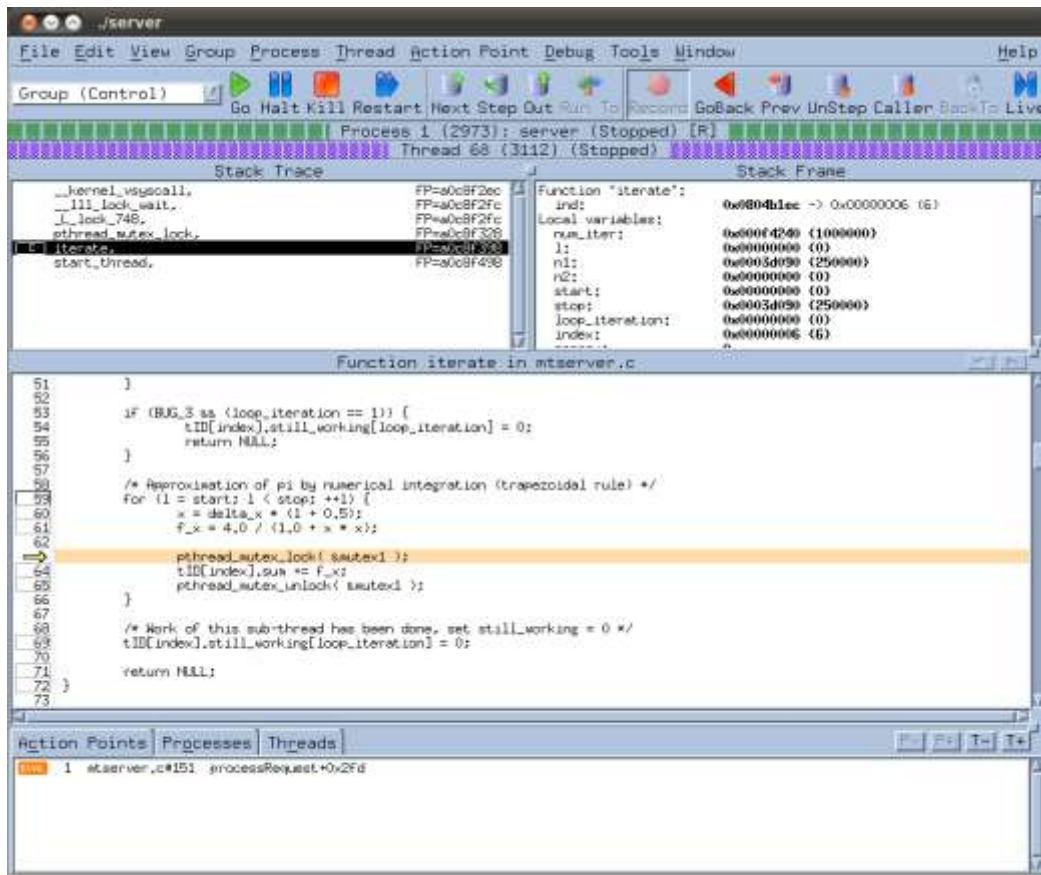


Figure 18: One sub-thread still executing iterate() in TotalView

The Stack Frame pane for thread 68 shows that the value of variable **loop\_iteration** is 0. This means the status of this thread is stored at **tID[6].still\_working[0]**. Using ReplayEngine and stepping the super-thread backwards through the barrier at lines 139 to 143 demonstrates the status of this thread will be disregarded. The issue has been solved.

```

for (i = BUG_2 ; i < SUBTHREADS; i++) {
    while (tID[index].still_working[i]) {
        usleep(1000);
    }
}

```

The value of i starts at 1 and increments upwards, so the super thread never checks to see if the thread for loop iteration 0 is done.

Notice how straightforwardly we could approach the solution to this issue; having the recorded execution history makes it really easy to check what the program is doing as we work backwards through the logic of the program from the incorrect result to the root cause of the bug.



## Thread Control Capabilities of GDB vs. TotalView

For the third debugging session, we rebuild the server with bug 3 enabled.

In this case, as can be seen in Figure 19, the results are identical for every request but the values are always lower than 3.14. The reason for this behaviour could be that one or more sub-threads do not finish the approximation properly.

```
PI = 2.26691687009408848752
This result was presented to you by thread no.: 3042986864

PI = 2.26691687009411158016
This result was presented to you by thread no.: 2975845232

PI = 2.26691687009401521280
This result was presented to you by thread no.: 2866740080

PI = 2.26691687009400677510
This result was presented to you by thread no.: 2883525488

PI = 2.26691687009401476871
This result was presented to you by thread no.: 2875132784

PI = 2.26691687009401965369
This result was presented to you by thread no.: 2757634928
```

**Figure 19: Results of a series of requests with BUG\_3 enabled**

In order to investigate this bug the following approach is planned:

1. Set a breakpoint at the very beginning of the `iterate()` function such that all sub-threads associated with a particular request are stopped.
2. Perform a few single-steps of these threads, in lockstep, until the calculation of the different `iteration()` sections start.
3. Move each individual thread in single-step mode in order to analyze its behaviour.

## Thread Control Capabilities of GDB

Again, the debug session will first be demonstrated using GDB.

The capability of controlling threads independently is limited in GDB. The ability to run a single thread is dependent on the referring operating system's ability to lock the scheduler. Furthermore, GDB generally stops the entire process when a single thread hits a breakpoint.

