



ATPESC 2018

Power Technologies and Techniques for Debugging HPC Applications

*Bill Burns – Sr. Director of Product Development &
Product Manager*



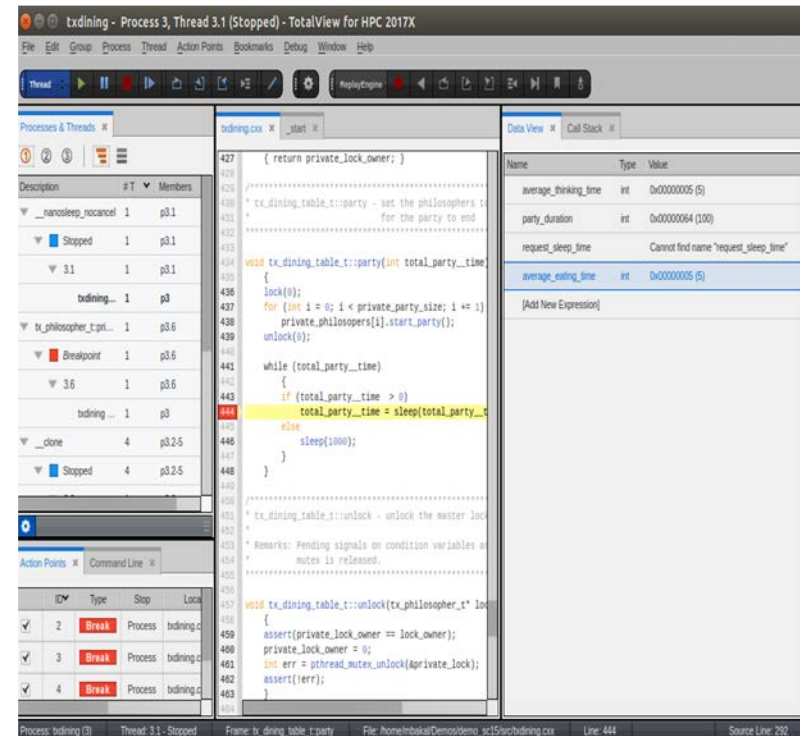
Agenda

- What is TotalView?
- Overview of TotalView features
- TotalView's new UI
- Python debugging
- Reverse debugging
- MPI and OpenMP parallel debugging
- Advanced C++ and Data debugging
- Memory debugging
- Unattended /Batch debugging
- Accelerator debugging
- Remote Display debugging
- Using TotalView
- TotalView resources and documentation
- Questions/Comments

TotalView Features

TotalView Features

- Comprehensive C, C++ and Fortran multi-core and multi-threaded analysis and debug environment
 - Thread specific breakpoints
 - Control individual thread execution
 - View thread specific stack and data
 - View complex data types easily
- CUDA and Xeon Phi debugging
- Integrated Reverse debugging
- Mixed Language - Python C/C++ debugging
- Track memory leaks in running applications
- Unattended debugging
- Linux, macOS and UNIX
- **More than just a tool to find bugs**
 - Understand complex code
 - Improve developer efficiency
 - Collaborate with team members
 - Improve code quality
 - Shorten development time



