



TotalView Training - NERSC

MAY 13, 2024

LBL/NERSC Agenda – May 2024

- Introduction
- Latest TotalView Features
- TotalView Roadmap
- Remote Debugging Techniques
- Review of General Debugging Features
- TotalView Debugging at NERSC Best Practices
- GPU Debugging with TotalView on Perlmutter (10:00am)
- MPI and OpenMP debugging
- Reverse Debugging
- Memory Debugging
- Common TotalView Usage Questions
- Q&A

Introductions

- Bill Burns (Senior Director of Software Engineering and Product Manager)

bburns@perforce.com

- John DeSignore (TotalView Chief Architect)

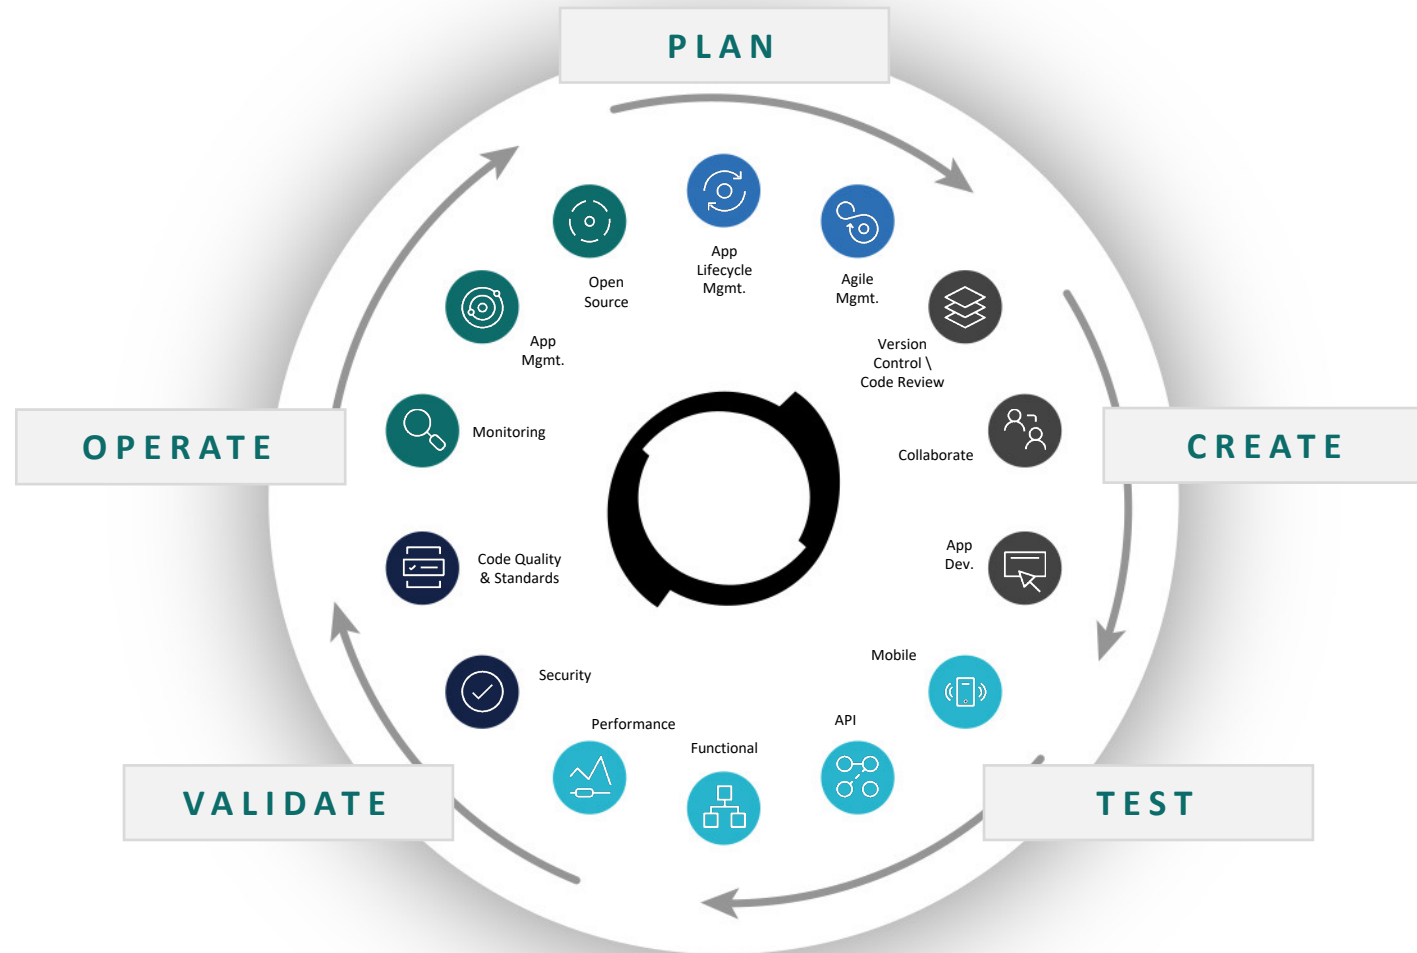
jdelsignore@perforce.com

- Larry Edelstein (Manager, Sales Engineering, TotalView)

ledelstein@perforce.com

PERFORMANCE

Solving the Hardest Challenges in DevOps



Perforce Products



Agile Management

Helix ALM

Hansoft

Gliffy



Code Management & Collaboration

Helix Core

Methodics

Helix4Git

JRebel

TotalView



Application Mgmt. & Components

Puppet

Akana

OpenLogic

Zend

Visualization

SourcePro

IMSL



Automated Testing

Helix QAC

Klocwork

Perfecto

BlazeMeter



Overview of TotalView Labs

Overview of TotalView Labs

Nine different labs and accompanying example programs


- Lab 1 - Debugger Basics: Startup, Basic Process Control, and Navigation
- Lab 2 - Viewing, Examining, Watching, and Editing Data
- Lab 3 - Examining and Controlling a Parallel Application
- Lab 4 - Exploring Heap Memory in an MPI Application
- Lab 5 - Debugging Memory Comparisons and Heap Baseline *
- Lab 6 - Memory Corruption discovery using Red Zones *
- Lab 7 - Batch Mode Debugging with TVScript
- Lab 8 - Reverse Debugging with ReplayEngine
- Lab 9 - Asynchronous Control Lab

Notes

- * Labs 5 and 6 require use of TotalView's Classic UI
- Sample program breakpoint files were created with GNU compilers. If a different compiler is used, they may not load and will need to be recreated.
- Several example programs use OpenMPI so you will need to configure your environment beforehand.
- We do not have a lab specific to Python Debugging yet. There are good examples and instructions in the TotalView *totalview.<version>/<linux-x86-64>/examples/PythonExamples* directory.
- Use this slide deck for GPU specific debugging information