1. To deal with these limitations of GDB, the debugging process will start by setting two breakpoints. The first one is placed at line 132 where the sub-threads are created in function `processRequest()`. The second breakpoint is placed on line 29 at the beginning of the function `iterate()`. The GDB commands `break 132 <return>` and `break 29 <return>` accomplish this.
2. Running the client by typing `./client localhost 25000 100000 <return>` or `./runclient.sh localhost 25000 100000 1 <return>` in a separate console window will cause the server to run to the first breakpoint, as seen in Figure 20. The GDB command `info threads <return>` helps identify the super-thread.

```
(gdb) break 132
Breakpoint 1 at 0x8048ca4: file mtserver.c, line 132.
(gdb) break 29
Breakpoint 2 at 0x8048957: file mtserver.c, line 29.
(gdb) run
Starting program: /home/luedtke/tests/listener/server
Missing separate debuginfo for /lib/ld-linux.so.2
Try: zypper install -C "debuginfo(build-id)=d7706cbaa0ca09319cb645eac789cb8399078797"
Missing separate debuginfo for /lib/libpthread.so.0
Try: zypper install -C "debuginfo(build-id)=964690b0ca2ed321e995340684e09981f5f986ad"
[Thread debugging using libthread_db enabled]
Missing separate debuginfo for /lib/libc.so.6
Try: zypper install -C "debuginfo(build-id)=7eb4e169e926464393ef2e98d99c37f56d5f5858"
No portnumber provided, listening on port no.: 25000
[New Thread 0xb7e4eb70 (LWP 2510)]
[Switching to Thread 0xb7e4eb70 (LWP 2510)]

Breakpoint 1, processRequest (sfd=0x804b0cc) at mtserver.c:132
132          for (i = 0; i < SUBTHREADS; i++) {
(gdb) info threads
* 2 Thread 0xb7e4eb70 (LWP 2510) processRequest (sfd=0x804b0cc)
   at mtserver.c:132
  1 Thread 0xb7e4f6c0 (LWP 2503) 0xffffe422 in __kernel_vsyscall ()
(gdb) █
```

Figure 20: GDB stopped at the first breakpoint

- As seen in Figure 21, typing **continue** **<return>**, or **c** **<return>**, will stop the first sub-thread at breakpoint 29 in function "iterate()."

```
Missing separate debuginfo for /lib/libc.so.6
Try: zypper install -C "debuginfo(build-id)=7eb4e169e926464393ef2e98d99c37f56d5f5858"
No portnumber provided, listening on port no.: 25000
[New Thread 0xb7e4eb70 (LWP 2510)]
[Switching to Thread 0xb7e4eb70 (LWP 2510)]

Breakpoint 1, processRequest (sfd=0x804b0cc) at mtserver.c:132
132         for (i = 0; i < SUBTHREADS; i++) {
(gdb) info threads
* 2 Thread 0xb7e4eb70 (LWP 2510)  processRequest (sfd=0x804b0cc)
   at mtserver.c:132
  1 Thread 0xb7e4f6c0 (LWP 2503)  0xffffe422 in __kernel_vsyscall ()
(gdb) c
Continuing.
[New Thread 0xb764db70 (LWP 2511)]
[Switching to Thread 0xb764db70 (LWP 2511)]

Breakpoint 2, iterate (ind=0x804b0cc) at mtserver.c:29
29         index = * ((int*) ind);
(gdb) info threads
[New Thread 0xb6e4cb70 (LWP 2512)]
  4 Thread 0xb6e4cb70 (LWP 2512)  0xb7f22848 in clone () from /lib/libc.so.6
* 3 Thread 0xb764db70 (LWP 2511)  iterate (ind=0x804b0cc) at mtserver.c:29
  2 Thread 0xb7e4eb70 (LWP 2510)  0xb7f22848 in clone () from /lib/libc.so.6
  1 Thread 0xb7e4f6c0 (LWP 2503)  0xffffe422 in __kernel_vsyscall ()
(gdb) █
```

