RogueWave by Perforce

# Techniques for Debugging HPC Applications

NIKOLAY PISKUN , DIRECTOR OF CONTINUING ENGINEERING, TOTALVIEW PRODUCTS

AUGUST 7 2019, ATRESC 2019

# Agenda

- What is debugging and why TotalView?
- Overview of TotalView  and TotalView's new UI
- Advanced C++ and Data debugging
- MPI and OpenMP parallel debugging
- Reverse debugging
- Memory debugging
- GPU debugging
- Python/C++  debugging
- Using TotalView on ANL
- TotalView resources and documentation
- Questions/Comments

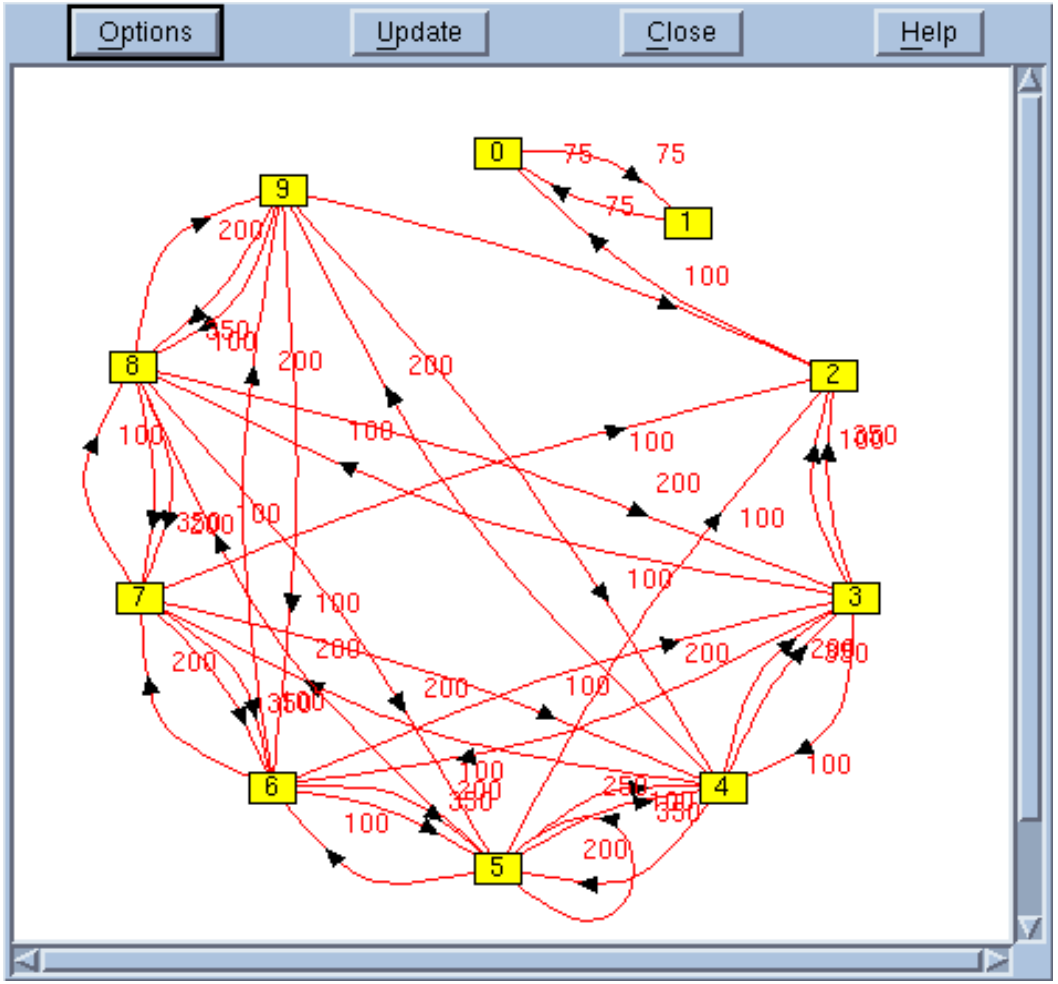


  

# What is Debugging  and
# Why do you need TotalView?

# What is Debugging?

- Debugging is the process of finding and resolving defects or problems within a computer program or a system.
  - Algorithm correctness
  - Data correctness
  - Scaling/Porting correctness

# TotalView debugger enables you to do:

- **Interactive debugging**
  - Live control of an executing program

- **Remote debugging**

  - Debug a program running on another computer

- **Post-mortem debugging (core files and reverse debugging)**
  - Debugging a program after it has crashed or exited


- **Memory debugging**

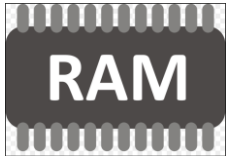


  - Find memory management problems (leaks, corruption …)
  - Comparing results between executions

- **Batch debugging (tvscript, CI environments)**
  - Unattended debugging

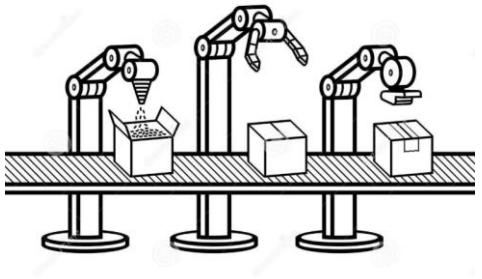

# TotalView for HPC and for All

- Leading debug environment for HPC users
  - Active development for 30+ years
  - Thread specific breakpoints
  - Control individual thread execution
  - View complex data types easily
  - From **MacBook** to **Top500** Supercomputers
- Track memory leaks in running applications
- Supports C/C++ and Fortran on Linux/Unix/Mac
- Support debugging mixed Python/C++
- Integrated Reverse debugging
- Batch non-interactive debugging.

- Allowing YOU to have
  - Predictable development schedules
  - Less time spent debugging

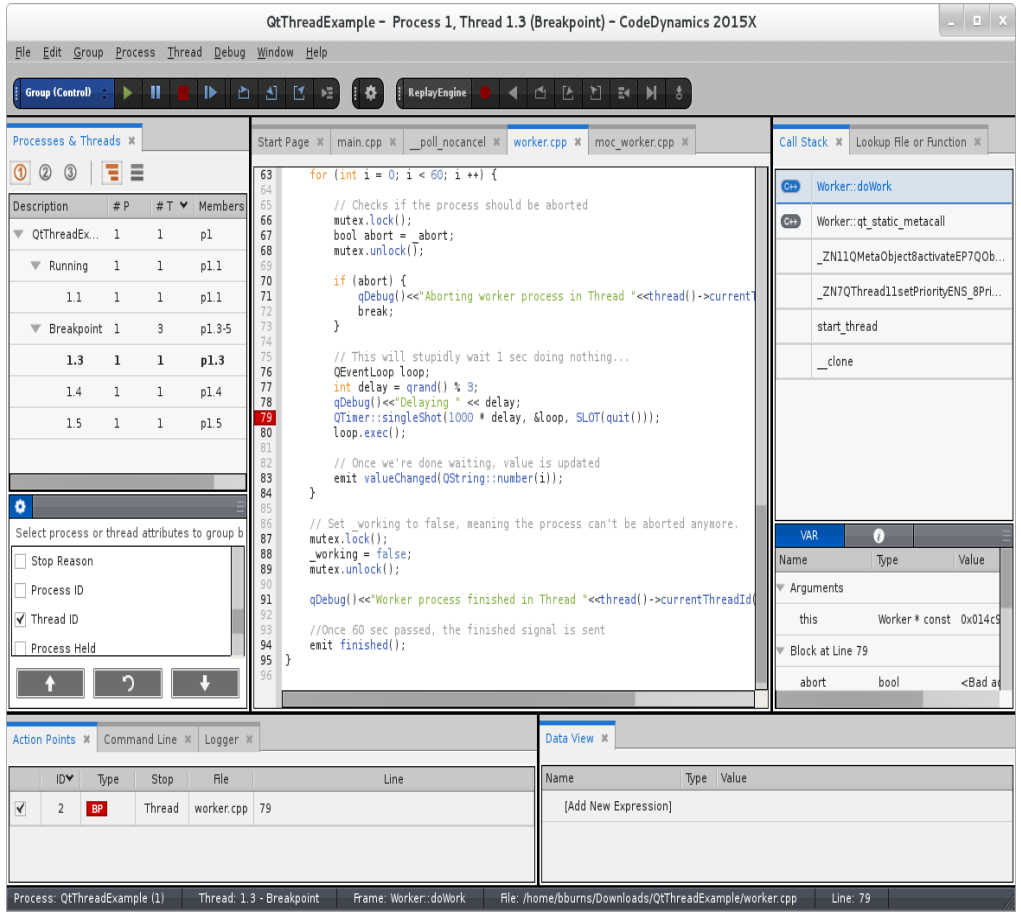

# TotalView’s GUI

# TotalView's Classic UI

- Original powerful design
- Better tested for high-scale MPI jobs
- Assembler support
- Better supported for Remote Display Client
- To use:
  - Set UI preference
  - Or command line argument

    `totalview -oldUI`

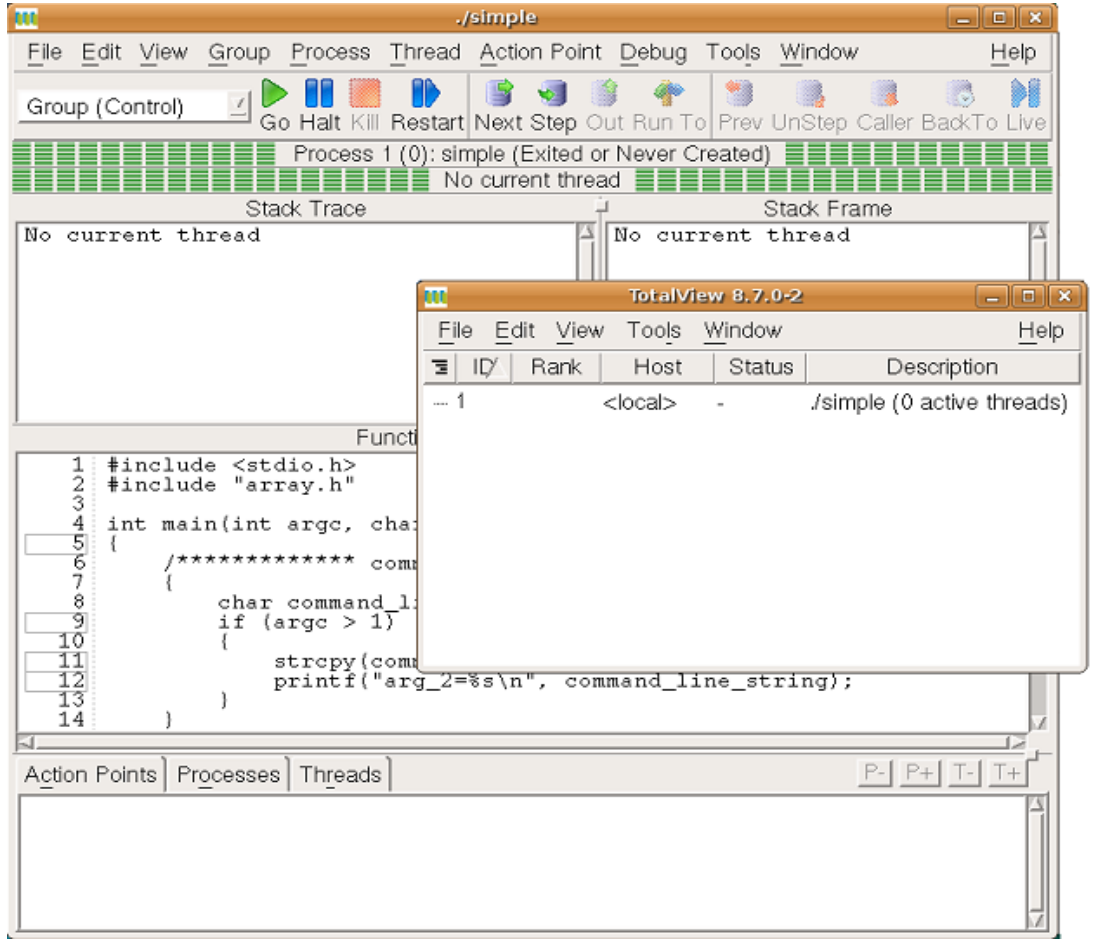

# TotalView's New UI (default)

- Provides a modern, dockable interface
- Easier to use, better workflows
- An architecture to grow
- To use:
  - Set UI preference
  - Or command line argument

```
totalview -newUI
```

- New UI gaps:
  - Missing array slicing and striding, view across, data visualization
  - No very high-scale support

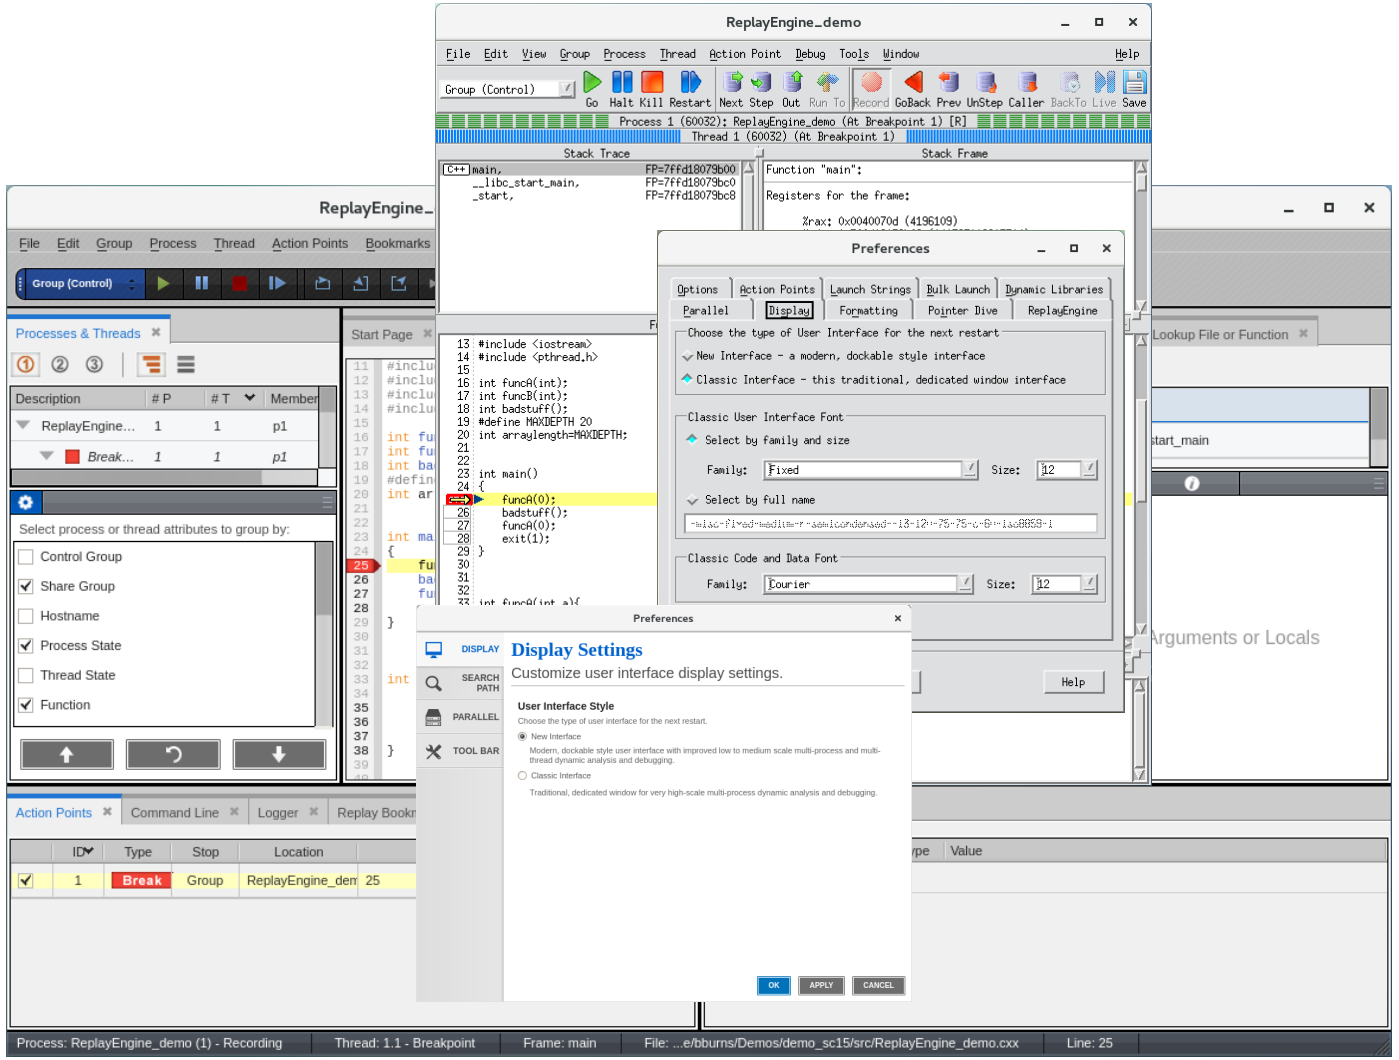

# Intro & Starting Up

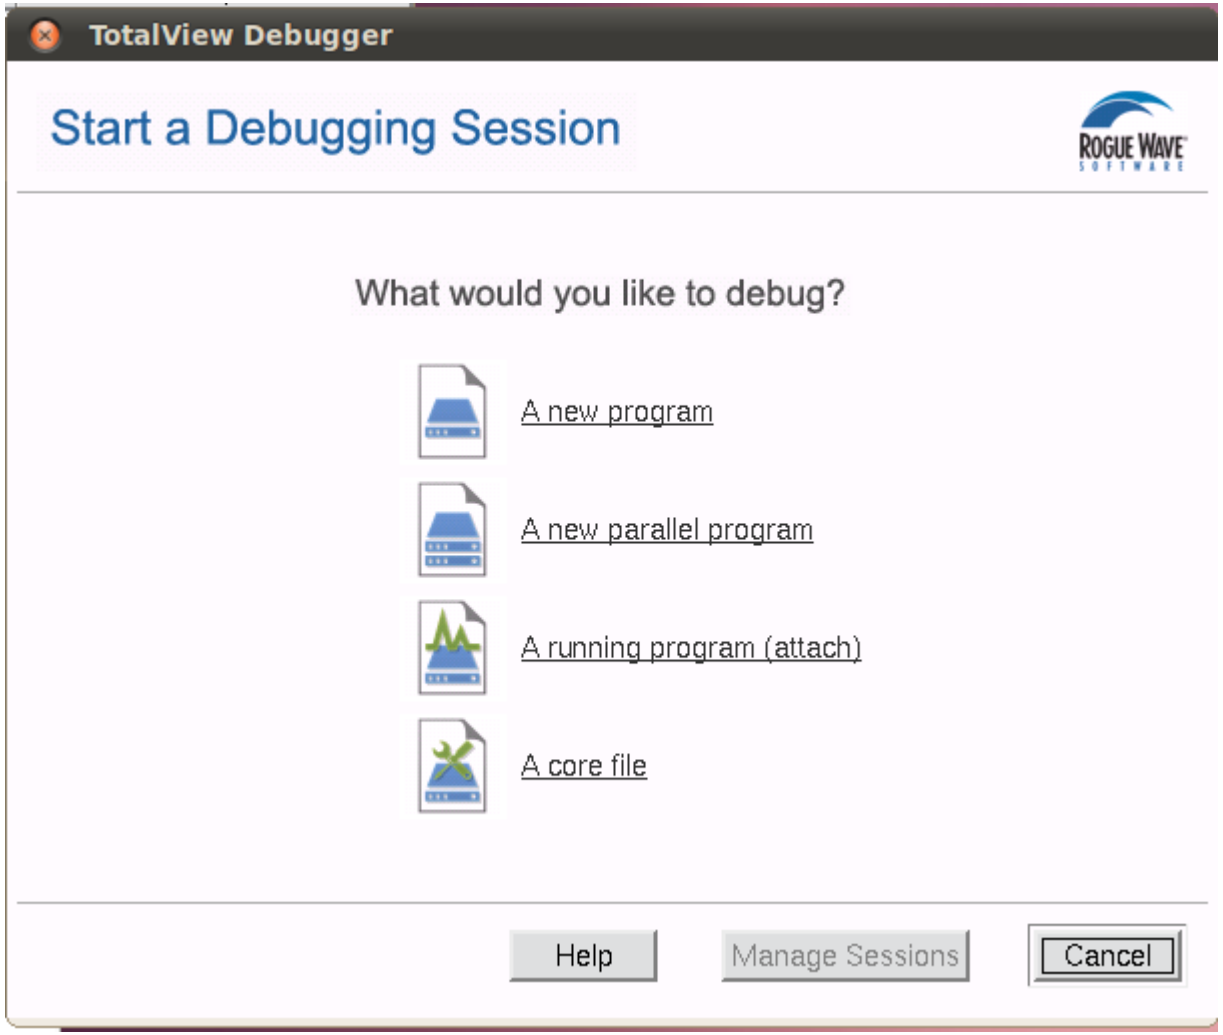
TotalView Debugger
Start a Debugging Session
ROGUE WAVE SOFTWARE
What would you like to debug?
A new program
A new parallel program
A running program (attach)
A core file
Help
Manage Sessions
Cancel