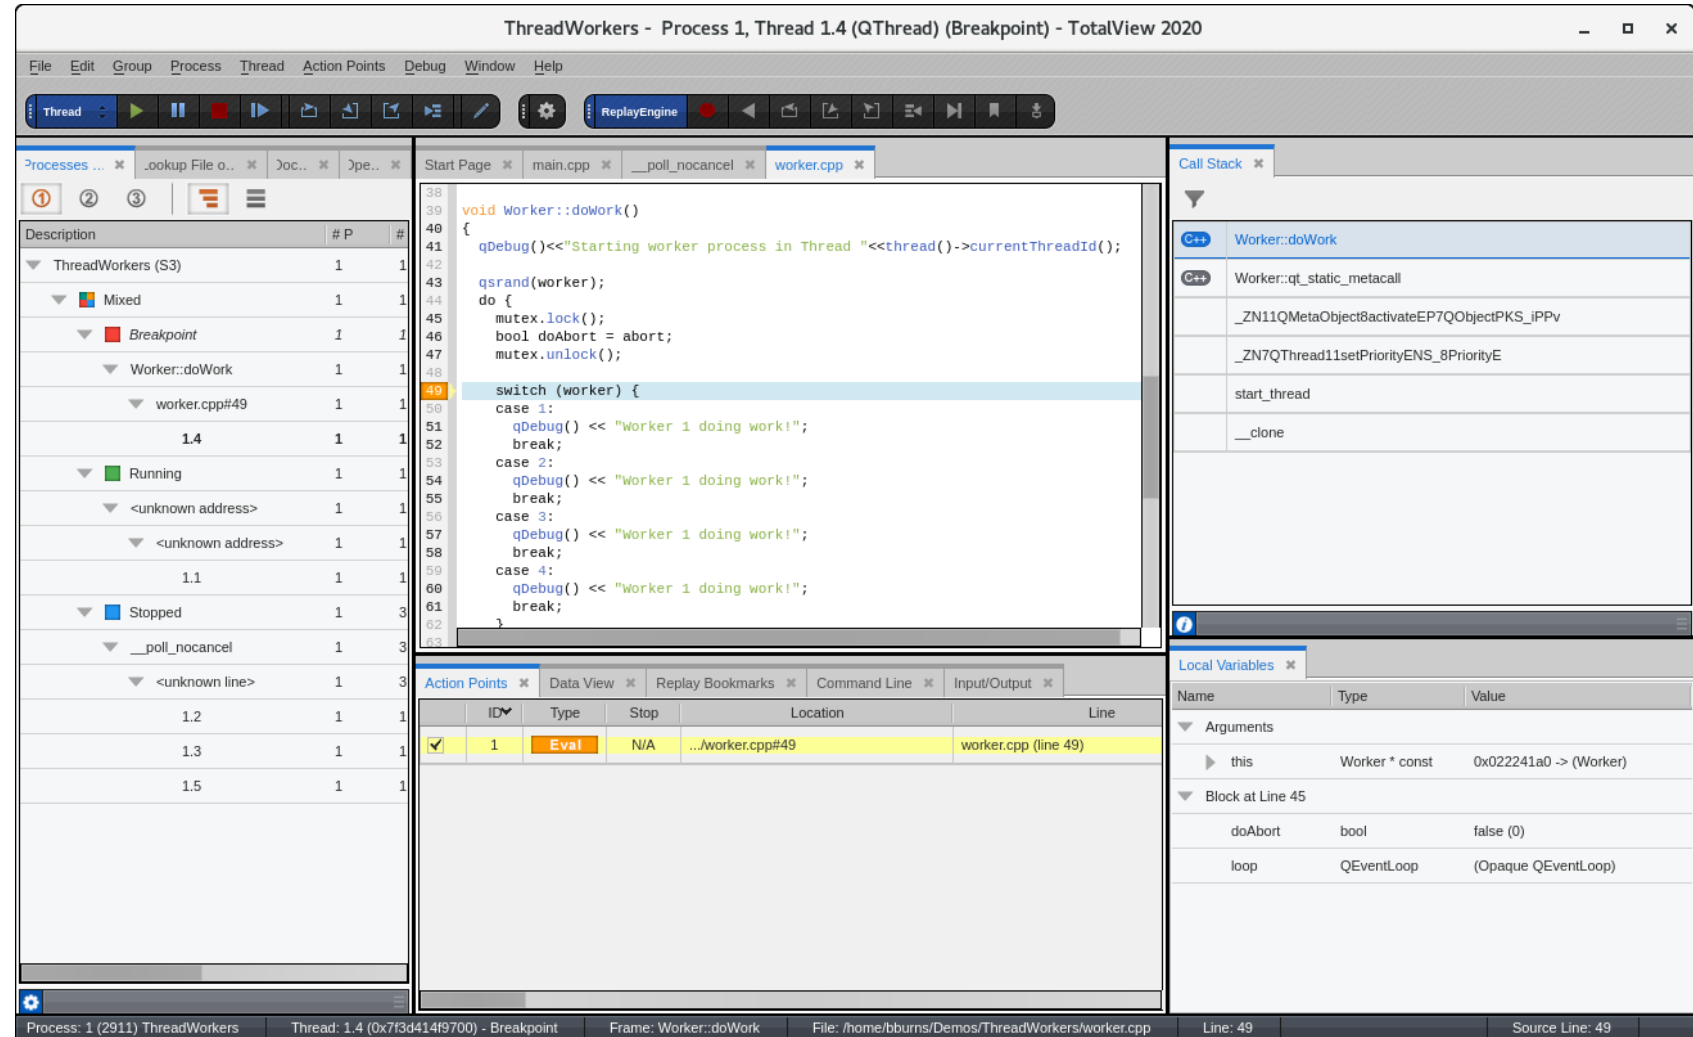
What is TotalView used for?

- Provides interactive Dynamic Analysis capabilities to help:
 - Understand complex code
 - Improve code quality
 - Collaborate with team members to resolve issues faster
 - Shorten development time
- Finds problems and bugs in applications including:
 - Program crash or incorrect behavior
 - Data issues
 - Application memory leaks and errors
 - Communication problems between processes and threads
 - CUDA application analysis and debugging
- Contains batch and Continuous Integration capabilities to:
 - Debug applications in an automated run/test environment