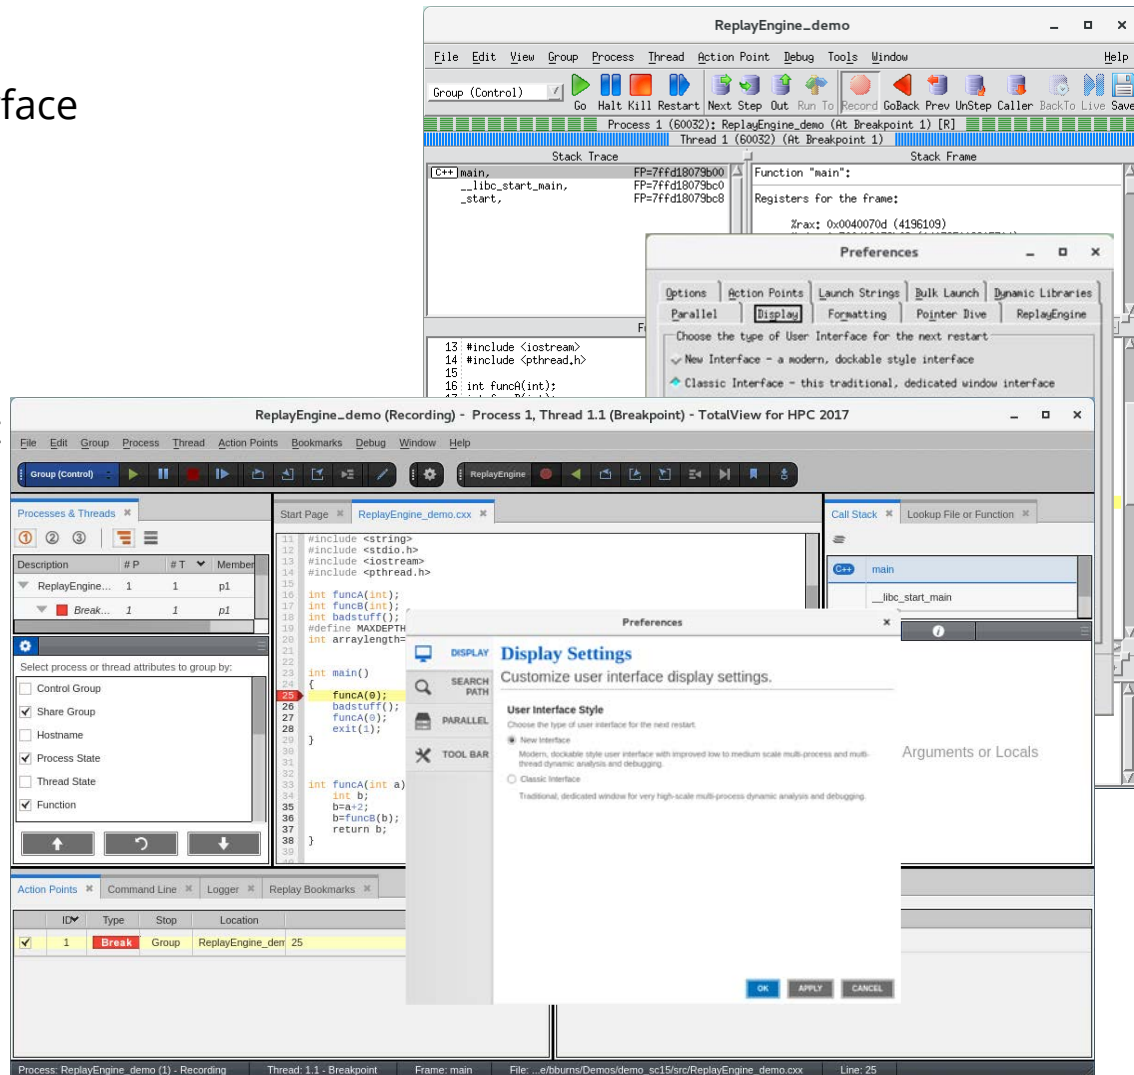
TotalView's New UI

TotalView's New UI

- 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
 - Memory debugging not integrated
 - No very high-scale support



Python Debugging

Python Development Trends

- Increased usage of Python to build applications that call out to C++
- Provides access to
 - High-performance routines
 - Leverage existing algorithms and libraries
 - Utilize advanced multi-threaded capabilities
 - Computational steering
- Calling between languages easily enabled using technologies such as SWIG, ctypes, Cython, CFFI, et al
- Debugging mixed language applications is not easy

Python Debugging

- Debugging one language is difficult enough
- Understanding the flow of execution across language barriers is hard
- Examining and comparing data in both languages is challenging
- 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++
 - Modest system requirements
 - Utilize reverse debugging and memory debugging
 - Python 2.7 (Python 3 coming soon)
- What TotalView does not provide (yet):
 - Setting breakpoints and stepping within Python code

Python Debugging Demo

python - Process 1, Thread 1.1 (Breakpoint) - TotalView for HPC 2017

File Edit Group Process Thread Action Points Bookmarks Debug Window Help

Group (Control) [Icons] [ReplayEngine] [Icons]

Processes & Threads

Description	# P	# T	Members
python (S3)	1	1	p1
Break...	1	1	p1
fact	1	1	p1.1
1.1	1	1	p1.1

Start Page python.c tv_python_example.cpp test_python_to_C.py

```
1 #!/usr/bin/python
2
3 def callFact(int_arg):
4     import tv_python_example as tp
5     # Test some locals
6     a = 3
7     b = 10
8     c = a+b
9     ch = "local string"
10    pi = 3.14159
11    long_var = 2.5
12    true_bool_var = True
13    false_bool_var = False
14    noType = None
15    cx = complex(2,-1)
16
17    return tp.fact(a)
18 if __name__ == '__main__':
19     b = 2
20     result = callFact(b)
21     print result
22
```

Call Stack

Language	Function
C++	fact
Py	callFact
Py	pySupportedTypes
Py	<module>
	__libc_start_main
	start

VAR

Name	Type	Value
Arguments		
int_arg	int	0x00000000
tp	module	0x7f9a284a
a	int	0x00000000

Action Points

ID	Type	Stop	Location	Line
1	Break	Group	tv_python_example	6

Data View

Name	Type	Value
[Add New Expression]		

Process: python (1) Thread: 1.1 - Breakpoint Frame: callFact File: ...me/bburns/Demos/PythonExamples/test_python_to_C.py Line: 17