Figure 21: GDB stopped at the second breakpoint

- In order to perform single-steps of only thread 3, set the scheduler locking mode to 'step': **set scheduler-locking step** **<return>**.
- Single-step thread 3 using either the command **step** **<return>** (stepping into functions) or **next** **<return>** (stepping over functions). After executing two single-steps, thread 4 will also enter the function `iterate()`, but it will stay at breakpoint 2, as seen in Figure 22.
- The focus has automatically been switched to thread 4, so it has to be manually switched back to thread 3 using the command **thread 3** **<return>**.

```

(gdb) info threads
[New Thread 0xb6e4cb70 (LWP 2512)]
  4 Thread 0xb6e4cb70 (LWP 2512) 0xb7f22848 in clone () from /lib/libc.so.6
* 3 Thread 0xb764db70 (LWP 2511) iterate (ind=0x804b0cc) at mtserver.c:29
  2 Thread 0xb7e4eb70 (LWP 2510) 0xb7f22848 in clone () from /lib/libc.so.6
  1 Thread 0xb7e4f6c0 (LWP 2503) 0xffffe422 in __kernel_vsyscall ()
(gdb) set scheduler-locking step
(gdb) step
32          pthread_detach(pthread_self());
(gdb) step
[Switching to Thread 0xb6e4cb70 (LWP 2512)]

Breakpoint 2, iterate (ind=0x804b0cc) at mtserver.c:29
29          index = * ((int*) ind);
(gdb) info threads
[New Thread 0xb664bb70 (LWP 2537)]
  5 Thread 0xb664bb70 (LWP 2537) 0xb7f22848 in clone () from /lib/libc.so.6
* 4 Thread 0xb6e4cb70 (LWP 2512) iterate (ind=0x804b0cc) at mtserver.c:29
  3 Thread 0xb764db70 (LWP 2511) iterate (ind=0x804b0cc) at mtserver.c:34
  2 Thread 0xb7e4eb70 (LWP 2510) 0xb7f22848 in clone () from /lib/libc.so.6
  1 Thread 0xb7e4f6c0 (LWP 2503) 0xffffe422 in __kernel_vsyscall ()
(gdb) thread 3
[Switching to thread 3 (Thread 0xb764db70 (LWP 2511))]#0 iterate (
  ind=0x804b0cc) at mtserver.c:34
34          num_iter = tID[index].lIterations;
(gdb) step
35          delta_x = 1.0 / num_iter;
(gdb) info threads
  5 Thread 0xb664bb70 (LWP 2537) 0xb7f22848 in clone () from /lib/libc.so.6
  4 Thread 0xb6e4cb70 (LWP 2512) iterate (ind=0x804b0cc) at mtserver.c:29
* 3 Thread 0xb764db70 (LWP 2511) iterate (ind=0x804b0cc) at mtserver.c:35
  2 Thread 0xb7e4eb70 (LWP 2510) 0xb7f22848 in clone () from /lib/libc.so.6
  1 Thread 0xb7e4f6c0 (LWP 2503) 0xffffe422 in __kernel_vsyscall ()
(gdb) █

```

**Figure 22: GDB after performing two single-steps**

- Single-stepping thread 3 until line 59, where the approximation starts, does not show any unusual behaviour. Therefore, switch focus to thread 4 in order to single-step this thread up to line 59, as seen in Figure 23.

```

35         delta_x = 1.0 / num_iter;
(gdb) info threads
  5 Thread 0xb664bb70 (LWP 2537) 0xb7f22848 in clone () from /lib/libc.so.6
  4 Thread 0xb6e4cb70 (LWP 2512) iterate (ind=0x804b0cc) at mtserver.c:29
* 3 Thread 0xb764db70 (LWP 2511) iterate (ind=0x804b0cc) at mtserver.c:35
  2 Thread 0xb7e4eb70 (LWP 2510) 0xb7f22848 in clone () from /lib/libc.so.6
  1 Thread 0xb7e4f6c0 (LWP 2503) 0xffffe422 in __kernel_vsyscall ()
(gdb) step
38         pthread_mutex_lock( &mutex0 );
(gdb) step
39         loop_iteration = tID[index].loop_count++;
(gdb) step
40         pthread_mutex_unlock( &mutex0 );
(gdb) step
42         n1 = num_iter / SUBTHREADS;
(gdb) step
43         n2 = num_iter % SUBTHREADS;
(gdb) step
45         if (loop_iteration == 0) {
(gdb) step
46             start = 0;
(gdb) step
47             stop = n1 + n2;
(gdb) step
53         if (BUG_3 && (loop_iteration == 1)) {
(gdb) step
59         for (l = start; l < stop; ++l) {
(gdb) info threads
  5 Thread 0xb664bb70 (LWP 2537) 0xffffe422 in __kernel_vsyscall ()
  4 Thread 0xb6e4cb70 (LWP 2512) iterate (ind=0x804b0cc) at mtserver.c:29
* 3 Thread 0xb764db70 (LWP 2511) iterate (ind=0x804b0cc) at mtserver.c:59
  2 Thread 0xb7e4eb70 (LWP 2510) 0xb7fb4970 in __nptl_create_event ()
    from /lib/libpthread.so.0
  1 Thread 0xb7e4f6c0 (LWP 2503) 0xffffe422 in __kernel_vsyscall ()
(gdb) █

```

**Figure 23: Single-stepping thread 3 to line 59 in GDB**

8. As shown in Figure 24, single-step thread 4, while taking care that it remains in focus shows that, in contrast to thread 3, it will hit lines 54 and 55, which will cause this thread to terminate.

The bug has been uncovered:

```

if (BUG_3 && (loop_iteration == 1)) {
    tID[index].still_working[loop_iteration] = 0;
    return NULL;
}

```



```

40      pthread_mutex_unlock( &mutex0 );
(gdb) step
[Switching to Thread 0xb5e4ab70 (LWP 2609)]
Breakpoint 2, iterate (ind=0x804b0cc) at mtserver.c:29
29      index = * ((int*) ind);
(gdb) step
32      pthread_detach(pthread_self());
(gdb) info threads
* 6 Thread 0xb5e4ab70 (LWP 2609)  iterate (ind=0x804b0cc) at mtserver.c:32
  5 Thread 0xb664bb70 (LWP 2537)  iterate (ind=0x804b0cc) at mtserver.c:29
  4 Thread 0xb6e4cb70 (LWP 2512)  iterate (ind=0x804b0cc) at mtserver.c:42
  3 Thread 0xb764db70 (LWP 2511)  (Exiting)  0xb7fb4980 in __nptl_death_event
  () from /lib/libpthread.so.0
  2 Thread 0xb7e4eb70 (LWP 2510)  0xffffe422 in __kernel_vsyscall ()
  1 Thread 0xb7e4f6c0 (LWP 2503)  0xffffe422 in __kernel_vsyscall ()
(gdb) thread 4
[Switching to thread 4 (Thread 0xb6e4cb70 (LWP 2512))]#0  iterate (
ind=0x804b0cc) at mtserver.c:42
42      n1 = num_iter / SUBTHREADS;
(gdb) step
43      n2 = num_iter % SUBTHREADS;
(gdb) step
45      if (loop_iteration == 0) {
(gdb) step
49          start = loop_iteration * n1 + n2;
(gdb) step
50          stop = (loop_iteration + 1) * n1 + n2;
(gdb) step
53      if (BUG_3 && (loop_iteration == 1)) {
(gdb) step
54          tID[index].still_working[loop_iteration] = 0;
(gdb) print loop_iteration
$1 = 1
(gdb) █