# Start New Process – Arguments

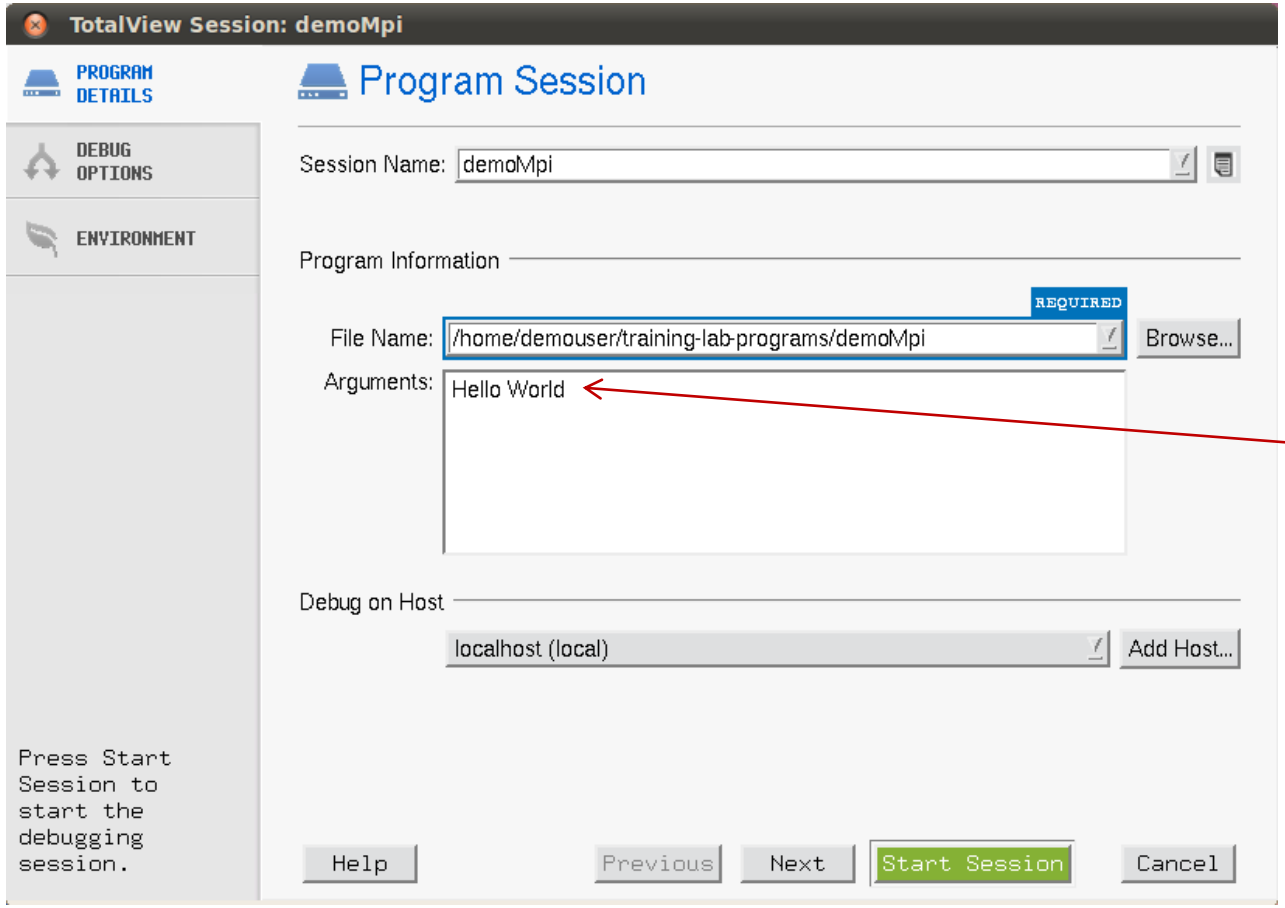

# Start New Process – Enable ReplayEngine

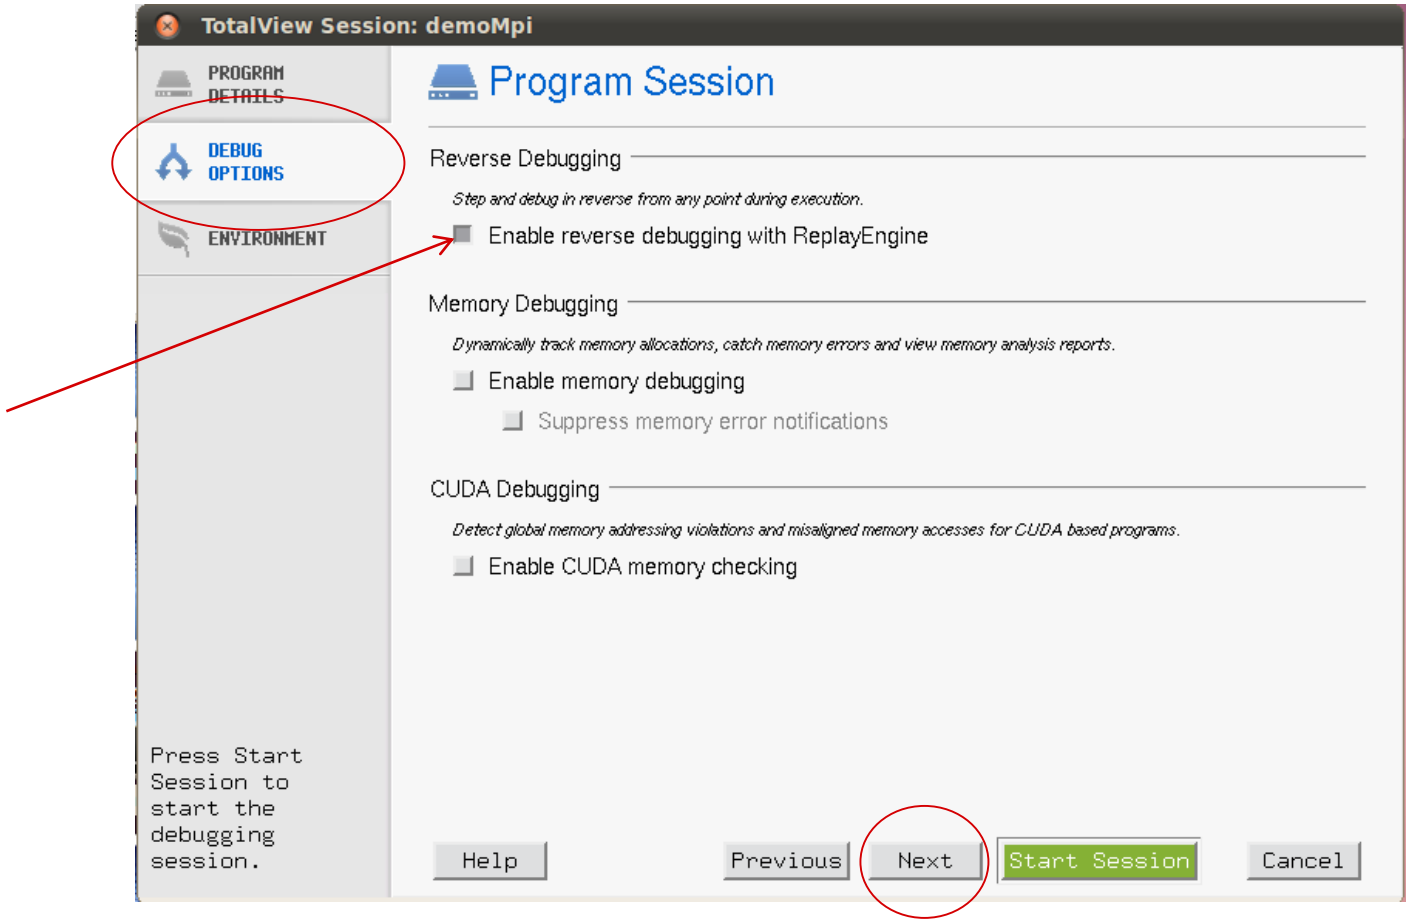

# Start New Process – Memory Debugging

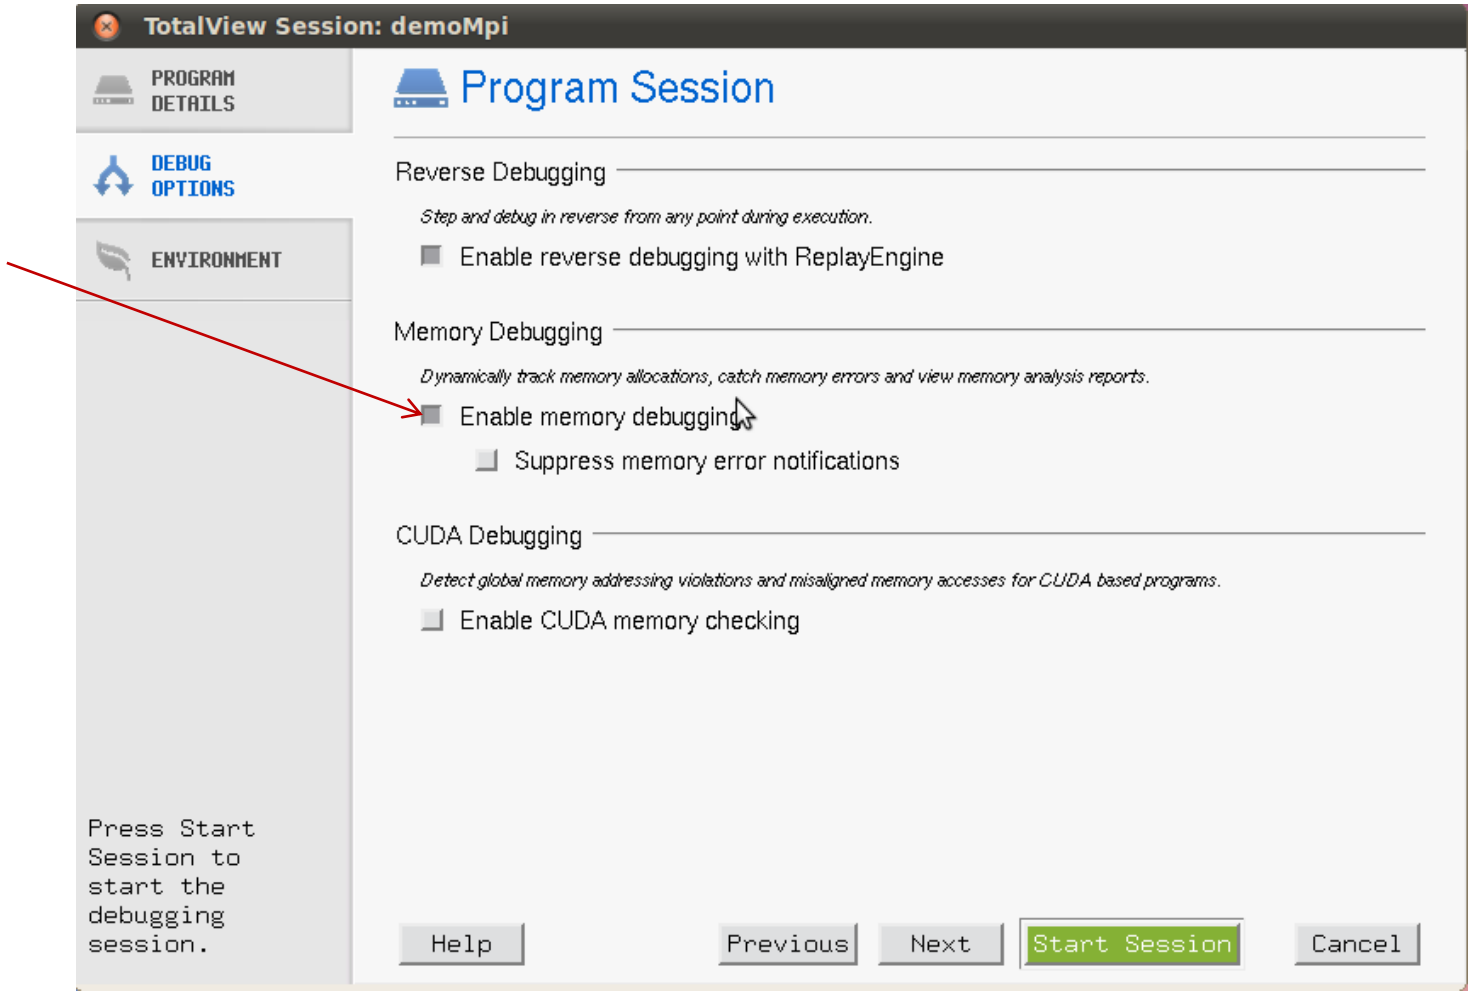

# CUDA memory checking

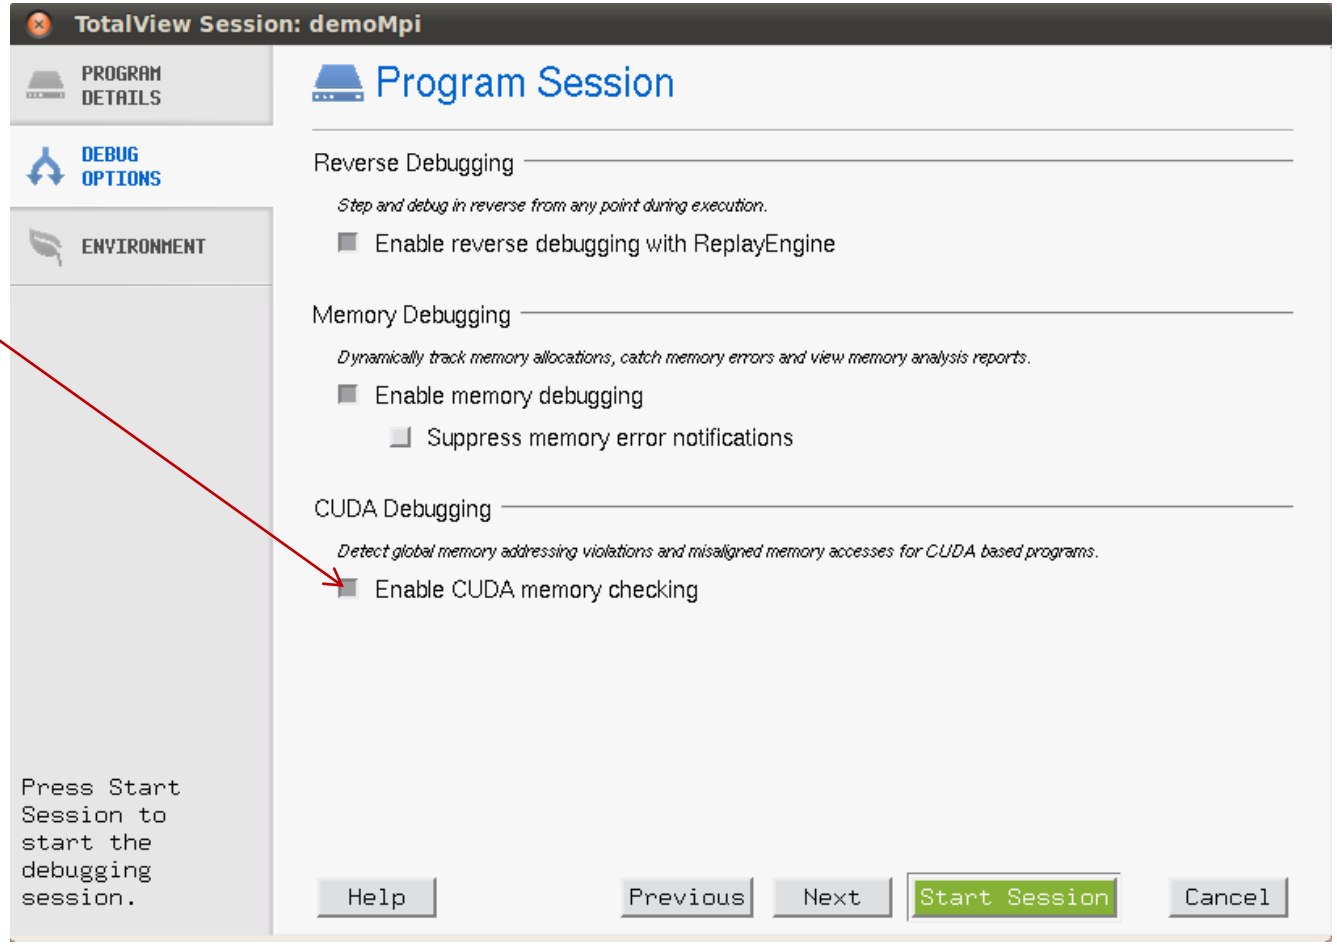

# Set environment variables

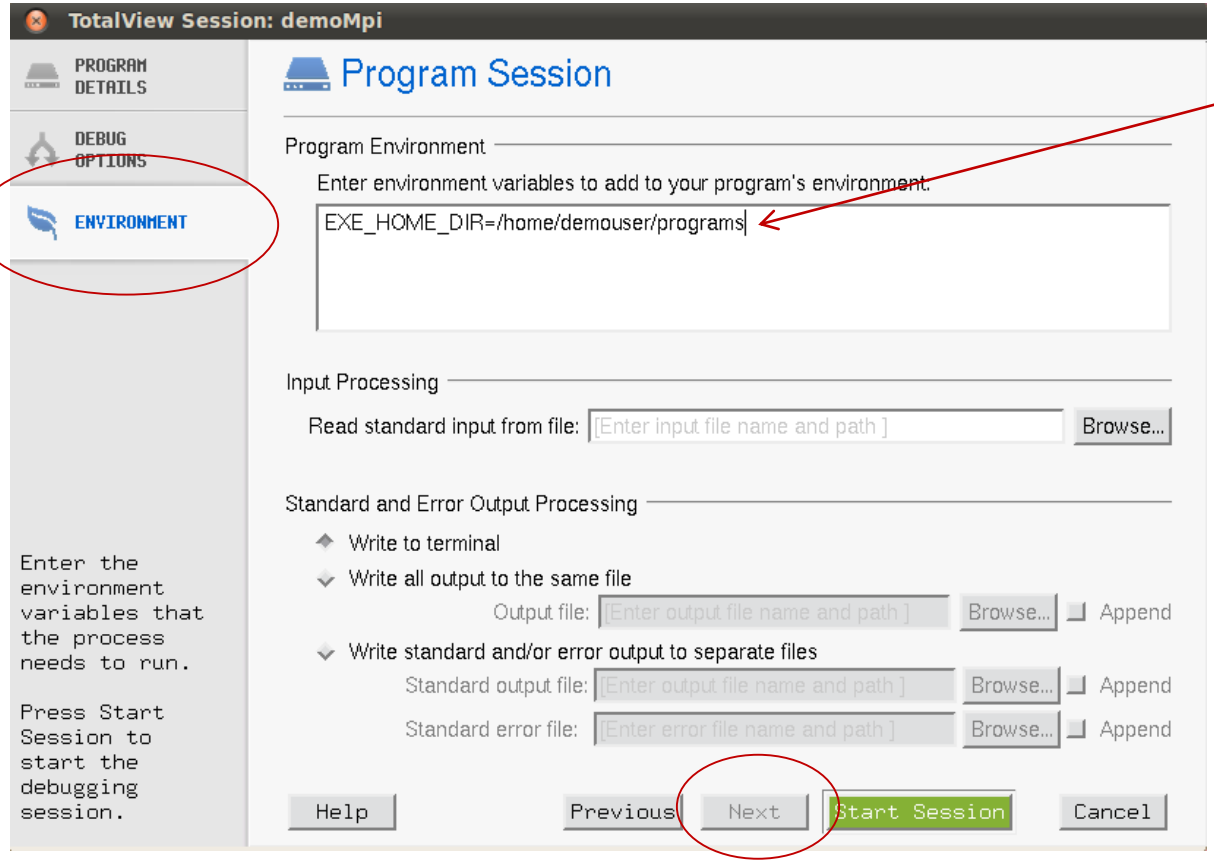

# Standard I/O redirection

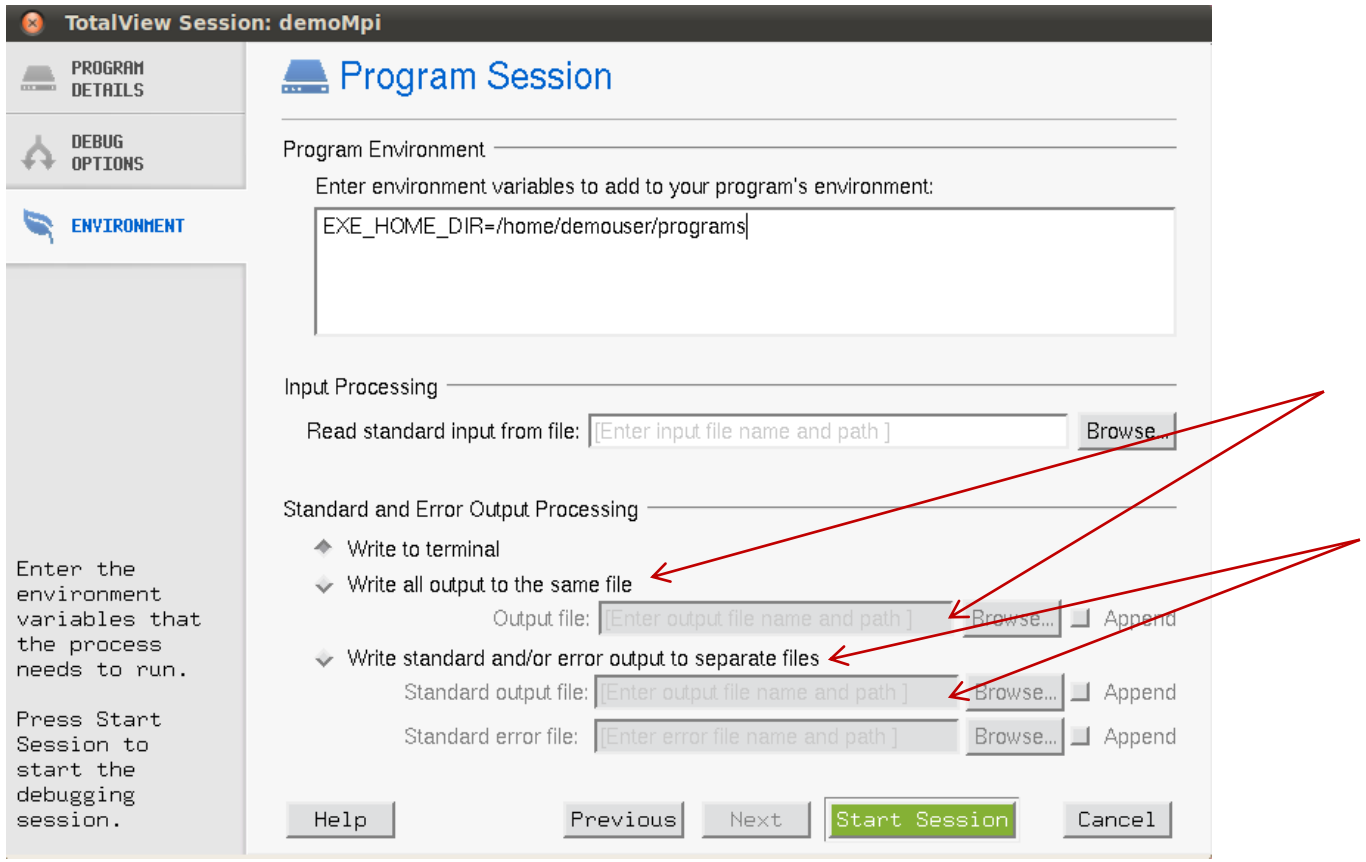

# Attach to Process

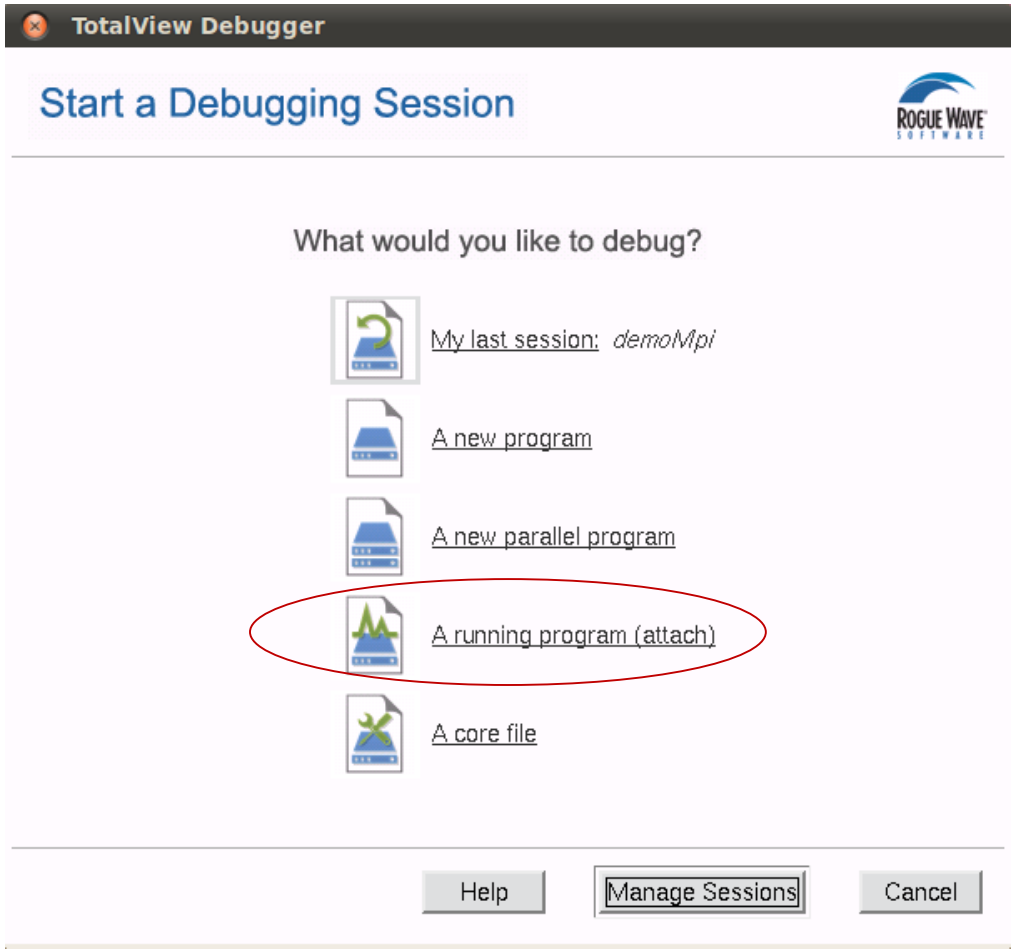

