

RED ESPAÑOLA DE
SUPERCOMPUTACIÓN



**Barcelona
Supercomputing
Center**
Centro Nacional de Supercomputación

Parallel Debugging with TotalView

-
BSC-CNS

AGENDA



- What debugging means?
- Debugging Tools in the RES
- Allinea DDT as alternative
- Introduction to TotalView (RogueWave Software)
 - ✓ What is TotalView
 - ✓ Compiling Your Program
 - ✓ Starting totalview
 - ✓ TotalView's Basic Look and Feel
 - ✓ Basic usage: Action points, groups, navigating the code...
 - ✓ Built-in variables and statements
 - ✓ Expression Evaluation and Code Fragments
 - ✓ Memory Debugging
 - ✓ Remote Display
 - ✓ Some notes on CLI
- Hands on

What debugging means?

Debugging 1.0: is a methodical process of finding and reducing the number of bugs.

Originally It literally meant the process to eliminate “bugs” ... like this:



What debugging means?



Debugging 2.0: Is the process to confirm all the things that you believe are true because there is, at least one, that is not

Things that you believe:

- This variable has been set before entering the loop
- This variable is only written by master process
- I am sending the right data type in all MPI communications



Serial Debugging:

- **Printf()**
- gdb and its frontends (DDD)
- Others ...

Parallel Debugging:

Both serial and parallel debuggers are useful.

Serial debuggers, like gdb, are what most programmers are used to, but parallel debuggers can attach to all the individual processes in an MPI job simultaneously, treating the MPI application as a single entity.



Linux Power TotalView 8.7.0-7



(c) Allinea Software 2002-2009

Version: 2.4.1
Build: Suse 10 ppc64
Build Date: Mar 20 2009

Debugging Tools



Session Control Search View Help

Current Group: User Defi | Focus on current: Group Process Thread Step Threads Together

All: 0 1 2 3 4 5 6 7 8 9 10 11 12 13 14 15

User Defined 1: 0 1 2 3 4 5

Create Group

Project Files: hello.c

```

80
81 for(x=0;x<12;x++)
82   for(y=0;y<12;y++)
83     tables[x][y] = (x+1)*(y+1);
84 MPI_Init(&argc, &argv);
85 MPI_Comm_rank(MPI_COMM_WORLD, &my_rank);
86 MPI_Comm_size(MPI_COMM_WORLD, &p);
87
88
89 alloc(sizeof(int)*100000);
90 for(x=0;x<10000;x++)
91 {
92   dynamicArray[x] = x*10;
93 }

```

Project Files

- Source Tree
 - home
 - usr
 - Header Files
 - Source Files
 - hello.c
 - func1(): void
 - func2(): int

File Edit View Group Process Thread Action Point Instrumentation Tools Window Help

Group (Control) | Go Halt Kill Restart Next Step Out Run To Prev UnStep Caller Back To Live

Process 1 (5274): combined (Stopped)

Thread 1 (5274) (Stopped) <Trace Trap>

Stack Trace

Circle::area,	FP=bf8d9fc8
Circle::not-in-charge Circle, FP=bf8d9fc8	
Cylinder::Cylinder,	FP=bf8da018
arrays,	FP=bf8da148
main,	FP=bf8da198
__libc_start_main,	FP=bf8delf8

Stack Frame

Function "Circle::area":
 this: 0xbf8da0a0 -> {s
 Block "\$b1":
 result: 1.52482386217018

Registers for the frame:
 %eax: 0xbf8da0a0 (-108123)
 %ecx: 0xbf8da0b4 (-108123)

Function Circle::area in combined.cpp

```

424 }
425
426 // Your basic circle area function
427
428 double Circle::area() {
429   double result;
430   result = PI * m_radius * m_radius; // Our old friend, pi r squared
431   return result;
432 }
433
434 // A Simple 3-D figure - class exercise: Do a cone figure in the
435 // fashion. - I forget how to calculate the surface area of a cone

```

Action Points | Processes | Threads

STOP	Process	Thread
STOP	9	combined.cpp#154 derived_class+0x09
STOP	2	combined.cpp#255 printArray<int>+0x49...
STOP	3	combined.cpp#357 user_templates+0x3e
STOP	6	combined.cpp#530 arrays+0x387...
STOP	6	combined.cpp#592 combine_waves_worker+0x17
STOP	5	combined.cpp#608 parallel_combine_waves+0x24
STOP	4	combined.cpp#670 nthreads_loop+0x81

DDT as alternative



DDT (Distributed Debugging Tool) from the Allinea Corporation

- Parallel debugger which provides many of the same basic features as Totalview, as well as some new elements.
- Totalview has a much larger feature set than DDT:
 - with debugging capability for more than one executable at a time
 - machine level language support
 - Tcl command line options
 - other advanced components ...

But..

These sometimes are not the primary reasons why scientists need a parallel debugger.

Why? ... Because ...

Most scientists need a simple and user friendly way to set breakpoints, step through code and halt execution while they examine variable values and code logic across different processors.



Some Features:

- DDT has an more intuitive user interface, especially for beginners
- Different ways of navigating through the processes
- None of both interfaces are suitable to debug apps with hundreds of processors (might become cumbersome)
- Both allow to dive into distributed multidimensionals arrays, subarrays, slices.
- And many more ...

<http://www.allinea.com/index.php?page=48>

<http://www.totalviewtech.com/support/documentation.html>



What is TotalView?

- TotalView is a sophisticated software debugger product of RogueWave Software
- Used for debugging and analyzing both serial and parallel programs.
 - Multi-threaded Debugging
 - Parallel Debugging: MPI, PVM, Others
- Especially designed for use with complex, multi-process and/or multi-threaded applications.
- Wide compiler & platform support
 - C, C++, Fortran 77 & 90, UPC
 - Unix, Linux, OS X
- Reverse debugging (Replay Engine)
- Long distance remote debugging
- Unattended Batch Debugging
- TVD along with DDT are the most popular HPC debuggers to date.



Compiling your program

Always compile with `-g` and `-O0`

- `O0` is important because with some optimizations, even when they not modify the code semantics, the source code may not reflect what is really happening.

- In the IBM compilers, some optimizations levels might alter the code semantics. That's why it is important to use `-qstrict` when using `-O3`

- TVD can debug code compiled without `-g` but assembler will be shown



Starting Totalview

TVD must be sent through the batch system

→ Connect to MareNostrum using -X option:

```
ssh -X bsc99704@mn4.bsc.es
```

→ Jump to a node above login4 (from login5 to login8)

```
ssh -Y login6
```

→ Submit the batch script:

```
mnsuubmit run.sh
```

```
1 #!/bin/bash
2 # @ job_name = simple2
3 # @ initialdir = ./
4 # @ output = mpi_%j.err
5 # @ error = mpi_%j.err
6 # @ total_tasks = 4
7 # @ cpus_per_task = 4
8 # @ wall_clock_limit = 01:10:00
9 # @ mining_level = 0
10# @ x11 = 1
11
12
13 /gpfs/apps/TOTALVIEW/totalview -mpi SLURM -np 4 ./simple -a "Prueba de TVD"
```

Introduction to TotalView



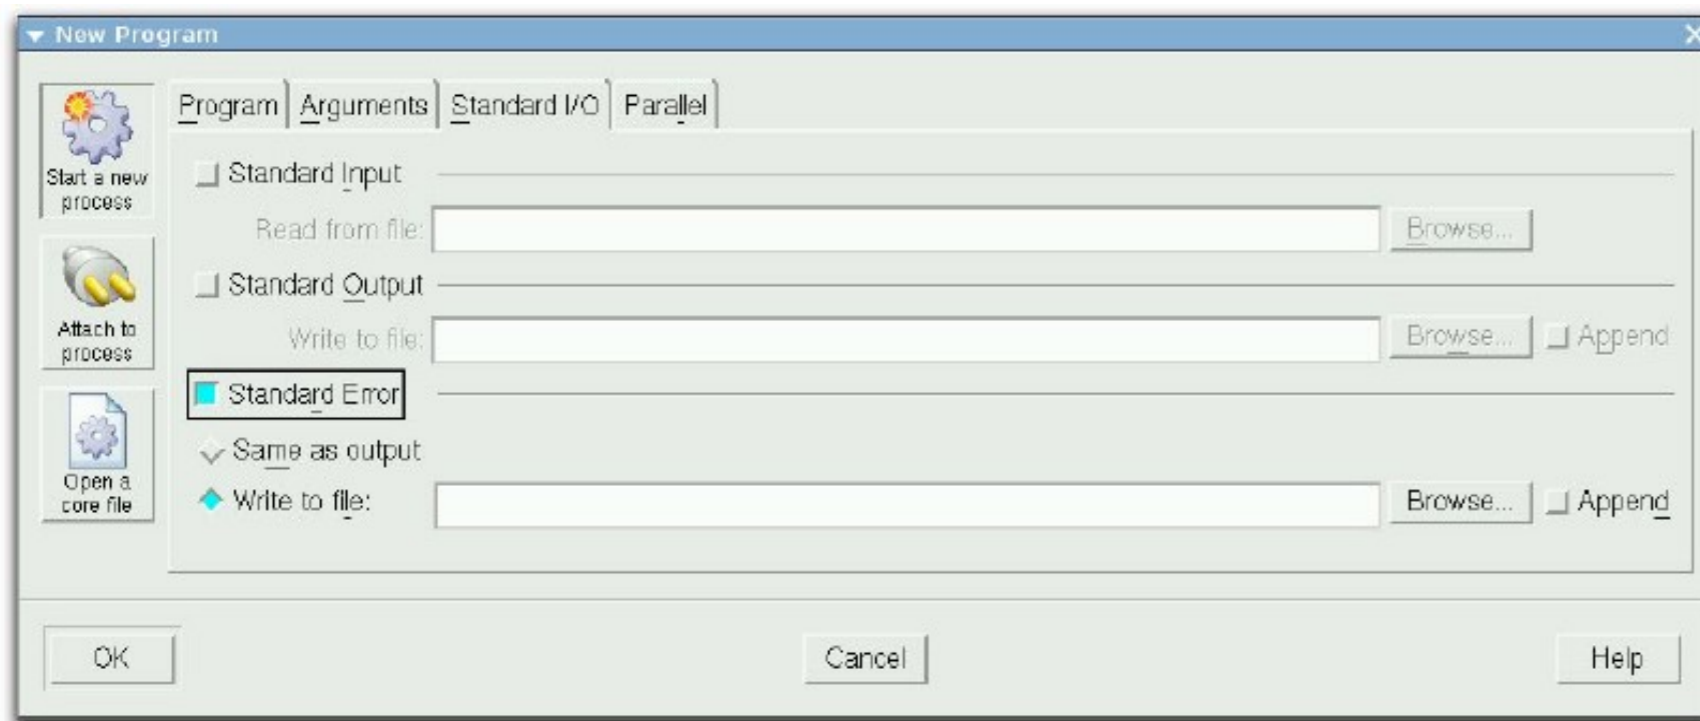
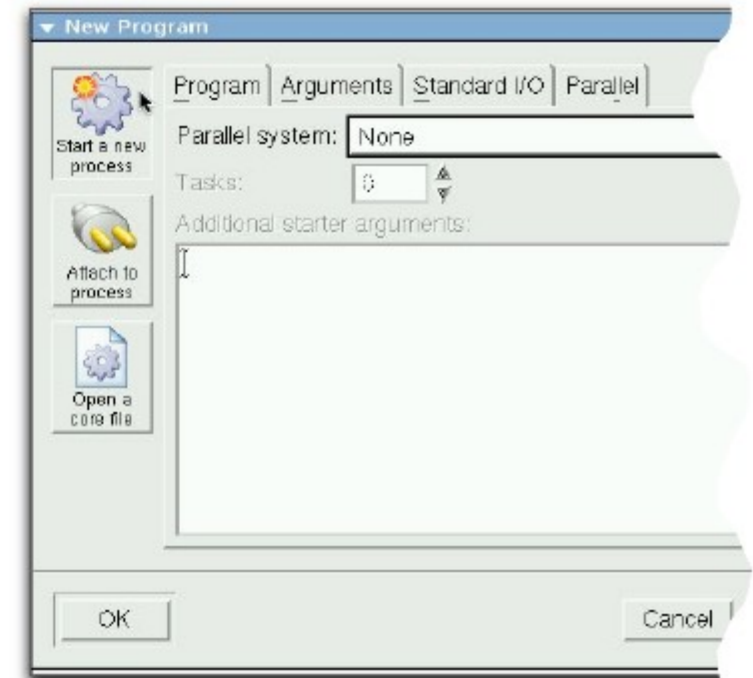
Starting Totalview

The New Program Screen lets you:

Start a New Process

Attach to an Existing Process

Open a Core File



Introduction to TotalView



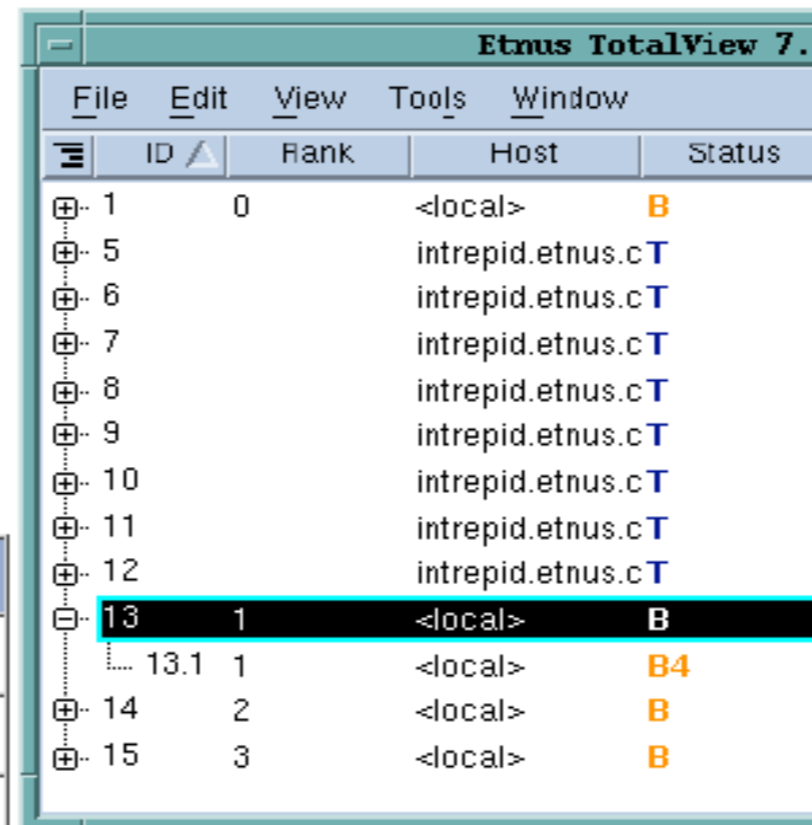
TotalView's Basic Look and Feel

TVD contains two kinds of windows:

- Root Window (Control)

- * States of processes
- * Processes and thread Status
- * Instant navigation access.