```

**Figure 24: GDB thread 4 terminates**

Here GDB does demonstrate the error, but you have to follow the lines as they come up quite carefully. In this case the difference between flow control in thread 3 and thread 4 leads thread 4's results to be disregarded.

## Thread Control Capabilities of TotalView

The following steps will show the advantages of TotalView's thread control capabilities:

- 1.) Start the server with the command **totalview ./server** and place a single breakpoint at line 29. Use the 'Action Point Properties' dialog to tell the breakpoint to act only on individual threads, as seen in Figure 25.

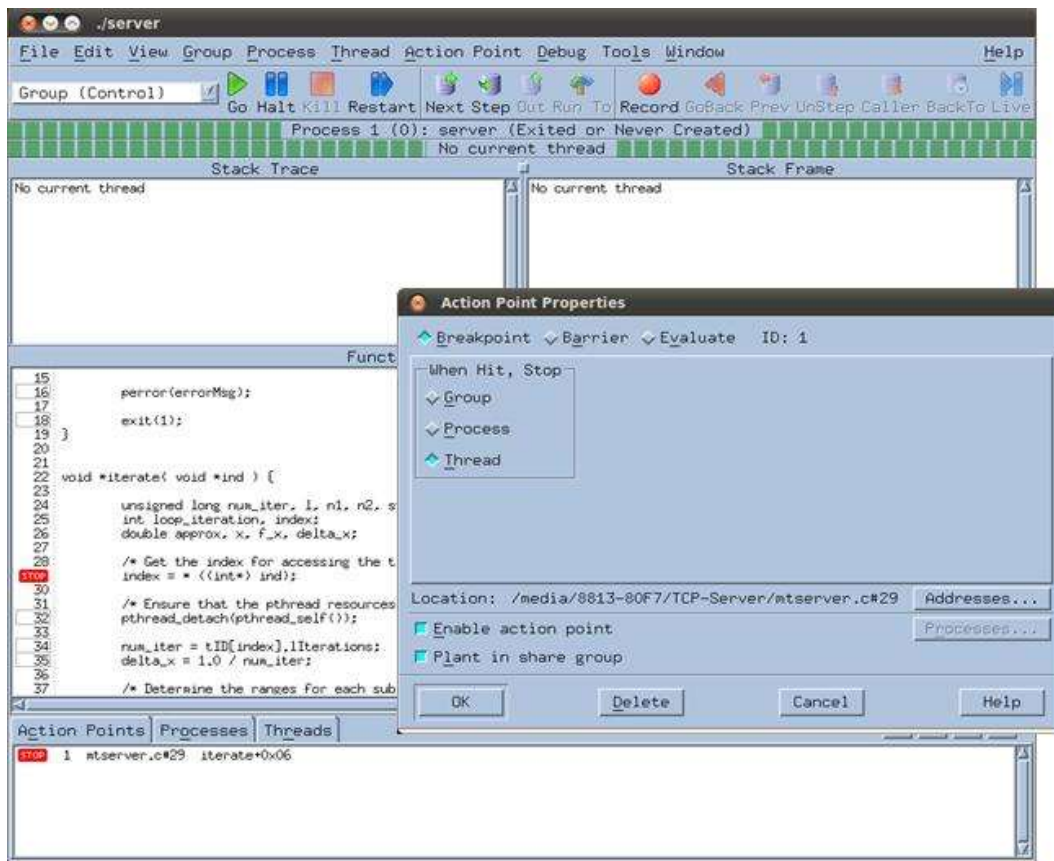


Figure 25: Setting Action Point properties in TotalView

- Run the server by pushing the 'Go' button and run the client one time as already described to force all four sub-threads to stop at the breakpoint, as seen in Figure 26.