# Attach to Process

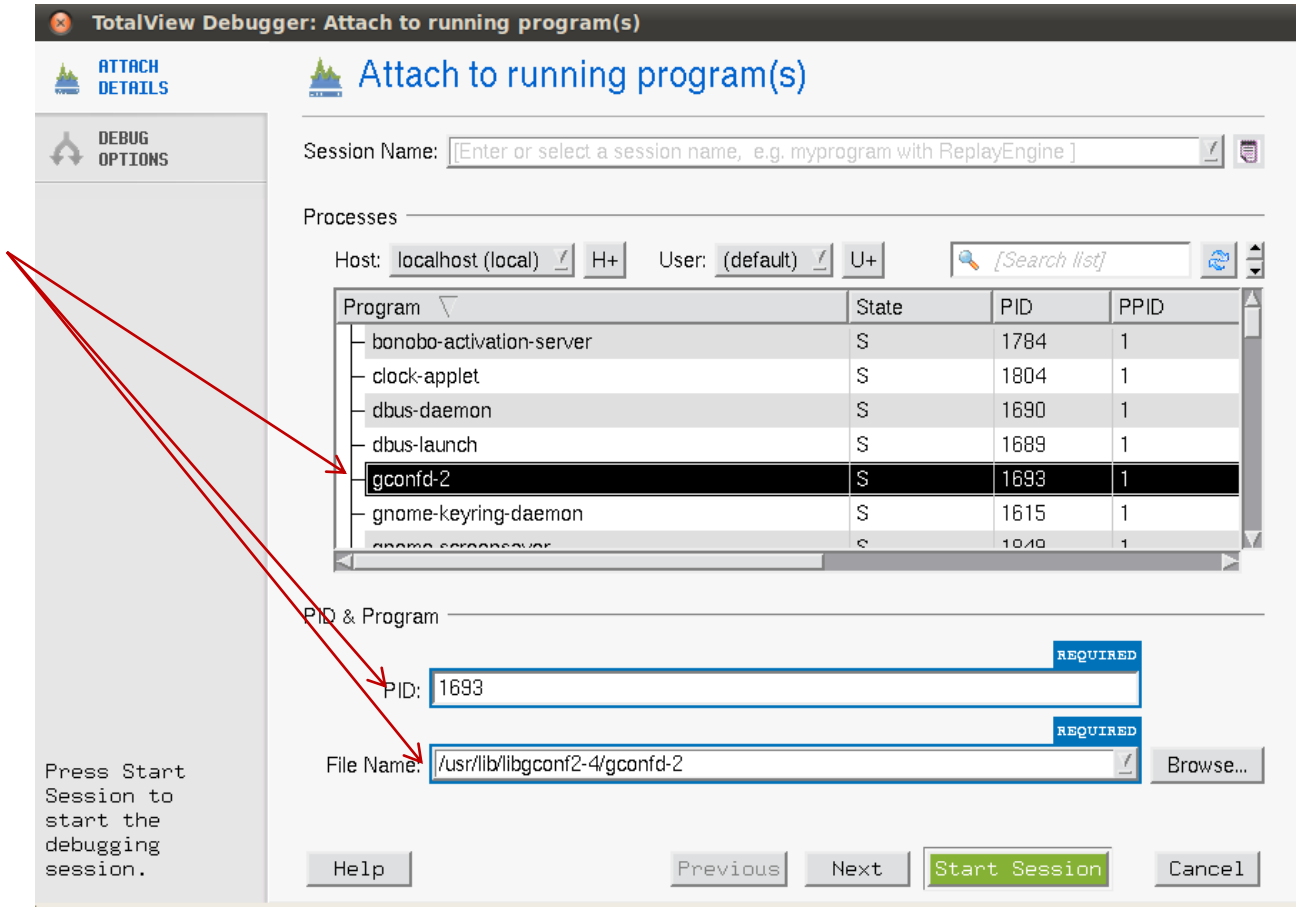

# Attach to Process – Enable Replay Engine

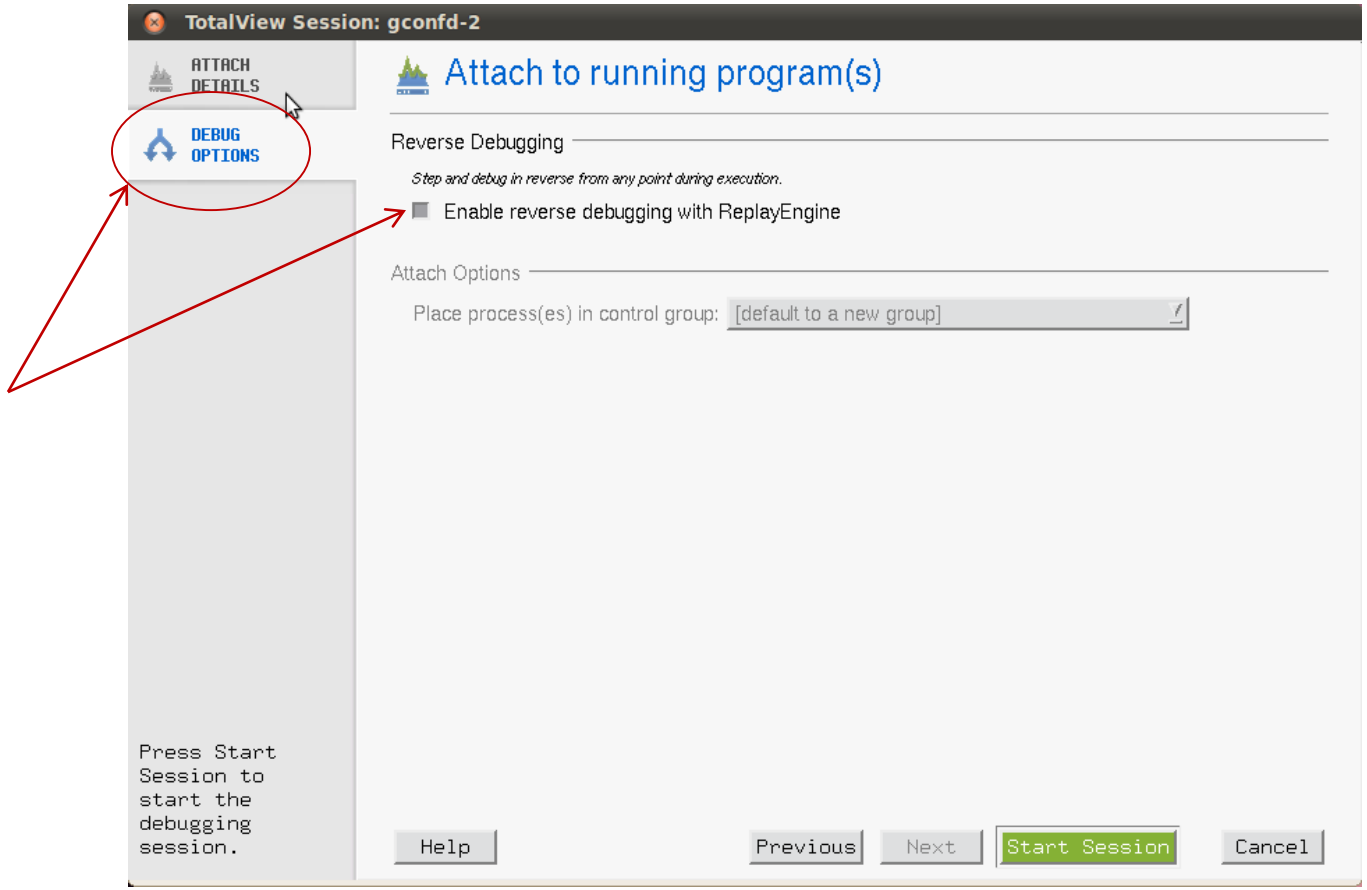

# Open a Core File

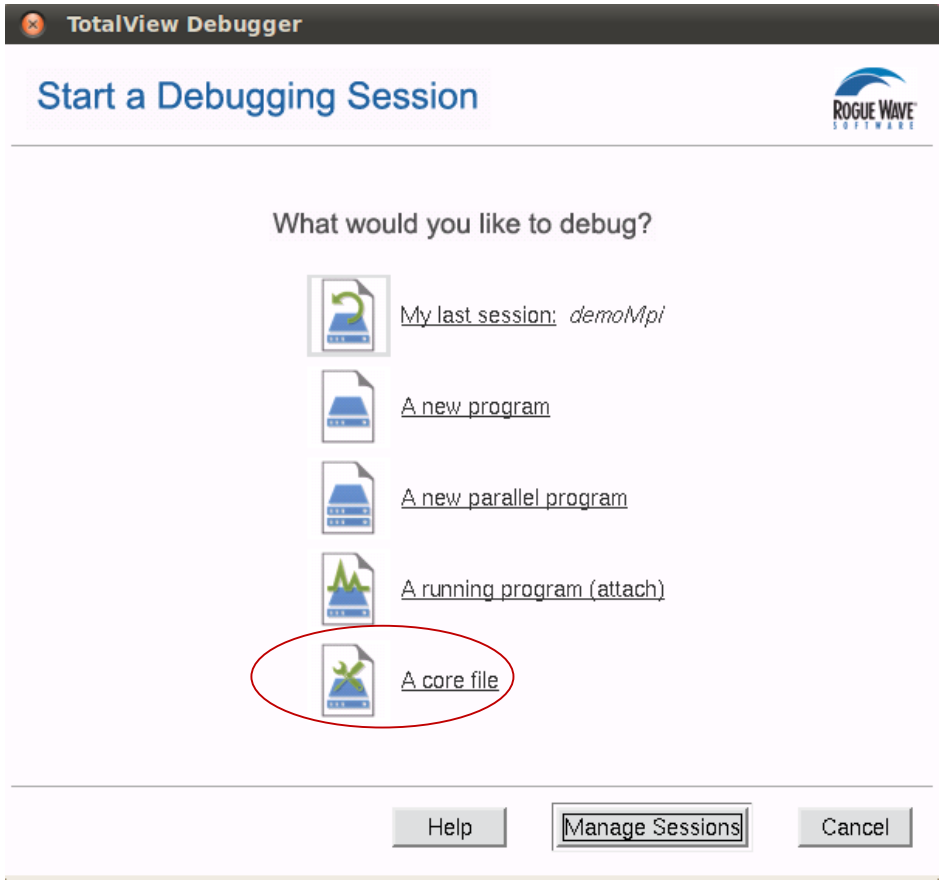

# Open a Core File

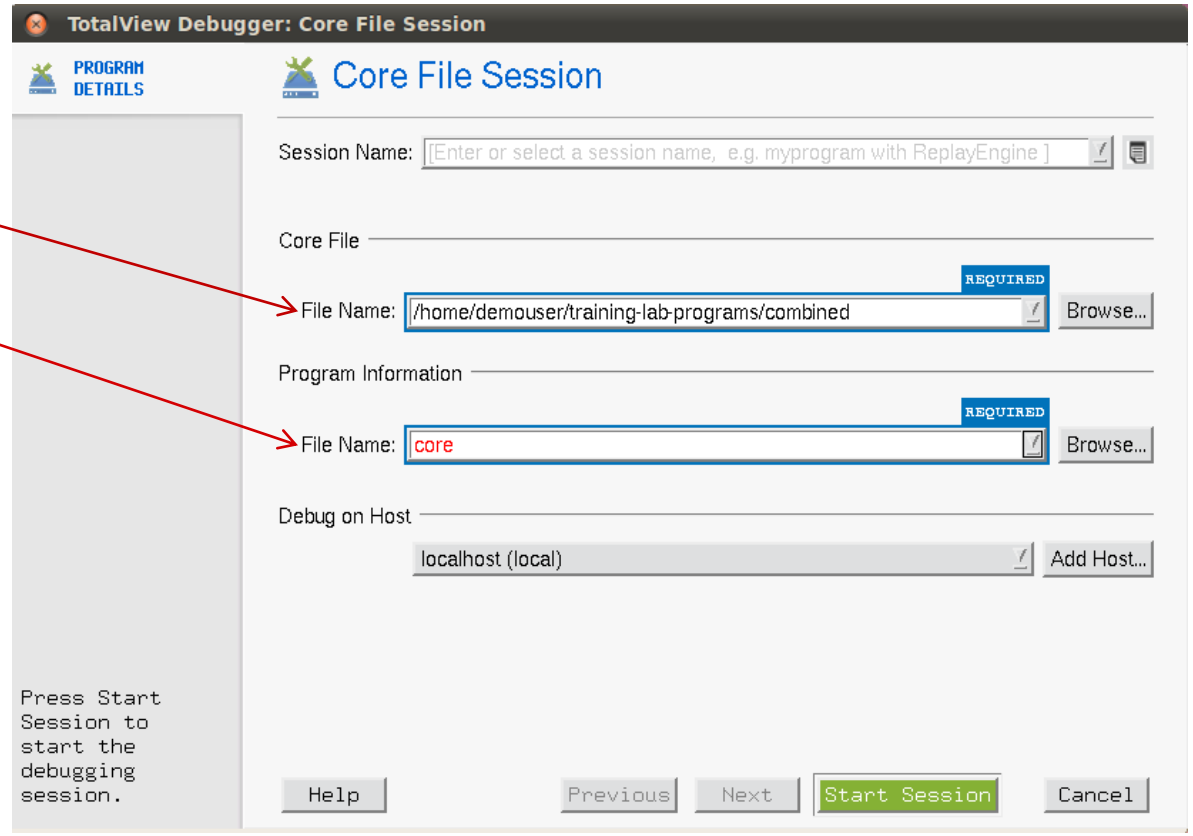

# Process Control & Navigation

# Interface Concepts

- State of all processes being debugged
- Process and Thread status
- Instant navigation access
- Sort and aggregate by status