State Code	Description
B	Stopped at a breakpoint
E	Stopped because of an error
H	In a Hold state
K	Thread is executing within the kernel
M	Mixed - some threads in a process are running and some not
R	Running
T	Thread is stopped
W	At a watchpoint



► Status Info

- T = stopped
- B = Breakpoint
- E = Error
- W = Watchpoint
- R = Running
- M = Mixed
- H = Held

Introduction to TotalView



TotalView's Basic Look and Feel

TVD contains two kinds of windows:

- *Process Window:*
Provides detailed view of one process

Toolbar

Stack Trace Pane

Stack Frame Pane

Source Pane

Tabbed Area



Basic usage

Controlling execution

Command	Description
Go	Start/resume execution
Halt	Stop execution
Kill	Terminate the job
Next	Run to next source line or instruction. If the next line/instruction calls a function, the entire function will be executed and control will return to the next source line or instruction (the function is "stepped over").
Step	Run to next source line or instruction. If the next line/instruction calls a function, the function will be "stepped into". Execution will stop within the function.
Out	Execute to the completion of a function. Returns to the instruction after the one which called the function.
Run To	Allows you to arbitrarily click on any source line and then run to that point
Next Instruction	Similar to Next, but applies only to machine instructions
Step Instruction	Similar to Step, but applies only to machine instructions
Hold/Release	Hold ignores other commands to resume execution Release allows other run commands to have effect
Restart	Restarts a running program, or one that has stopped without exiting
Set PC	Sets the Program Counter to a desired source line, machine instruction, or absolute address

Introduction to TotalView



Basic usage

Controlling execution

Based on
PC location

```
2 int sub2(int);
3 int sub3(int);
4
5 int main ()
6 {
7     int j, k, i = 0;
8
9     j = sub1(i); k = sub3(j);
10    printf ("The value of k is %d\n", k);
11 }
12
13 int sub1(int x)
14 {
15     x = sub2(2);
16     return (x++);
17 }
18
19 int sub2(int y)
20 {
21     y = y + 10;
22     y++;
23     return (y+10);
24 }
25
26 int sub3(int z)
27 {
28     return (z*z);
```




Basic usage

“Diving” : In TVD this concept is widely used to refer the way user navigates through the application in a debugging session in order to:

- Obtain more information
- Refocus the process window
- Open variables
- ...

You can “dive” by:

- Double-clicking the left mouse button
- Selecting “Dive” in the context menu

You can dive on:

- Variables names to open a variable window (viewing data)
- Function names to open the source
- Process and threads in the root window to open a process window

Introduction to TotalView

Basic usage

“Diving” example over a variable in the common block in the stack frame

The image shows a sequence of four debugger windows illustrating the 'Dive' operation:

- window 1:** Stack Frame view showing variables: SKEWED_ARRAY, INT3_ARRAY, INT2_ARRAY, MASTER_ARRAY, LOCAL_VAR, COMP16_ARRAY_2, COMP_ARRAY_1, and Common blocks: ieee.
- window 2:** Variable inspection window for 'ieee' (Type: \$void, Address: 0x1406214a0). It shows a table with variables 'denoms' (integer(1000)) and 'ieee_array' (integer(6)).
- window 3:** Variable inspection window for 'denoms' (Type: integer(1000), Address: 0x1408214b8). It shows a list of values: (1) 0, (2) 0, (3) 0, (4) 0, (5) 0, (6) 0, (7) 0, (8) 0, (9) 0, (10) 0, (11) 0.
- window 4:** Variable inspection window for 'ieee_array' (Type: integer(6), Address: 0x1406214a0). It shows a list of values: (1) 0, (2) 0, (3) 0, (4) 0, (5) 0, (6) 0.

Arrows labeled 'Dive' indicate the progression from the stack frame to the variable, then to its elements, and finally to the array elements.



Basic usage

Viewing Data:

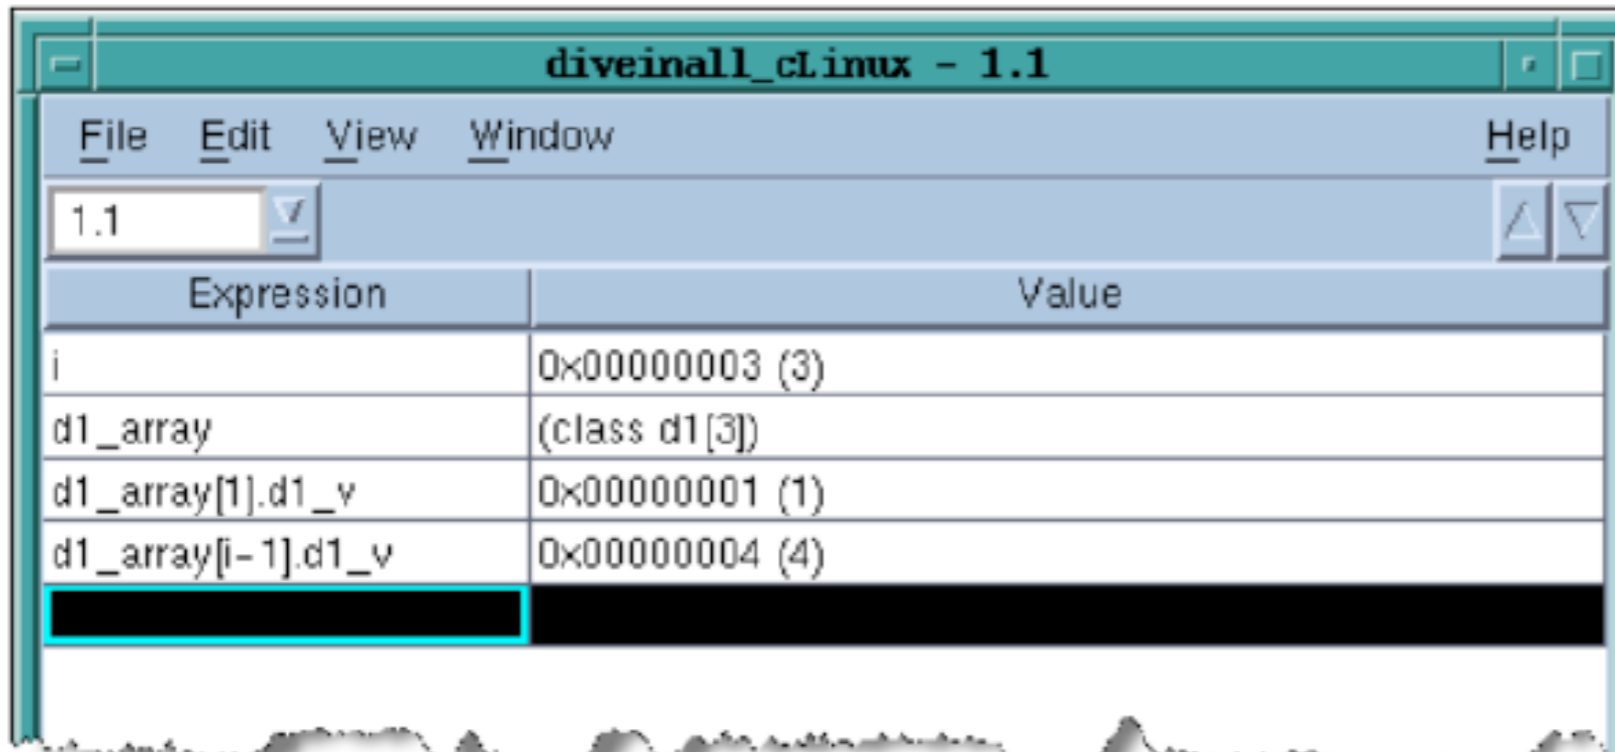
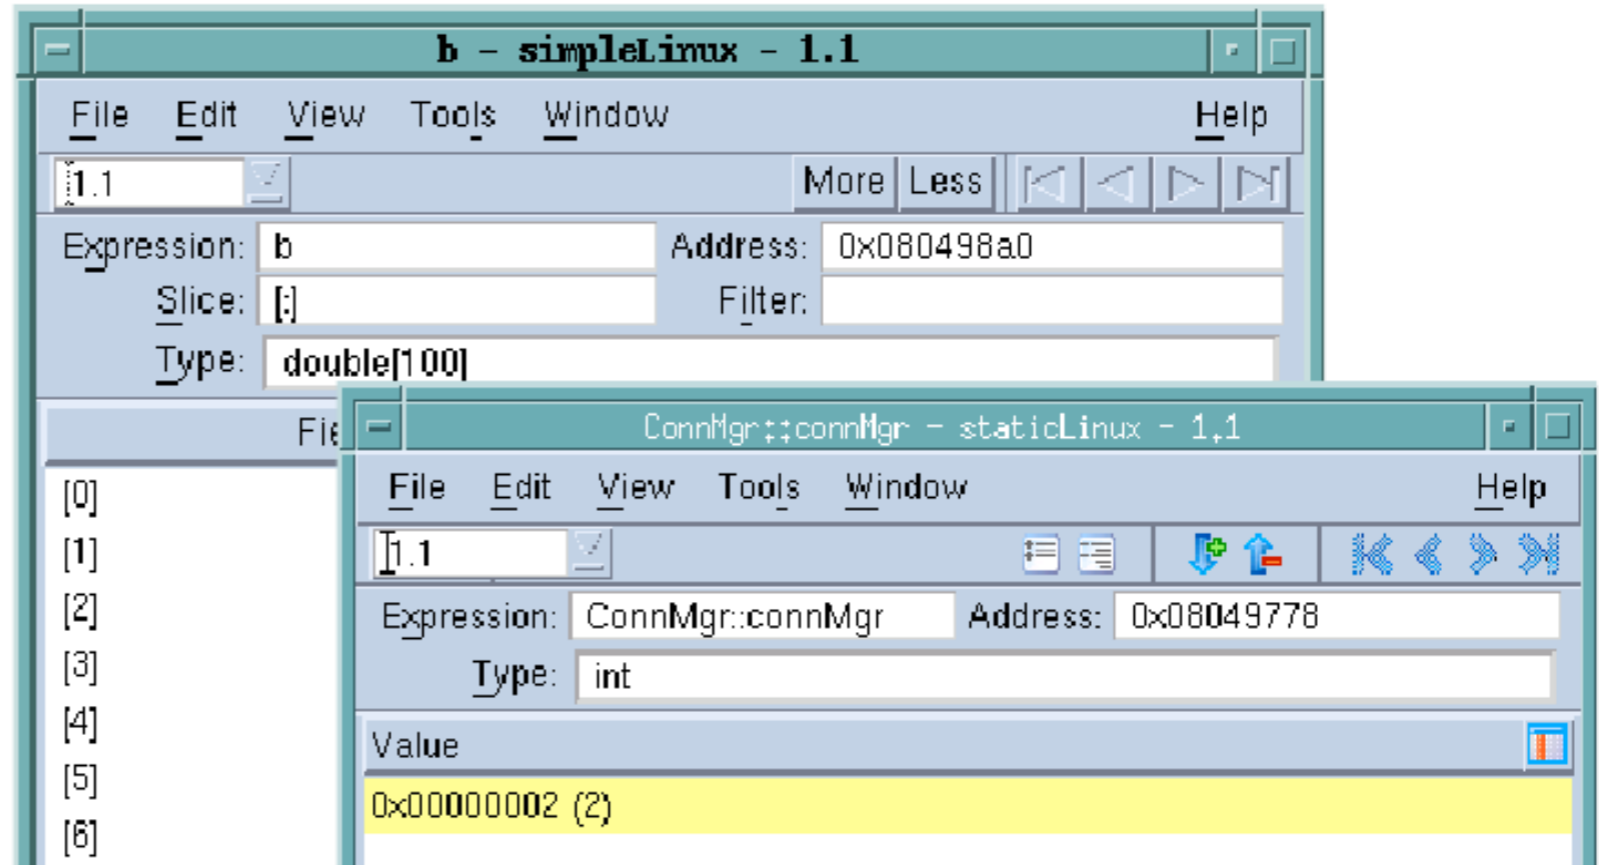
- Diving on a variable opens the Variable Window
 - Contents are updated automatically
 - Changed values are highlighted
 - “Last value” column
 - Clicking on the variable let the user to edit it:
 - Editing values changes the memory of the program
 - “Enter” to commit changes
 - “Esc” to cancel changes
- Using the Expression List Window
 - Variables can be added using right click on the variable
 - Adding expression directly in the window