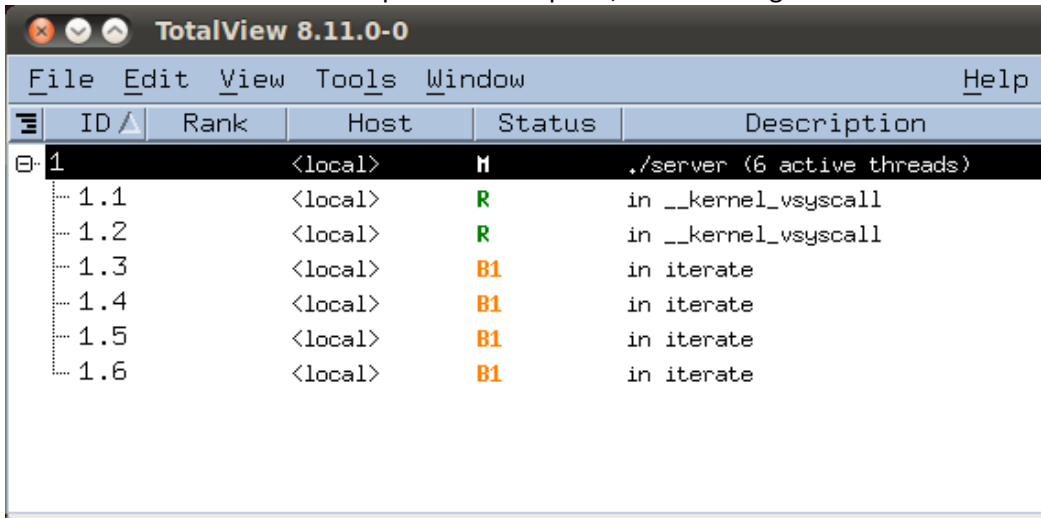


Figure 26: TotalView Root Window showing all threads have been stopped at the breakpoint

- Move your cursor down to highlight one of the sub-threads.
- As seen in Figure 27, select the scope modifier 'Group (Lockstep)' in the toolbar's P/T Set Control pull down menu to allow single-stepping of all four sub-threads in sync. It makes sense to move all threads together until line 38 because the mutex will only allow one thread at a time to pass this barrier. Single-stepping in TotalView is done by pushing either the 'Next' (stepping over functions) or 'Step' (stepping into functions) buttons.

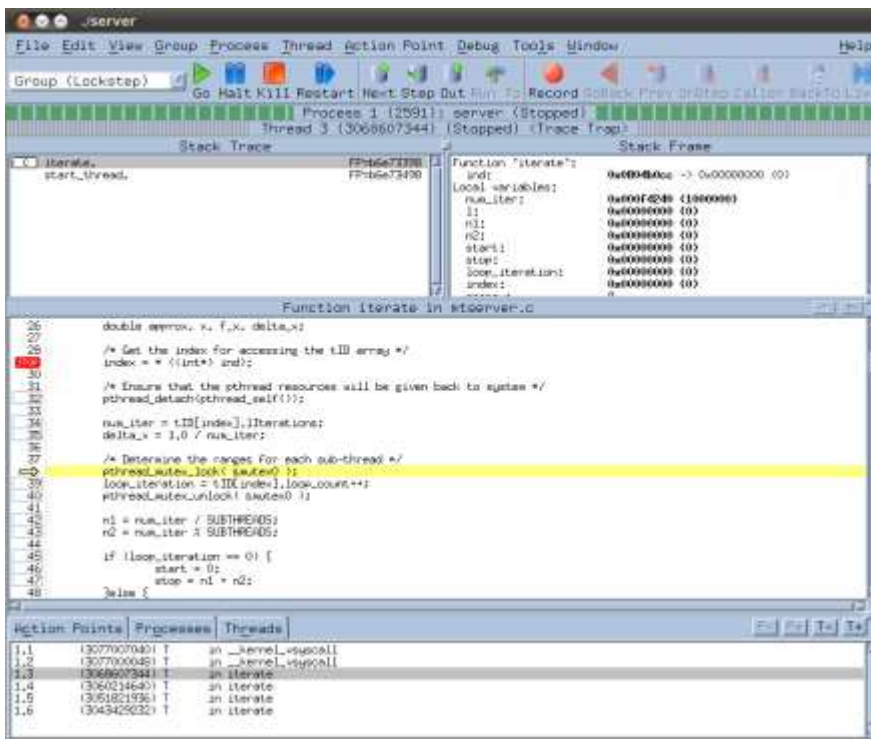
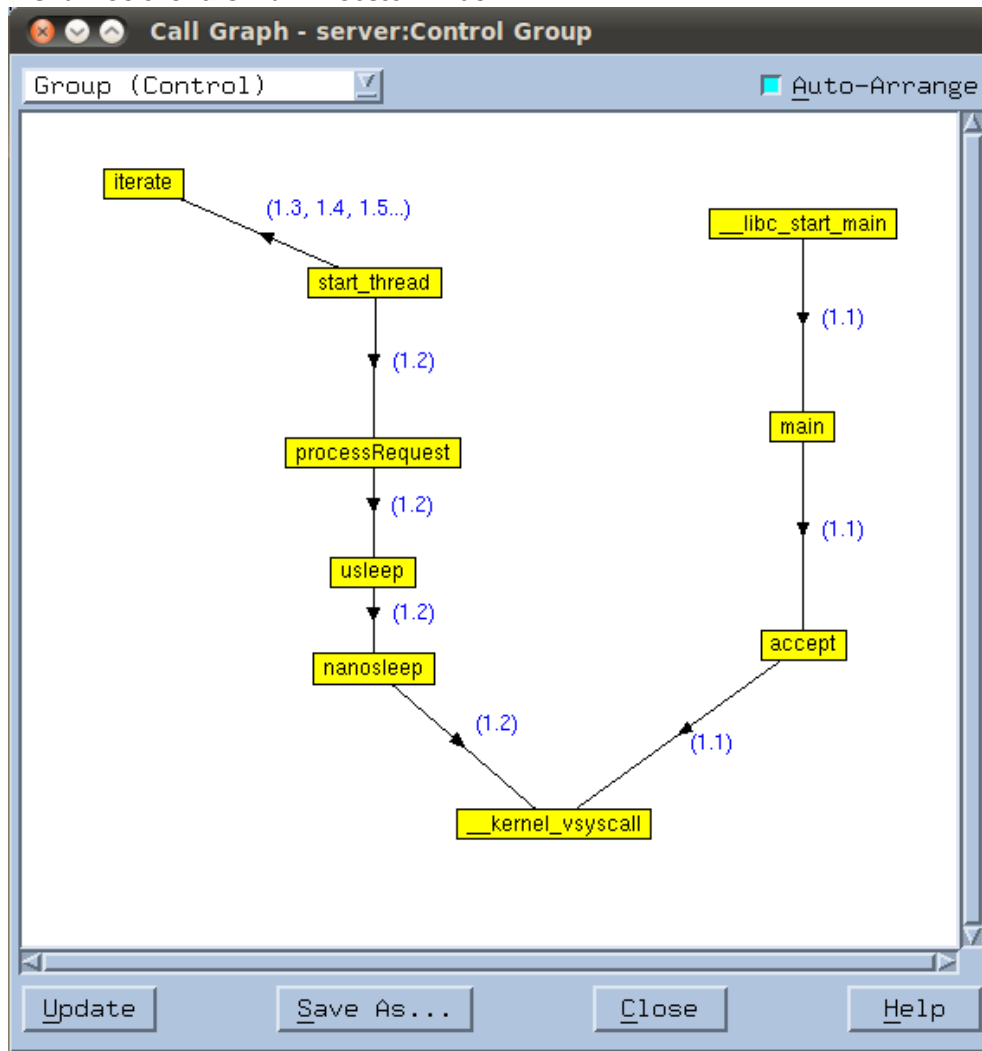