TotalView Features

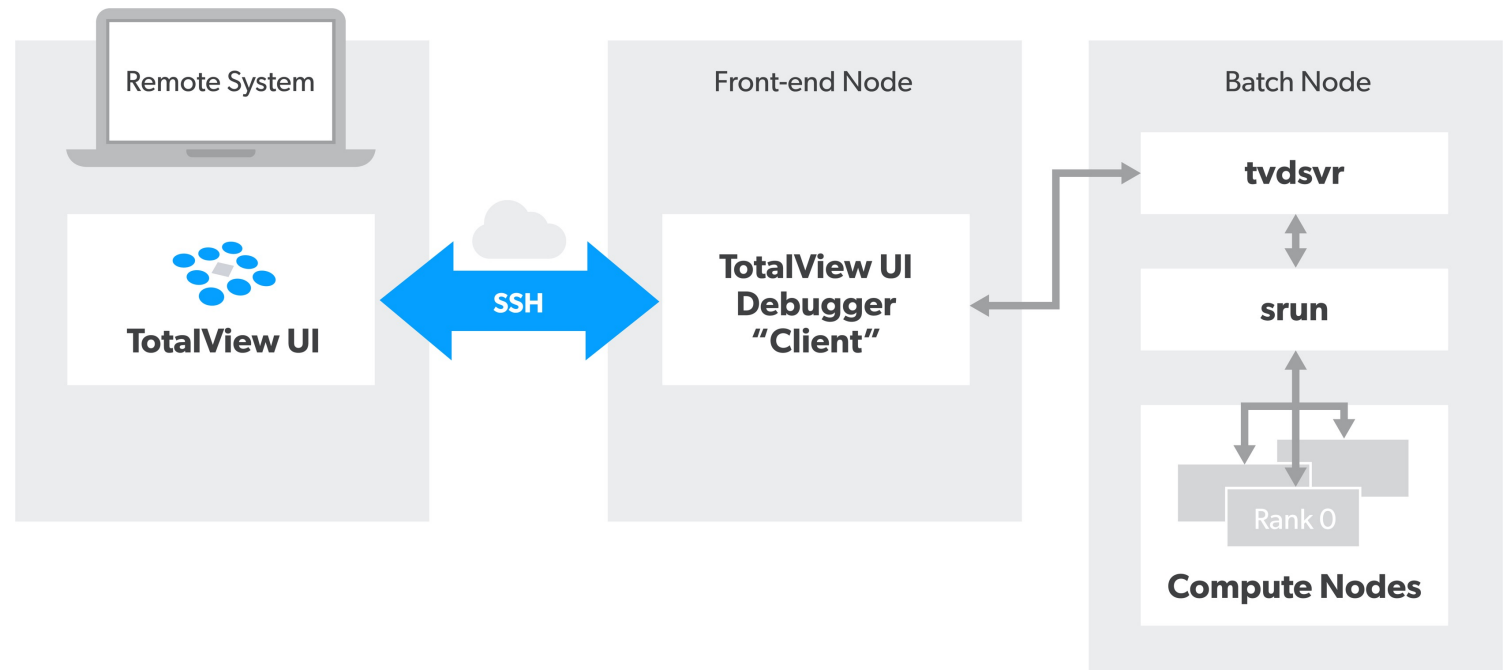
- Multi-process/thread dynamic analysis and debugging
- Comprehensive C, C++ and Fortran Support
- Thread specific breakpoints with individual thread control
 - View thread specific stack and data
- View complex data types easily
- MPI, OpenMP, Hybrid support
- NVIDIA (CUDA) and AMD (HIP) GPU support
- Convenient remote debugging
- Integrated Reverse debugging
- Mixed Language - Python C/C++ debugging
- Memory debugging
- Script 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



Recent TotalView Features

TotalView Remote Client for Windows

- TotalView 2024.1 adds native **Windows remote client support**
- Combine the convenience of establishing a remote connection to a cluster and the ability to run the TotalView GUI locally
- Front-end GUI architecture does not need to match back-end target architecture (macOS front-end -> Linux back-end)
- Secure communications
- Convenient saved sessions
- Once connected, debug as normal with access to all TotalView features
- **Windows, macOS and Linux native front-ends**



TotalView 2024.1 Platform Updates

Platform / Compiler Updates

- macOS Sonoma
- AMD GPU ROCm 6.0 and MI300

Other Updates

- Various bug fixes and other minor enhancements
- Third-party open-source package updates — security



Other Recent TotalView Updates

- Assembly and Register View
- C++ Type Transformations
 - Iterator support
 - Additional container classes
- Array Debugging
 - Array View
 - Array Visualization
- Apple ARM M1/2/3 support
- Memory Debugging Additions
 - Hoarding and Painting
 - Buffer overwrite detection

Category	Class	Iterator kind	Transformed type
Sequence	array (C++11)	Random	Array (Dense)
	vector	Random	Array (Dense)
	deque	Random	Array (Sparse)
	forward_list (C++11)	Forward	List
	list	Bidirectional	List
Associative	set	Bidirectional	Tree
	multiset	Bidirectional	Tree
	map	Bidirectional	Tree
	multimap	Bidirectional	Tree
Unordered associative	unordered_set	Forward, Local	Hashtable
	unordered_multiset	Forward, Local	Hashtable
	unordered_map	Forward, Local	Hashtable
	unordered_multimap	Forward, Local	Hashtable
Adaptors	stack (deque,list,vector)		Struct
	queue (deque,list)		Struct
	priority_queue (deque,vector)		Struct
General utilities	pair		Struct
	tuple (C++11)		Struct
Memory management	unique_ptr (C++11)		Struct
	shared_ptr (C++11)		Struct
	weak_ptr (C++11)		Struct
Numeric	complex		Struct
Strings	string		Struct