## Root Window

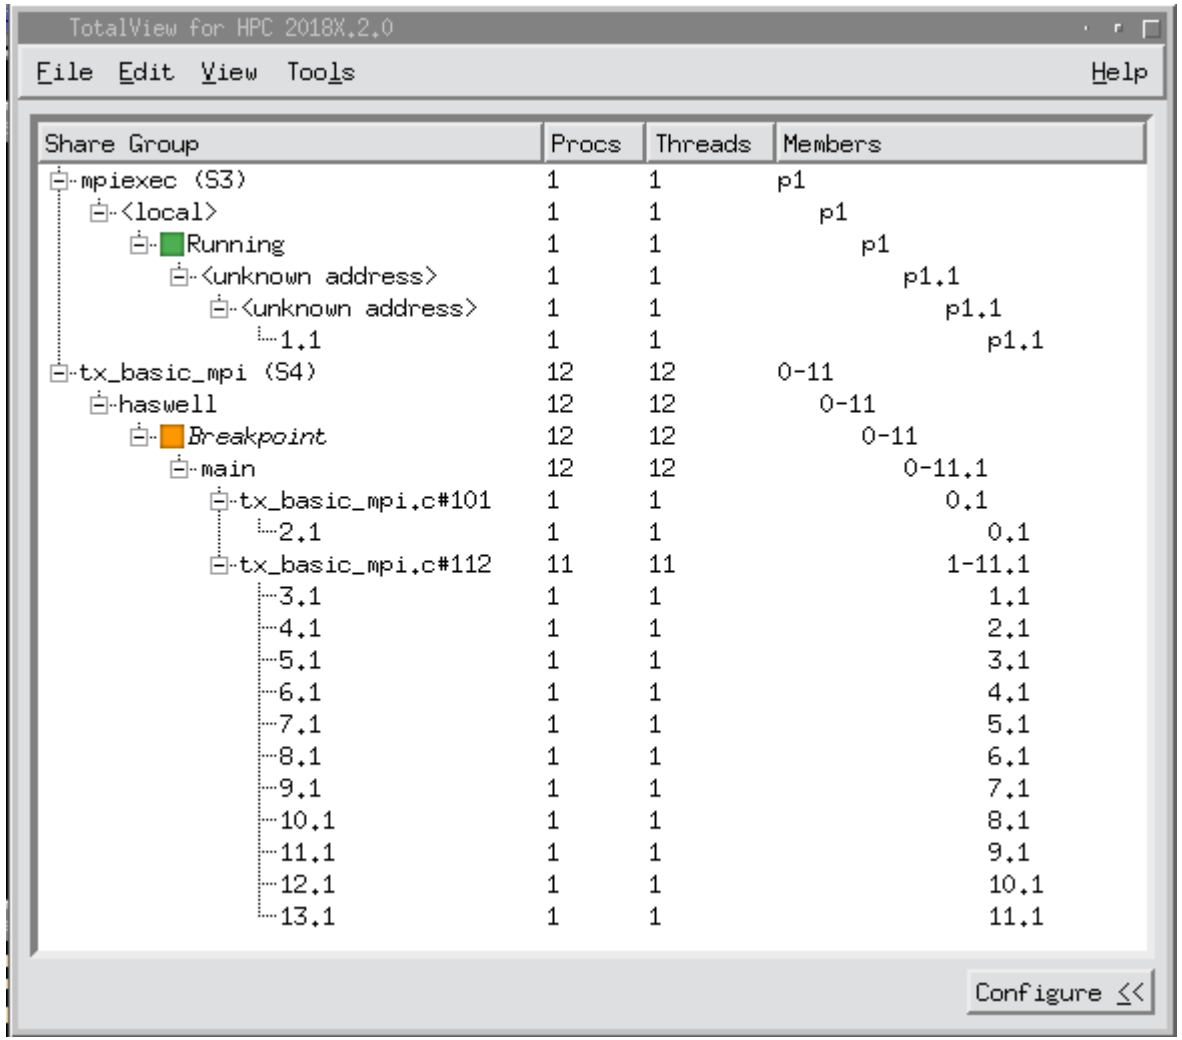

# Process Window Overview

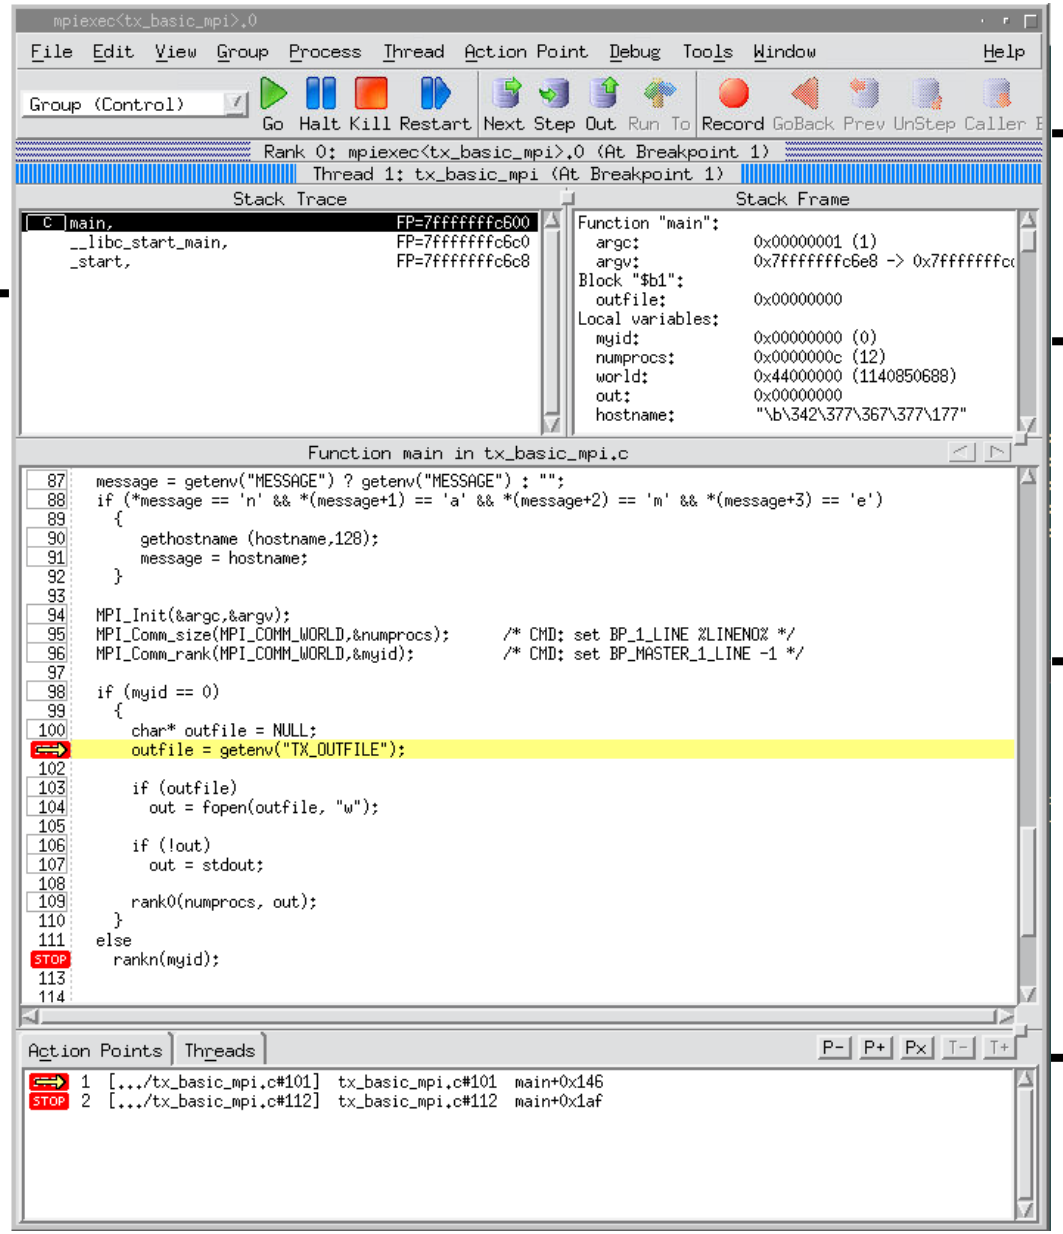


Stack Trace Pane

Toolbar

Stack Frame Pane

Source Pane

Tabbed Area

**Provides detailed state of one process, or a single thread within a process**

**A single point of control for the process and other related processes**

# Stepping Commands


Based on

PC location

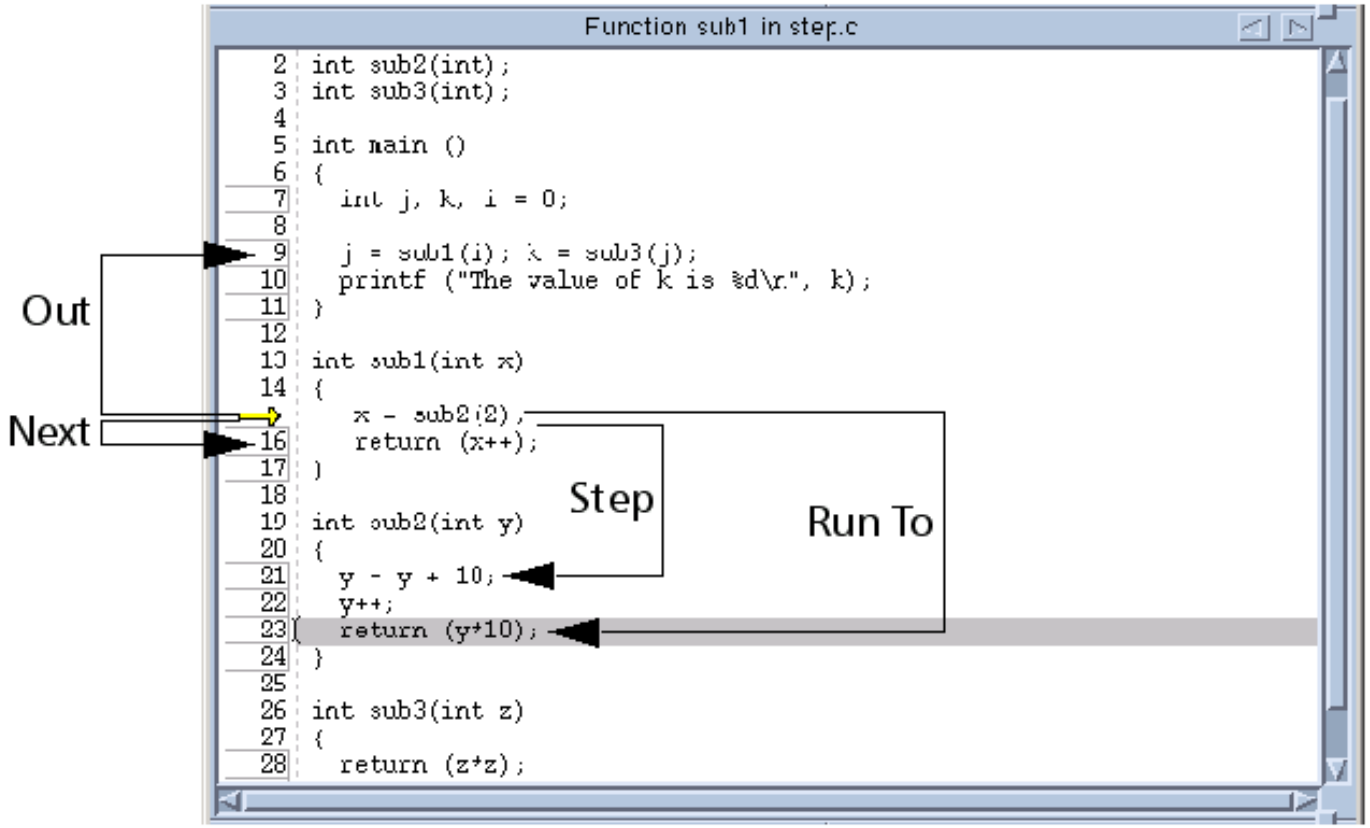

# Action Points

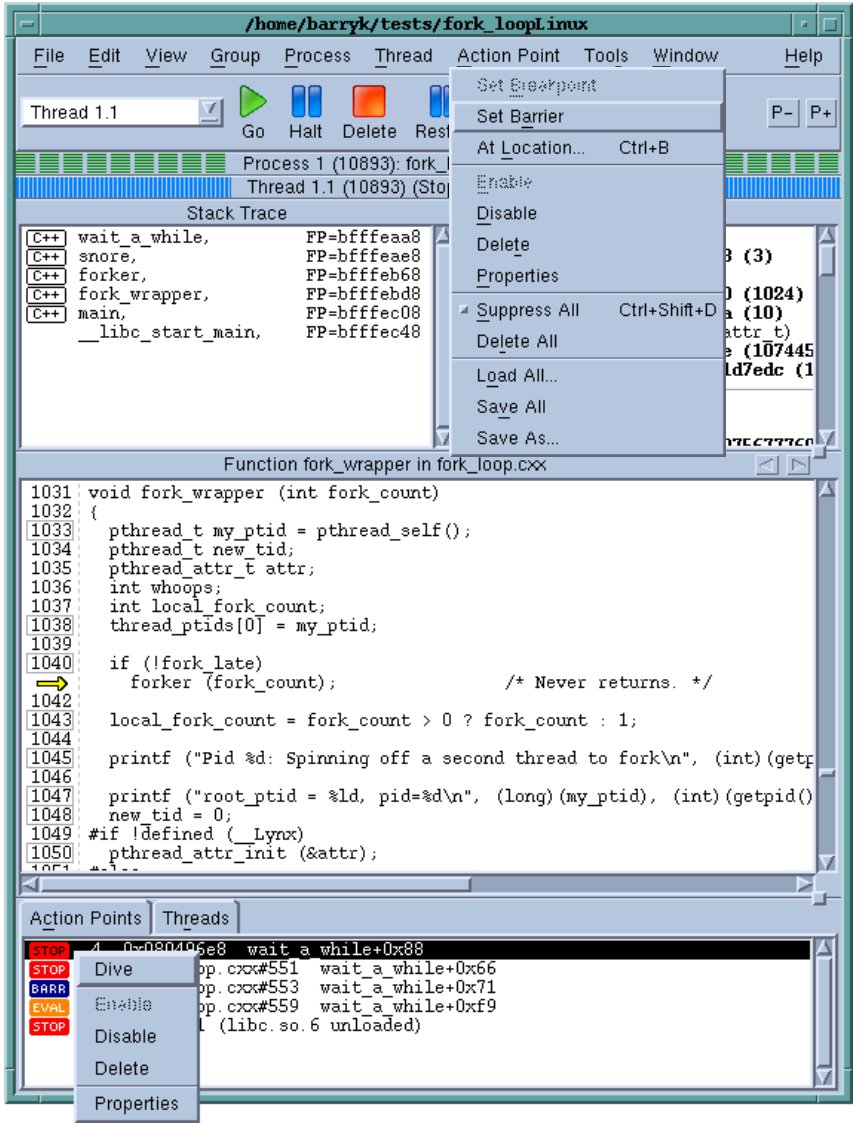


***Breakpoints***

***Barrier Points***

***Conditional Breakpoints***

***Evaluation Points***

***Watchpoints***

# Conditional Breakpoint

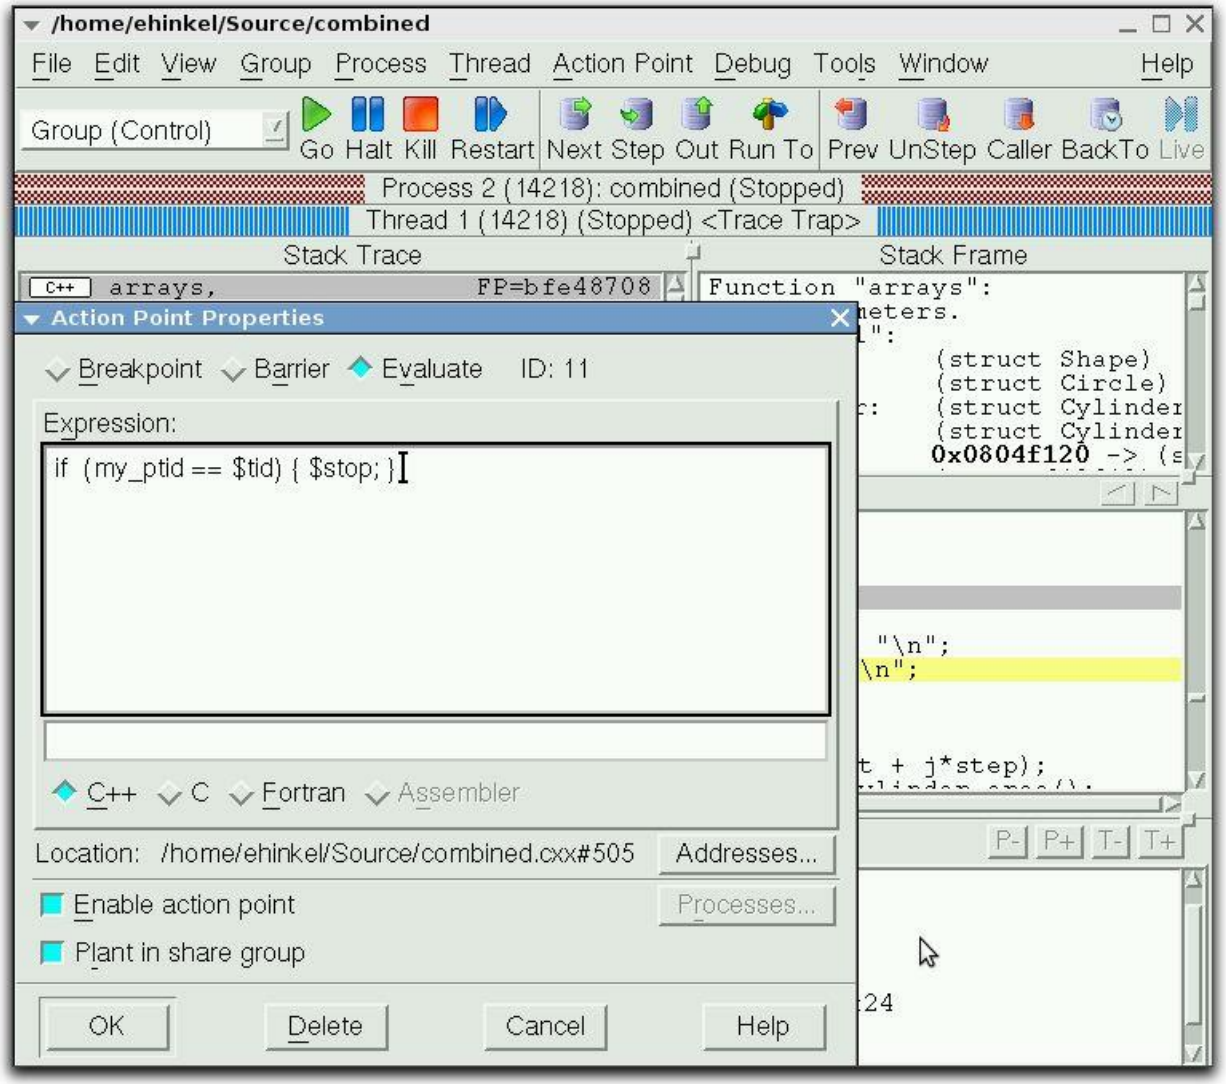

# Evaluation Point – Test Fixes on the Fly!

- Test small source code patches
- Call functions
- Set variables
- Test conditions
- C/C++ or Fortran
- Some limitations:
  - Can't use C++ constructors

```
item 0;99.099998  range 99.099998 - 99.099998  sum 99.099998
item 1;42.000000  range 42.000000 - 99.099998  sum 141.100006
item 2;58.299999  range 42.000000 - 99.099998  sum 199.400009
item 3;39.000000  range 39.000000 - 99.099998  sum 238.400009
item 4;77.699997  range 39.000000 - 99.099998  sum 316.100006
item 5;78.000000  range 39.000000 - 99.099998  sum 394.100006
item 6;92.099998  range 39.000000 - 99.099998  sum 486.200012
item 7;58.500000  range 39.000000 - 99.099998  sum 544.700012
item 8;58.000000  range 39.000000 - 99.099998  sum 602.700012
item 9;91.000000  range 39.000000 - 99.099998  sum 693.700012
N 10 min 39.000000 max 99.099998
total 693.700012 mean 69.370001
```

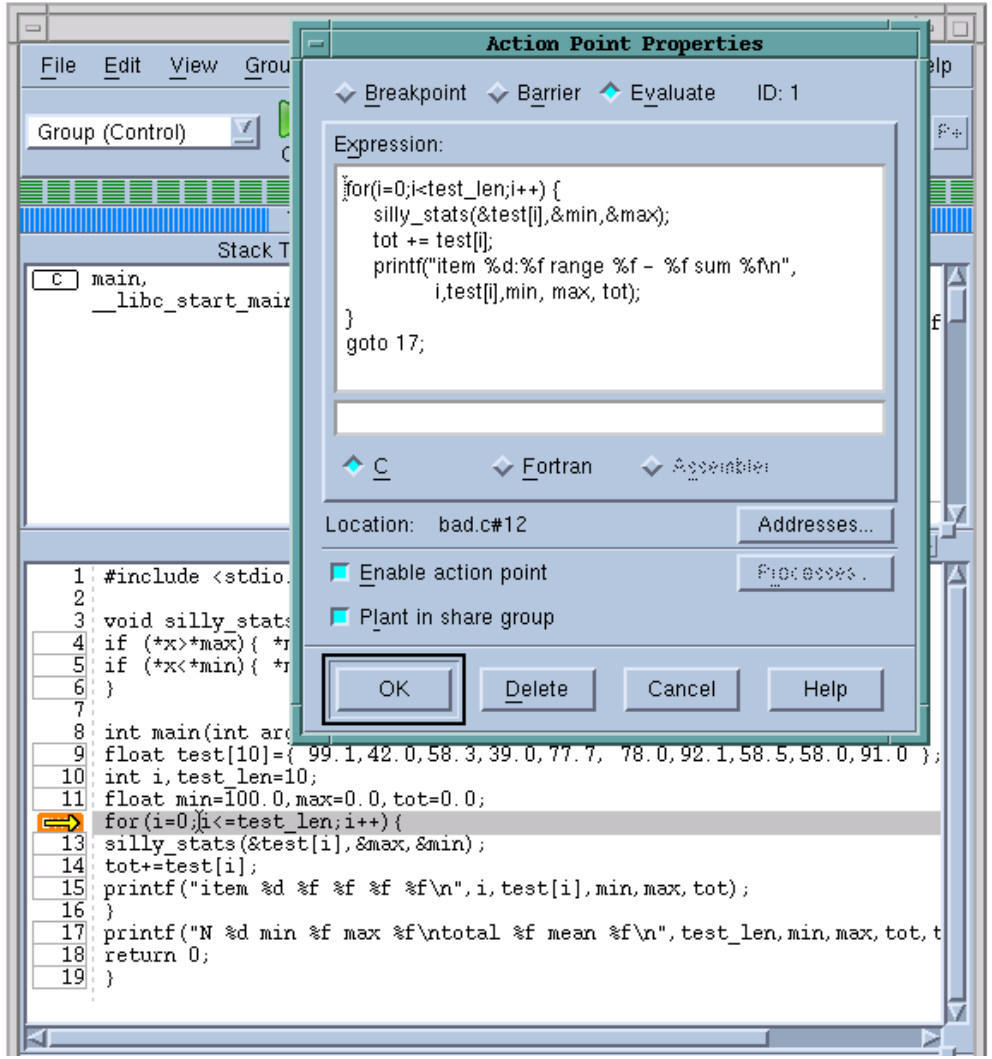

# Watchpoints

Watchpoints are set on a fixed memory region

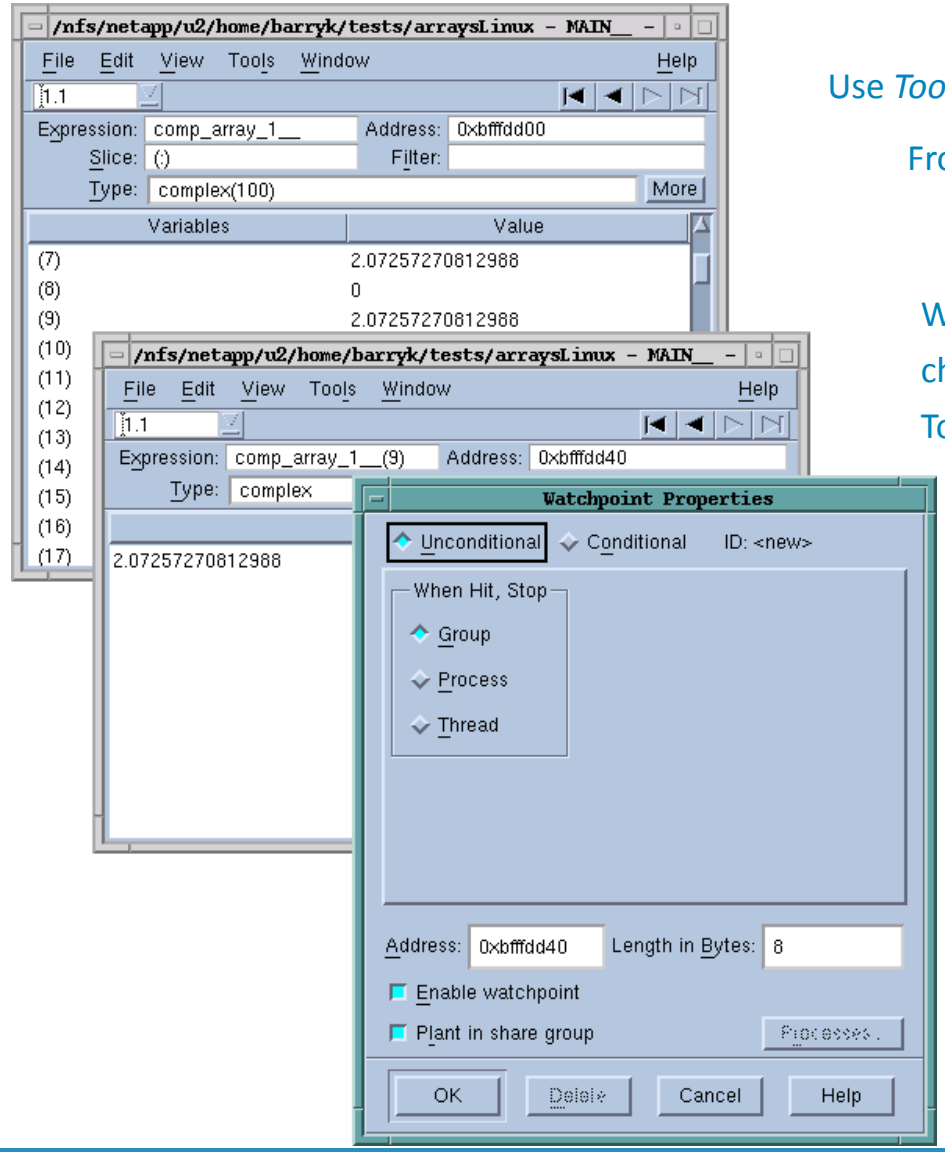


Use *Tools > Watchpoint* from a Variable Window or

From source pane with contextual menu

When the contents of watched memory change, the watchpoint is triggered and TotalView stops the program.

Watchpoints are not set on a variable. You you need to be aware of the variable scope.

Watchpoints can be conditional or unconditional

Uses Hardware Watchpoints with various limitations based on architecture

# Advanced C++ and Data Debugging

# Advanced C++ and Data Debugging

```
1  #include <functional>
2  #include <vector>
3  #include <iostream>
4  double eval(std::function<double(double)> f, double x = 2.0){
5    return f(x);}
6
7  int main(){
8  //   // One line lambdas
9     auto glambda1 = [](int a, float b) { return a < b; };
10    // Two line lambda
11    auto glambda2 = [](int a, float && b) {
12      if (a < b)
13      return 1;
14      if (b>a)
15      return -1;
16      return 0;
17    };
18
19    bool b = glambda1(3, 3.14);
20    int i = glambda2(3, 3.14);
21    for (int i=0; i<10;i++)
22      b = glambda1(i, 3.14+i);
23
24
25    std::function<double(double)> f0    = [](double x){
26      return 1;};
27    auto                          f1    = [](double x){
28      return x;};
29    decltype(f0)                  fa[3] = {f0,f1,[](double x){
```

- TotalView supports debugging the latest C++11/14 features including:
  - lambdas, transformations for smart pointers, auto types, R-Value references, range-based loops, strongly-typed enums, initializer lists, user defined literals

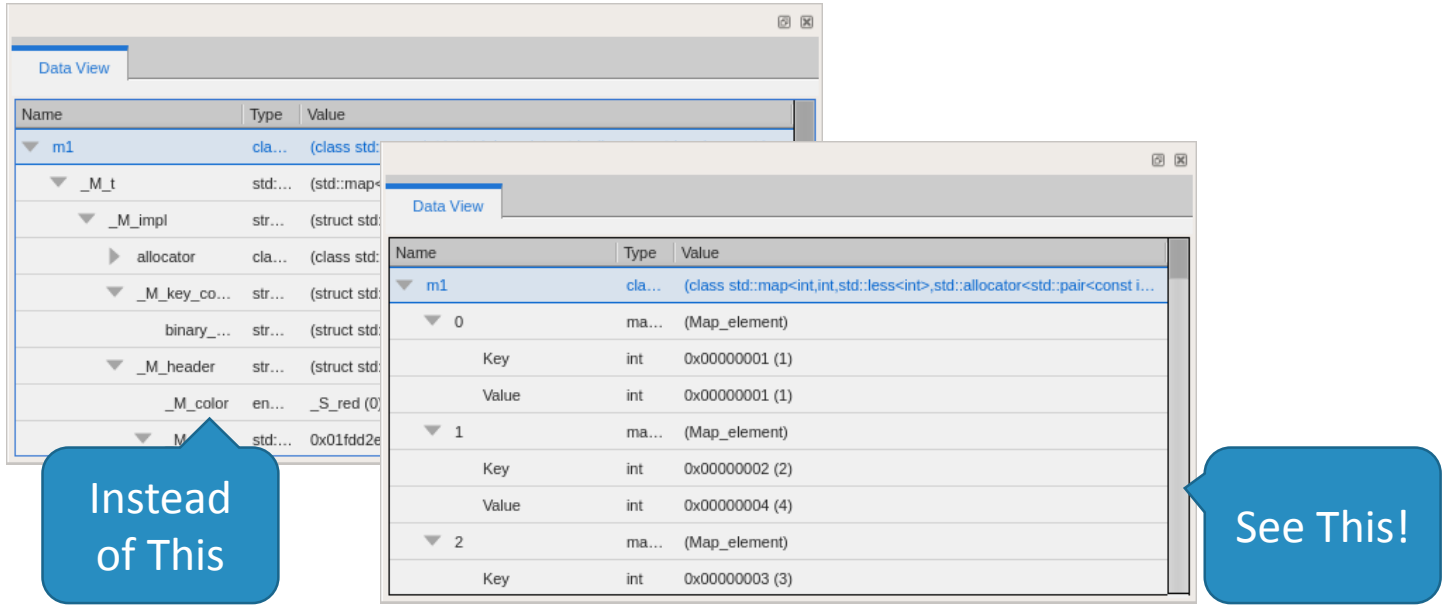


- TotalView transforms many of the C++ and STL containers such as:
  - array, forward_list, tuple, map, set, vector and others.

# Array Slicing, Striding and Filtering (classic UI)

- Slicing – reduce display to a portion of the array
  - [lower_bound:upper_bound]
  - [5:10]
- Striding – Skip over elements
  - [::stride]
  - [::5], [5:10:-1]
  - Filtering
    - Comparison: ==, !=, <, <=, >, >=
    - Range of values: [**>**] *low-value* **:** [**<**] *high-value*
    - IEEE values: $nan, $inf, $denorm

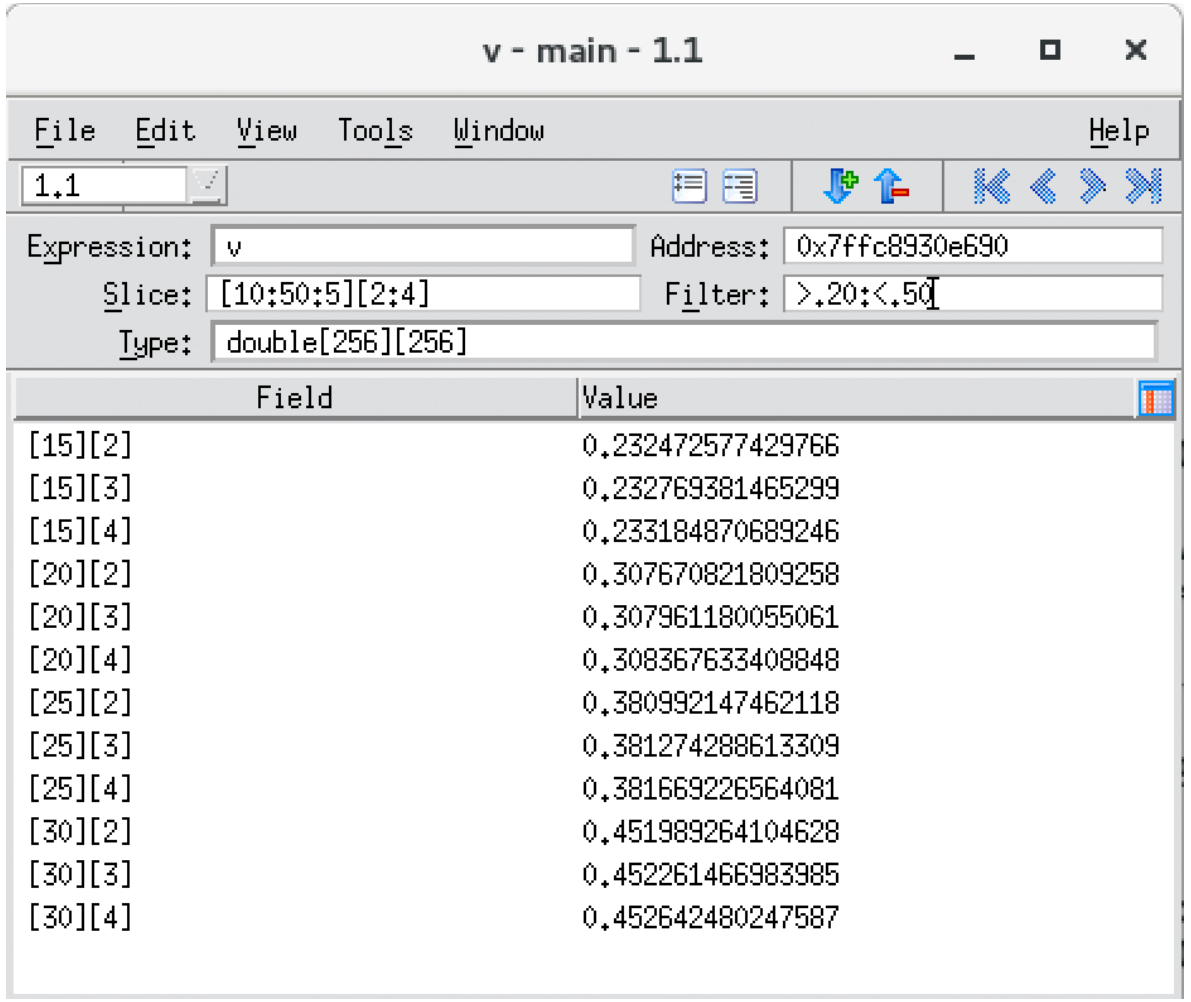

# Array Statistics

- Easily display a set of statistics for the filtered portion of your array

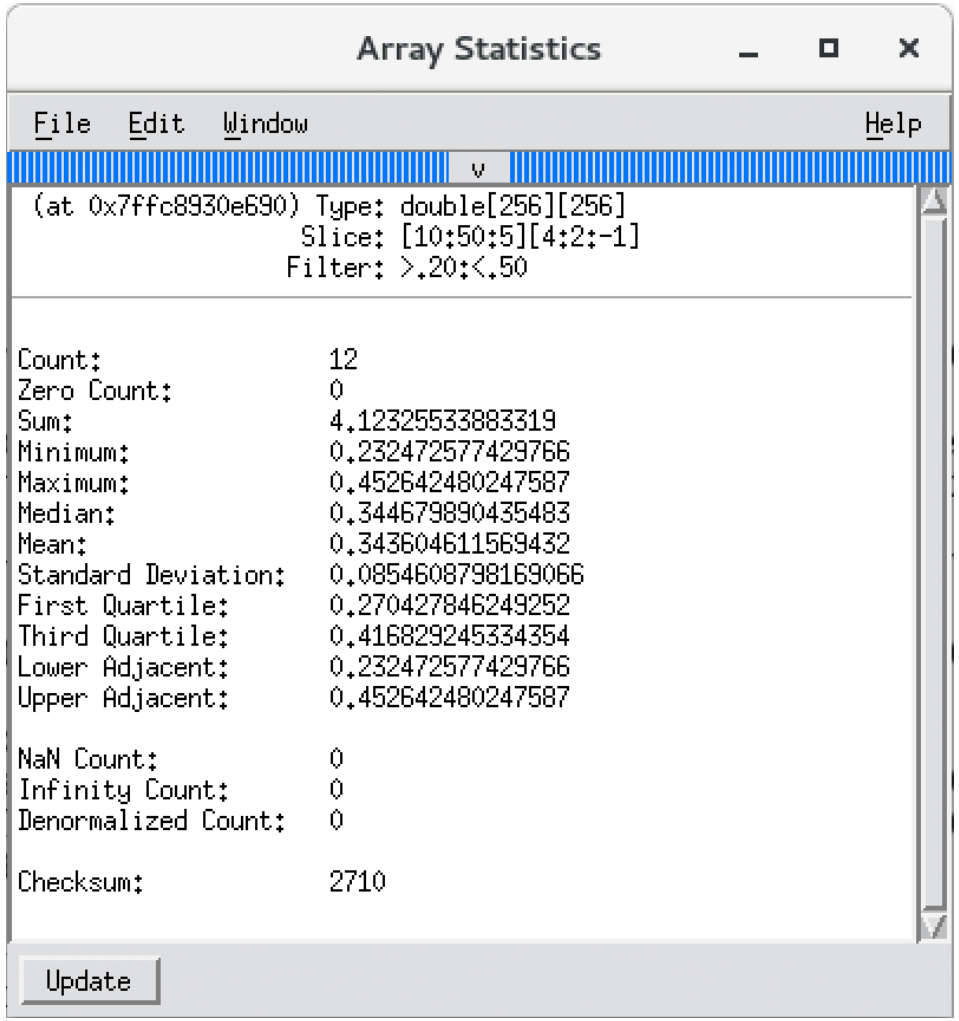

# Visualizing Array Data

- Visualizer creates graphic images of your program's array data.
- Visualize one or two dimensional arrays
- View data manually through the Window > Visualize command on the Data Window
- Visualize data programmatically using the $visualize function