Introduction to TotalView



Basic usage

Viewing Data:



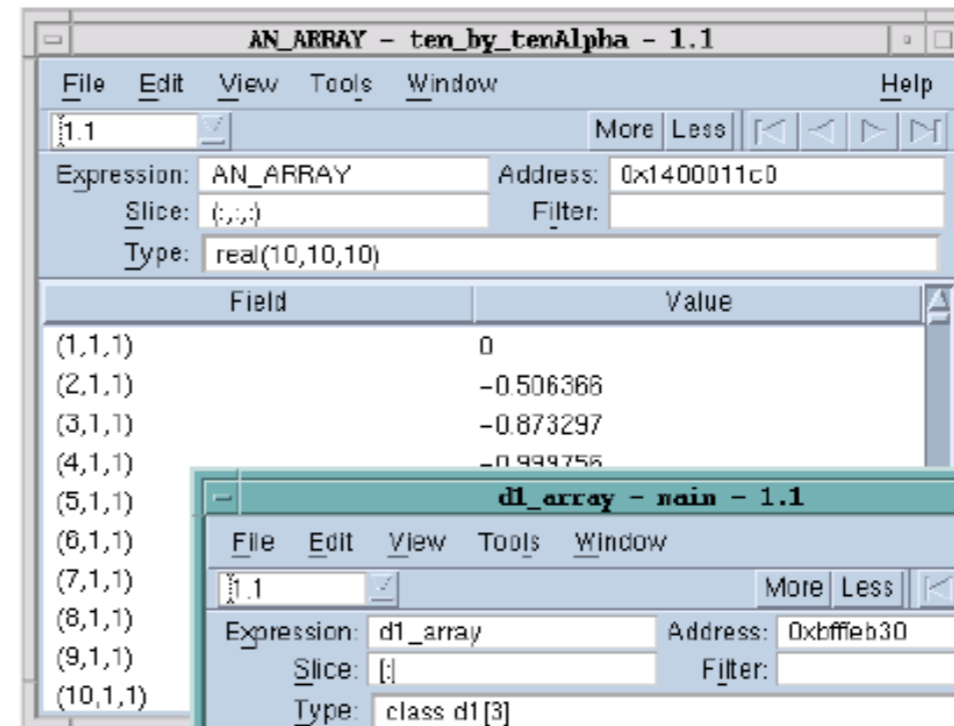


Basic usage

Viewing Data: Viewing arrays

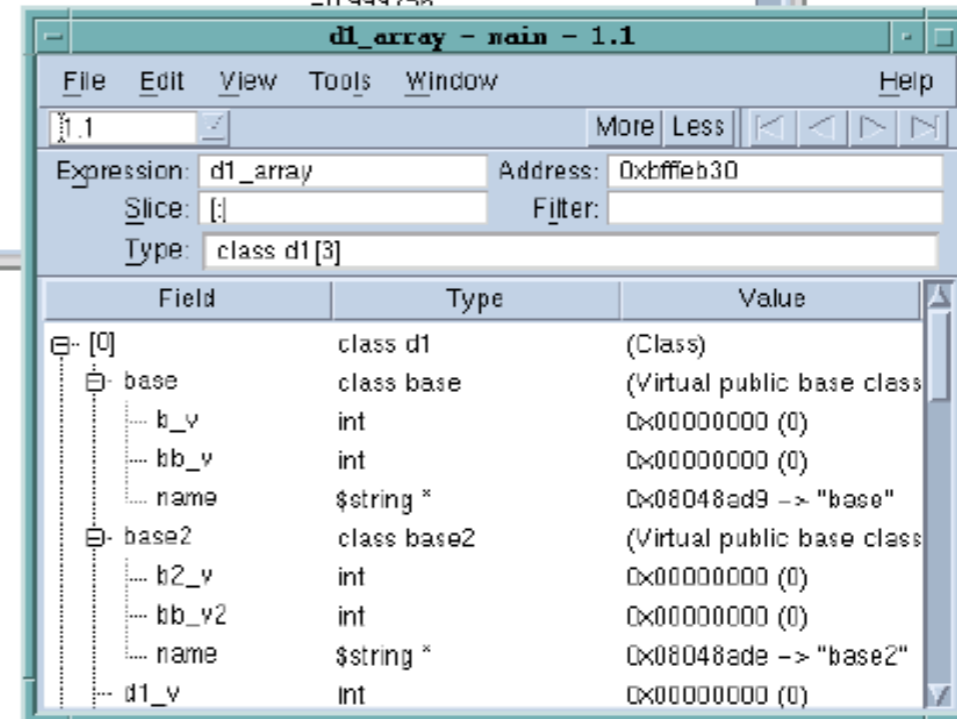
For array data, TotalView provides several additional features:

- Displaying array slices
- Data filtering
- Data Sorting
- Array statistics



Data Arrays

Structure Arrays



Introduction to TotalView



Basic usage

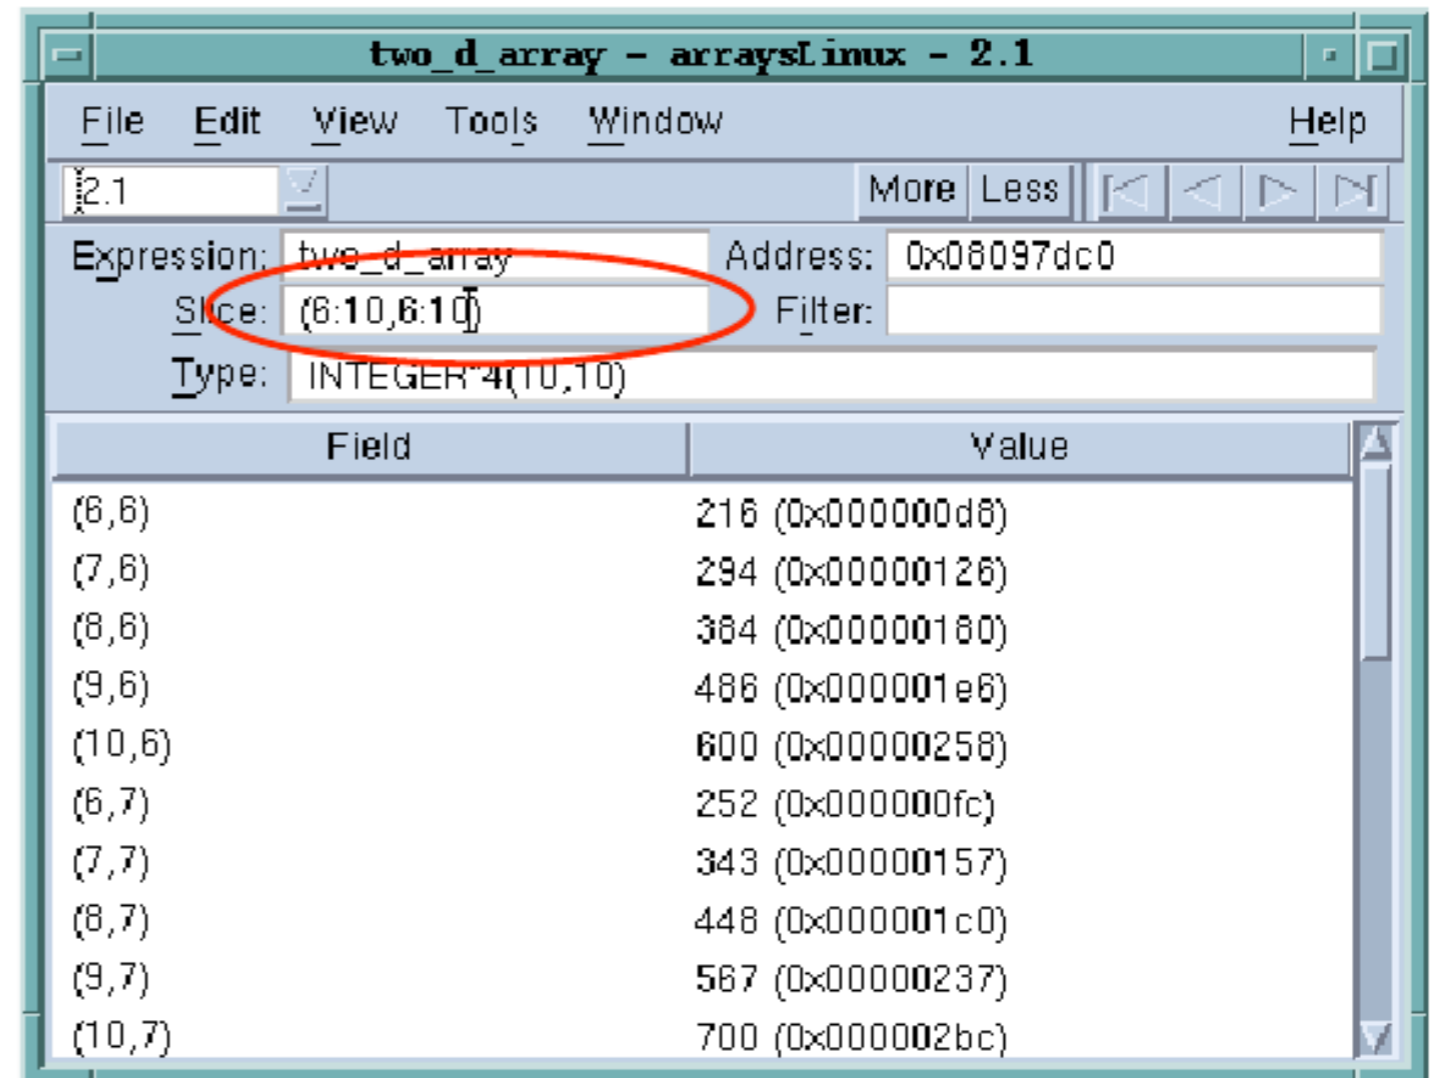
Viewing Data: Slicing arrays

Used to display subsections of an array.
Particularly useful if only a small section
of a large array is of interest.

lower_bound:upper_bound:stride

Fortran Slice: (1:5, 3:8)

C/C++ Slice: [::2][1:20]



Slice notation is [start:end:stride]



Basic usage

Viewing Data: Filtering arrays

Arrays containing data types of character, integer or floating point can be filtered to display only desired data.

Filtering can be:

- By arithmetic comparison
- For IEEE values
- By a range of values
- Within an expression

See the TotalView documentation for additional examples, syntax options and other important information.

Examples:

Fortran

Filter: `.gt. 250`

Filter: `.eq. $nan`

Filter: `7:512`

C/C++

Filter: `>= 100`

Filter: `!= $inf`

Filter: `128:<1024`



Basic usage

Viewing Data: Filtering arrays

The first screenshot shows a TotalView window for 'ieee_array' at address 0x1406214a0. The filter is set to '.eq. \$inf'. The data table shows two entries: (1) INF and (2) -INF.

Field	Value
(1)	INF
(2)	-INF

The second screenshot shows the same 'ieee_array' window. The filter is changed to '.eq. \$denorm'. The data table shows two entries: (5) 1.4013e-45 <denormalized> and (6) -1.4013e-45 <denormalized>.

Field	Value
(5)	1.4013e-45 <denormalized>
(6)	-1.4013e-45 <denormalized>

The third screenshot shows a TotalView window for 'int2_array__' at address 0xbfff0450. The filter is set to '\$value > 20 .and. \$value < 100'. The data table shows entries from (16) to (25) with values ranging from 22 to 40.