TotalView Memory Debugging

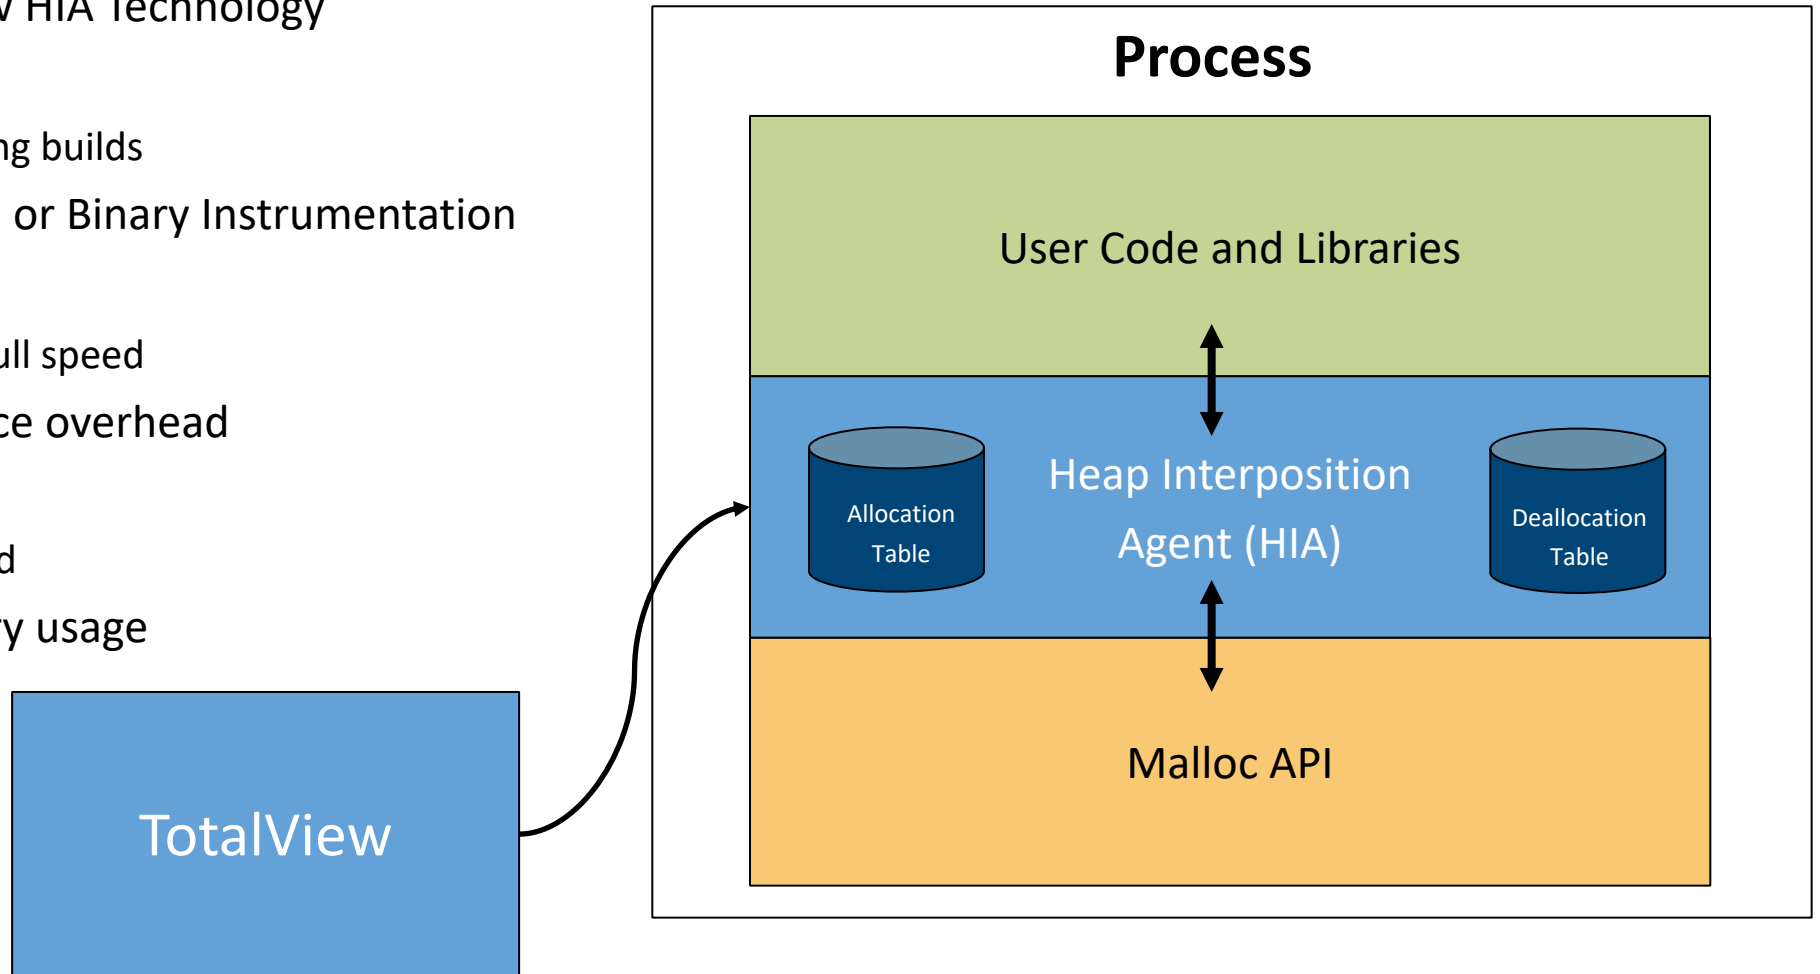
What is a Memory Bug?

- A Memory Bug is a mistake in the management of heap memory
 - Leaking: Failure to free memory
 - Dangling references: Failure to clear pointers
 - Failure to check for error conditions
 - Memory Corruption
 - Writing to memory not allocated
 - Overrunning array bounds