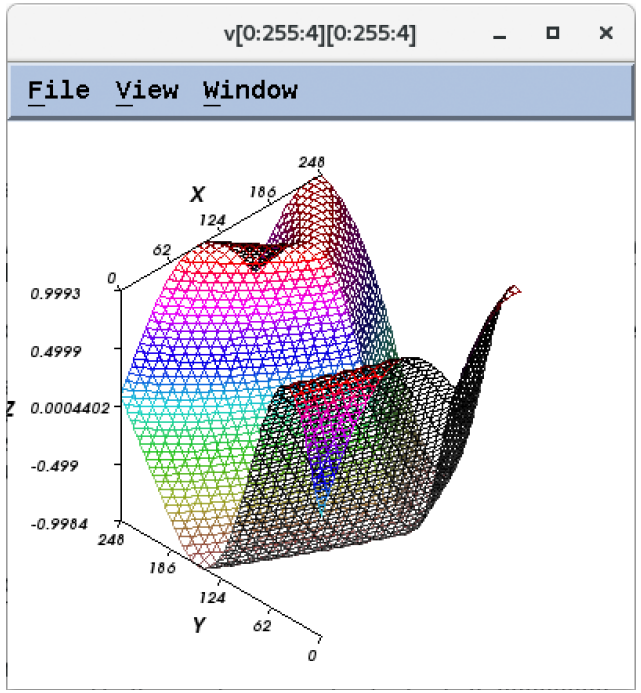


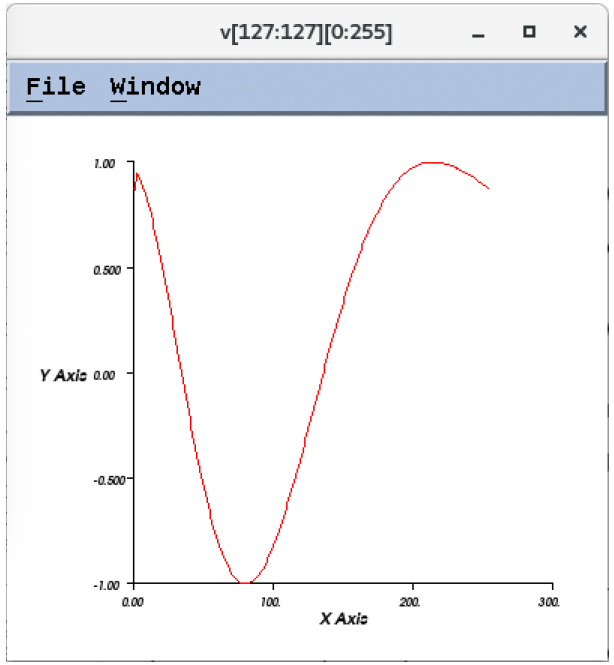

# Dive in All

- Dive in All
  - Use Dive in All to easily see each member of a data structure from an array of structures

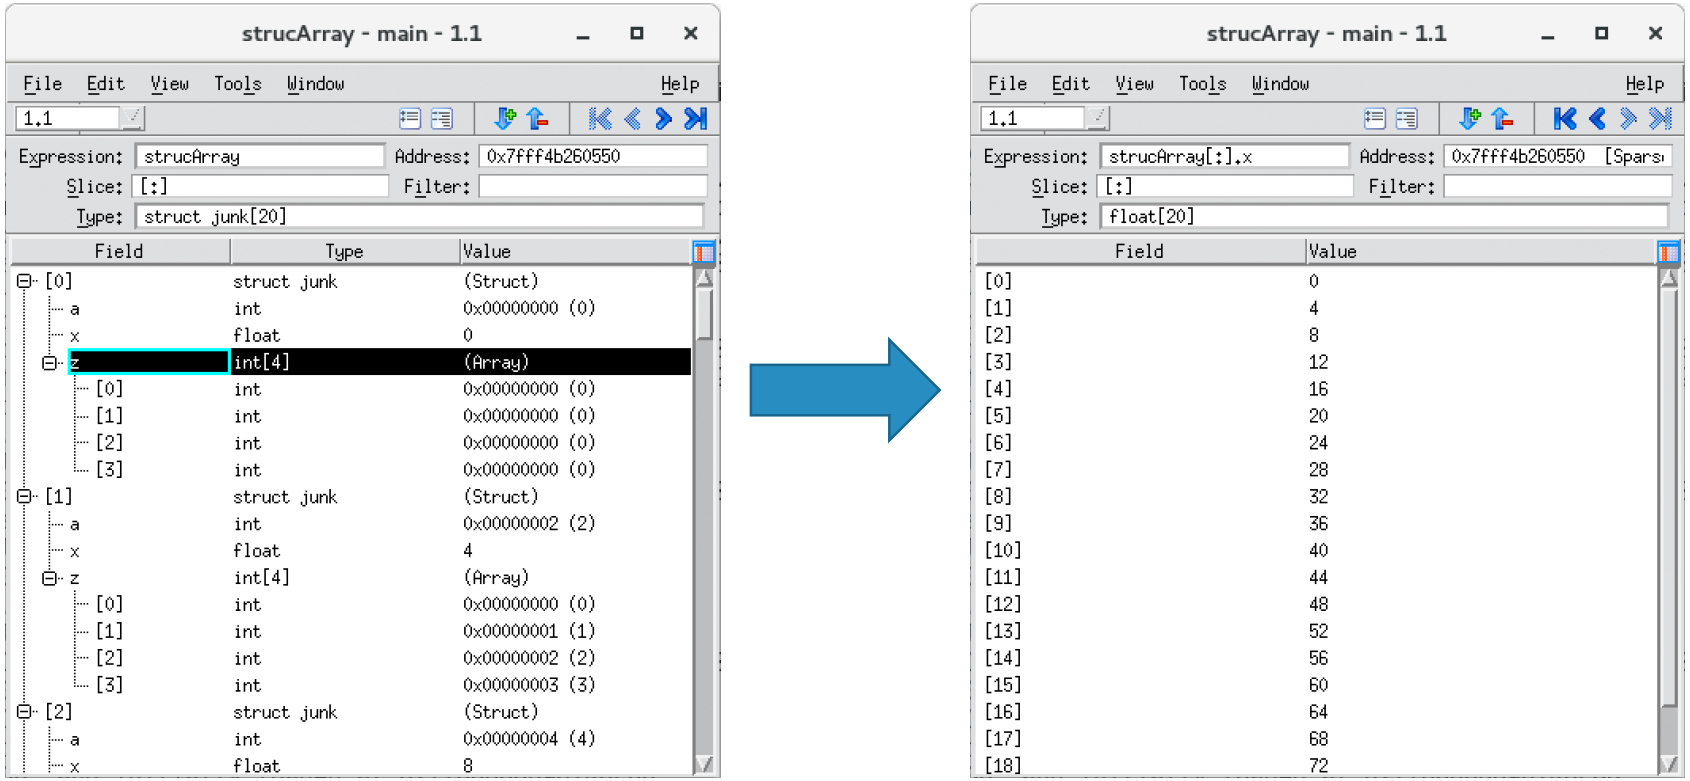

# Looking at Variables Across Processes

- TotalView allows you to look at the value of a variable in all MPI processes
  - Right Click on the variable
  - Select the View > View Across
- TotalView creates an array indexed by process
- You can filter and visualize
- Use for viewing distributed arrays as well.
- You can also View Across Threads

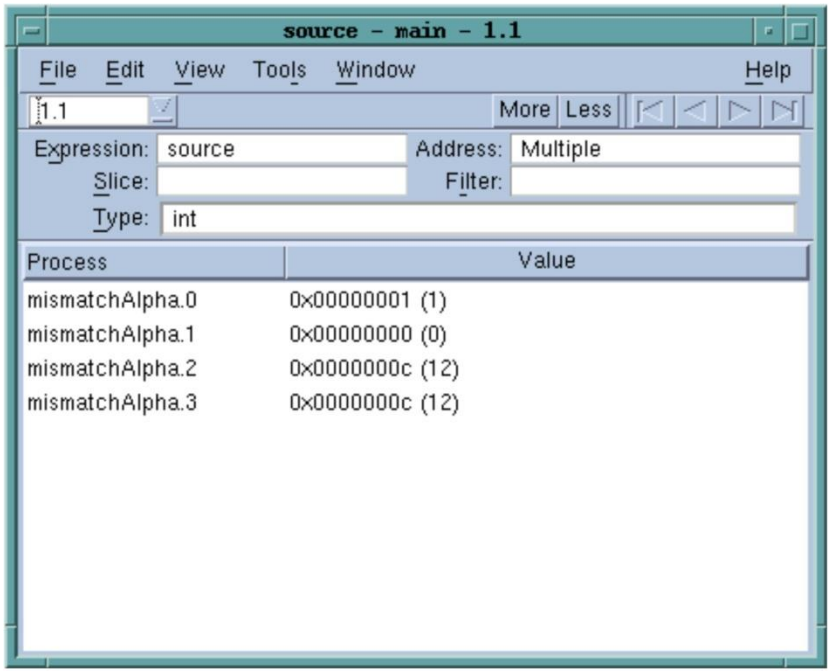

# Multi-Thread and Multi-Process Parallel Debugging

# In the Parallel Program Session select:

Select:

- MPI preference
- number of tasks
- number of nodes
- starter arguments

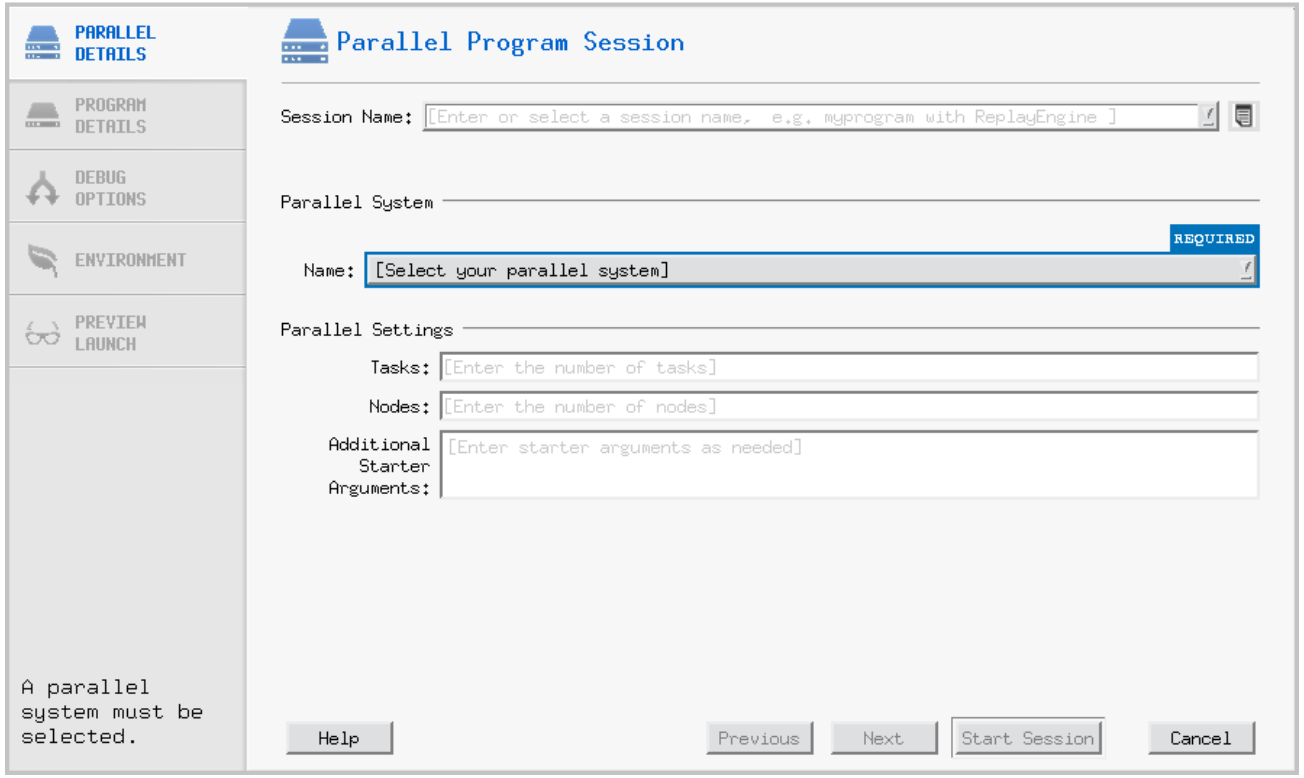


… then save all this in Session

# Stepping Commands

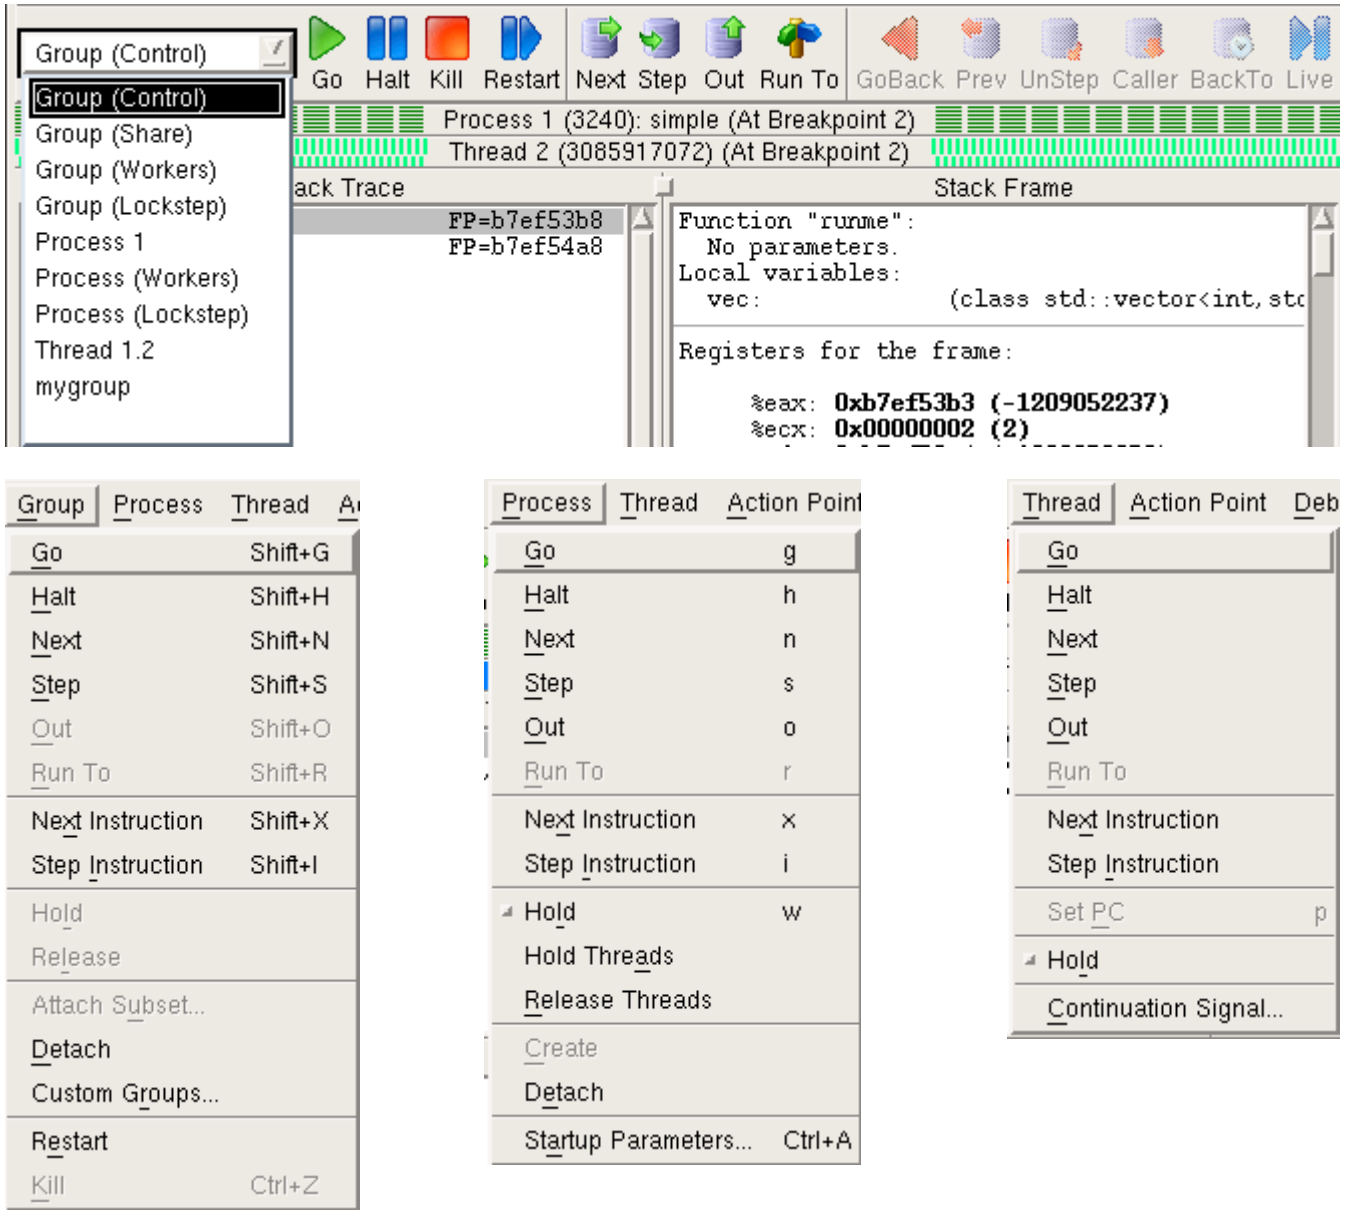

# Message Queue Graph

- Hangs & Deadlocks
- Pending Messages
  - Receives
  - Sends
  - Unexpected
- Inspect
  - Individual entries
- Patterns

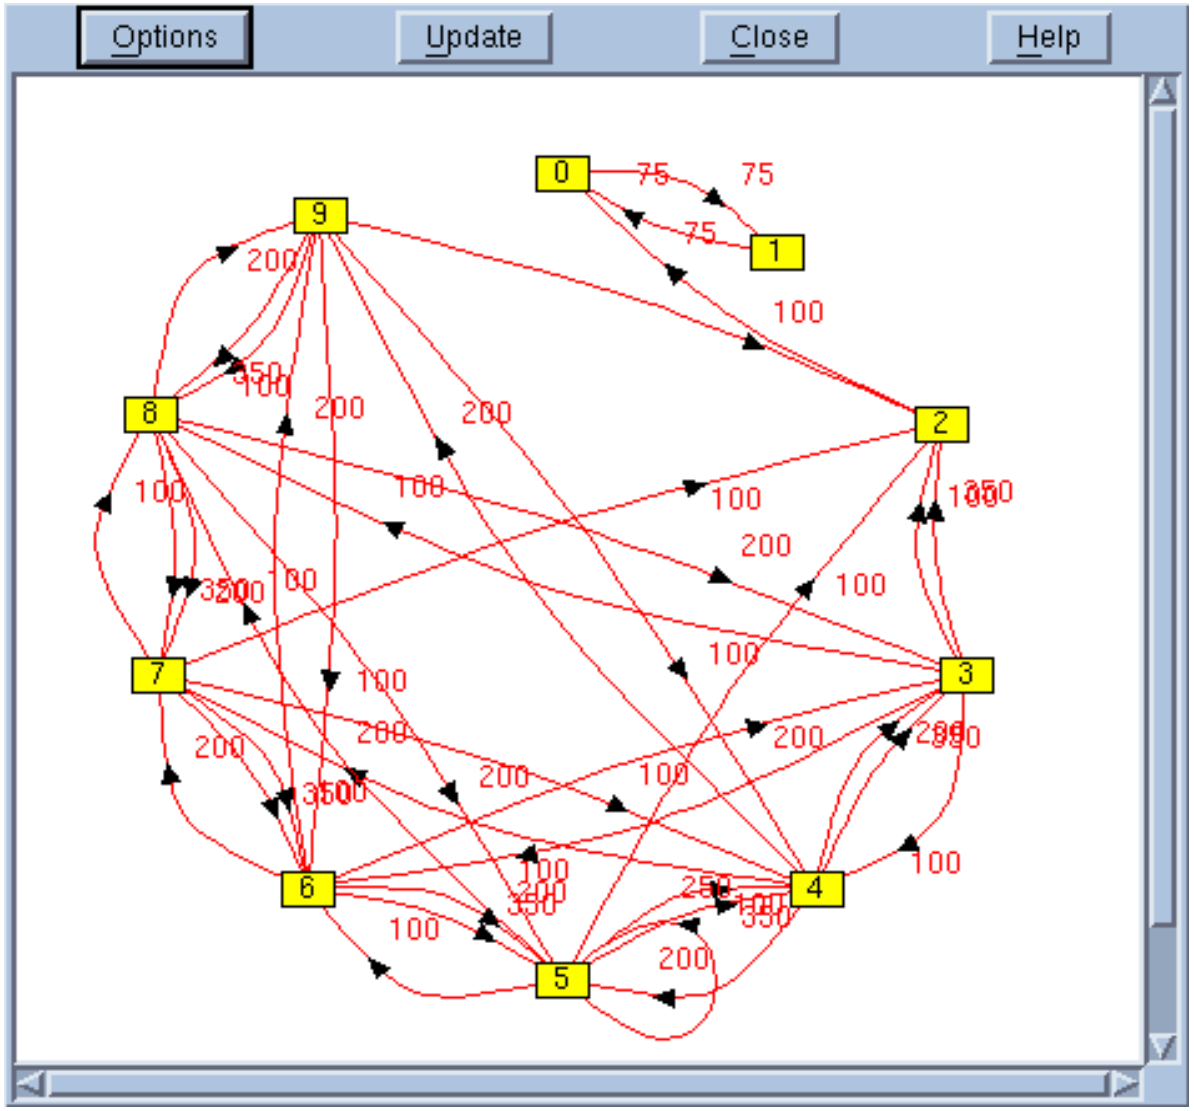

# Find Deadlocks and Performance Sinks

- Filtering
  - Choose messages to track
  - Choose MPI Communicators
- Cycle detection →