Field	Value
(16)	22 (0x0016)
(17)	24 (0x0018)
(18)	26 (0x001a)
(19)	28 (0x001c)
(20)	30 (0x001e)
(21)	32 (0x0020)
(22)	34 (0x0022)
(23)	36 (0x0024)
(24)	38 (0x0026)
(25)	40 (0x0028)

Introduction to TotalView



Basic usage

Sorting Array Data

Field	Value
(62,7)	4.4520000000000000e+04
(61,7)	4.3890000000000000e+04
(60,7)	4.3260000000000000e+04
(59,7)	4.2630000000000000e+04
(58,7)	4.2000000000000000e+04

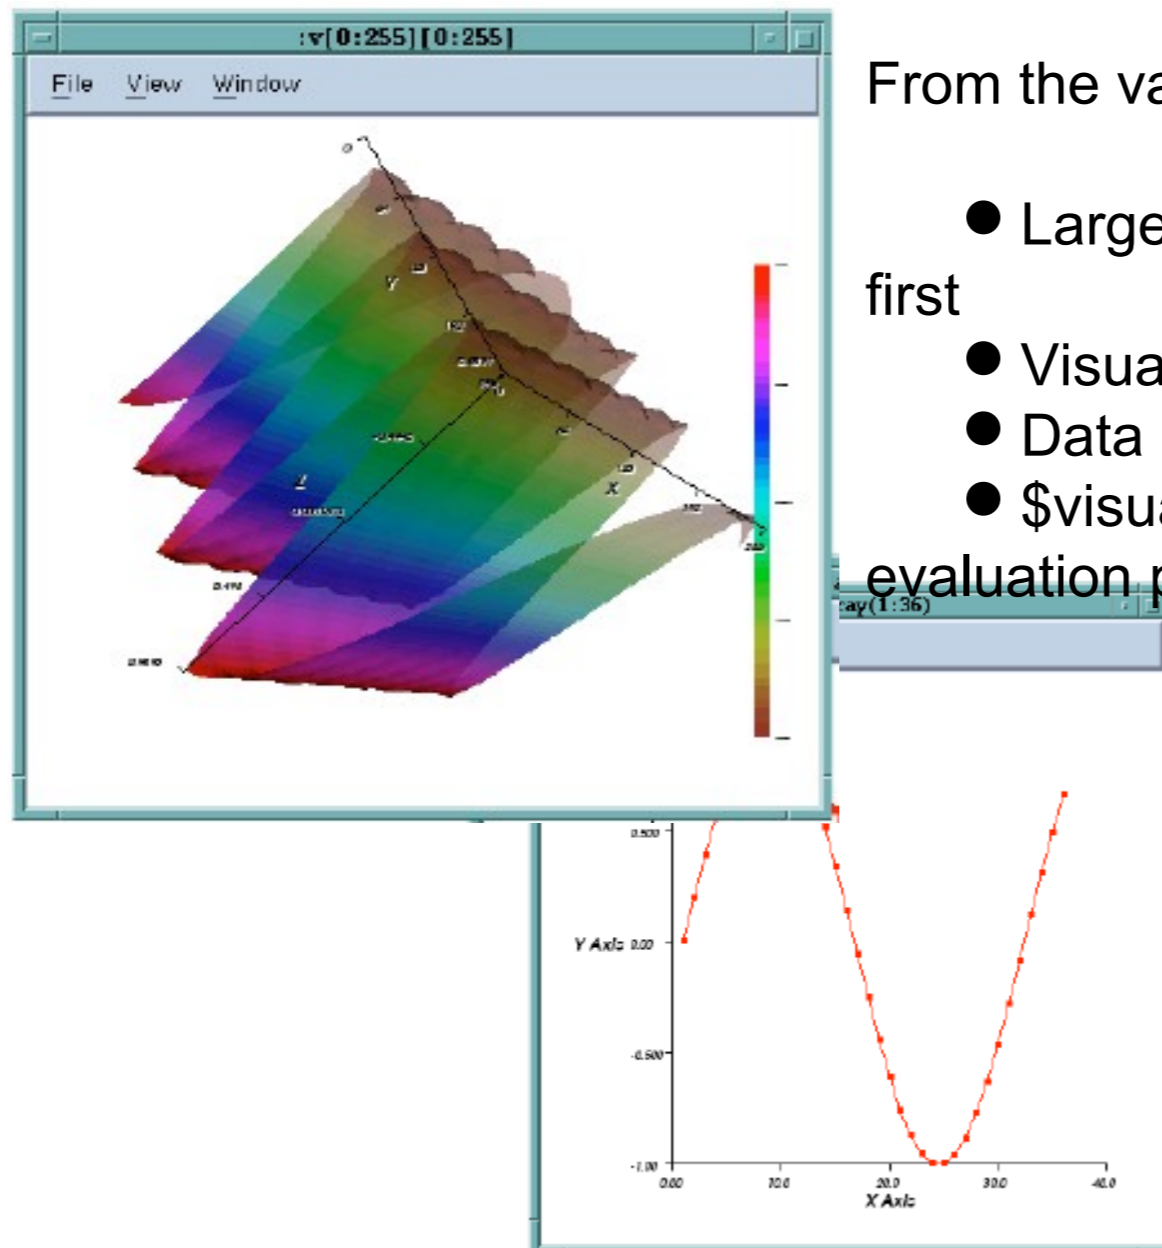
Clicking on the Value bar toggles the sort direction

Triangle icon shows sort direction



Basic usage

Visualizing Array Data



From the variable window click on Tools → Visualize

- Large arrays can be sliced down to a reasonable size first
- Visualize allows to spin, zoom, etc ...
- Data is not updated. You must revisualize
- `$visualize()` is a directive and can be used in evaluation points

Introduction to TotalView



Basic usage

Viewing Data: Laminating arrays (view across processes)

Totalview allows you to look at the values of a variable in all MPI processes

- Right Click on the variable
- Select the variable window → view across

You can Filter, visualize, explore distributed arrays....



Basic usage

Viewing Data: Viewing STL

TVD transform templates into understandable information:

–STLView supports `std::vector`, `std::list`, `std::map`, `std::string`

The screenshot shows a TotalView window titled 'x - main - 1.1'. The menu bar includes File, Edit, View, Tools, Window, and Help. The expression 'x' is entered, with address '0xbfffdba0'. The 'Actual Type' is 'float[3]' and the 'Type' is 'class vector<float, allocator<float> >'. Below this, a table displays the array elements:

Field	Value
[0]	1.3
[1]	2.2
[2]	3.1

The screenshot shows a TotalView window titled 'x - main - 1.1'. The expression 'x' is entered, with address '0xbffe1a0'. The 'Type' is 'class vector<float, allocator<float> >'. Below this, a tree view shows the internal structure of the vector:

- ⊖ _Vector_base class _Vector_base<float> (Private base class)
 - ⊖ _Vector_alloc_base class _Vector_alloc_base (Public base class)
 - ... _M_start float * 0x08052368 -> 1.3
 - ... _M_finish float * 0x08052374 -> 0**
 - ... _M_end_of_storage float * 0x08052378 -> 9.80909e-

Introduction to TotalView



Basic usage

Call graph: Allows a quick view of the program state

- Functions are nodes
 - Calls are edges
- Labels are MPI rank

The screenshot displays the TotalView IDE interface. The main window shows the source code for a C program named 'pingpong'. The code includes a function definition for 'pingpong' that uses MPI. A red arrow points to the line: `MPI_Comm_size(MPI_COMM_WORLD, &numprocs);`. The 'Stack Trace' window shows the current call stack, with 'pingpong' at the top. The 'Stack Frame' window shows the local variables for the 'pingpong' function, including 'myid', 'numprocs', 'left', 'right', 'buffer', 'buffer2', and 'request'. The 'Call Graph' window shows a hierarchical view of the program's execution flow, with nodes representing functions and edges representing calls. The nodes are labeled with MPI ranks, such as '(1.1, 7.1, 8.1...)' and '(1.2, 7.2, 8.2...)', indicating the rank of the process that called the function. The call graph shows the flow from the main function to the 'pingpong' function, and from the 'clone' function to the 'start_thread' function, which then calls 'mx__progress_thread', 'mx__wait', 'mx__ioctl', and finally '.ioctl'.



Basic usage

MPI Messages queue and graph:

- Provides information from the MPI layer
 - pending messages
 - unexpected messages

- Messages can be filtered by tags, MPI Communicators.

- Useful in deadlock situations and load balancing studies.

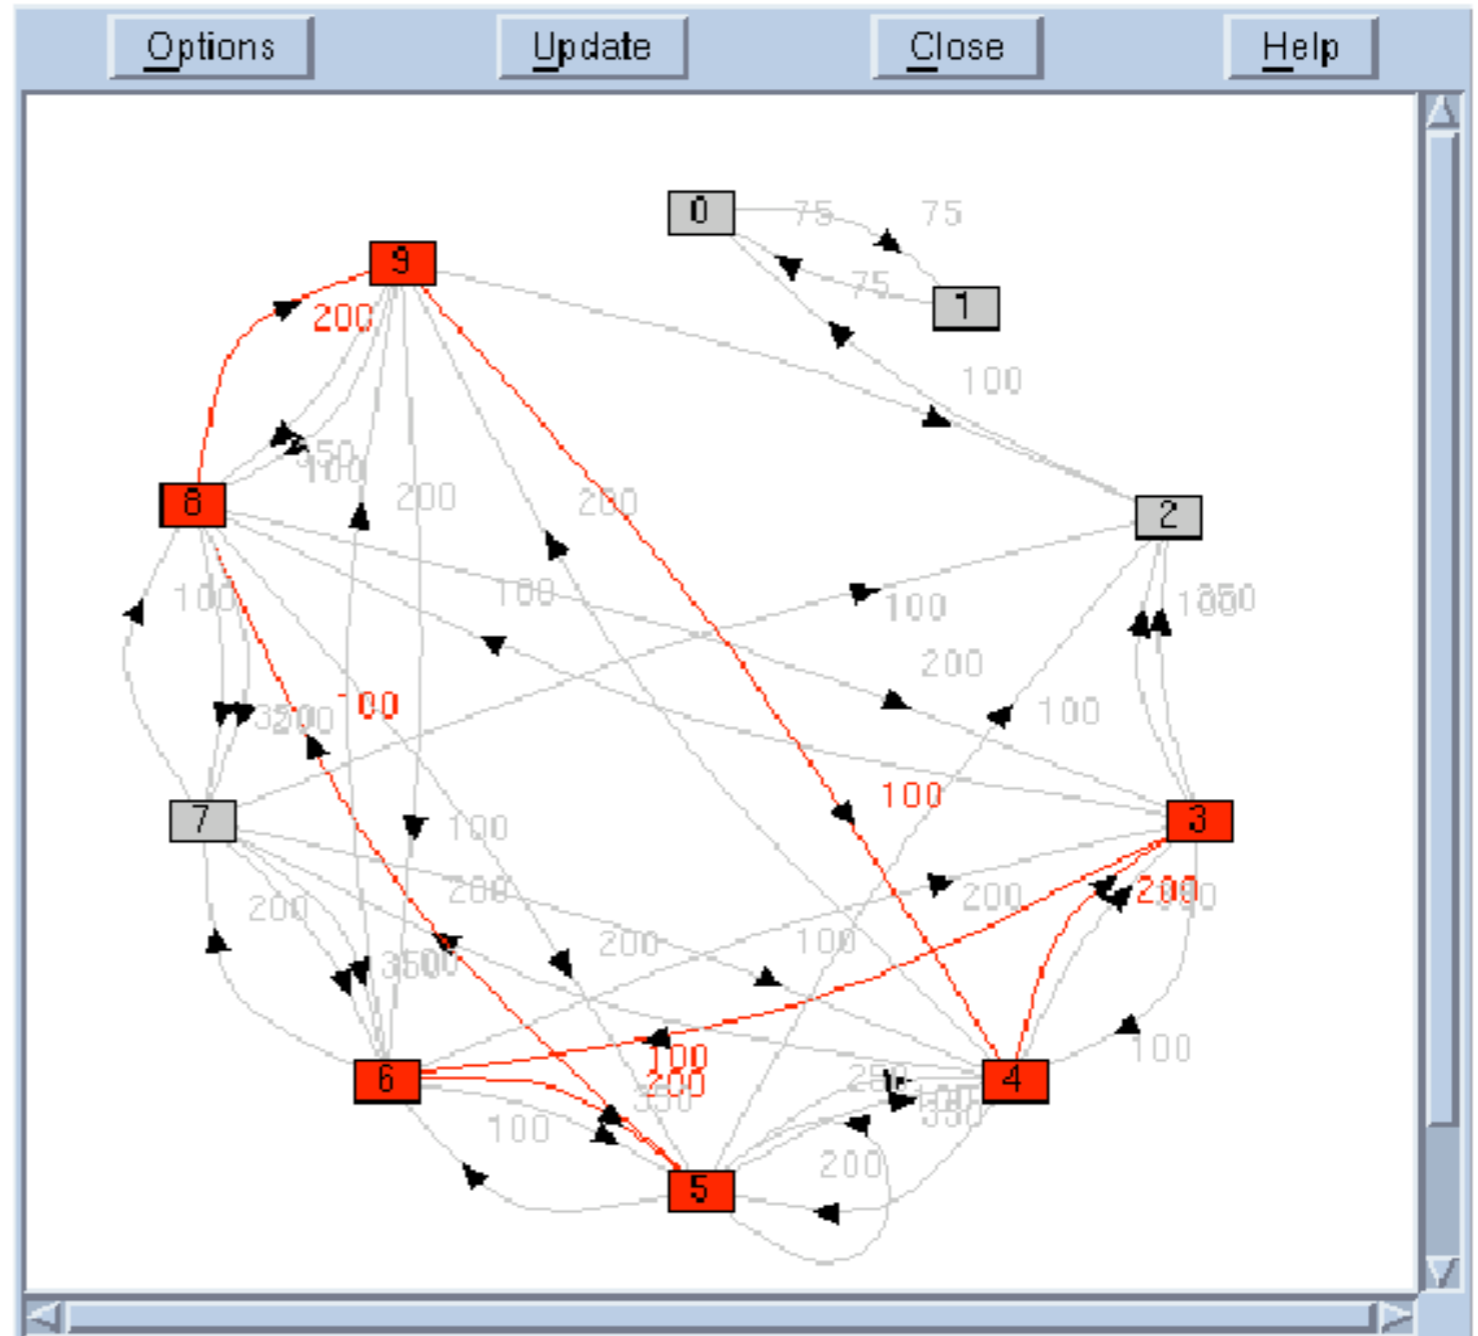
- May be to be enabled in the MPI Library
 - --enable debug

Introduction to TotalView



Basic usage

```
File Edit View Window Help
Message State - 1.1 "springs.0"
MPI_COMM_WORLD_collective
Comm_size 4
Comm_rank 0
Pending receives
[0]
  Status Pending
  Source 1 (springs.1)
  Tag 3 (0x00000003)
  User Buffer 0x0809d028 -> 0x00000000 (0)
  Buffer Length 100 (0x00000064)
Unexpected messages
[0]
  Status Pending
  Source 2 (springs.2)
```