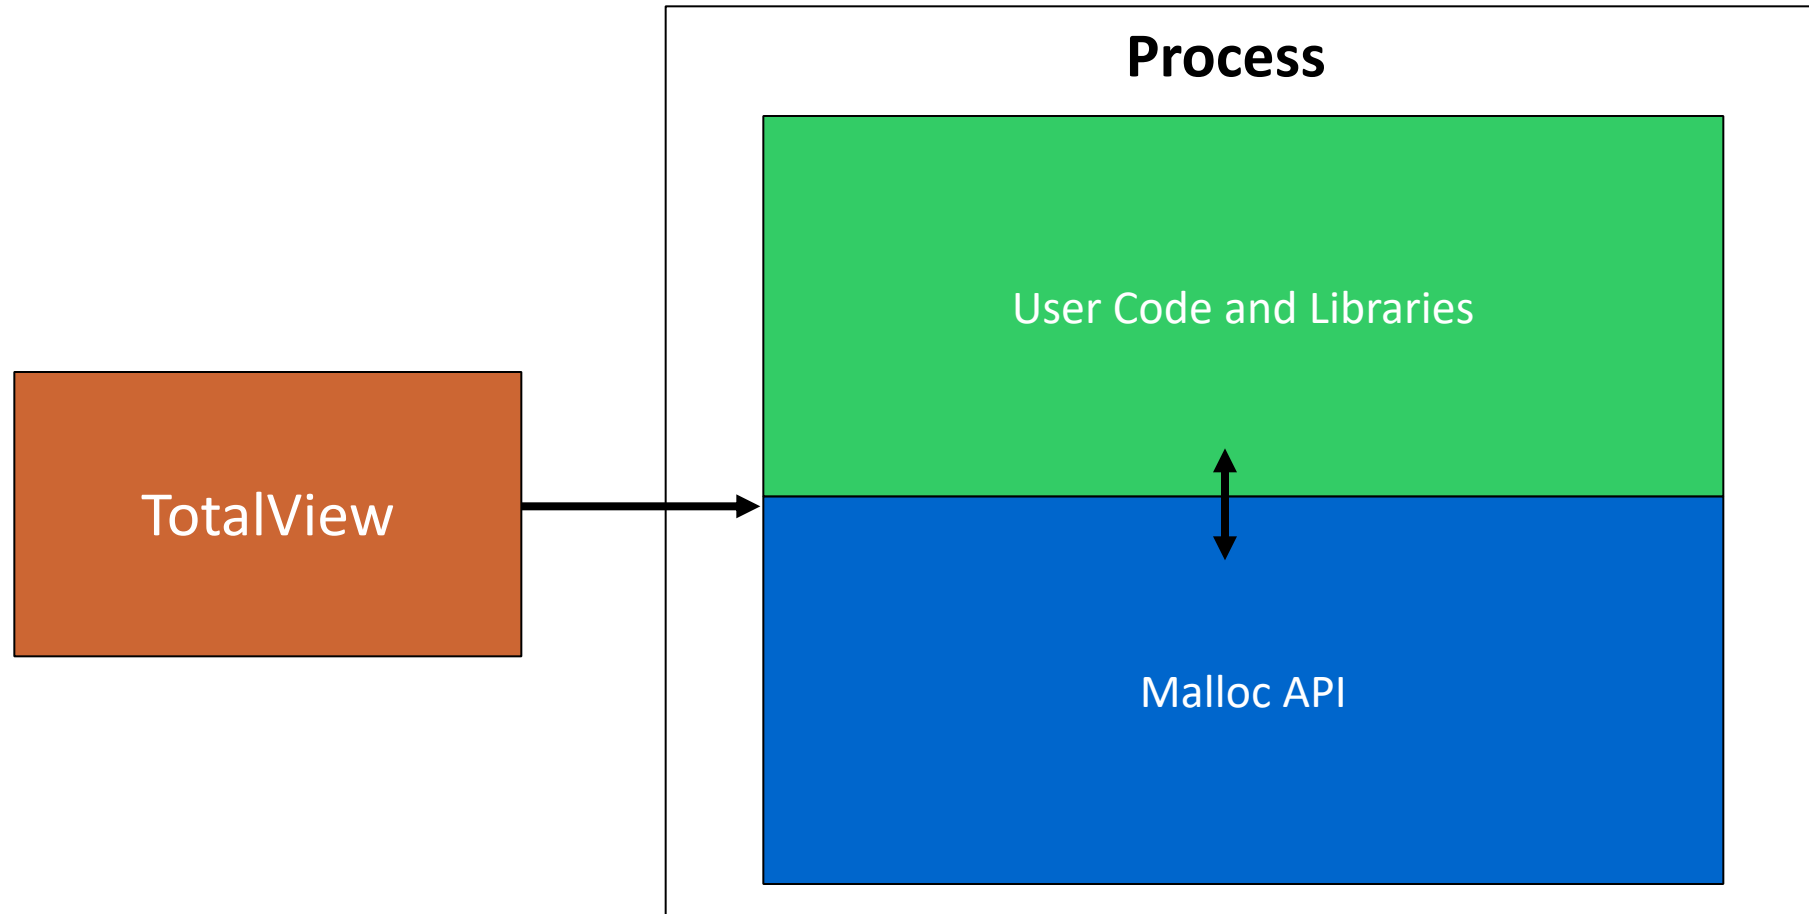


TotalView Heap Interposition Agent (HIA) Technology

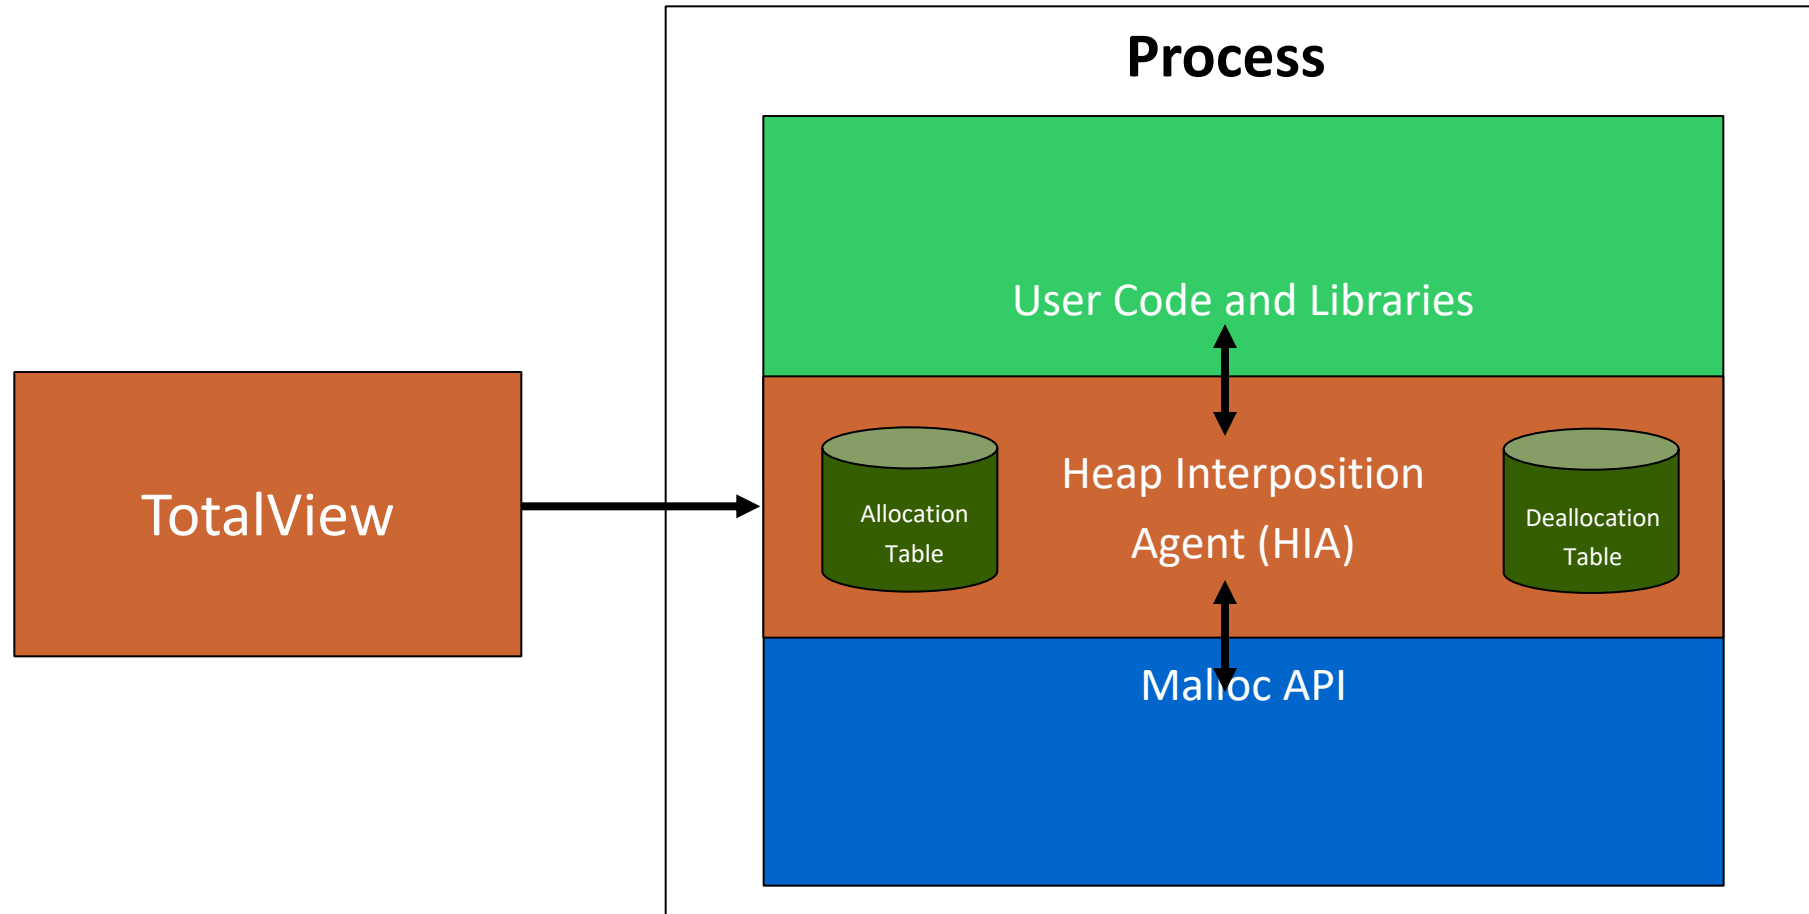
- Advantages of TotalView HIA Technology
 - Use it with your existing builds
 - No Source Code or Binary Instrumentation
 - Programs run nearly full speed
 - Low performance overhead
 - Low memory overhead
 - Efficient memory usage