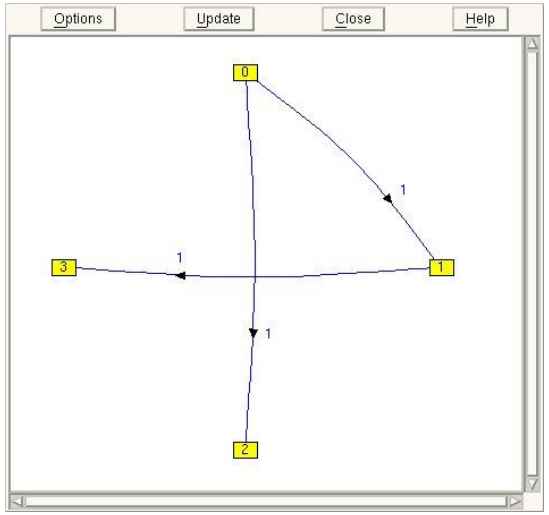


← Sink

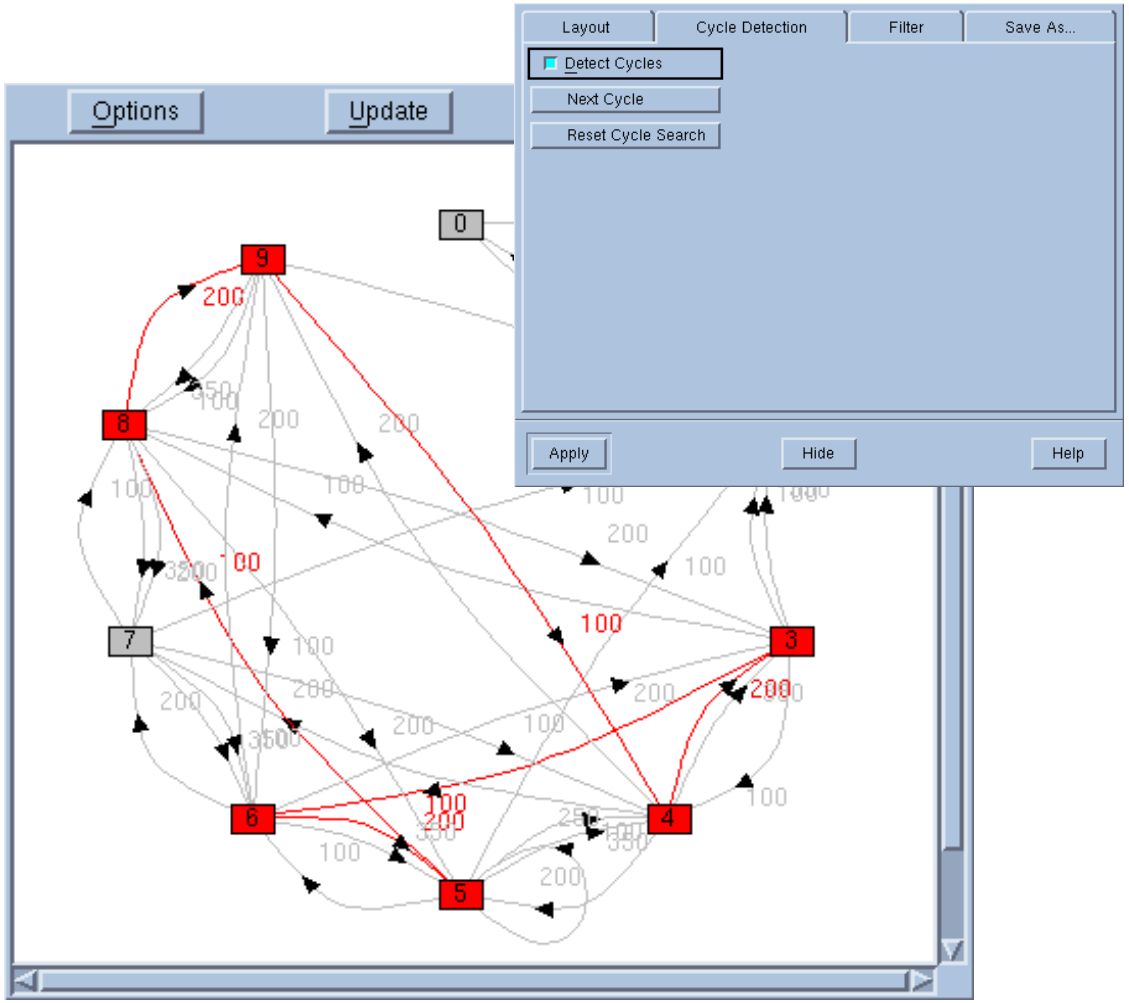

# Reverse Debugging

# Reverse Debugging

## Replay Engine – The right way to debug


Step forward over functions


Step *backward* over functions


Step forward into functions


Step *backward* into functions


Advance forward out of current Function, after the call


Advance backward out of current Function, to before the call


Advance forward to selected line

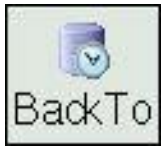


Advance backward to selected line

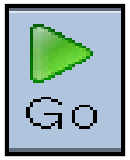


Run forward


Run *backward*


Advance forward to “live” session

# ReplayEngine

- **Captures execution history**
  - Records all external input to program
  - Records internal sources of non-determinism
- **Replays execution history**
  - Examine any part of the execution history
  - Step back as easily as forward
  - Jump to points of interest
- **An add-on product to TotalView**
  - Support for
    - Linux/x86
    - Linux x86- 64



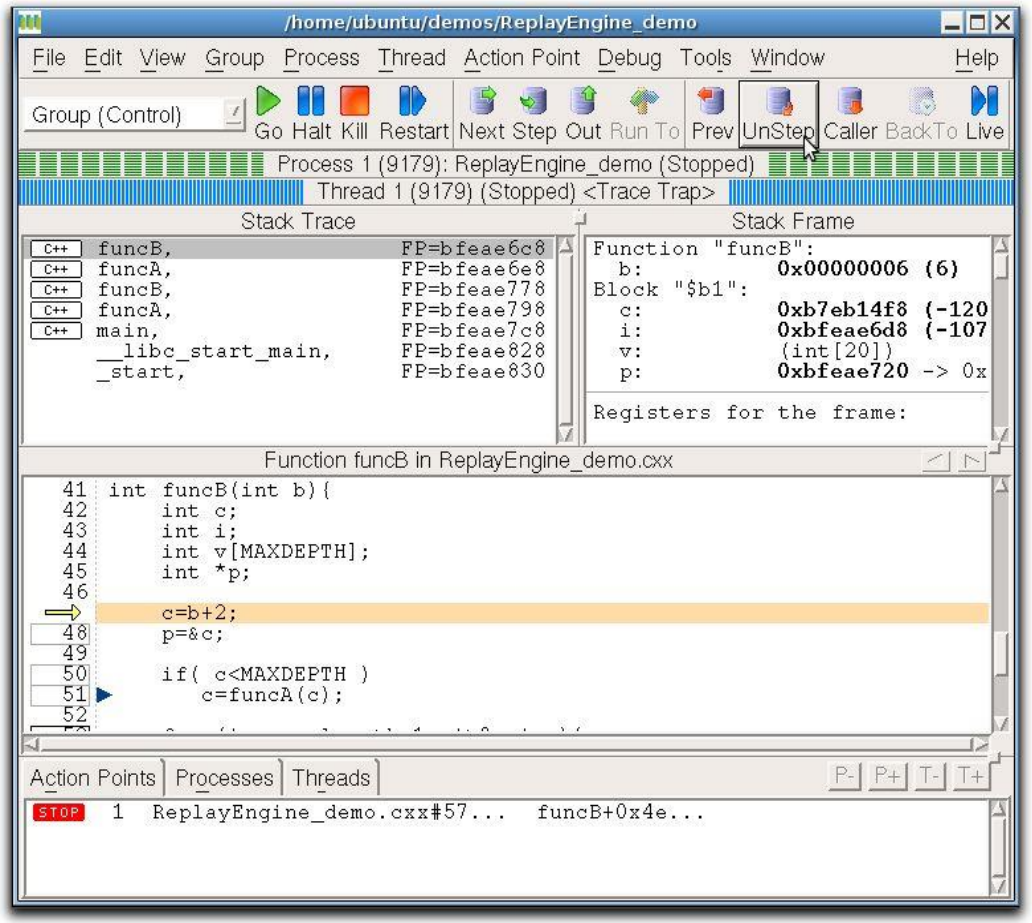



# Memory Debugging

# Memory Debugging

- TotalView's memory debugging technology allows you to
  - Easily find memory leaks and other memory errors
  - Detect malloc/free new/delete API misuse
  - Dangling pointer detection
  - Detect buffer overruns
  - Paint memory blocks on allocation and deallocation
- Memory debugging results can be easily shared as
  - HTML reports or raw memory debugging files.
- Compare memory results between runs to verify elimination of leaks
- Supports parallel applications
- Low overhead and does not require recompilation or instrumentation

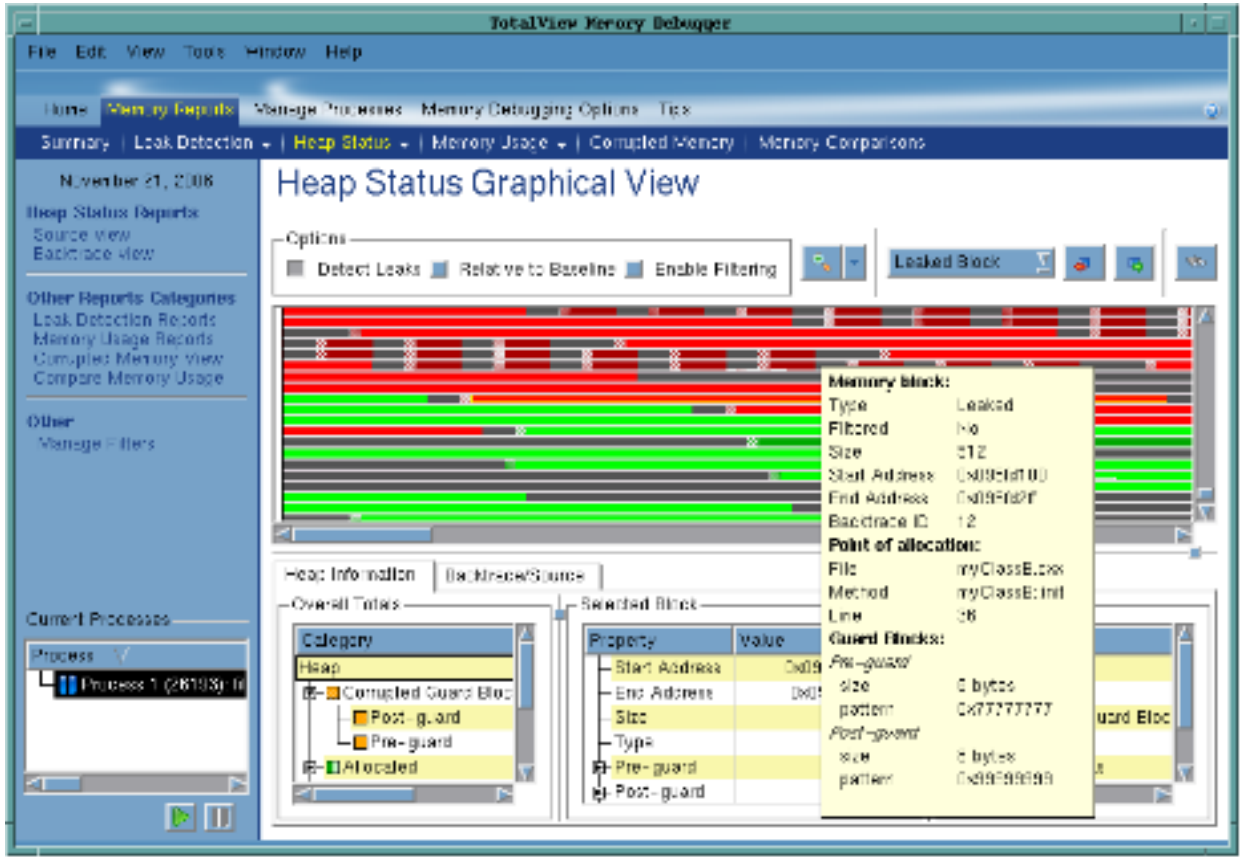

# Strategies for Parallel Memory Debugging

- Run the application and see if memory events are detected
- View memory usage across the MPI job
  - Compare memory footprint of the processes
    - Are there any outliers? Are they expected?
- Gather heap information in all processes of the MPI job
  - Select and examine individually
    - Look at the allocation pattern.
      Does it make sense?
    - Look for leaks
  - Compare with the 'diff' mechanism
    - Are there any major differences?
      Are they expected?

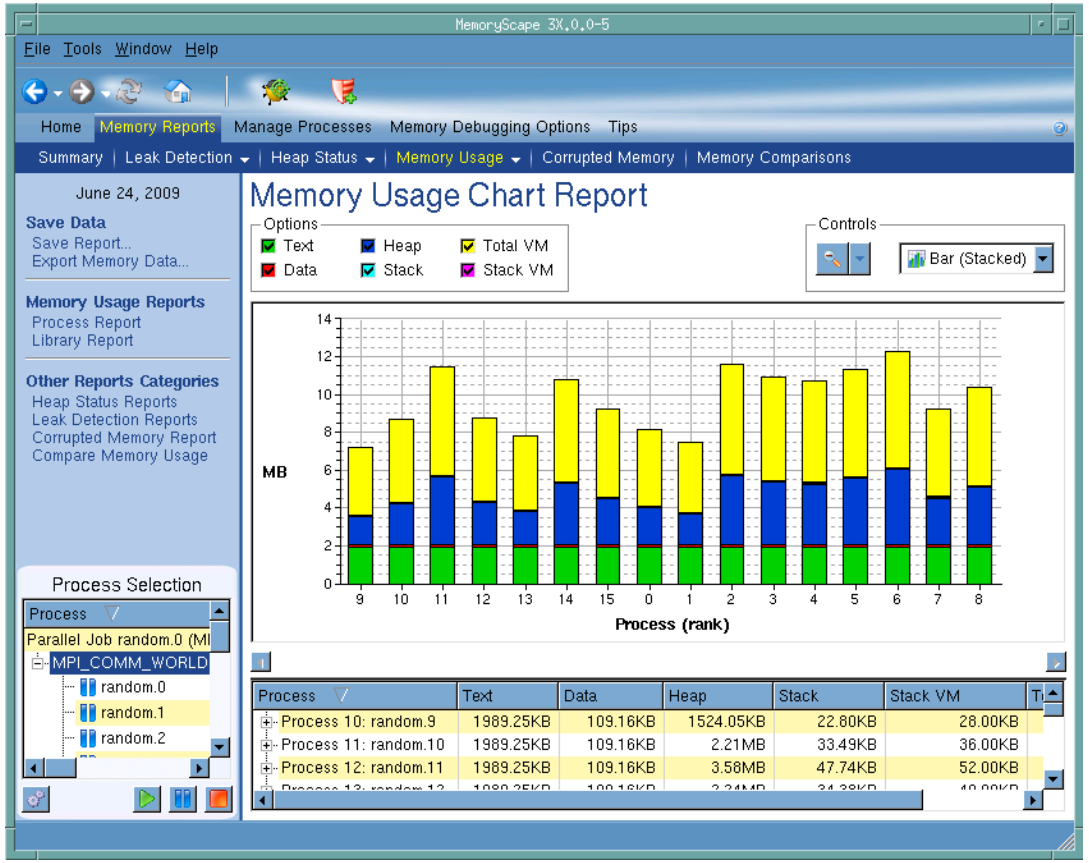

# GPU Debugging

# GPU debugging with TotalView

- NVIDIA CUDA support
  - Multiple platforms : X86-64,PowerLE, ARM64
  - Multiple cards:  from Jetson to Turing
- Features and capabilities include
  - Support for dynamic parallelism
  - Support for MPI based clusters and multi-card configurations
  - Flexible Display and Navigation on the CUDA device
    - Physical (device, SM, Warp, Lane)
    - Logical (Grid, Block) tuples
  - CUDA device window reveals what is running where
  - Support for CUDA Core debugging
  - Leverages CUDA memcheck
  - Support for OpenACC

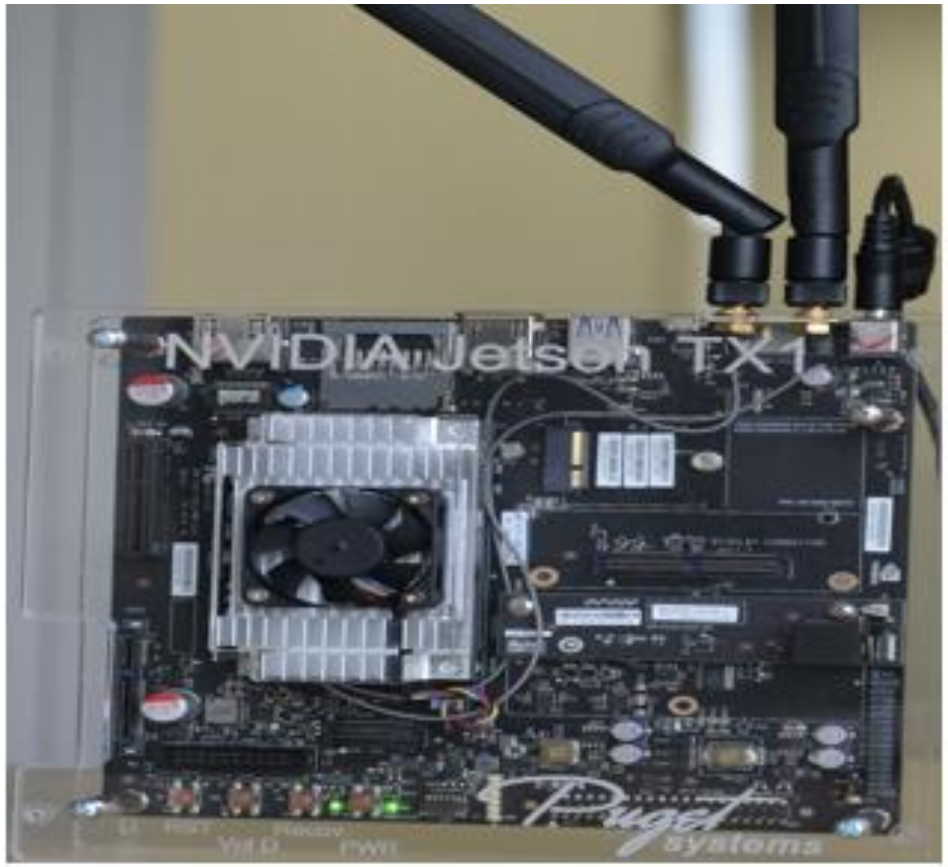

# GPU Debugging Model Improvements

- First in class Unified Source debugging
- Improves and streamlines debugging CUDA applications

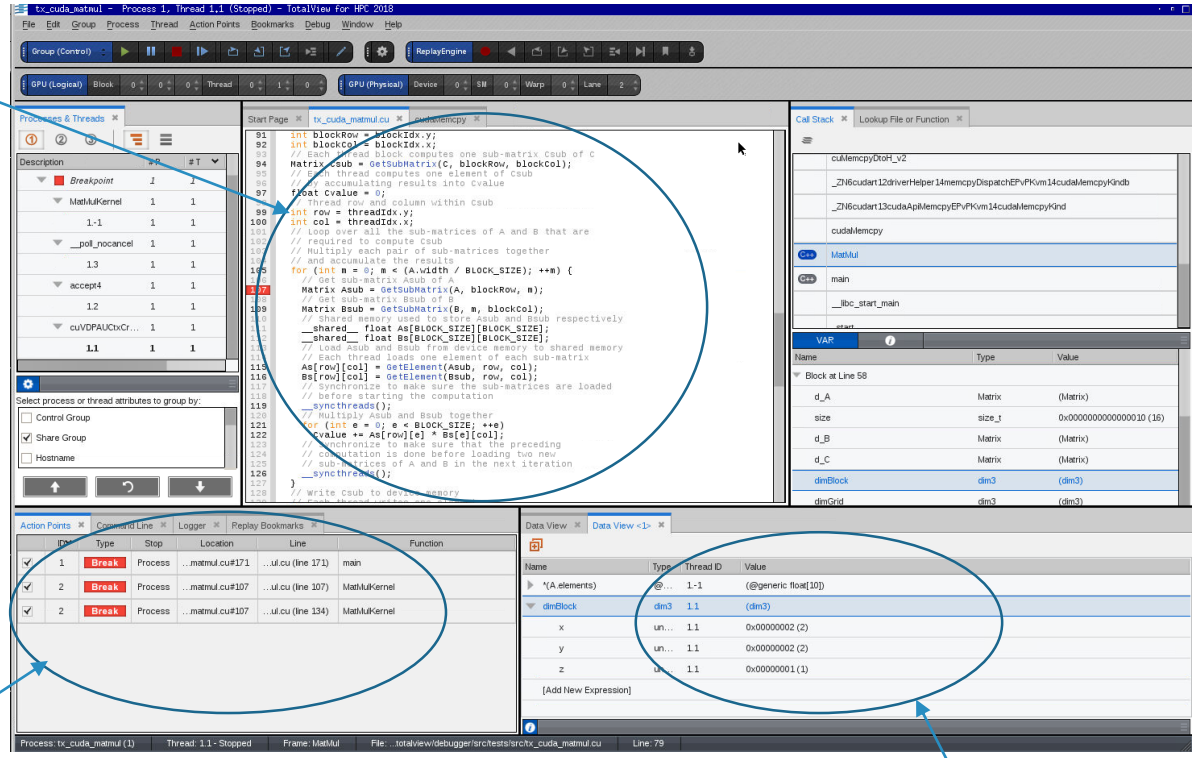

- Set breakpoints in CPU **and** GPU kernel code before it is launched on the GPU
- Compare variables in CPU and GPU code together