Figure 27: Controlling the lockstep group in TotalView

Another convenient way to control a defined number of threads using TotalView is to create a Call Graph group. As seen in Figure 28, the Call Graph is a graphical representation of the recent stack trace of specified processes and threads. It can be opened by choosing 'Call Graph' on the menu 'Tools' of the main Process Window.



**Figure 28: The TotalView Call Graph**

The Call Graph is useful for understanding the functionality of the application by showing which threads have a displayed routine on their call stack. This graph shows three different thread behaviours. Thread 1.1 includes main and is currently in a routine called `accept()`, thread 1.2 is sleeping in a routine called `processRequest()`, and threads 1.3 – 1.5 are in `iterate()`.

Additionally, double clicking on a routine in the graph automatically creates a thread group called "call\_graph," containing all of the threads which call that routine, as seen in Figure 29.

The new group will immediately appear in the Process Window's pull down menu and can be selected for actions affecting the entire thread group, for example, moving this entire group of threads in single step mode to line 38.






Figure 29: TotalView Process Window with the new call\_graph group selected

5.) Single-step each thread starting from line 38 to easily uncover the bug, as displayed in Figure 30.

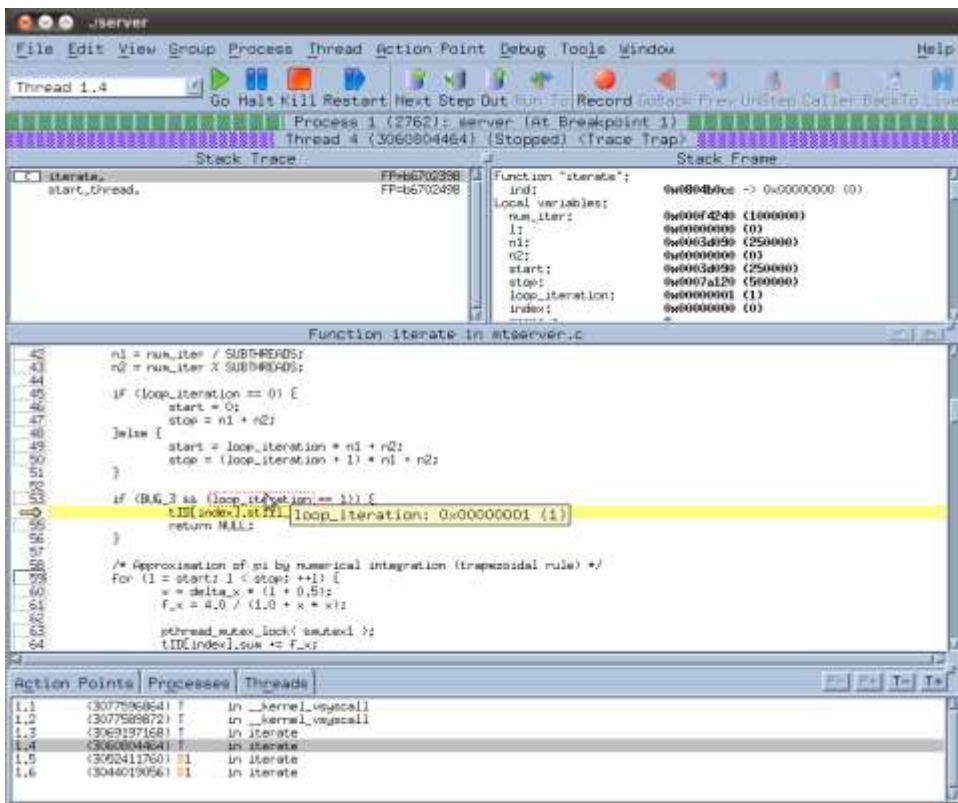


Figure 30: The bug has been found in TotalView

Stepping each thread through the section shows that one thread takes a divergent path. Hovering over the variables shows why the program is taking that path.