Reverse Debugging

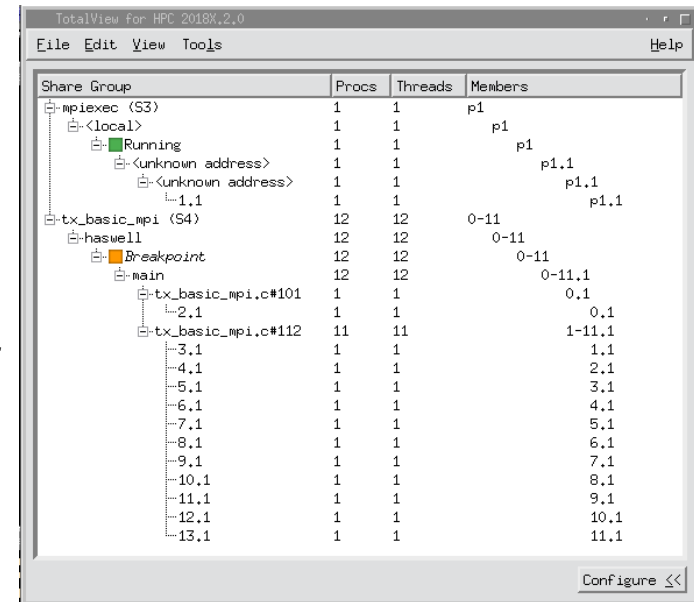
Reverse Debugging

- Reverse debugging is an amazing feature that
 - Saves developers time by finding issues in one debugging session
 - Allows developers to quickly learn new or complex code
 - Enables collaboration and sharing of run sessions
- By
 - Capturing and deterministically replay execution
 - Enables stepping backwards and forward by function, line or instruction
 - Recording from the start of execution or on demand
 - Saving recording files for later analysis or collaboration
 - Linux x86/x86-64
- **Demo!**

MPI and OpenMP Parallel Debugging

Parallel Debugging

- TotalView provides the power to
 - Simultaneously debug many MPI processes and threads in a single debugging session
 - Help locate deadlocks and race conditions
 - Examine message queues between MPI processes
 - Debug all or a subset of the processes in your job
 - Understand complex parallel applications
- By
 - Providing control of entire groups processes, individual processes or even down to individual threads within a process
 - Enabling thread level breakpoints and barrier controls
 - Showing aggregated process and thread state display