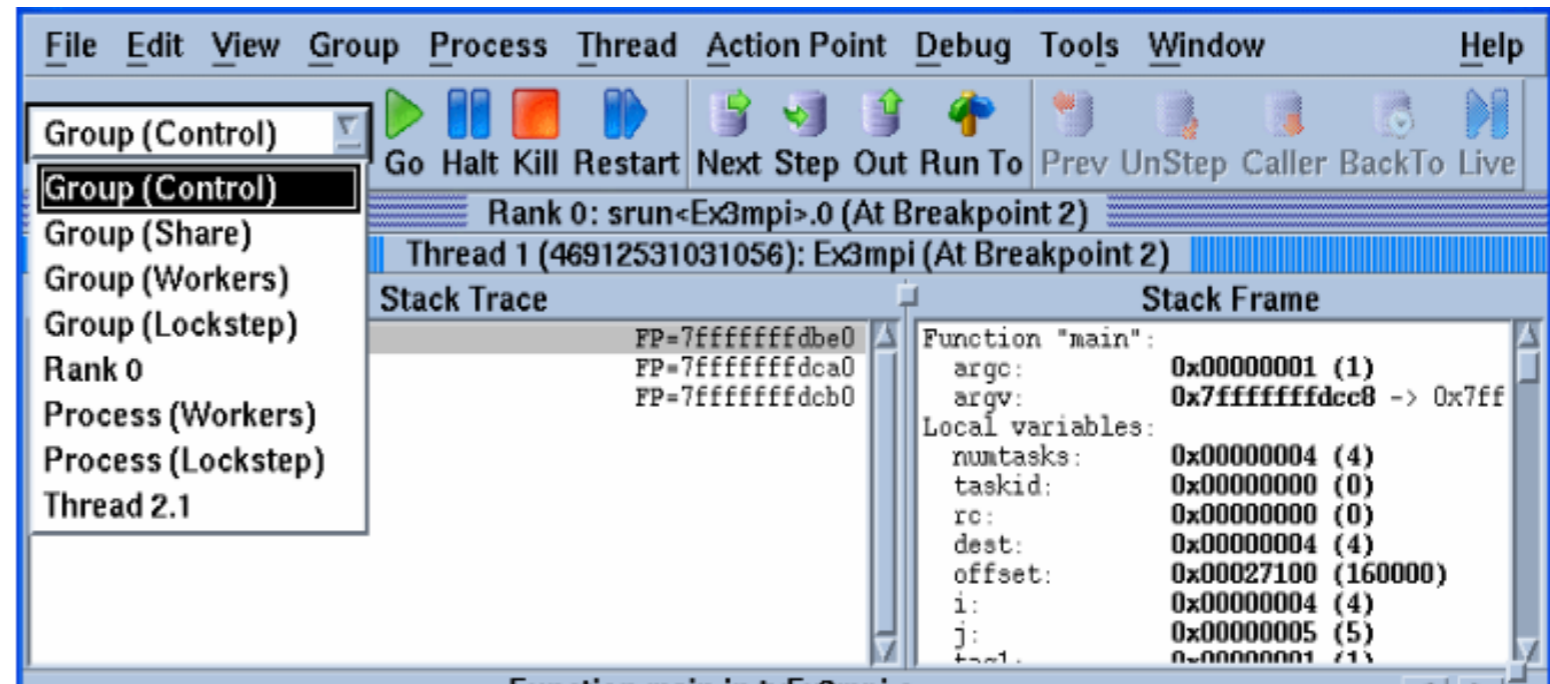


Introduction to TotalView



Basic usage

Working with groups



P/T Selection	What is affected by any execution Command
Group (Control)	Default. All processes and their threads.
Group (Share)	All processes and their threads that are in the same share group as the POI (process-of-interest)
Group (Workers)	All threads that are executing user code
Group (Lockstep)	All user threads that are stopped at the same PC
Rank 0	Only the POI and its threads. In the above example, the POI happens to have an MPI rank of 0
Process (Workers)	User threads in the POI
Process (Lockstep)	User threads stopped at the same PC in the POI
Thread 2.1	Only the TOI (thread-of-interest). In the above example, the TOI happens to be 2.1

Introduction to TotalView



Basic usage

Creating custom groups

Group → Custom Groups ...

The screenshot shows the TotalView debugger interface for a process named 'simple'. The main window displays the source code of 'simple.c' with the following lines visible:

```
39     printf("Pi using %i elements = %f\n",N,pi*4);
40
41     //printf("rank = %i\tpi = %f\n", rank, pi_0);
42
43     //sleep(60);
```

The 'Custom Groups - simple:Control Group' dialog is open, showing a list of custom groups with 'odd' selected. The membership of the selected group is listed as p1, p3, p4, and p5. The dialog includes 'Add...', 'Remove', 'OK', and 'Help' buttons.

The bottom status bar shows the current thread is stopped at line 36 of 'simple.c' in the 'main+0x17c' function.



Basic usage

Action Points

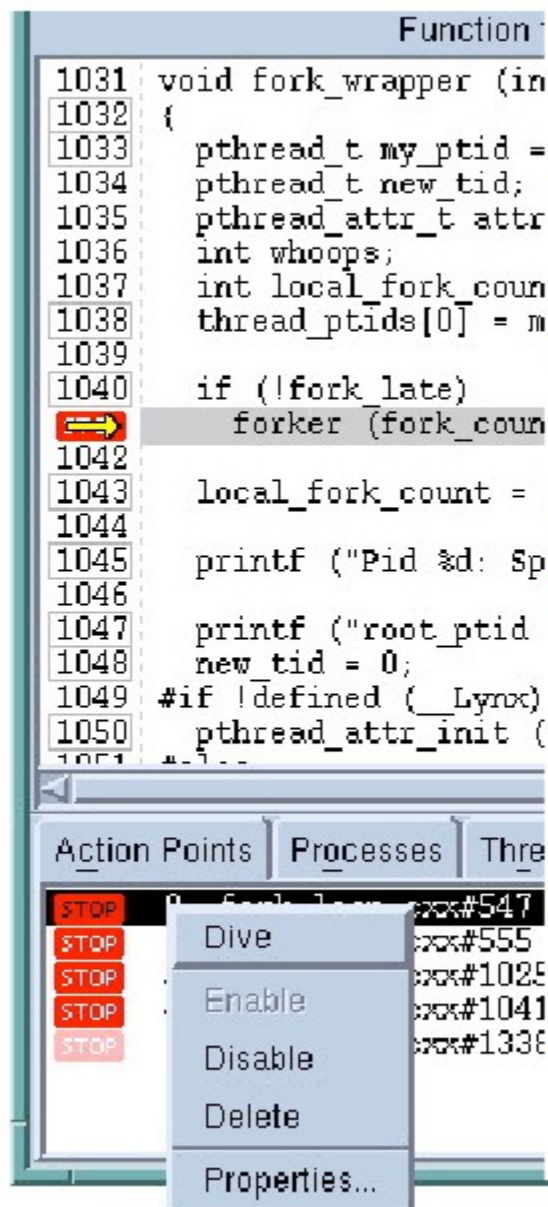
TotalView supports four different types of action points:

- **Breakpoint** - stops execution of the processes or threads that reach it. Note that breakpoints apply to the entire process - if any thread executing a process reaches a breakpoint, TotalView will stop the entire process.
- **Process Barrier Point** - holds each process when it reaches the barrier point until all processes in the group have reached the barrier point. Primarily for MPI programs.
- **Evaluation Point** - causes a code fragment to execute when it is reached. Enables you to set "**conditional breakpoints**" and perform conditional execution.
- **Watchpoint** - enables you to monitor a location in memory and either stop execution or evaluate an expression when the value stored in memory is modified.



Basic usage

Action Points: managing breakpoints

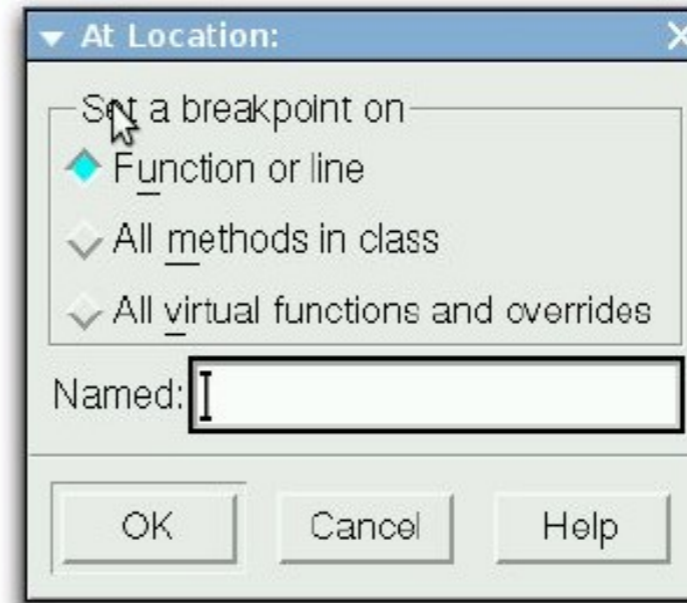
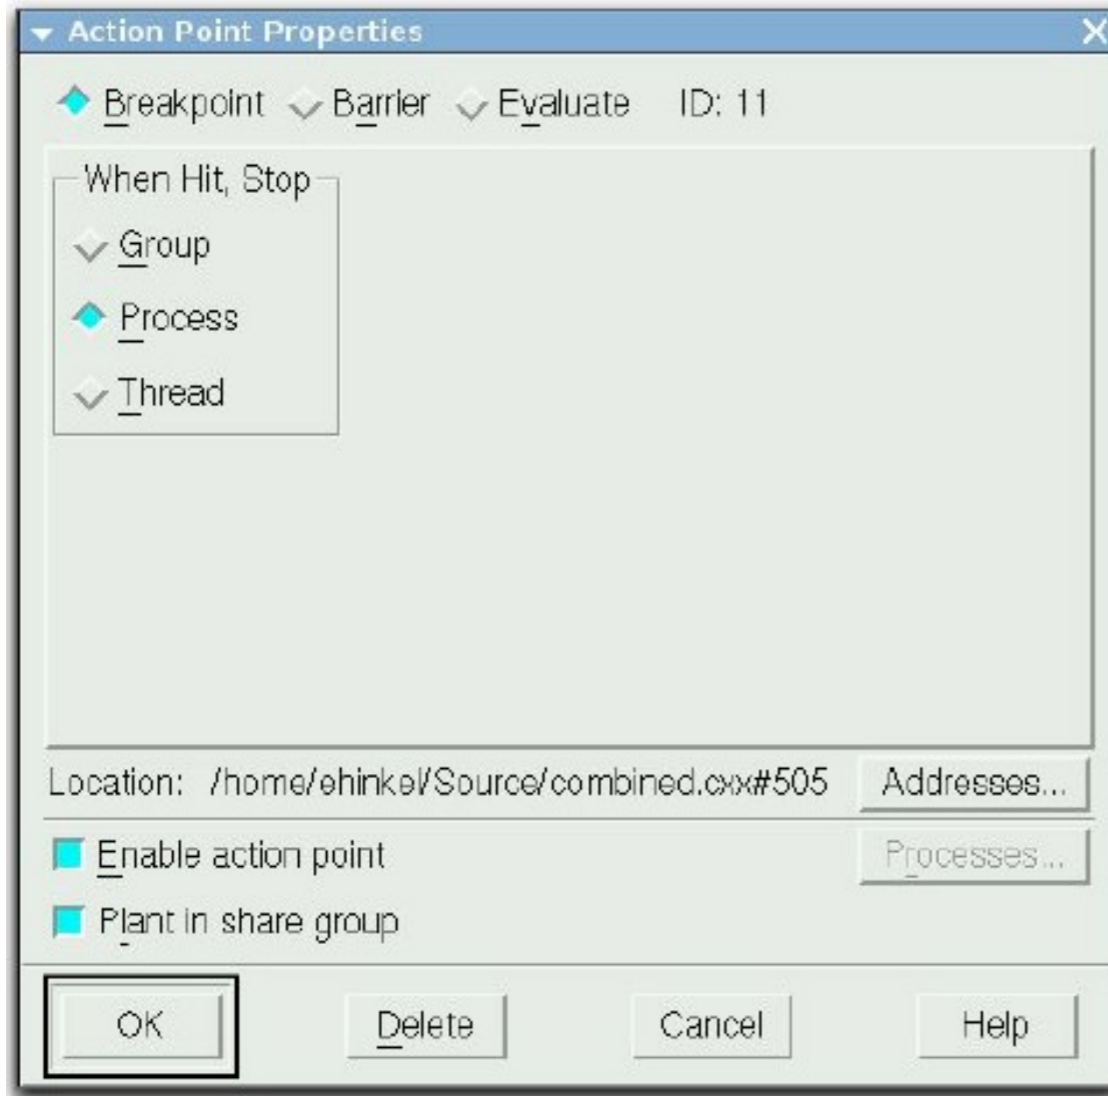


- **Setting action points**
 - Single-click line number
- **Deleting action points**
 - Single-click action point line
- **Disabling action points**
 - Single-click in Action Points Tab Pane
- **Optional contextual menu access for all functions**
- **Action Points Tab**
 - Lists all action points
 - Dive on an action point to focus it in source pane
- **Action point properties**
 - In Context menu
- **Saving all action points**
 - Action Point > Save All