# CUDA Debugging Demo

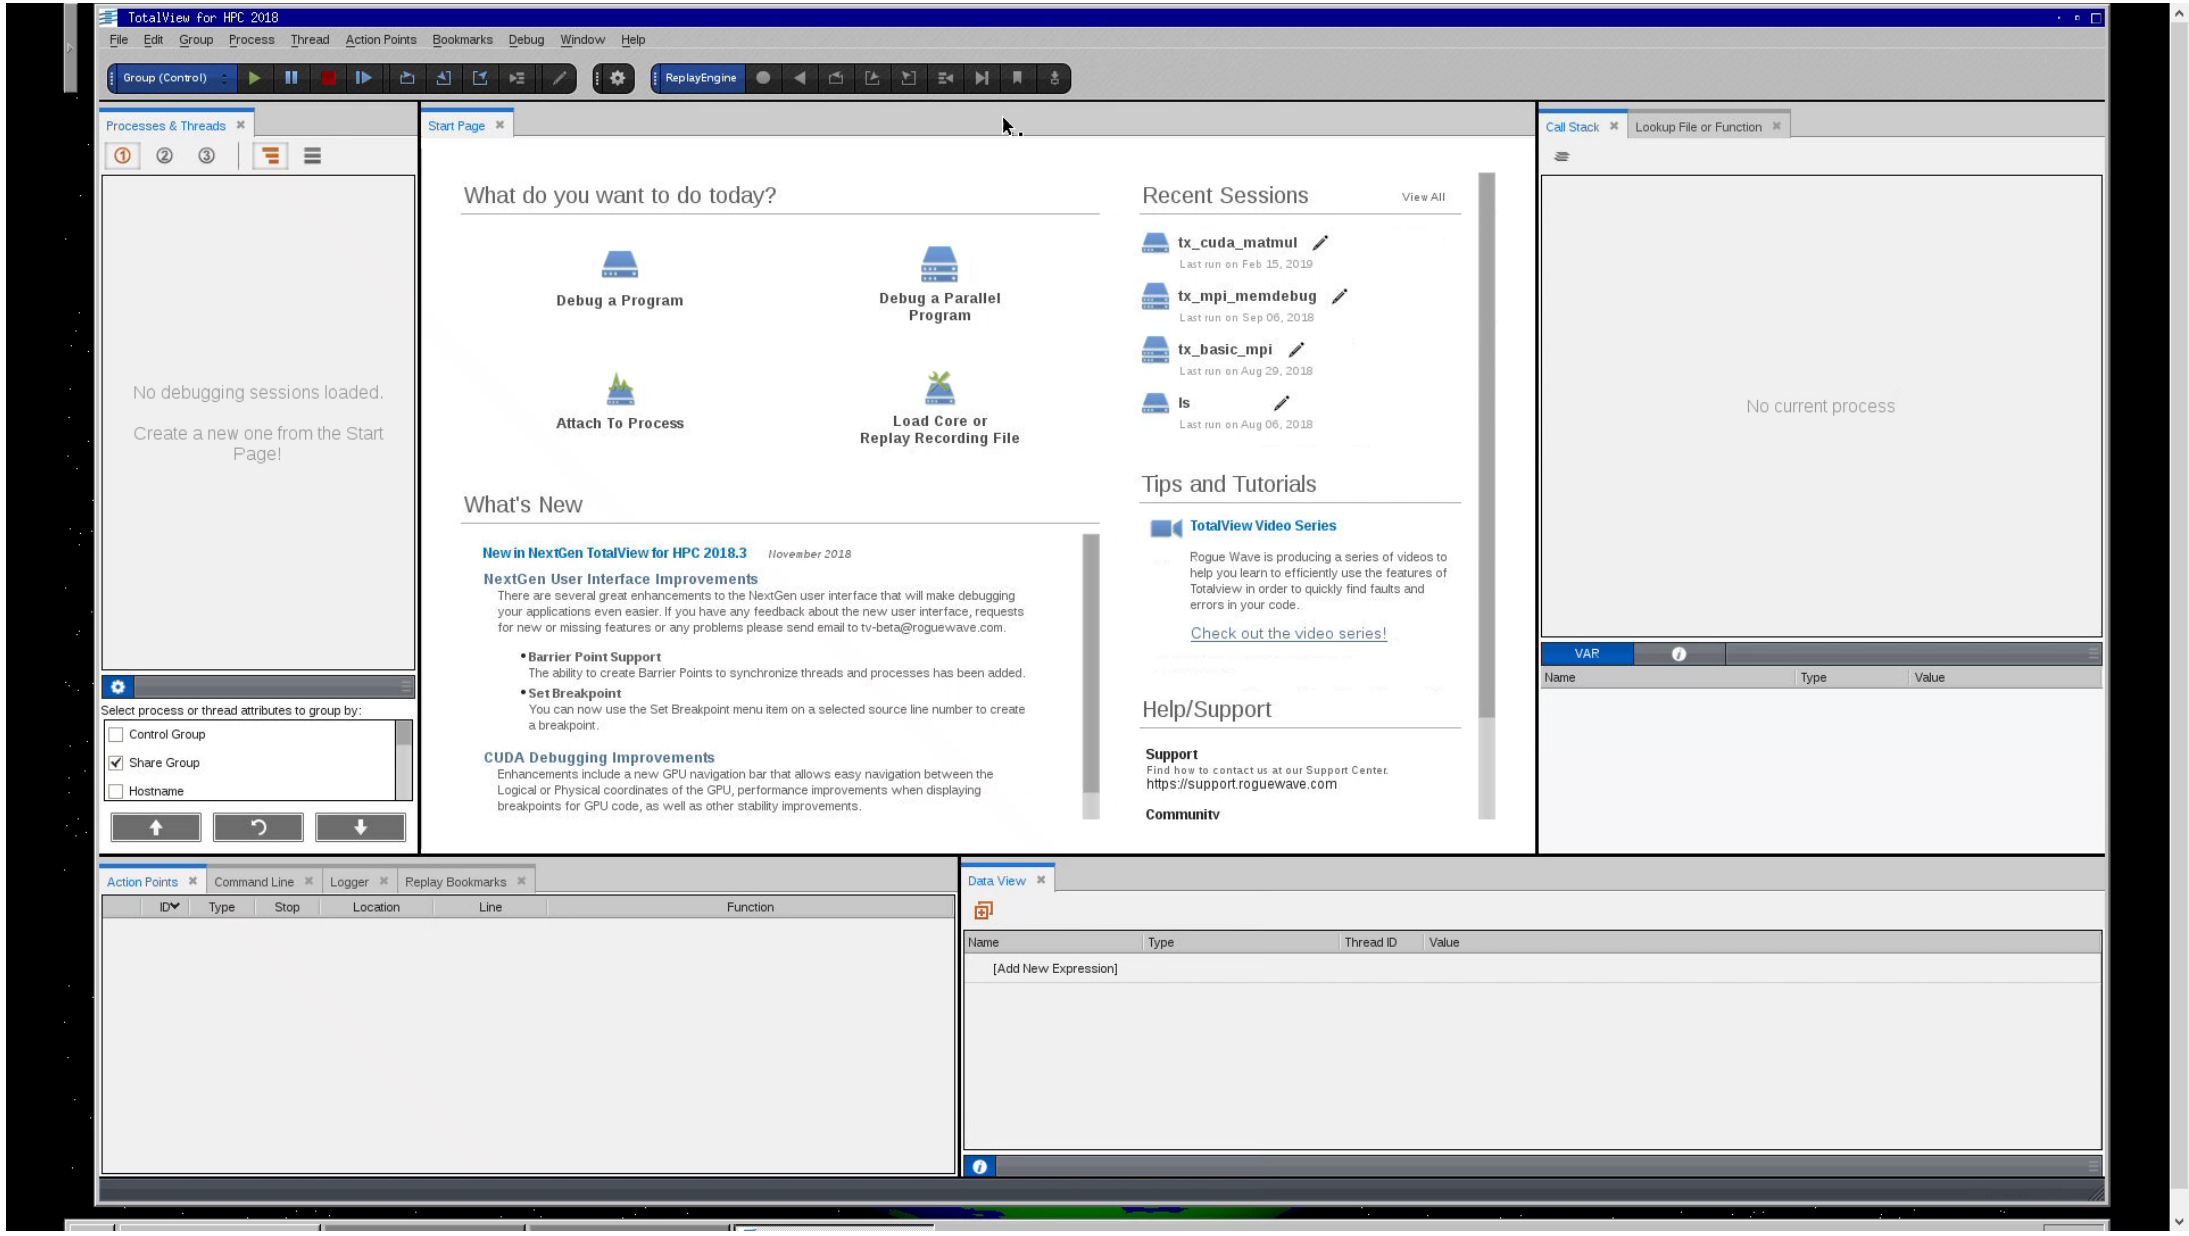
TotalView for HPC 2018
File Edit Group Process Thread Action Points Bookmarks Debug Window Help
Group (Control)
ReplayEngine
Processes & Threads
No debugging sessions loaded.
Create a new one from the Start Page!
Select process or thread attributes to group by:
Control Group
Share Group
Hostname
Start Page
What do you want to do today?
Debug a Program
Debug a Parallel Program
Attach To Process
Load Core or Replay Recording File
What's New
New in NextGen TotalView for HPC 2018.3
November 2018
NextGen User Interface Improvements
There are several great enhancements to the NextGen user interface that will make debugging your applications even easier. If you have any feedback about the new user interface, requests for new or missing features or any problems please send email to tv-beta@roguewave.com.
Barrier Point Support
The ability to create Barrier Points to synchronize threads and processes has been added.
Set Breakpoint
You can now use the Set Breakpoint menu item on a selected source line number to create a breakpoint.
CUDA Debugging Improvements
Enhancements include a new GPU navigation bar that allows easy navigation between the Logical or Physical coordinates of the GPU, performance improvements when displaying breakpoints for GPU code, as well as other stability improvements.
Recent Sessions
View All
tx_cuda_matmul
Last run on Feb 15, 2019
tx_mpi_memdebug
Last run on Sep 06, 2018
tx_basic_mpi
Last run on Aug 29, 2018
ls
Last run on Aug 06, 2018
Tips and Tutorials
TotalView Video Series
Rogue Wave is producing a series of videos to help you learn to efficiently use the features of Totalview in order to quickly find faults and errors in your code.
Check out the video series!
Help/Support
Support
Find how to contact us at our Support Center.
https://support.roguewave.com
Community
Call Stack
Lookup File or Function
No current process
VAR
Name Type Value
Action Points
Command Line
Logger
Replay Bookmarks
ID Type Stop Location Line Function
Data View
Name Type Thread ID Value
[Add New Expression]

# Extending Debugging Capabilities: How to Debug (AI) Mixed Python/C++ Code

# Debugging multiple languages

- Debugging one language is difficult enough
  - Especially with many threads/processes
- The language intersection is tougher
  - Data comparison
  - Glue code
- Issues are:
  - Type mismatches
  - Extraneous stack frames

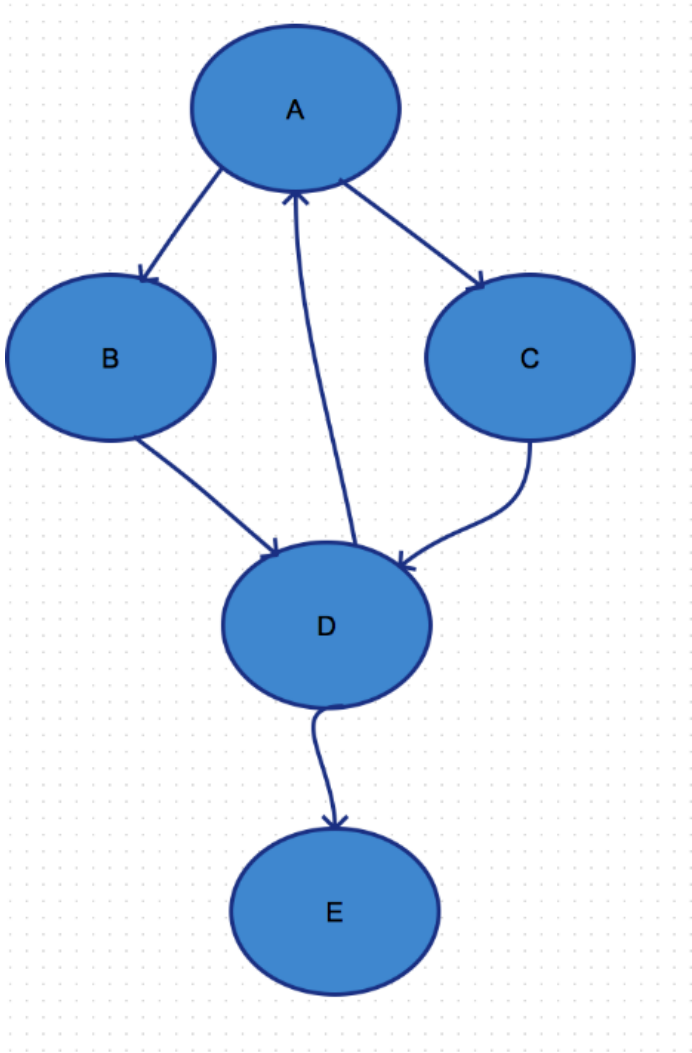

# Python debugging with TotalView (New GUI only)

- What TotalView provides:
  - Easy Python debugging session setup
  - Fully integrated Python and C/C++ call stack
    - ”Glue” layers between the languages removed
  - Easily examine and compare variables in Python and C++
  - Utilize reverse debugging and memory debugging

- What TotalView does not provide (yet):
  - Setting breakpoints and stepping within Python code

# Demo

```
#!/usr/bin/python

def callFact():
  import tv_python_example as tp
  a = 3
  b = 10
  c = a+b
  ch = “local string”
   ……
  return tp.fact(a)
if __name__ == '__main__’:
  b = 2
  result = callFact()
  print result
```

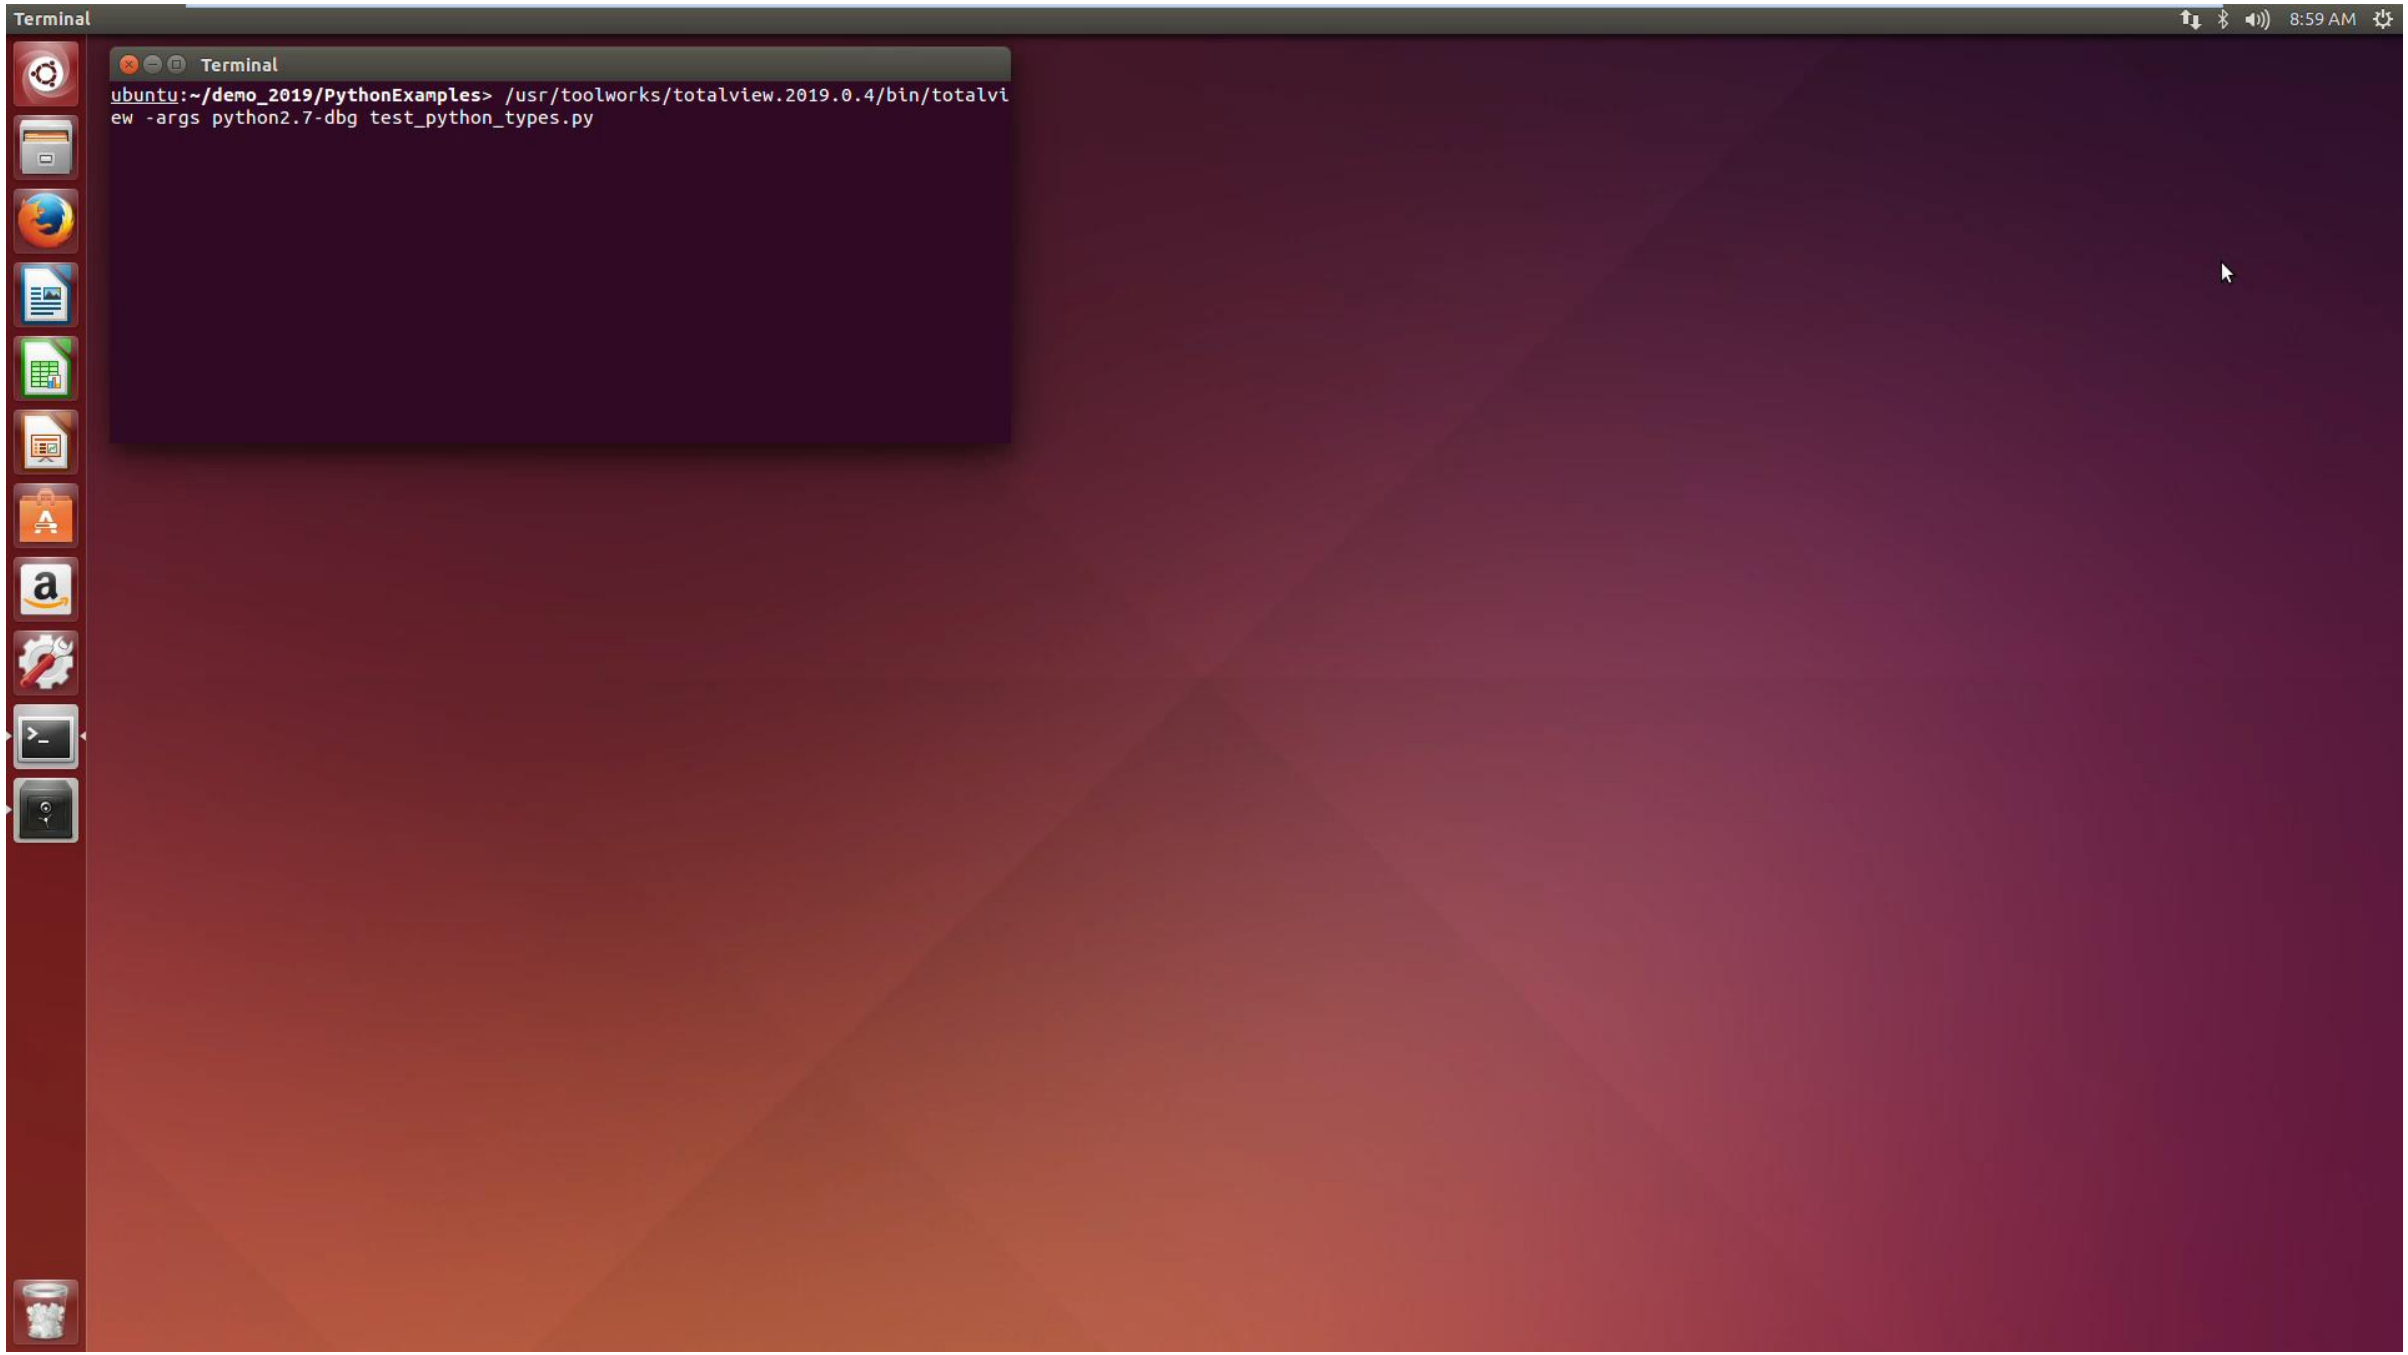
Terminal
8:59 AM
Terminal
ubuntu:~/demo_2019/PythonExamples> /usr/toolworks/totalview.2019.0.4/bin/totalvi
ew -args python2.7-dbg test_python_types.py