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
8 int main(){
9     // One line lambdas
10    auto glambda1 = [](int a, float b) { return a < b; };
11    // Two line lambda
12    auto glambda2 = [](int a, float && b) {
13        if (a < b)
14            return 1;
15        if (b > a)
16            return -1;
17        return 0;
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    std::function<double(double)> f0 = [](double x){
25        return 1;};
26    auto f1 = [](double x){
27        return x;};
28    decltype(f0) fa[3] = {f0, f1, [](double x){
29
```

- 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

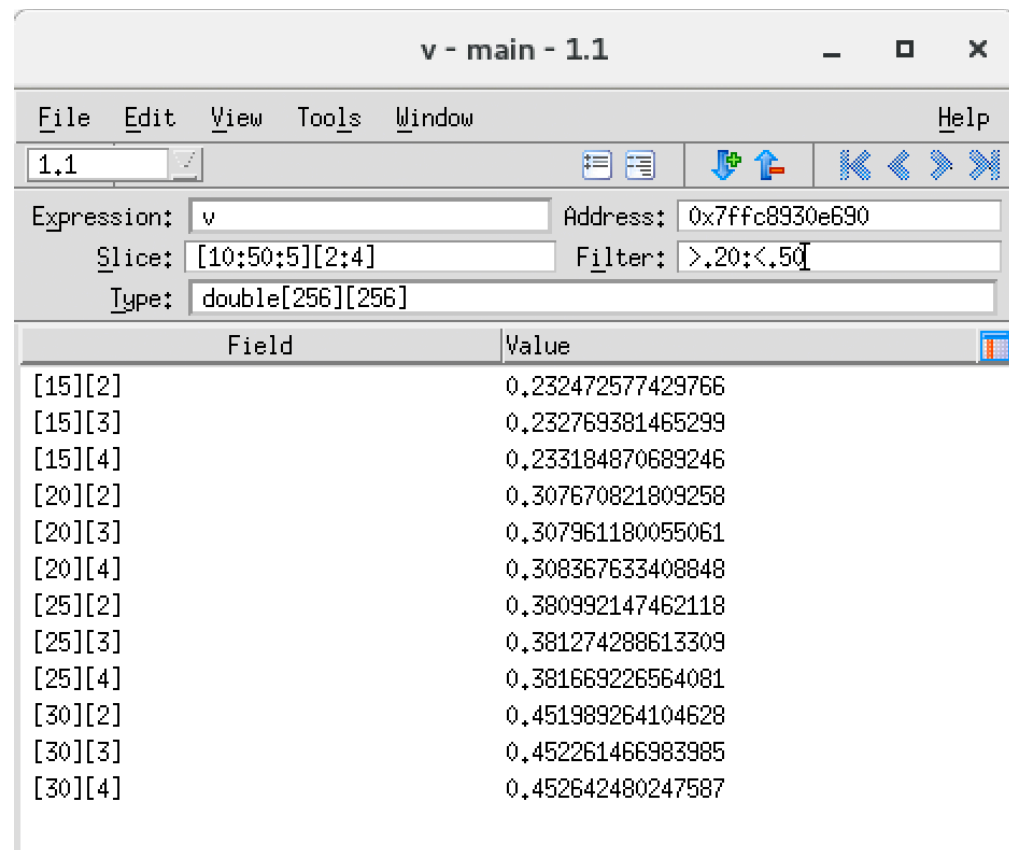
Instead of This

See This!

- 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

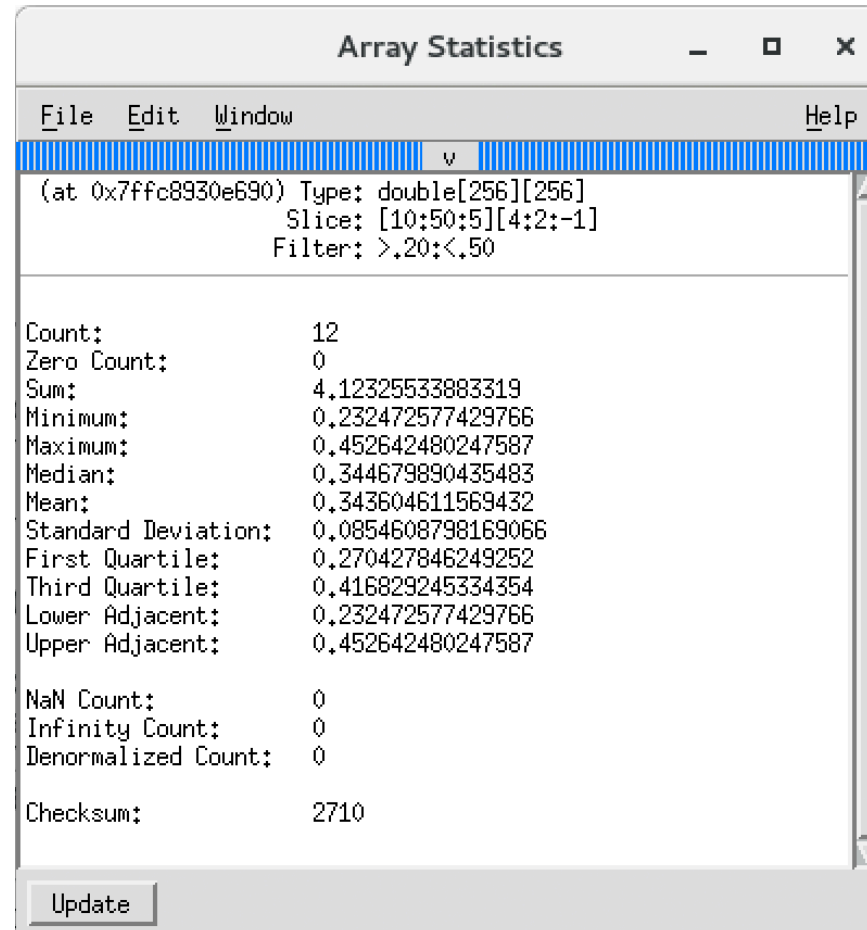
- Slicing – reduce display to a portion of the array
 - [lower_bound:upper_bound]
 - [5:10]
- Striding – Skip over elements
 - [::stride]
 - [::5], [5:10:-1]
- Filtering
 - Comparison: ==, !=, <, <=, >, >=
 - Range of values: [>] *low-value* : [<] *high-value*
 - IEEE values: \$nan, \$inf, \$denorm



Field	Value
[15][2]	0.232472577429766
[15][3]	0.232769381465299
[15][4]	0.233184870689246
[20][2]	0.307670821809258
[20][3]	0.307961180055061
[20][4]	0.308367633408848
[25][2]	0.380992147462118
[25][3]	0.381274288613309
[25][4]	0.381669226564081
[30][2]	0.451989264104628
[30][3]	0.452261466983985
[30][4]	0.452642480247587

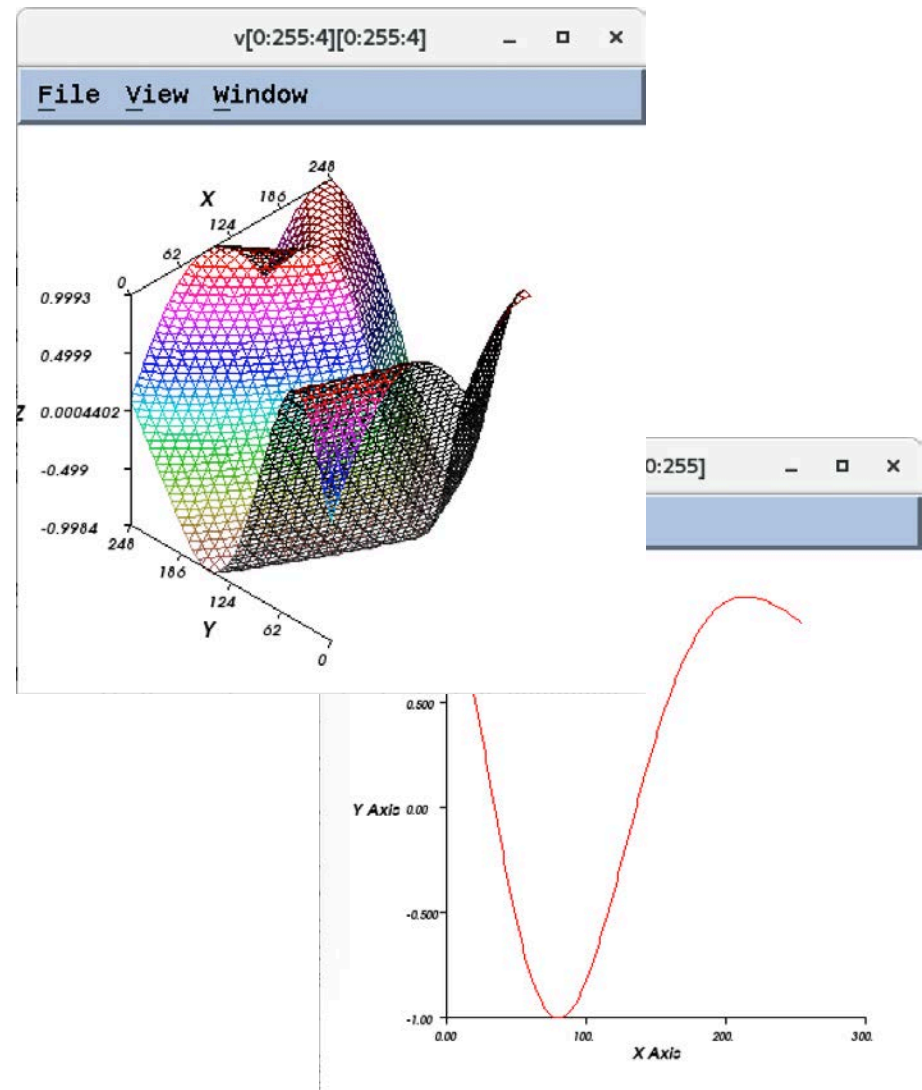
Array Statistics

- 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

strucArray - main - 1.1

File Edit View Tools Window Help

1.1

Expression: strucArray Address: 0x7fff4b260550

Slice: [:] Filter:

Type: struct_junk[20]

Field	Type	Value
[0]	struct junk	(Struct)
a	int	0x00000000 (0)
x	float	0
z	int[4]	(Array)
[0]	int	0x00000000 (0)
[1]	int	0x00000000 (0)
[2]	int	0x00000000 (0)
[3]	int	0x00000000 (0)
[1]	struct junk	(Struct)
a	int	0x00000002 (2)
x	float	4
z	int[4]	(Array)