Compared to GDB, the TotalView debugger has far superior thread capabilities. Not only does it offer a convenient interface designed to allow complete control over individual threads, it also supports asynchronous thread control, thread groups, stepping commands that operate on thread groups, and thread width breakpoints.



## Displaying Local Variables of Different Threads

It is usually necessary to monitor the local variables belonging to each thread in order to check whether a program is working as expected. In order to demonstrate such a case, rebuild the server with all of the bugs turned off and the number of sub-threads set to 10.

In this debugging session, monitor changes to the local variable `f_x` for each of the 10 sub-threads while in single-step mode. To accomplish this, place a breakpoint at line 60.

As with the sessions above, GDB will be demonstrated first.

### *Displaying Local Variables of Different Threads Using GDB*

1. Start the server by typing `gdb ./server <return>`, set the breakpoint by typing `break 60 <return>`, run the server by typing `run <return>` and send a client request to the server by typing `./client localhost 25000 100000 <return>` in a separate console. The first of the 10 sub-threads will hit the breakpoint and the process will stop.
2. Because GDB always focuses on the thread that last hit the breakpoint, you can serially display the value of `f_x` for each thread by alternating the commands `print f_x <return>` and `c <return>`.



As shown in Figure 31, the nature of the GDB output makes it very difficult to compare the values stored by the various threads, especially if hundreds or thousands of threads are running simultaneously.

```
Starting program: /home/luetke/tests/listener/server
Missing separate debuginfo for /lib/ld-linux.so.2
Try: zypper install -C "debuginfo(build-id)=d7706cbaa0ca09319cb645eac789cb8399078797"
Missing separate debuginfo for /lib/libpthread.so.0
Try: zypper install -C "debuginfo(build-id)=964690b0ca2ed321e995340684e09981f5f986ad"
[Thread debugging using libthread_db enabled]
Missing separate debuginfo for /lib/libc.so.6
Try: zypper install -C "debuginfo(build-id)=7eb4e169e926464393ef2e98d99c37f56d5f5858"
No portnumber provided, listening on port no.: 25000
[New Thread 0xb7e4eb70 (LWP 5068)]
[New Thread 0xb764db70 (LWP 5069)]
[Switching to Thread 0xb764db70 (LWP 5069)]

Breakpoint 1, iterate (ind=0x804b0cc) at mtserver.c:60
60          x = delta_x * (1 + 0.5);
(gdb) print f_x
$1 = 0
(gdb) c
Continuing.

Breakpoint 1, iterate (ind=0x804b0cc) at mtserver.c:60
60          x = delta_x * (1 + 0.5);
(gdb) print f_x
$2 = 3.9999999999999899999
(gdb) c
Continuing.

Breakpoint 1, iterate (ind=0x804b0cc) at mtserver.c:60
60          x = delta_x * (1 + 0.5);
(gdb) print f_x
$3 = 3.99999999999910001
(gdb) █
```

Figure 31: GDB displaying the value of a local variable for every thread

Of course, it might be nice for GDB to tell you which thread you are seeing each data element from; you can inquire into it with `info threads`.

### ***Displaying Local Variables of Different Threads Using TotalView***

The TotalView session is set-up similar to the GDB session above.

- 1.) Start the server within TotalView by typing `totalview ./server <return>` and set a breakpoint in front of line 60 by clicking on the rectangle around the line number with the left mouse button.
- 2.) Using the right mouse button, click on the breakpoint to bring up its properties menu. Set the breakpoint to act on threads only.



- 3.) Start the server by pushing the 'Go' button and send it a client request by typing `./client localhost 25000 100000 <return>` in a separate console. The 10 sub-threads stopped at the breakpoint, as seen in Figure 32.

ID/	Rank	Host	Status	Description
1		<local>	M	./server (12 active threads)
- 1.1		<local>	R	in __kernel_vsyscall
- 1.2		<local>	R	in __kernel_vsyscall
- 1.3		<local>	B1	in iterate
- 1.4		<local>	B1	in iterate
- 1.5		<local>	B1	in iterate
- 1.6		<local>	B1	in iterate
- 1.7		<local>	B1	in iterate
- 1.8		<local>	B1	in iterate
- 1.9		<local>	B1	in iterate
- 1.10		<local>	B1	in iterate
- 1.11		<local>	B1	in iterate
- 1.12		<local>	B1	in iterate

**Figure 32: The TotalView Root Window showing all sub-threads stopped at breakpoint 1**

- 4.) In the Root Window, double-click on the line marked 1.3 to change the focus to thread 3. The current position of thread 3 is highlighted in the Process Window's source code pane.
- 5.) Drill down into the variable `f_x` by double-clicking it in the source code pane. This opens a variable window that displays the value of `f_x` for the currently selected thread, in this case thread 3.



- 6.) The 'View' menu of the variable window includes the option 'Show Across ->', as seen in Figure 33.' Select 'Show Across ->Threads' to create an array-like view of  $f_x$ , showing the values for every thread.