Basic usage

Action Points: Setting



Introduction to TotalView

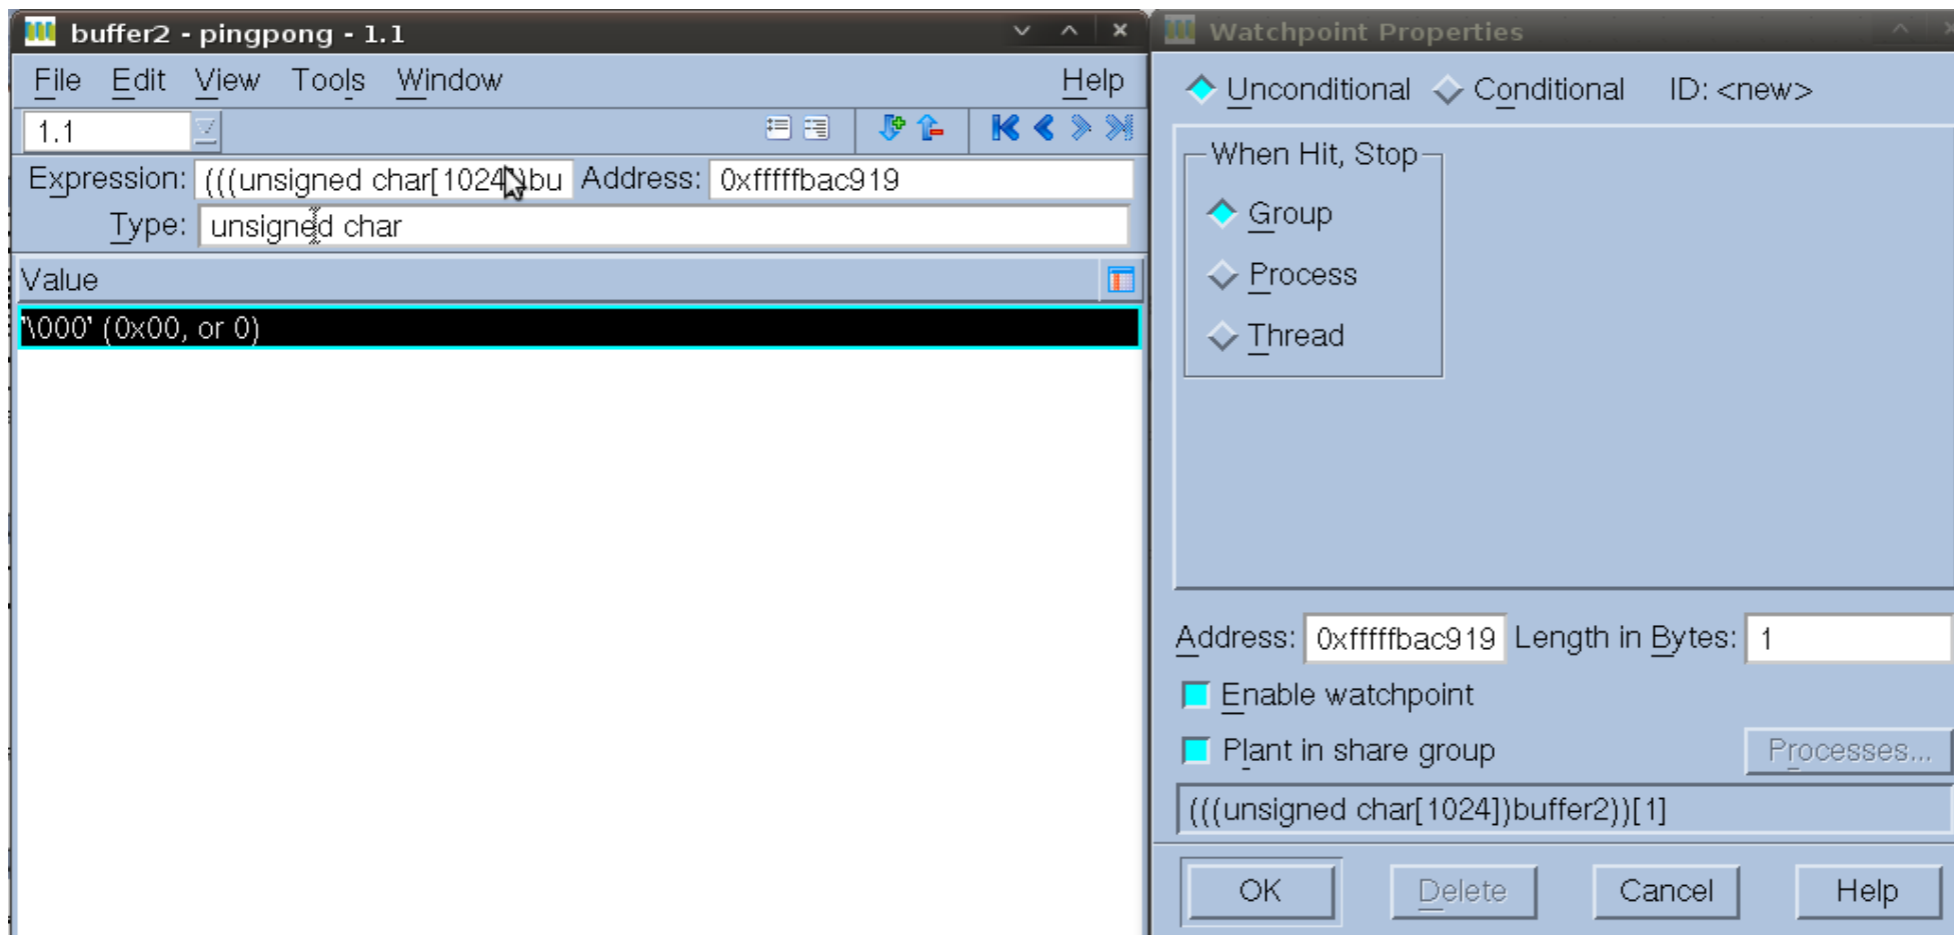


Basic usage

Action Points: Watchpoint

When the contents of a watched variable change, TVD stops the program

Watchpoints are set from the Variable Window: **Tools** → **Watchpoint**



Watchpoints are NOT set on a variable but on a memory region as well
So, user must be aware of the scope of the variable

Introduction to TotalView



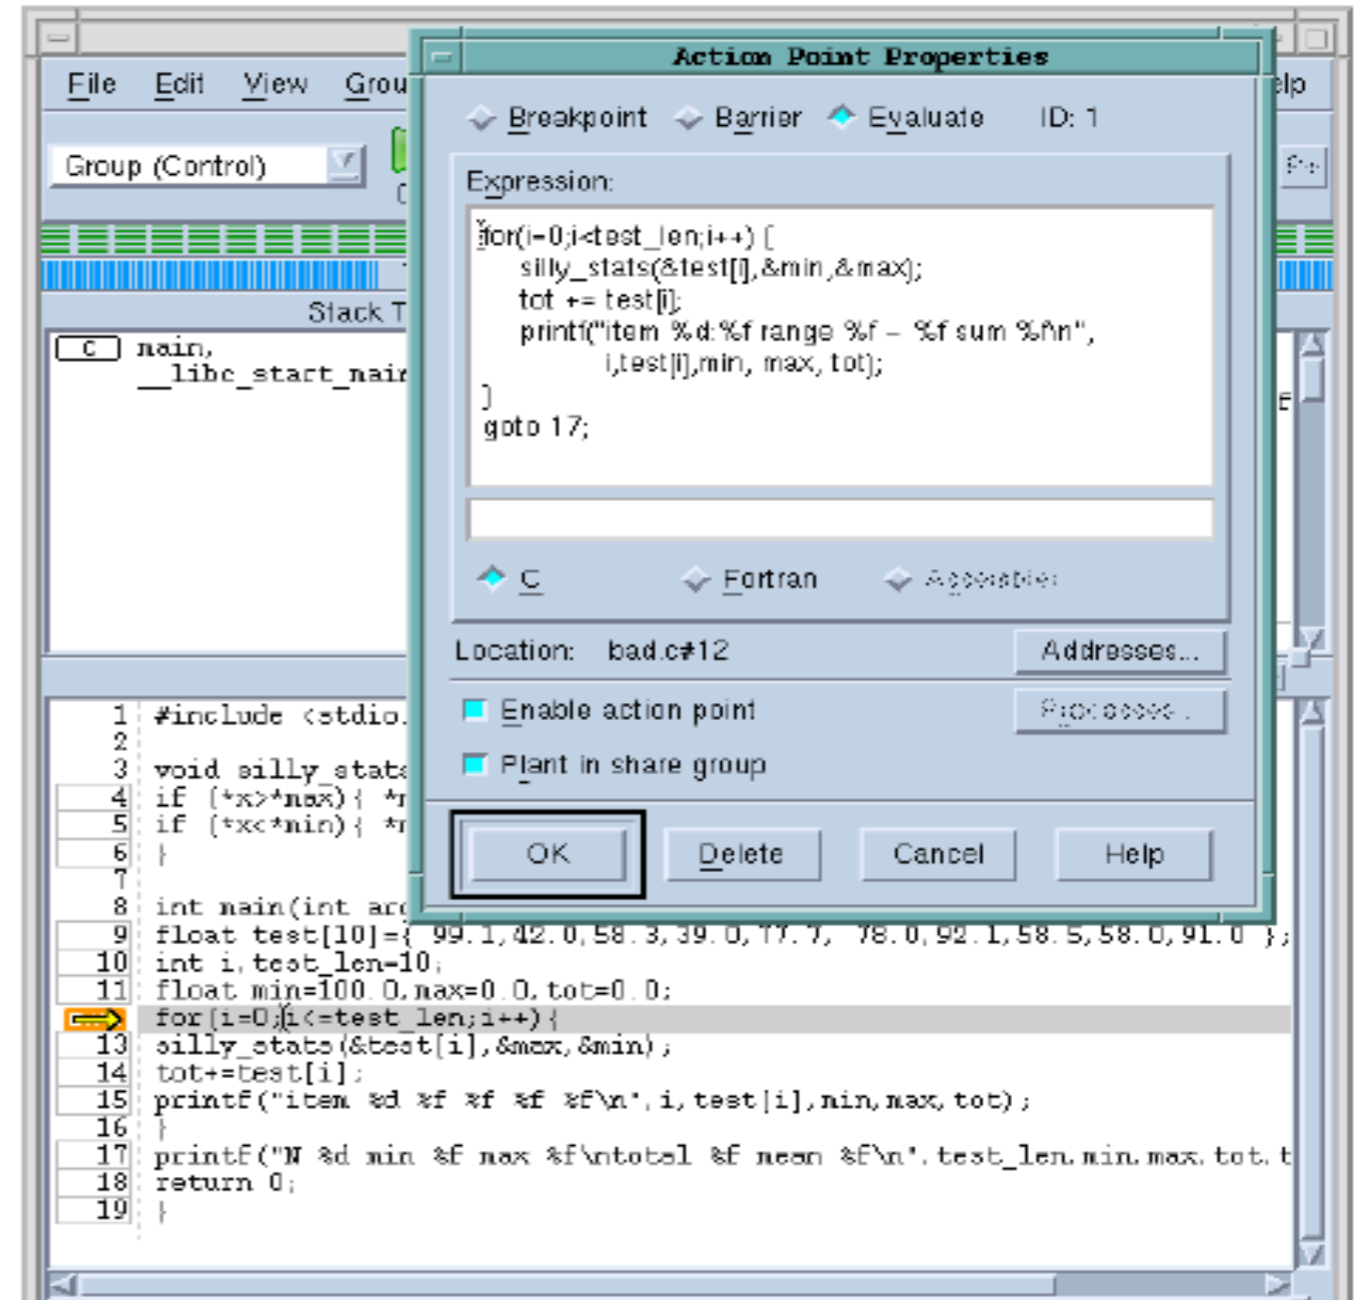
Basic usage

Action Points: *Evaluation and conditional breakpoints*

It is a cool feature that allows:

- Testing small source code patching
- Call functions
- Set variables
- Test conditions
- Use program variables