[0]	int	0x00000000 (0)
[1]	int	0x00000001 (1)
[2]	int	0x00000002 (2)
[3]	int	0x00000003 (3)
[2]	struct junk	(Struct)
a	int	0x00000004 (4)
x	float	8



strucArray - main - 1.1

File Edit View Tools Window Help

1.1

Expression: strucArray[:].x Address: 0x7fff4b260550 [Sparse]

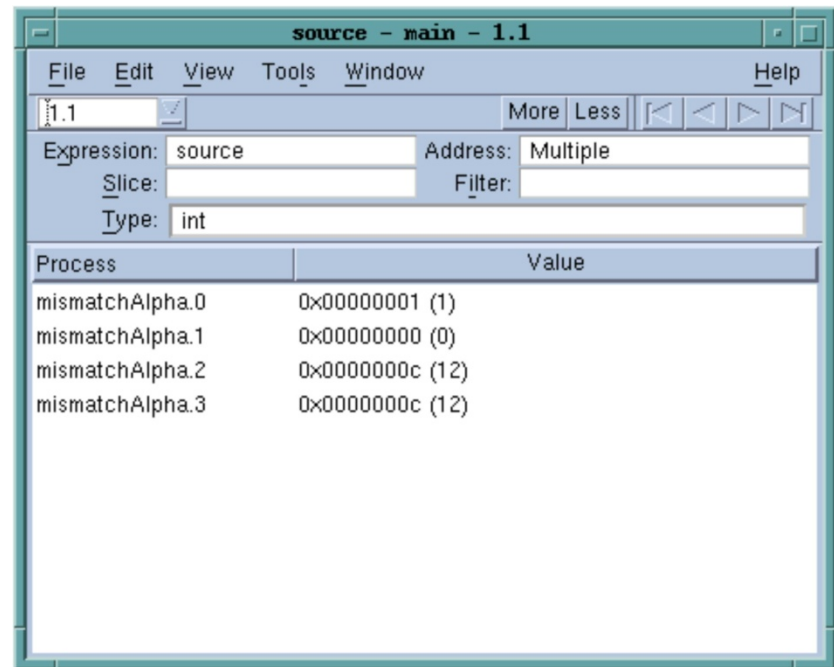
Slice: [:] Filter:

Type: float[20]

Field	Value
[0]	0
[1]	4
[2]	8
[3]	12
[4]	16
[5]	20
[6]	24
[7]	28
[8]	32
[9]	36
[10]	40
[11]	44
[12]	48
[13]	52
[14]	56
[15]	60
[16]	64
[17]	68
[18]	72

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



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

TotalView Heap Interposition Agent (HIA)

- Advantages of TotalView HIA Technology
 - Use it with your existing builds
 - No Source Code or Binary Instrumentation
 - Programs run nearly full speed
 - Low performance overhead
 - Low memory overhead
 - Efficient memory usage
 - Support wide range of platforms and compilers

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?



Unattended/Batch Debugging

Unattended Debugging

- tvscript provides for unattended, straightforward TotalView batch debugging
 - As an alternative to interactive debugging
 - Usable whenever jobs need to be submitted or batched
 - Provides a tool more powerful and flexible than Printf-style debugging
 - Can be used to automate test/verify environments
- Complete documentation in Chapter 4 of the TotalView Reference Guide:
 - http://docs.roguewave.com/totalview/current/html/index.html#page/Reference_Guide%2FBatchDebuggingUsingScripts.html

Think of tvscript as “Printf on steroids”!

Unattended debugging

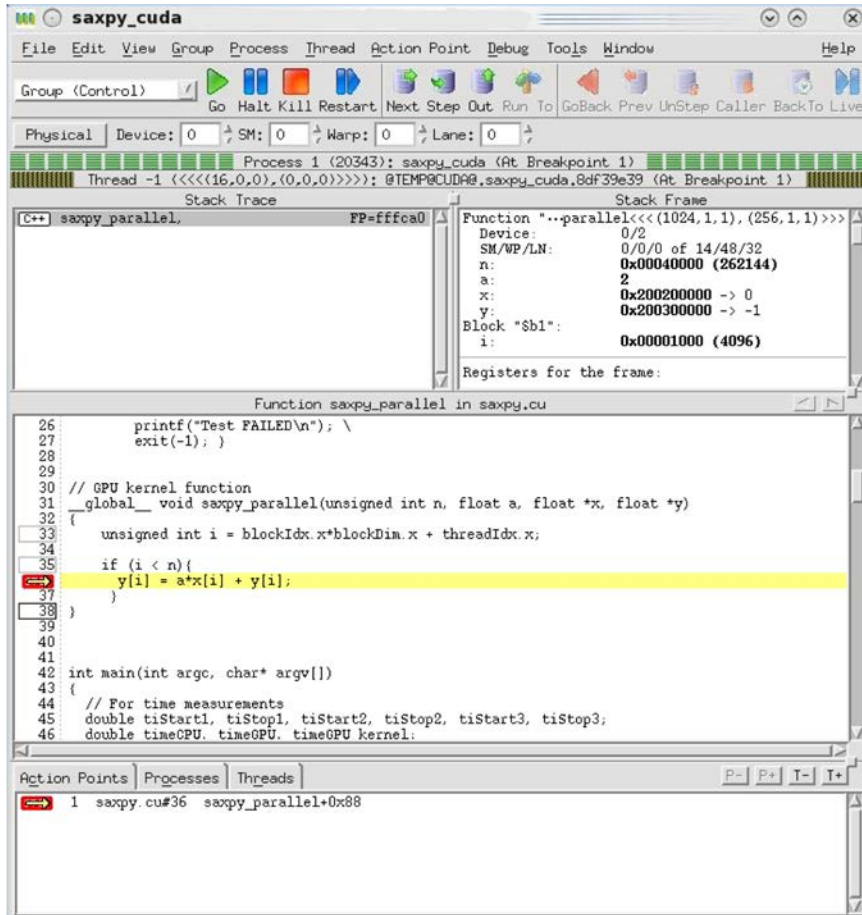
- Command line invocation to run TotalView and MemoryScape unattended
- tvscript can be used to set breakpoints, take actions at those breakpoints and have the results logged to a file. It can also do memory debugging
 - `tvscript -create_actionpoint "method1=>display_backtrace show_arguments" \`