The Agent and Interposition



The Agent and Interposition



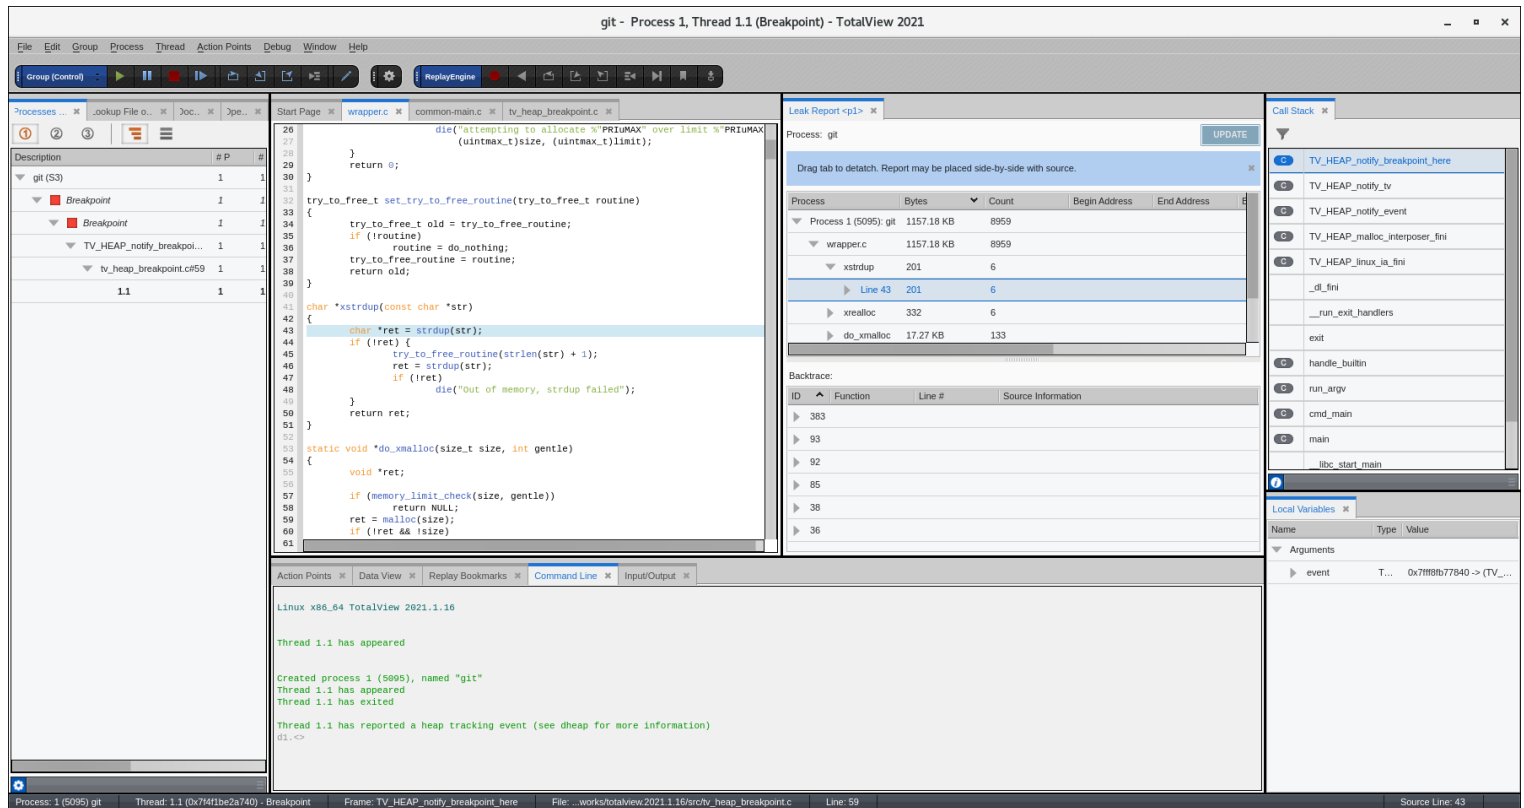
Memory Debugging in TotalView's New UI

TotalView 2024.1 Features

- Leak detection
- Dangling pointer detection
- Heap allocation overview
- Automatically detect allocation problems
- Memory Corruption Detection - Guard Blocks
- Memory Block Painting
- Memory Hoarding

Coming Features

- Graphical heap view
- Memory Corruption Detection - Red Zones
- Memory Comparisons between processes



Memory Debugging Demo

Demo

- Memory Debugging Demo

Debugging OpenMP Applications

TotalView Support for OpenMP

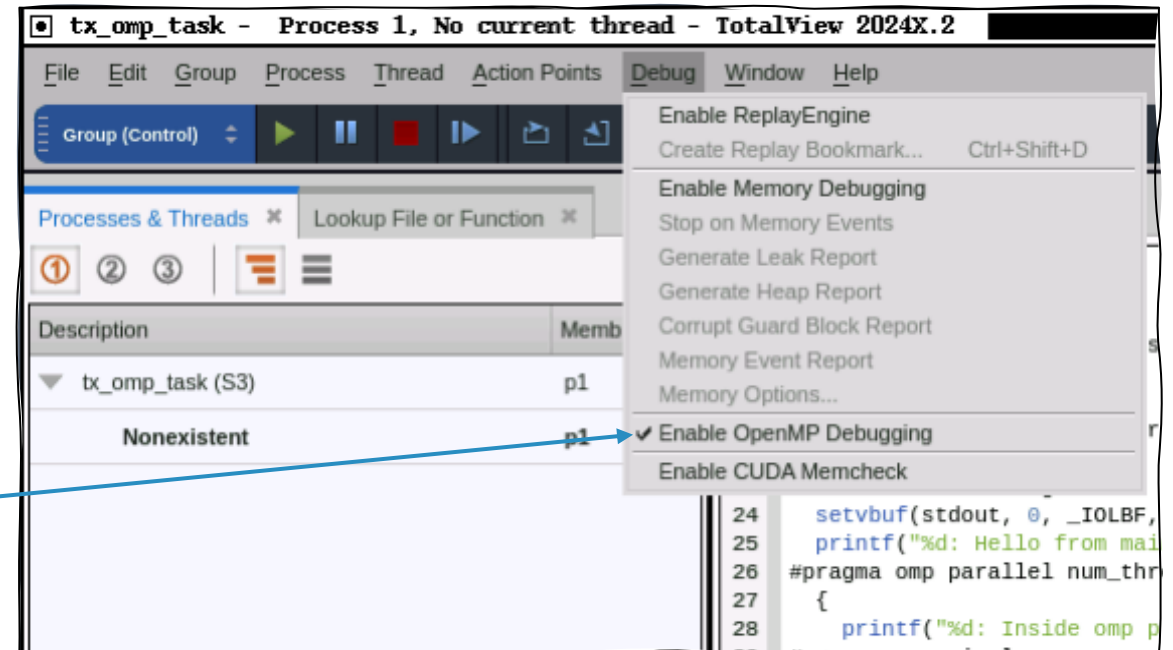
- Source-level debugging of the original OpenMP code (C, C++, and Fortran)
- Debug code inside of OMP **parallel** and **task** regions
 - Set stop-thread breakpoints, single-step, etc.
 - View OMP **shared**, **private**, and **threadprivate** variables
- Debug code inside OMP **target** regions
 - On NVIDIA GPUs, similar to debugging CUDA code
 - On AMD GPUs, similar to debugging HIP code
- CORAL-2 OMP/OMPD support (scheduled for TotalView 2024.2 release)
 - Focuses on Clang, AMD Clang/Flang, and HPE CCE compilers
 - Evaluates quality DWARF debug information produced by the compiler
 - Adds TotalView OpenMP views to display the runtime state of regions, threads, control variables, and ICVs
 - OpenMP thread-stack transformations: filter-out OMP runtime frames, annotate parallel/task regions, and insert parent links
 - Demangles OpenMP outlined function names