Can't be used with Replay Engine

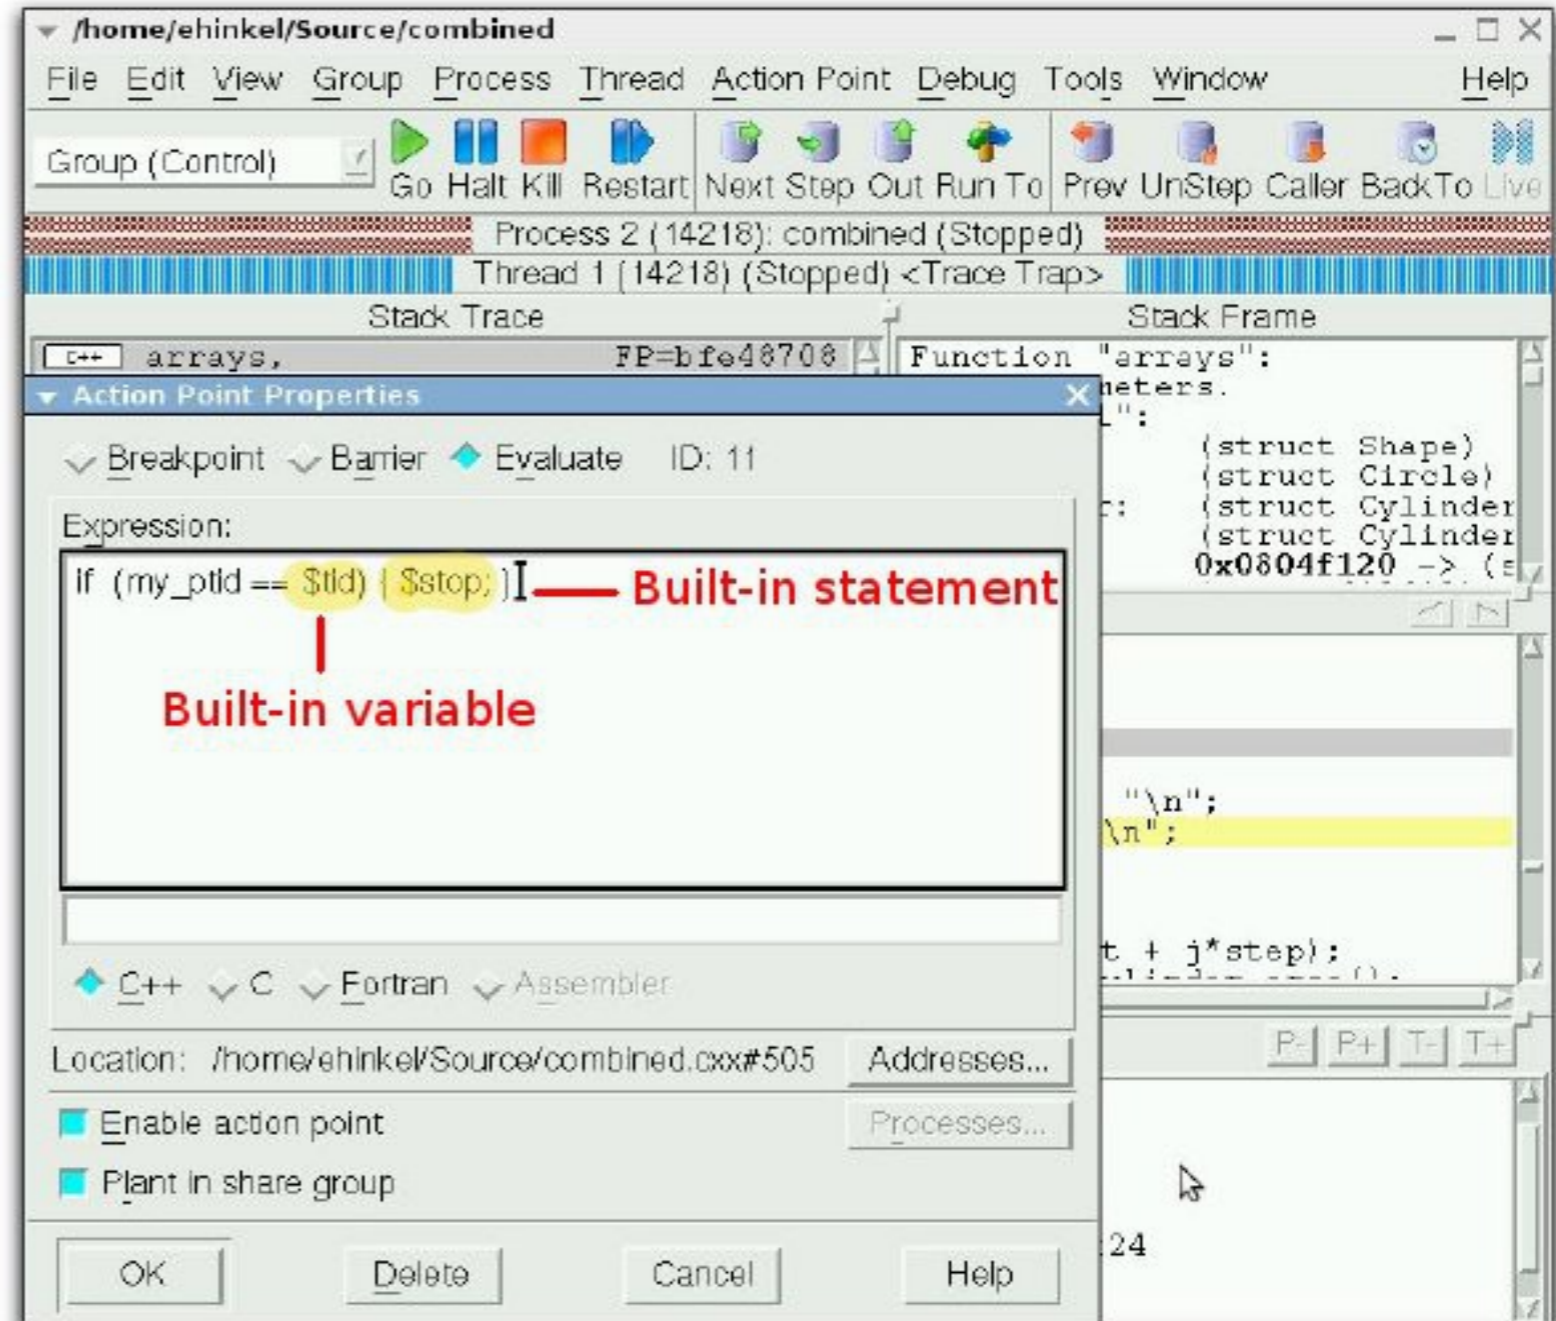


Introduction to TotalView



Basic usage

Action Points:
Evaluation
And
conditional breakpoints





Built-in variables and statements

Built-in Variable	Description
\$clid	Returns the cluster ID.
\$duid	Returns the TotalView-assigned Debugger Unique ID (DUID).
\$newval (watchpoints only).	Returns the value just assigned to a watched memory location
\$nid	Returns the node ID.
\$oldval	Returns the value that existed in a watched memory location before a new value modified it (watchpoints only).
\$pid	Returns the process ID.
\$processduid	Returns the DUID of the process.
\$systid	Returns the system-assigned thread ID. When referenced from a process, generates an error.
\$tid	Returns the TotalView-assigned thread ID. When referenced from a process, generates an error.



Built-in variables and statements

TotalView provides a set of built-in statements that you can use when writing code fragments. The statements are available in all languages, and are shown in the table below.

Built-In Statement	Description
\$count expression	Sets a process-level countdown breakpoint.
\$countprocess expression	When any thread in a process executes this statement for the number of times specified by expression, the process stops. The other processes in the program group continue to execute.
\$countall expression	Sets a program-group-level countdown breakpoint. All processes in the program group stop when any process in the group executes this statement for the number of times specified by expression.
\$countthread expression	Sets a thread-level countdown breakpoint. When any thread in a process executes this statement for the number of times specified by expression 1, it stops. The other threads in the process continue to execute. If the target system does not support asynchronous stop, this executes as a \$countprocess.
\$hold \$holdprocess	Holds the current process. If all other processes in the group are already held in breakpoint state at this eval point, then all will be released. If other processes in the group are running, they continue to run.
\$holdstopall \$holdprocessstopall	Exactly like \$hold, except any processes in the group which are running are stopped. Note that the other processes in the group are not automatically held by this call -- they are just stopped.
\$holdthread	Freezes the current thread leaving other threads running.
\$holdthreadstop \$holdthreadstopprocess	Exactly like \$holdthread except it stops the process. The other processes in the group are left running.
\$holdthreadstopall	Exactly like \$holdthreadstop except it stops the entire group.
\$stop \$stopprocess	Sets a process-level breakpoint. The process that executes this statement stops, but other processes in the program group continue to execute.
\$stopall	Sets a program-group-level breakpoint. All processes in the program group stop when any thread or process in the group executes this statement.
\$stopthread	Sets a thread-level breakpoint. The thread that executes this statement stops, but all other threads in the process continue to execute. If the target system does not support asynchronous stop, this executes as a \$stopprocess.
\$visualize(expression[,slice])	Visualizes the data specified by expression and modified by the optional slice. Expression and slice must be written in the syntax of the code fragment's language. The expression can be any valid expression that yields a data-set (after modification by slice) that can be visualized. The slice is a quoted string containing a slice expression. For more information on how to use \$visualize in an expression, see "Visualizing Data in Expressions" in the TotalView User Guide.



Expression Evaluation and Code Fragments

Code fragments interact with your program, and are evaluated within its runtime context.

They can therefore be used for a variety of purposes, such as:

- Setting conditional breakpoints
- Program patching - branching around code and/or adding new code
- Effecting conditional execution
- Displaying program data
- Modifying program data

TotalView enables you to enter "code fragments" during a debugging session.

Code fragments can include a mixture of:

- C, Fortran or Assembler language code
- TotalView built-in variables (\$tid, \$pid, \$systid ...)
- TotalView built-in statements (\$stop, \$hold, \$stopall ...)



Memory Debugging

Beginning with TotalView version 8.7, the memory debugging functions of TotalView are packaged as a separate, but integrated, client called **MemoryScape**.

Prior to 8.7, the memory functionality was launched from an integrated Memory Debugging Window.

Key features include:

- Memory usage reports
- Leak detection
- Heap status
- Corrupted memory detection
- Dangling pointers

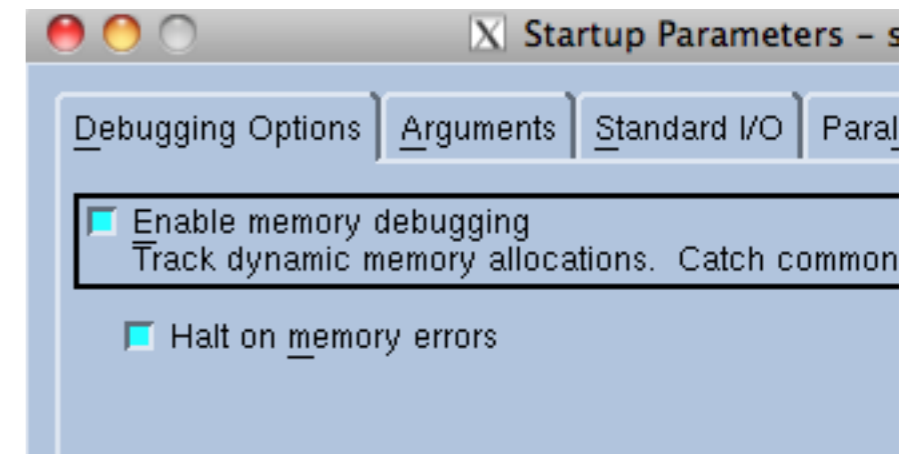
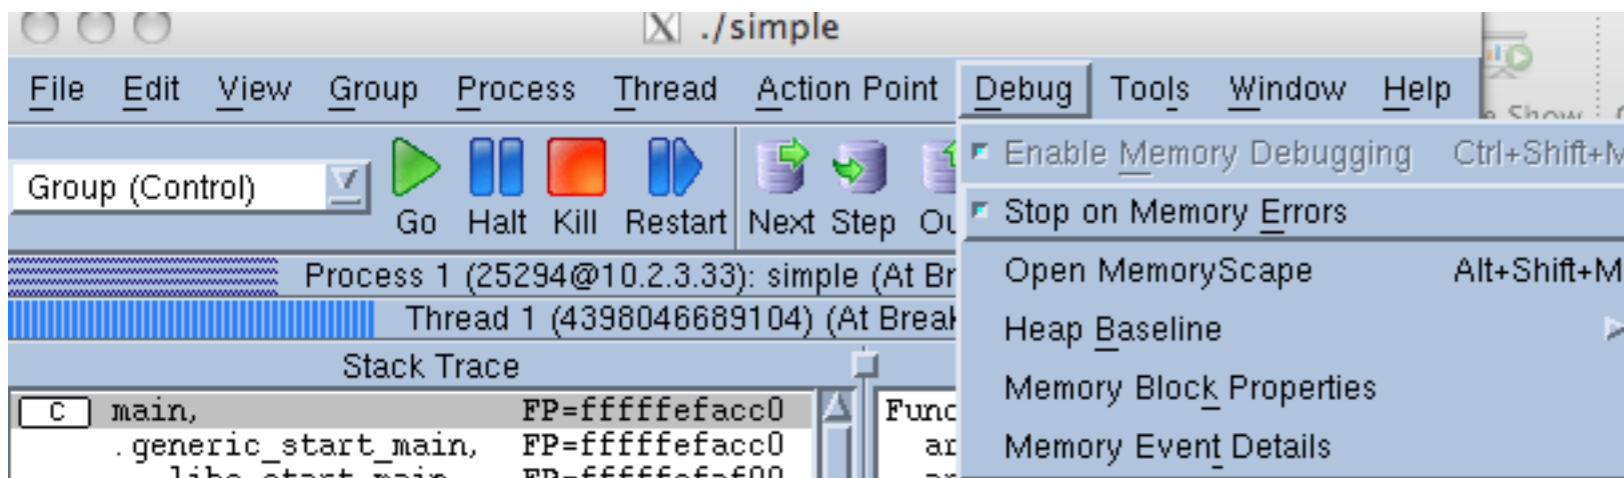
MemoryScape would require a separate tutorial