`-create_actionpoint "method.c#342=>print x" myprog -a dataset 1`
- memscript can be used to run memory debugging on processes and display data when a memory event takes place. Exit is ALWAYS an event
 - `memscript -event_action`
`"alloc_null=list_allocations,any_event=check_guard_blocks" \`
`-guard_blocks -maxruntime "00:30:00" \`
`-display_specifiers "noshow_pc,noshow_block_address,show_image" \`
`myProgram -a myProgramArg1`
- Memscript data can be saved in html, memory debug file, text heap status file

Accelerator Debugging

Accelerator Debugging

- High-level debugging support for
 - CUDA
 - Up-to SDK 9.2 and Volta
 - Xeon Phi
 - Native and offloads
 - OpenACC
 - PGI, Cray compilers

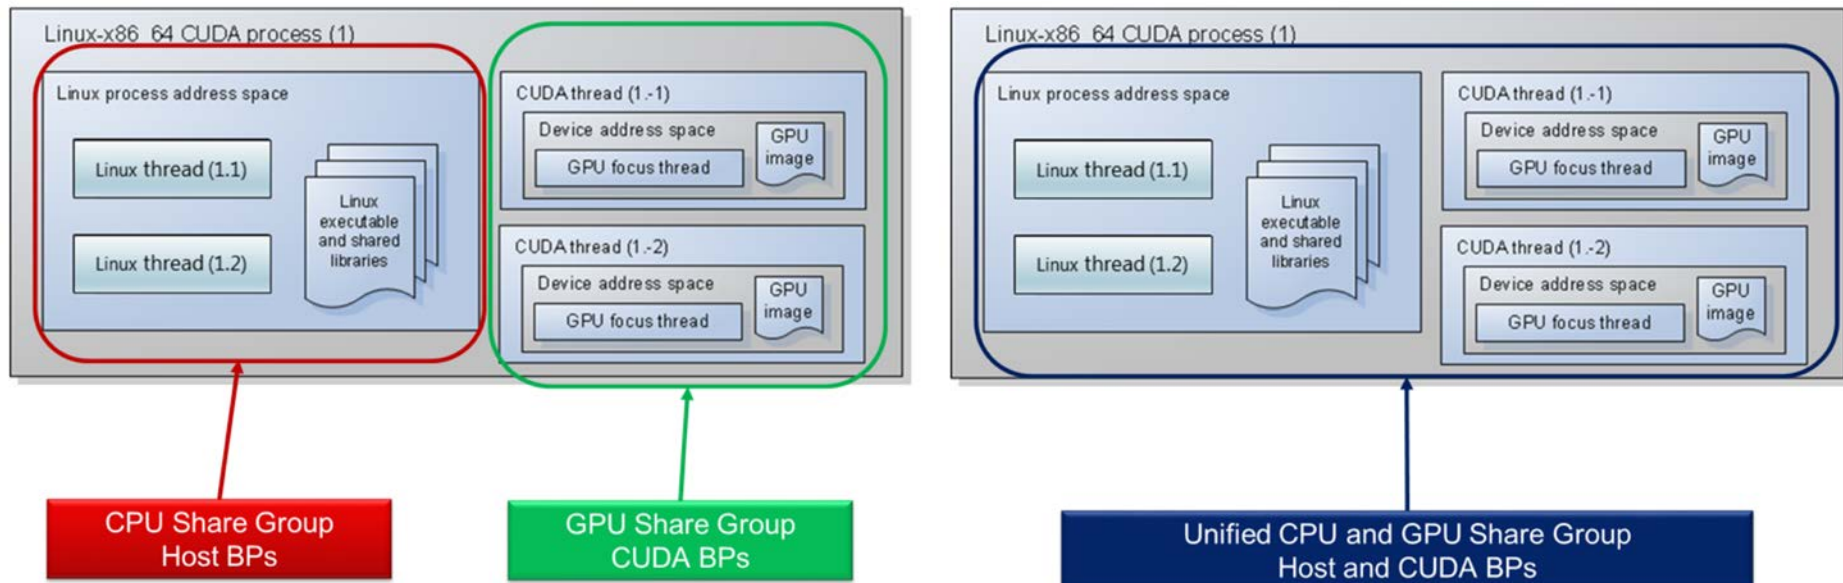
CUDA GPU Debugging with TotalView



- NVIDIA CUDA 7.5, 8.0, 9.2
- 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**

CUDA Debugging Model Improvements

- Coming in TotalView 2018.2 in August
- Set breakpoints in GPU kernel code before it is launched on the GPU
- Unification of source code locations across images (GPU kernel code and shared libraries)
- Improves and streamlines debugging CUDA applications



Intel Xeon Phi Debugging

Supports All Major Intel Xeon Phi Coprocessor Configurations

- Native Mode
 - With or without MPI
- Offload Directives
 - Incremental adoption, similar to GPU
- Symmetric Mode
 - Host and Coprocessor
- **Multi-device, Multi-node**
- Clusters
- KNL Support – Just works like a normal node
- Supports AVX2

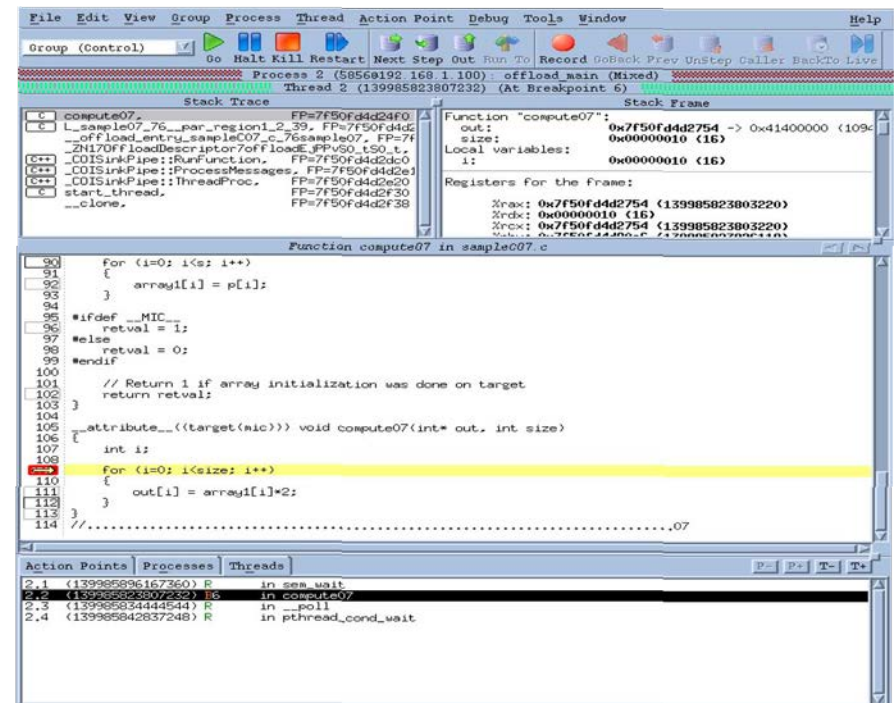
User Interface

- MPI Debugging Features
 - Process Control, View Across, Shared Breakpoints
- Heterogeneous Debugging
 - Debug Both Xeon and Intel Xeon Phi Processes

Memory Debugging

- Both native and symmetric mode

ID	Rank	Host	Status	Description
1		<local>	R	/opt/intel/composerxe/Sample
1.1		<local>	R	in main
1.2		<local>	R	in __poll
1.3		<local>	R	in __poll
1.4		<local>	R	in pthread_cond_wait