# Python without special debugger support

No viewing of Python data and code

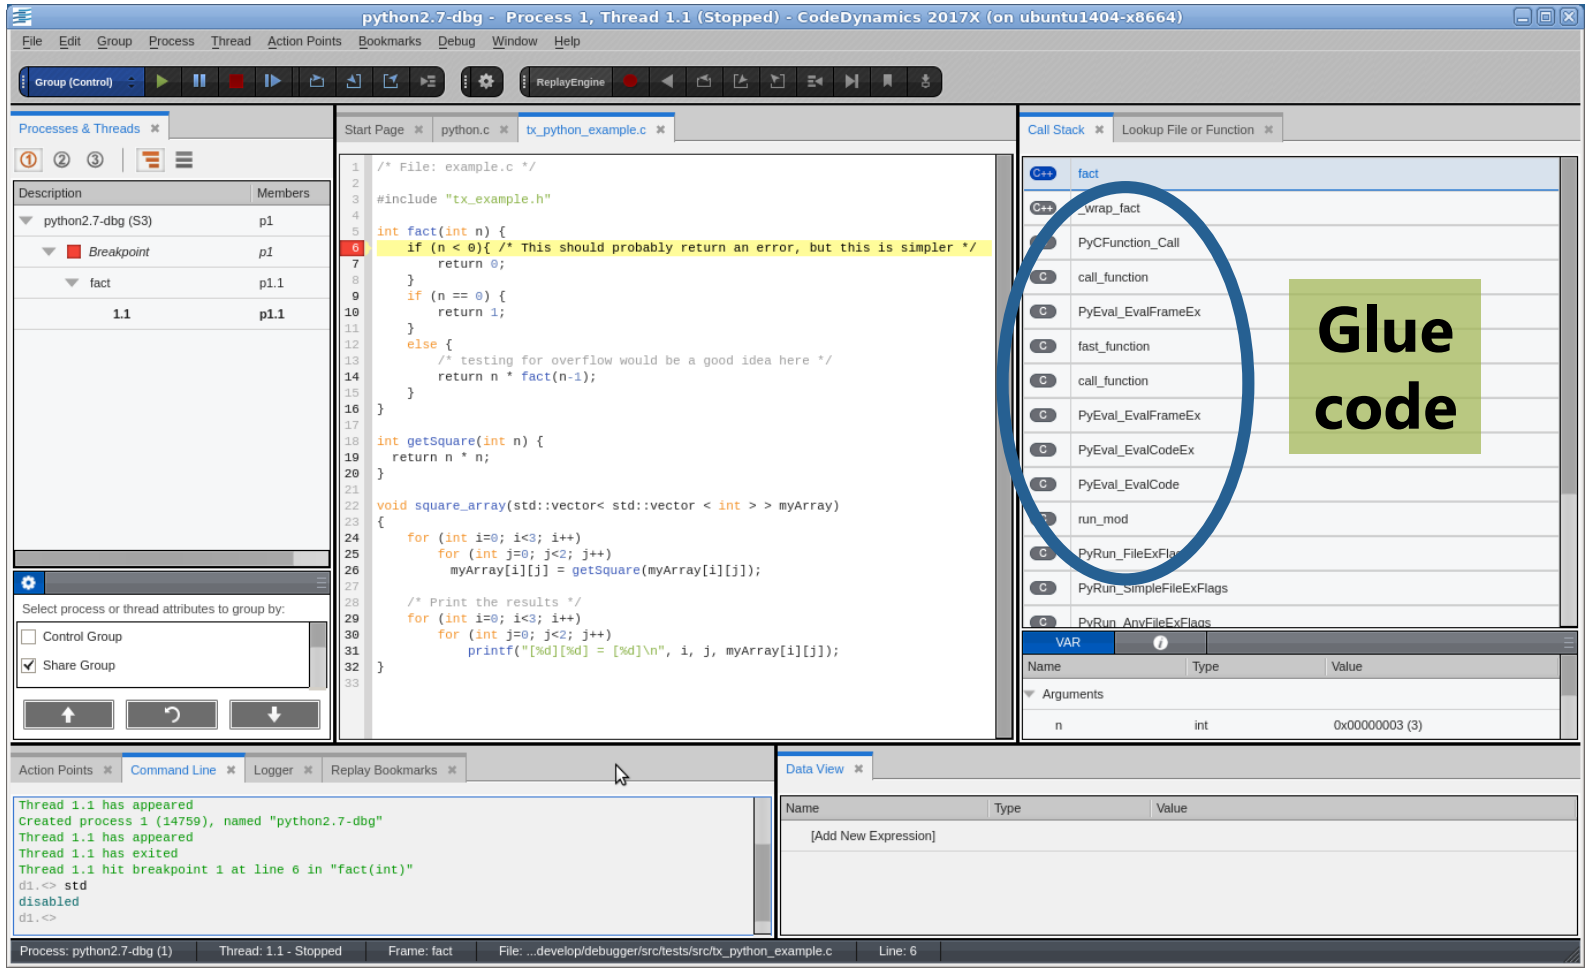

# Showing C code with mixed data

- Glue code filtered out
- Python data and code available for viewing

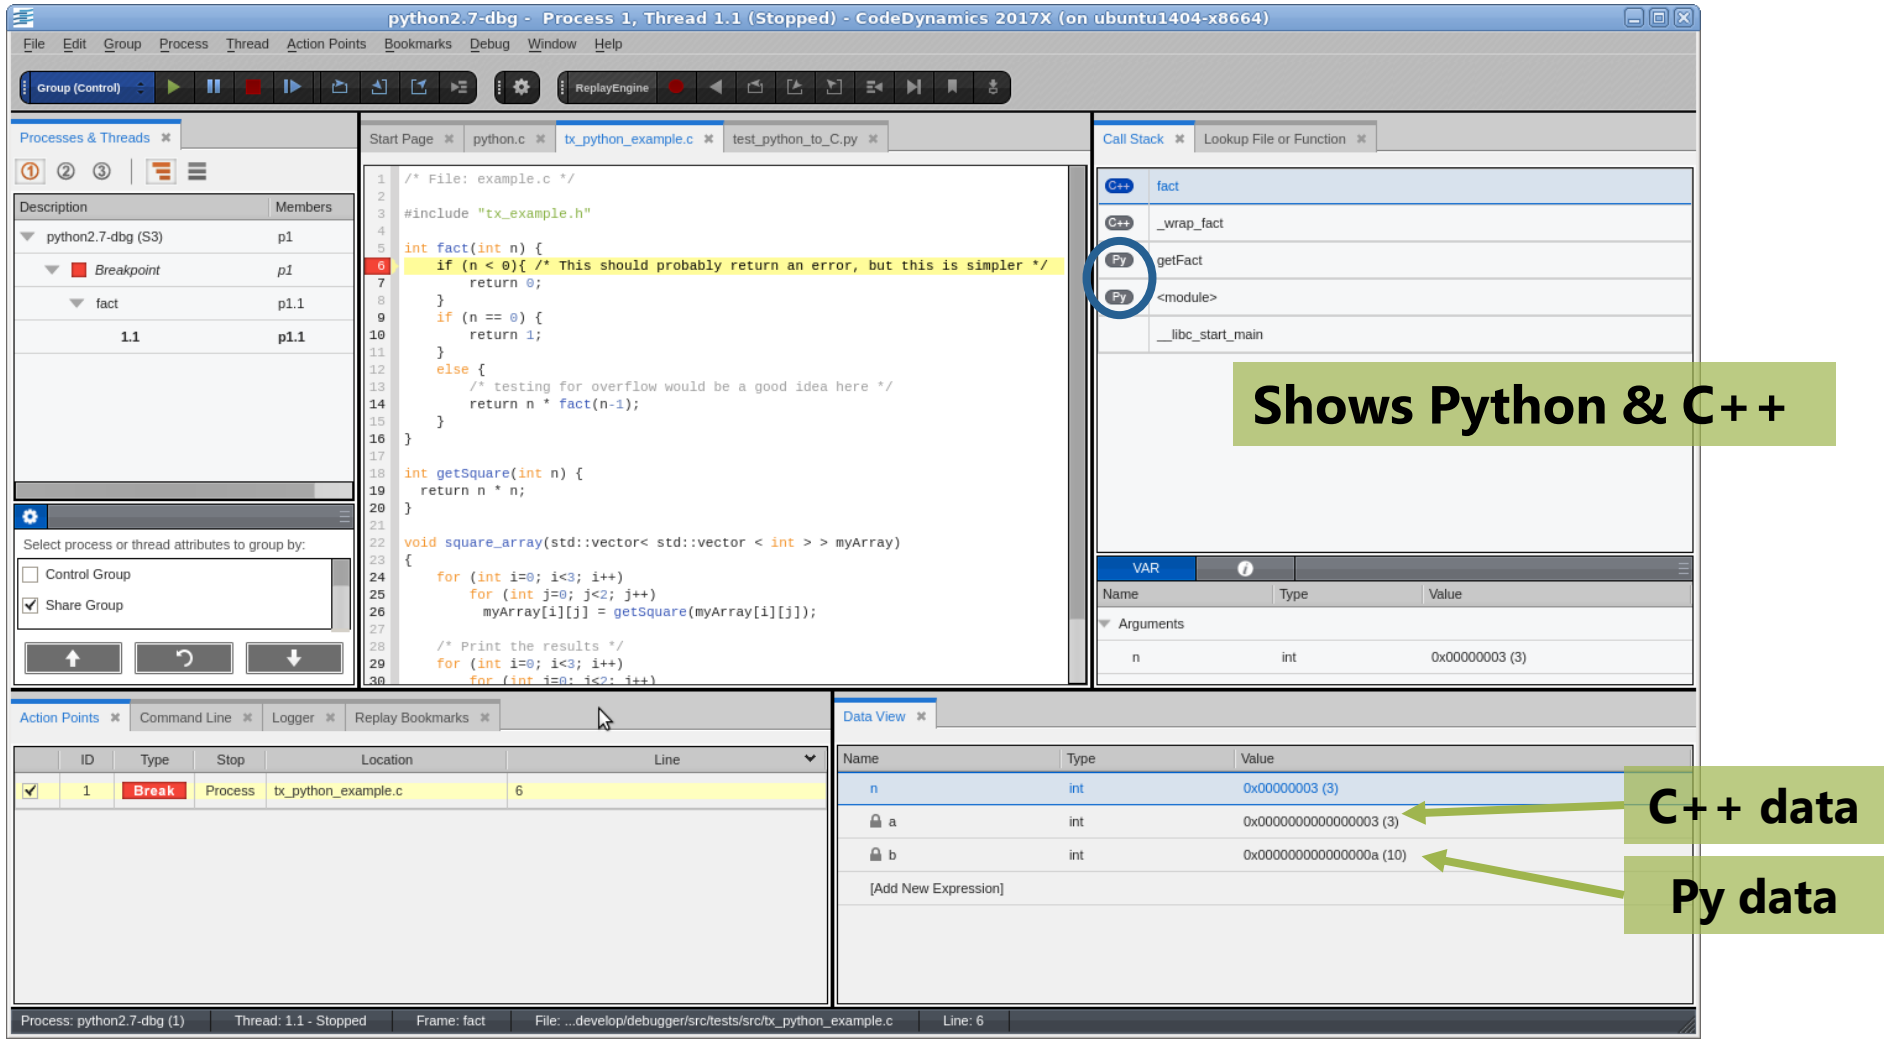

# Remote Display Debugging

# Remote Display Client (RDC)

- Offers users the ability to easily set up and operate a TotalView debug session that is running on another system
- Consists of two components
  - Client – runs on local machine
  - Server – runs on any system supported by TotalView and “invisibly” manages the secure connection between host and client
- Free to install on as many clients as needed
- Remote Display Client is available for:
  - Linux x86, x86-64
  - Windows
  - Mac OS X

# Remote Display Client

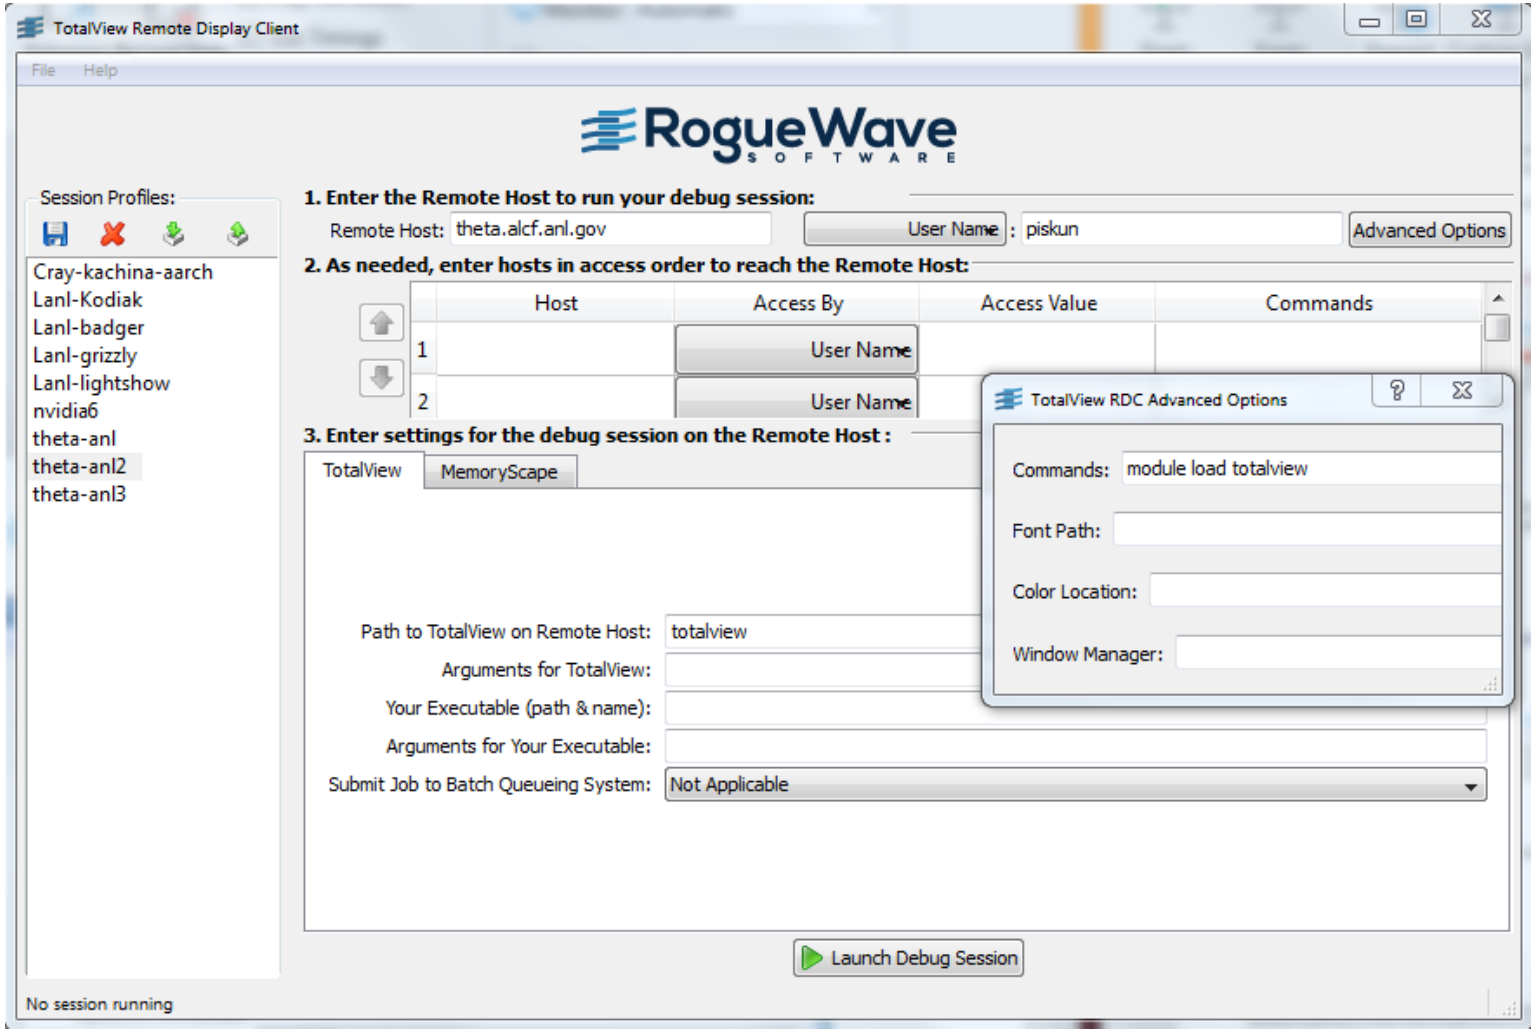

# Remote Display Client (Argonne NL)

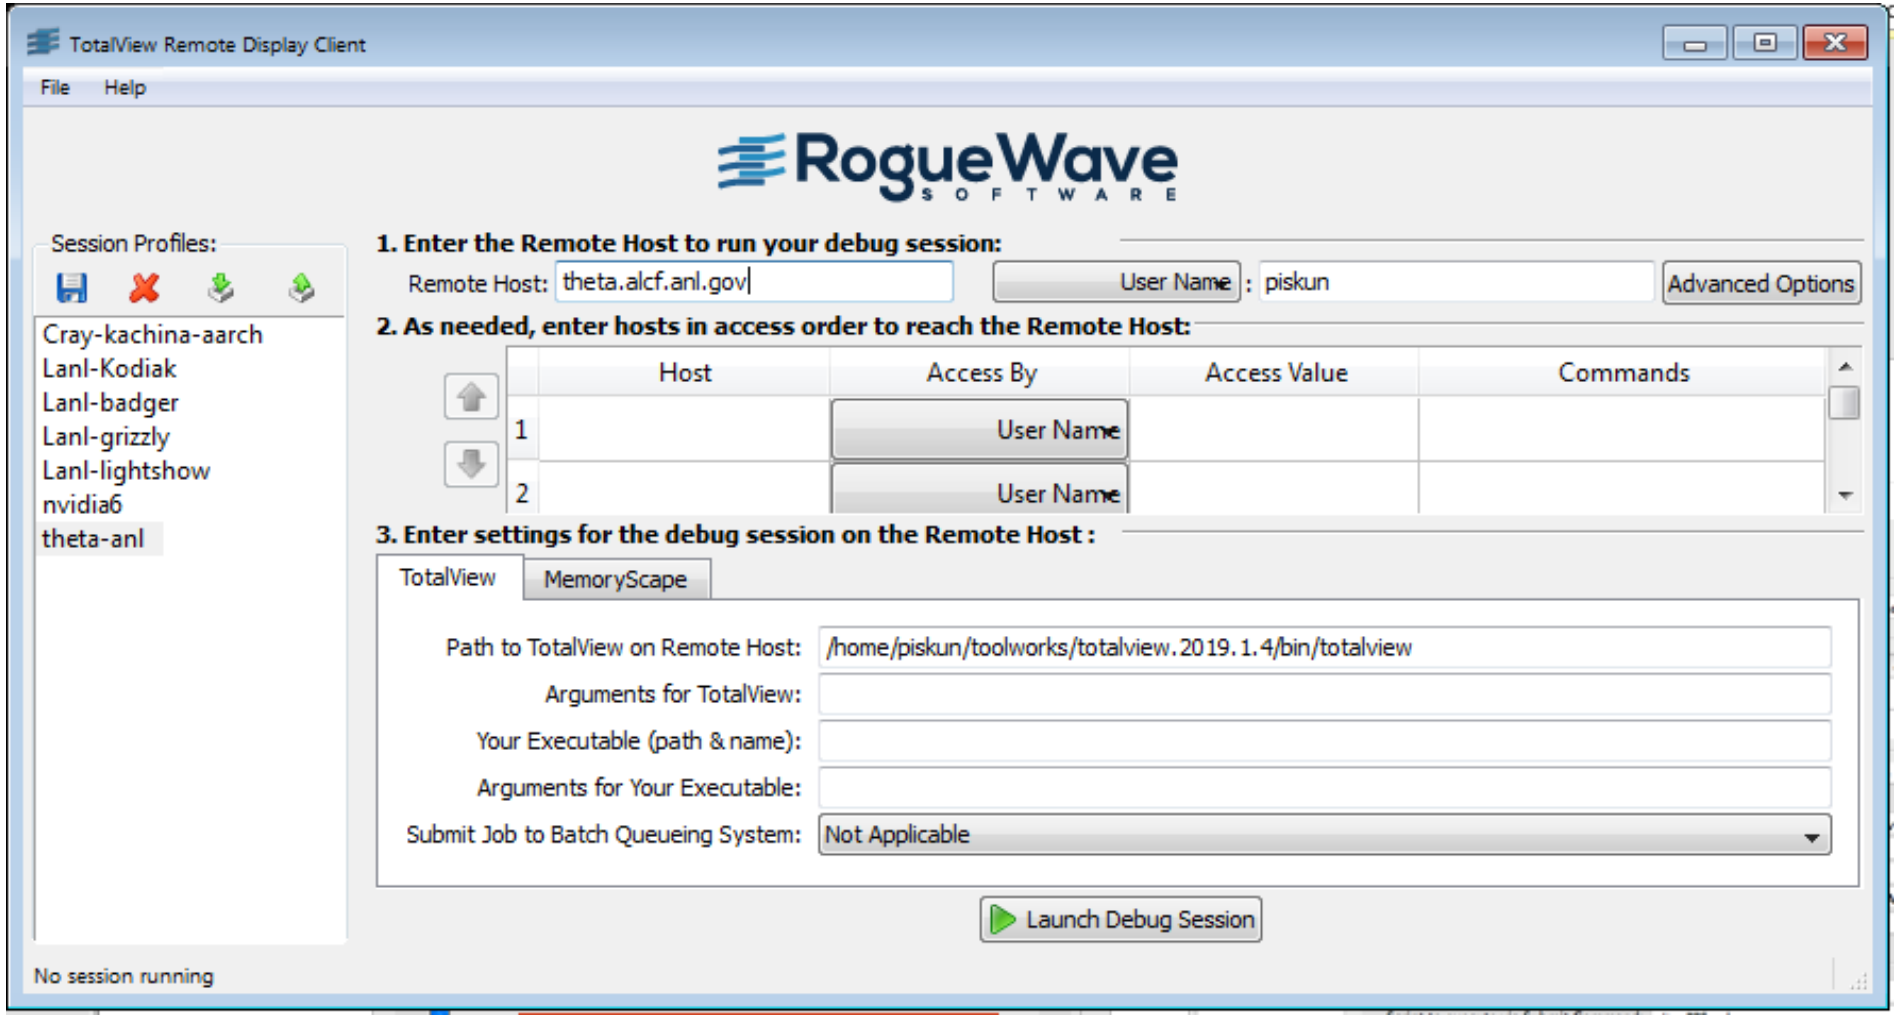

# Summary

- Use of modern debugger saves you time.

- TotalView can help you because:
  - It's **cross-platform** (the only debugger you ever need)
  - Allow you to debug accelerators (GPU) and CPU in **one session**
  - Allow you to debug **multiple languages** (C++/Python/Fortran)

# Using TotalView for Parallel Debugging on ANL

# Starting a MPI job – method 1

For HPC we have two methods to start the debugger

The ‘classic’ method

- **`totalview -args mpiexec -np 512 ./myMPIprog myarg1 myarg2`**
- This will start up TotalView on the parallel starter (mpiexec, srun, runjob, etc) and when you hit ‘Go’ the job will start up and the processes will be automatically attached. At that point you will see your source and can set breakpoints.

- Some points to consider…
  - You don’t see your source at first, since we’re ‘debugging’ the mpi starter
  - Some MPI’s don’t support the process acquistion method (most do, but might be stripped of symbols we need when packaging)
  - In general more scalable than the next method...

# Starting a MPI job – method 2

The ‘indirect’ method

- Simply ‘totalview’ or ‘totalview myMPIprog’ and then you can choose a parallel system, number of tasks, nodes, and arguments to the program.

- With this method the program source is available immediately
- Less dependent on MPI starter symbols
- May not be as scalable as some ‘indirect’ methods launch a debug server per process

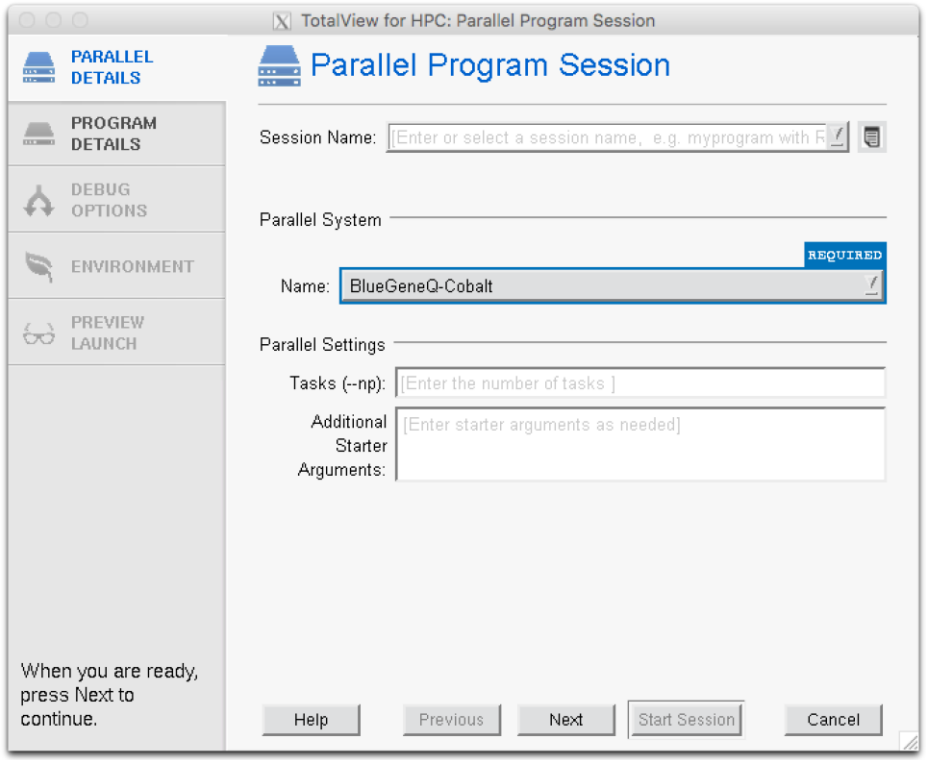

# Using TotalView at Argonne

- TotalView available on Theta, Vesta, Mira, Cooley
  - Installed at:
    /soft/debuggers/totalview-2019-08-01/toolworks/totalview.2019T.2.7/bin/totalview
  - module load totalview
- Download and install RDC from https://www.roguewave.com/products-services/features/remote-display-client
- Connect to Theta
- Get allocation first
  - On Theta : qsub -A ATPESC2019 –n <N> –q debug-flat-quad –I
  - Module load totalview
  - totalview -args aprun –np <N> ….

# TotalView Resources and Documentation

# TotalView Resources & Documentation

- TotalView documentation:
  - https://support.roguewave.com/documentation/tvdocs/en/current/
  - User Guides:  Debugging, Memory Debugging and Reverse Debugging
  - Reference Guides:  Using the CLI, Transformations, Running TotalView
- TotalView online HTML doc:
  - http://docs.roguewave.com/totalview/current/html/index.html
- Other Resources (Blogs, videos, white papers, etc):
  - https://www.roguewave.com/resources?tagid=18
- New UI resources:
  - Reference CodeDynamics Help
    https://www.roguewave.com/help-support/documentation/codedynamics
- New UI videos:
  - https://www.roguewave.com/products-services/codedynamics/videos
- Python Debugging blog:
  - http://blog.klocwork.com/dynamic-analysis/the-challenge-debugging-python-and-cc-applications/

# Questions/Comments

- Any questions or comments?
  - Don’t hesitate to reach out to me directly with any problems or suggestions!
  - **Email:** nikolay.piskun@roguewave.com

- **Thank you for your time today!**