Introduction to TotalView



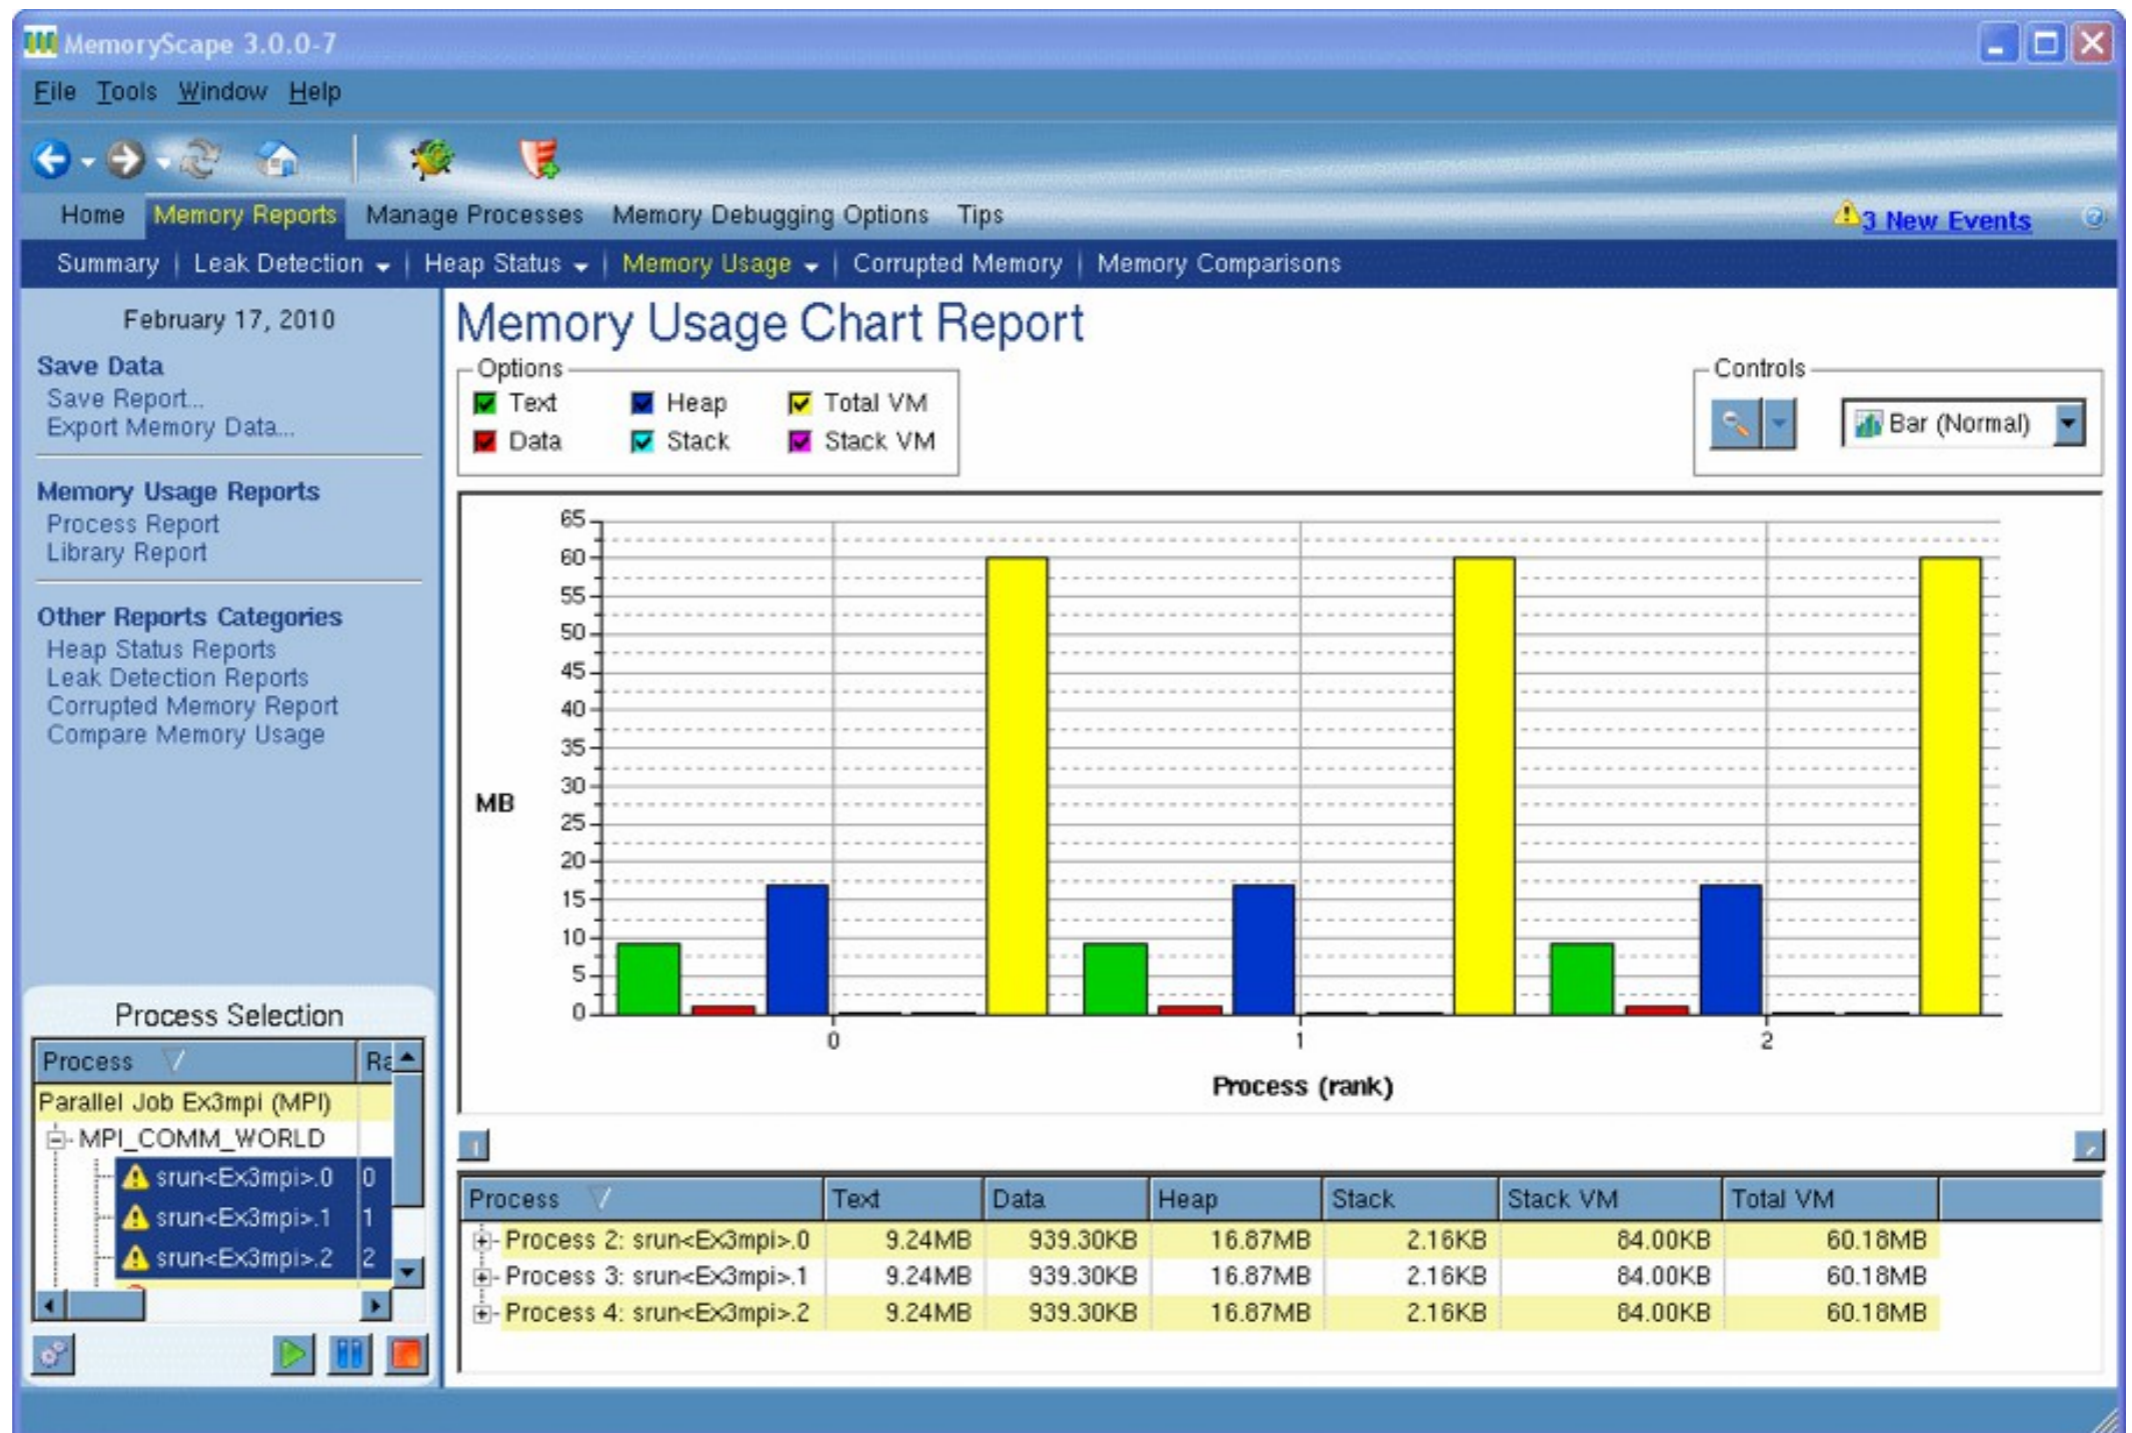
Memory Debugging

MemoryScape can be launched as a standalone application or from within TotalView
→ It is important to click on the checkbox “Enabling Memory Debugging”
when TVD is started.



Introduction to TotalView

Memory Debugging





Memory Debugging

The screenshot shows the MemoryScope 3.0.0-7 application window. The title bar reads "MemoryScope 3.0.0-7". The menu bar includes "File", "Tools", "Window", and "Help". The toolbar contains navigation icons and a "3 New Events" notification. The breadcrumb trail shows "Home", "Memory Reports", "Manage Processes", "Memory Debugging Options", and "Tips".

The main content area is titled "Memory Debugging Options" and includes a "Basic Options" button. Below the title, there is a section for "Enable memory debugging" which is currently unchecked. Underneath, there are several expandable sections:

- Halt execution on memory event or error**: Includes a description and an "Advanced..." button.
- Guard allocated memory**: Contains controls for "Pre-Guard Size" (8 bytes), "Pattern" (0x77777777), "Post-Guard Size" (8 bytes), "Pattern" (0x99999999), and "Maximum Guard Size" (0 bytes).
- Use Red Zones to find memory access violations**: Currently checked.
- Paint memory**: Includes options for "Paint allocations" (Pattern: 0xa110ca7f) and "Paint deallocations" (Pattern: 0xdea110cf).

At the bottom of the main area, there is explanatory text about guard blocks and their use in detecting memory access violations.

On the left side, there is a "Process Selection" panel with a table:

Process	Rank	Event
Parallel Job simple (MPI)		
M_PI_COMM_WORLD		
simple	-1	About
simple	-1	About
simple	-1	About
simple	-1	About

At the bottom right of the application window, there is a "Restore Defaults" button.

Introduction to TotalView

Memory Debugging

MemoryScope 3.0.0-7

File Tools Window Help

Home Memory Reports Manage Processes Memory Debugging Options Tips

Summary Leak Detection Heap Status Memory Usage Corrupted Memory Memory Comparisons

November 3, 2010

Save Data
Save Report...
Export Memory Data...

Leak Detection Reports
Backtrace Report

Other Reports Categories
Heap Status Reports
Memory Usage Reports
Corrupted Memory Report
Compare Memory Usage

Other
Manage Filters

Leak Detection Source Report

Options
 Relative to Baseline Enable Filtering

Process	Bytes	Count	Begin Address	End Address	Backtrace ID	Allocator	Owner
Process 1: simple	4.32KB	5					
pingpong.h	2.00KB	2					
pingpong	2.00KB	2					
Line 22	1024	1					
Line 23	1024	1					
libmyriexpress.so	2.32KB	3					
Process 3: simple	4.32KB	5					
pingpong.h	2.00KB	2					
libmyriexpress.so	2.32KB	3					
Process 4: simple	4.32KB	5					
pingpong.h	2.00KB	2					
libmyriexpress.so	2.32KB	3					

Process Selection

Process	Rank	Event
Parallel Job simple (MPI)		
MPI_COMM_WORLD		
simple	-1	About
simple	-1	About
simple	-1	About
simple	-1	About

Backtrace

ID	Function	Line #	Source Information
94	malloc	166	malloc_wrappers_dlopen.c
	pingpong	22	pingpong.h
	main	56	simple.c
	generic_start_main		libc.so.6
	__libc_start_main		libc.so.6
95	malloc	166	malloc_wrappers_dlopen.c
	pingpong	23	pingpong.h
	main	56	simple.c
	generic_start_main		libc.so.6
	libc_start_main		libc.so.6

Source

```

/gpfs/home/bsc99/bsc99704/EXAMPLES/debug_test/pingpong.h
16 MPI_Status status;
17
18 MPI_Comm_size(MPI_COMM_WORLD, &numprocs);
19 MPI_Comm_rank(MPI_COMM_WORLD, &myid);
20
21 // reserve memory for buffers
22 buffer = (char*)malloc(UNKB*sizeof(char));
23 buffer2 = (char*)malloc(UNKB*sizeof(char));
24
25 // fill buffer data with rank

```

43%



Remote Display

This feature is not currently available in MareNostrum but it is worth to mention it here

TotalView Remote Display lets you start and then view TotalView as it executes on another system.

For example, debugging in MS Windows from home in a PC which is outside a firewall

Some notes on the CLI



One interesting TotalView feature is the CLI (Command Line Interpreter)

- The TotalView Command Line Interpreter (CLI) provides a command line debugger interface
- CLI commands can be integrated into user-written Tcl programs/scripts for "automated" debugging (Advanced)

CLI is useful when:

- a program takes several days to execute
- the program must be run under a batch scheduling system or network conditions that inhibit GUI interaction.
- network traffic between the executing program and the person debugging is not permitted or limits the use of the GUI.

For details see the TotalView documentation located at
www.totalviewtech.com


Hands on



The End



Thanks for your attention



RED ESPAÑOLA DE
SUPERCOMPUTACIÓN