2		192.168.1.1	M	/tmp/coi_proc/1/5856/offloa
2.1		192.168.1.1	R	in sem_wait
2.2		192.168.1.1	B	in compute07
2.3		192.168.1.1	R	in __poll
2.4		192.168.1.1	R	in pthread_cond_wait



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 (RDC)

- User must provide information necessary to connect to remote host
- Connection info can be saved for reuse
- Information required includes:
 - User name, public key file, other ssh information
 - Directory where TotalView/MemoryScape is located
 - Path and name of executable to be debugged
 - If using indirect connection with host jump, each host
 - Host name
 - Access type (User name, public key, other ssh information)
 - Access value
- Client also allows for batch submission via PBS Pro or LoadLeveler

Remote Display Client



Session Profiles:





perseid

vesta

1. Enter the Remote Host to run your debug session:

Remote Host: User Name

2. As needed, enter hosts in access order to reach the Remote Host:



	Host	Access By	Access Value	Commands
1		<input type="text" value="User Name"/>		
2		<input type="text" value="User Name"/>		

3. Enter settings for the debug session on the Remote Host :

Path to TotalView on Remote Host:

Arguments for TotalView:

Your Executable (path & name):

Arguments for Your Executable:

Submit Job to Batch Queueing System:

4. Enter batch submission settings for the Remote Host :

Submit Command:

Script to execute via Submit Command:

Additional Submit Command Options:

Using TotalView for Parallel Debugging

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 acquisition 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

The screenshot shows the 'TotalView for HPC: Parallel Program Session' dialog box. On the left is a sidebar with four tabs: 'PARALLEL DETAILS' (selected), 'PROGRAM DETAILS', 'DEBUG OPTIONS', and 'ENVIRONMENT'. Below these is a 'PREVIEW LAUNCH' section. The main area is titled 'Parallel Program Session' and contains the following fields:

- Session Name:** A text field with a placeholder: '[Enter or select a session name, e.g. myprogram with R]'. There is a small icon to the right of the field.
- Parallel System:** A dropdown menu. The selected value is 'BlueGeneQ-Cobalt'. A blue 'REQUIRED' label is to the right of the dropdown.
- Parallel Settings:** A section containing:
 - Tasks (-np):** A text field with a placeholder: '[Enter the number of tasks]'.
 - Additional Starter Arguments:** A text field with a placeholder: '[Enter starter arguments as needed]'.

At the bottom of the dialog, there is a message: 'When you are ready, press Next to continue.' and a row of buttons: 'Help', 'Previous', 'Next', 'Start Session', and 'Cancel'.

Using TotalView at Argonne

- TotalView available on Theta, Vesta, Mira, Cooley
 - Installed at:
/soft/debuggers/totalview-2018-07-26/toolworks/totalview.2018.1.12/bin
- Memory Debugging on BG/Q and Cray should link against the agent, either static or dynamically
 - BG/Q:
 - -L<path> -Wl,@<path>/tvheap_bgqs.ld #static
 - -L<path> -ltvheap_64 -Wl,-rpath,<path> #dynamic
 - Cray:
 - -L<path> -ltvheap_cnl # static
 - -L<path> -ltvheap_cnl -Wl,-rpath,<path> #dynamic
 - <path> = Path to platform specific TV lib
 - export TVLIB=/soft/debuggers/totalview-2018-07-26/toolworks/totalview.2018.1.12/linux-x86-64/lib
 - Substitute linux-power on BlueGene

Job Control at Argonne

- TotalView can be run on simple serial programs on login nodes (though maybe not the preferred method)
- MPI jobs require an allocation, either an interactive session (`qsub -I`) or through a batch script that creates an interactive session.
- `tvscript` and `memscript` can be run totally in batch.

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

Questions/Comments

- Any questions or comments?
 - Don't hesitate to reach out to me directly with any problems or suggestions!
 - **Email:** bill.burns@roguewave.com
- **Thank you for your time today!**