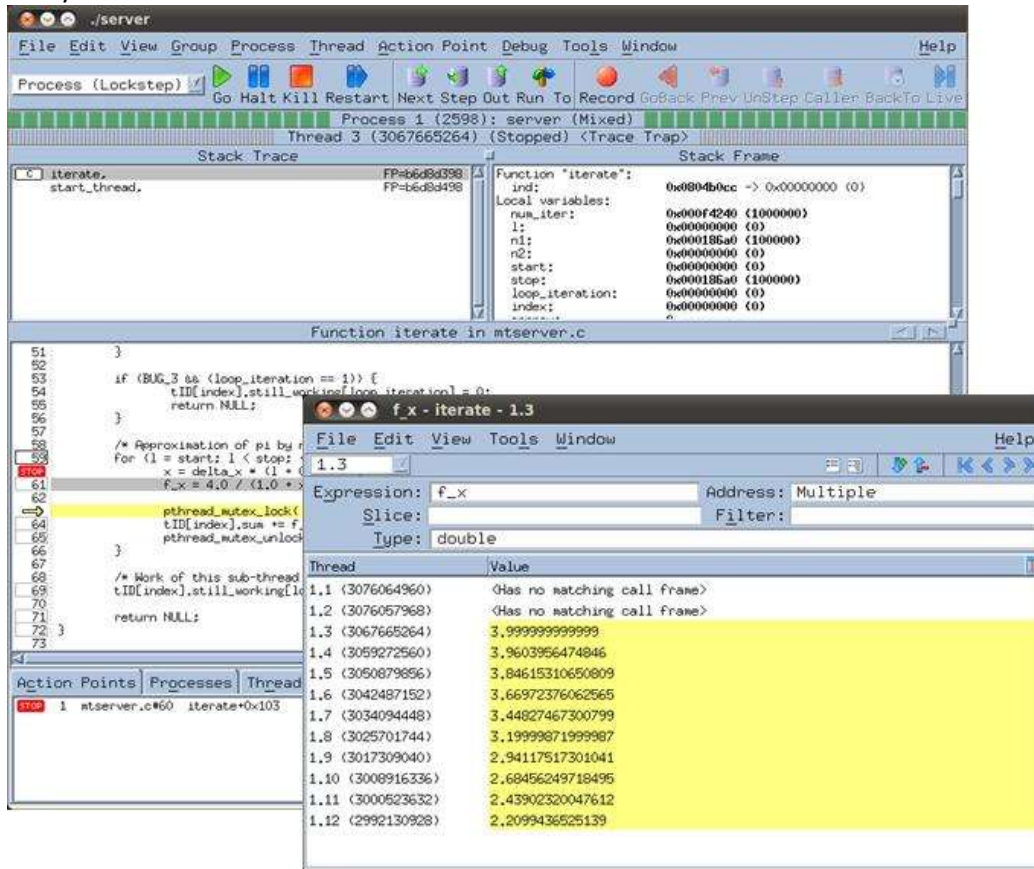


Figure 33: TotalView's Show Across capability

Variables displayed in a variable window are updated and highlighted as changes to them occur.

View across makes it much easier to do thread vs. thread data comparisons. With TotalView you can synchronize your threads so they are all at the same well-understood point in their execution, and do an apples-to-apples comparison of the data each thread contains at that same point. Alternately, you can run the threads in any execution sequence and see what happens, not just in that thread, but in all the other threads.

## Conclusion

This paper shows what it is like to debug several common problems in a multithreaded application using basic scenarios. It demonstrates how this debugging can be done with the open source GDB and with the TotalView debugger. TotalView provides a variety of useful features that make it easy to see and control what is going on in a multithreaded program. As you debug and explore the behavior of a multithreaded application, features such as asynchronous thread control, thread groups, thread width breakpoints, the Call Graph display, the ability to easily navigate data, and the ability to view variables across multiple threads will help to provide a clear indication of what the threads are doing.

Multithreaded programs can manifest intermittent defects. Tracking and resolving these defects without the right tools can be a tedious and frustrating process of running and re-running the program, only to have the problem go away. TotalView gives developers the ability to capture and deterministically replay the execution history of the program. This radically simplifies the troubleshooting process. Once an intermittent error is captured in the recording, it is easy to go back and forth as many times as needed to understand what has happened, making errors easier to solve.

## About TotalView

[TotalView](#) is a scalable and intuitive debugger for parallel applications written in C, C++, and Fortran. Designed to improve developer productivity, TotalView simplifies and shortens the process of developing, debugging, and optimizing complex applications. TotalView provides a powerful combination of capabilities for pinpointing and fixing hard-to-find bugs, such as race conditions, memory leaks, and memory overruns. Providing developers the ability to step freely, both forwards and backwards, through program execution, TotalView's unique reverse debugging capabilities drastically reduce the amount of time invested in troubleshooting code. To help developers maximize hardware capabilities, TotalView also provides debugging support for NVIDIA® CUDA™, OpenACC®, and the Intel® Xeon® Phi™ coprocessor.

## About Rogue Wave Software

Rogue Wave Software, Inc. is the largest independent provider of cross-platform software development tools and embedded components for the next generation of HPC applications. Rogue Wave marries High Performance Computing with High Productivity Computing to enable developers to harness the power of parallel applications and multicore computing. [Rogue Wave products](#) reduce the complexity of prototyping, developing, debugging, and optimizing multi-processor and data-intensive applications. Rogue Wave customers are industry leaders in the Global 2000, ISVs, OEMs, government laboratories and research institutions that leverage computationally-complex and data-intensive applications to enable innovation and outperform competitors. For more information, visit <http://www.roguewave.com>.