Debugging OpenMP Applications

- OpenMP programs are multi-threaded programs
 - Use normal debugging techniques for multi-threaded programs
 - Use stop-thread breakpoints inside parallel regions
 - Use thread-level execution controls: hold/unhold thread, single-step threads / thread groups
- OpenMP **parallel** and **task** regions are in outlined functions
 - A single line of source code can generate multiple block of machine code
 - Outlined function names are mangled by the compiler, but TV 2024.2 will support OMP name demangling
 - TV 2024.2 will support step into and out of OMP parallel region, if the compiler supports OMPD properly
 - Otherwise, set a breakpoint inside the parallel region and let the process run to it
- OpenMP **target** regions are offloaded to the GPU
 - Use normal debugging techniques for CUDA / HIP
 - Unfortunately, OMPD information is not available for target regions

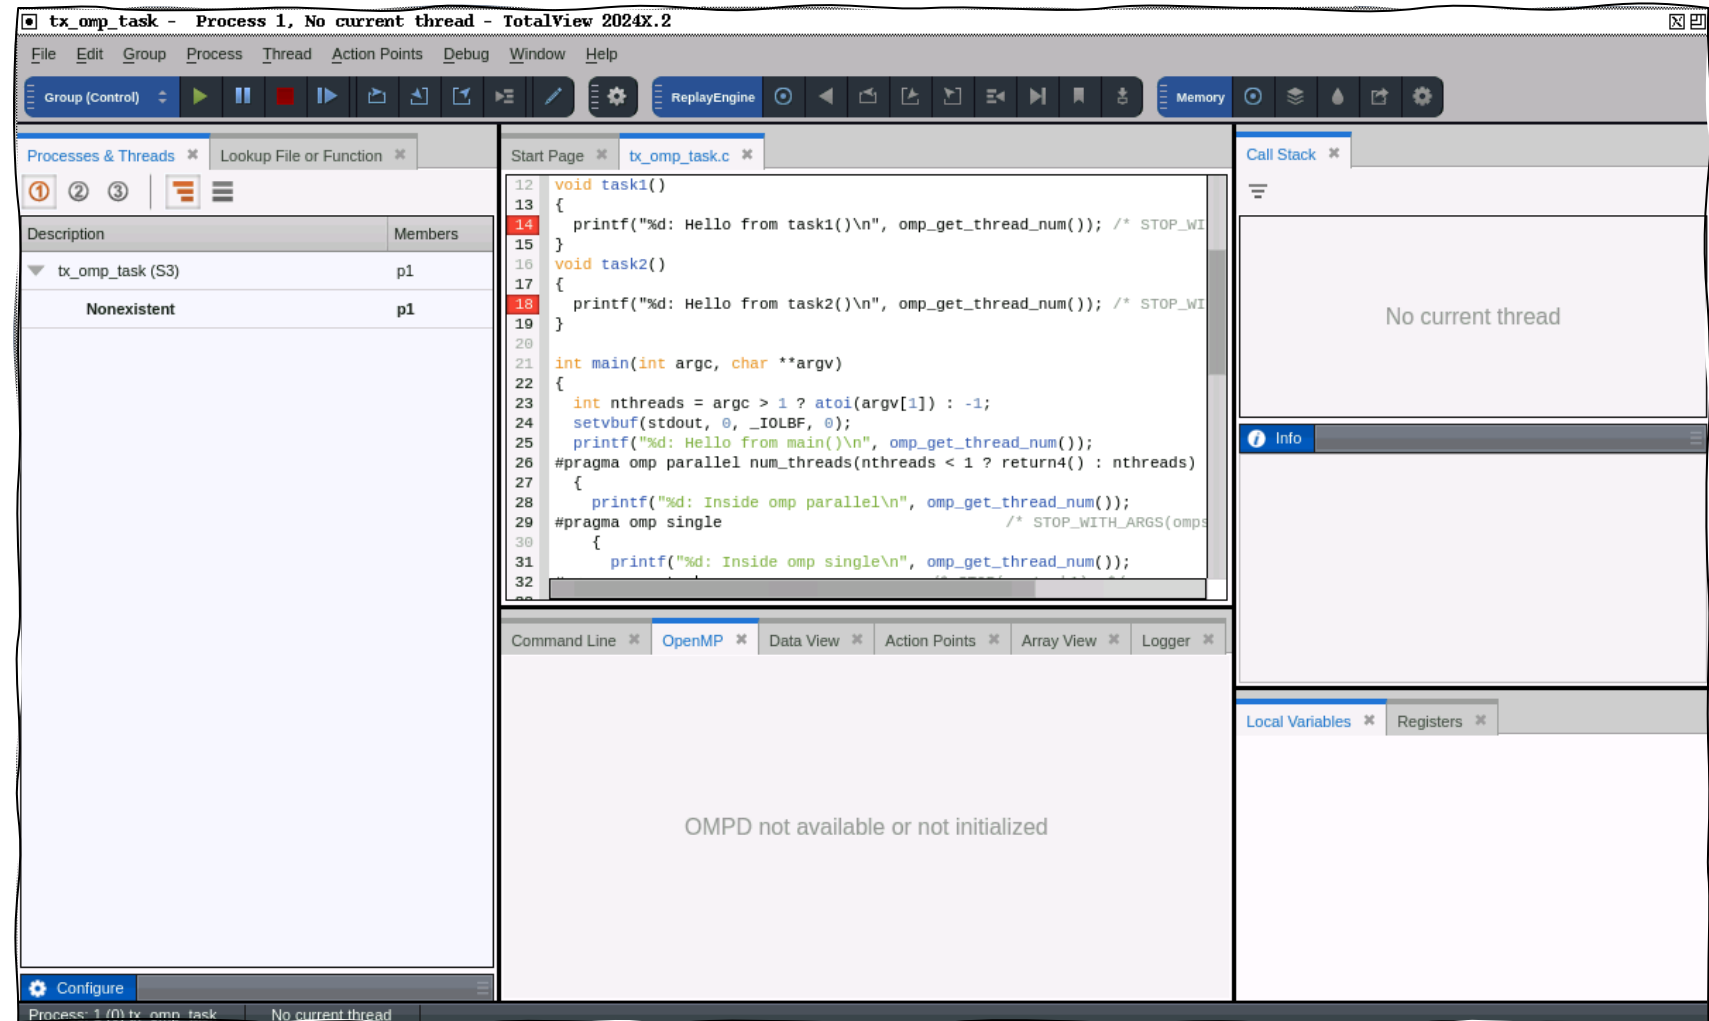
Enabling OMPD Support (TV 2024.2)

- TotalView OMPD support requires compiler support
 - Clang 15+, HPE CCE 17+, AMD Clang/Flang 17+
 - OMPD support is still maturing
 - Special linking rules might apply (check the doc)
- Set OMPD environment variable
 - **OMP_DEBUG=enabled**
 - MPI + OpenMP codes require setting the environment variable in the MPI processes
 - Use to propagate **OMP_DEBUG** setting
- Select “Enable OpenMP Debugging”
 - Good for non-MPI codes launched by the debugger
 - Does not work for processes that are attached to



Example OpenMP Debugging Session (TV 2024.2)

- Example using AMD Clang 17
 - LLVM Clang works but has some issues
 - HPE CCE works but has some issues
 - gcc / gfortran works but OMPD support not tested
- Set stop-thread breakpoints inside parallel / task region
- Make sure the whole process is stopped so that the OpenMP runtime is in a consistent state



OpenMP > Regions (TV 2024.2)

- Displays **parallel** and **task** regions
 - Aggregated view of all OpenMP threads
- “Regions” tab shows
 - Source-code line-number of OMP region
 - OMP implicit or explicit task function name
 - OpenMP threads that are in the region

Task Line	Task Function	Members
▼ / (Parallel Regions)		1:4[p1.1-4]
tx_omp_task.c#26	.omp_outlined..7	1:4[p1.1-4]
▼ / (Task Regions)		1:4[p1.1-4]
▼ tx_omp_task.c#26	.omp_outlined..7	1:4[p1.1-4]
tx_omp_task.c#32	.omp_task_entry.	1:1[p1.1]
tx_omp_task.c#34	.omp_task_entry..6	1:1[p1.3]

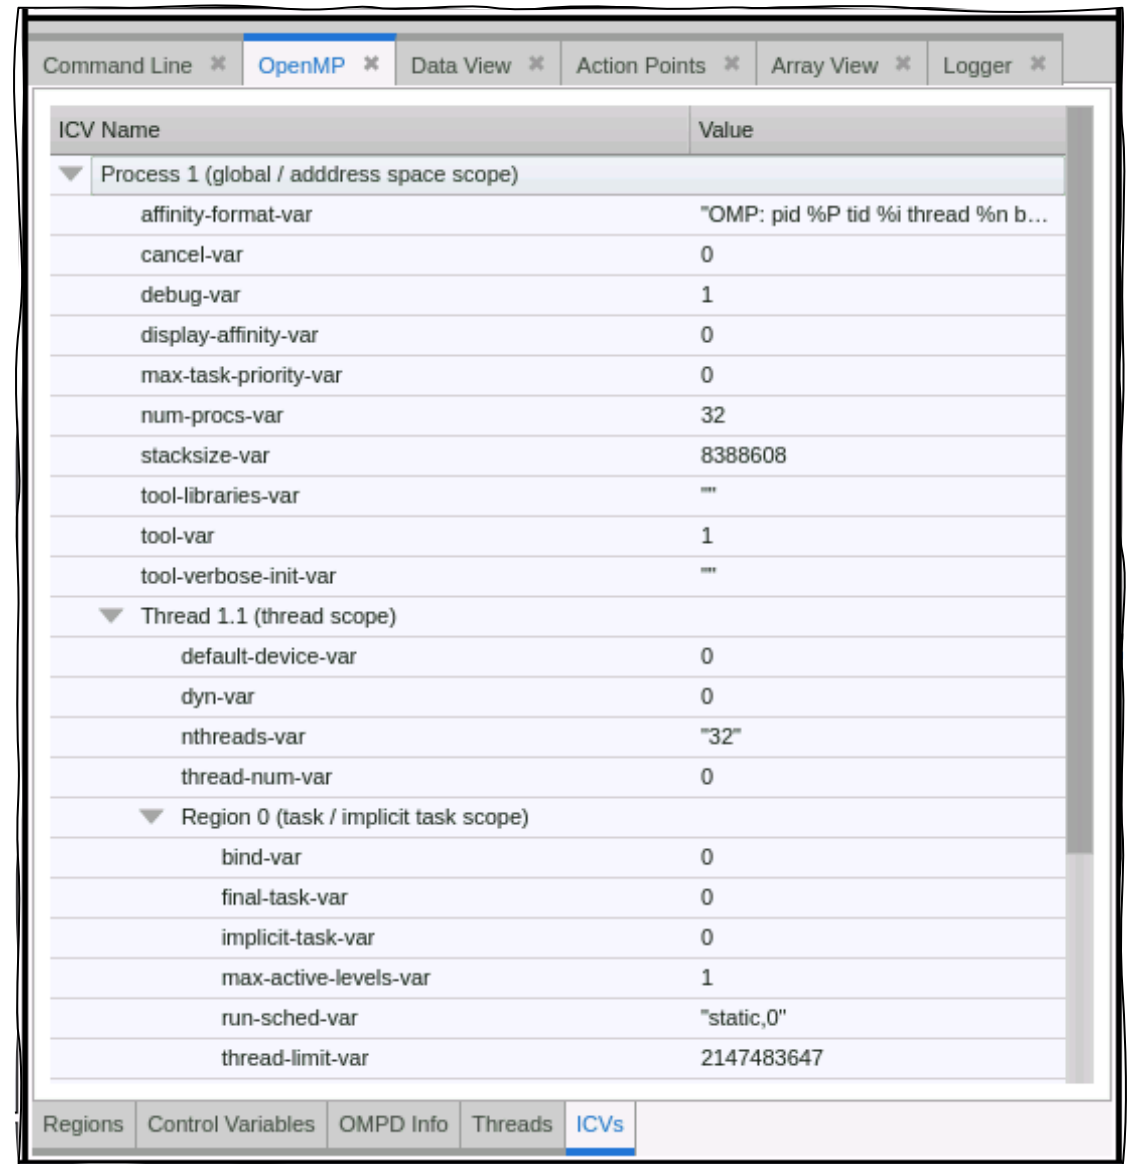
OpenMP Threads (TV 2024.2)

- Thread-oriented view of OMP threads
 - For the focus process
 - Current state plus nest of OMP generating task regions
- “Threads” tab shows
 - Debugger process/thread ID and OMP thread-num
 - Current state of OMP thread / region #
 - Wait ID / Parent (encountering thread) ID
 - Region flags
 - “i” implicit vs. explicit task
 - “p” active parallel region
 - “f” final task
 - Task function and source-code line-number
 - Runtime frame information (not shown)

Thread ID	OMP	State/Region#	Wait Id/Pare	Flags	Task Function	Task Line
▼ p1.1	0	work_parallel	0x0	-p-	.omp_task_entry.	[tx_omp_task.c#32]
		region 0		-p-	.omp_task_entry.	[tx_omp_task.c#32]
		region 1	p1.4	ip-	.omp_outlined..7	[tx_omp_task.c#26]
		region 2	p1.1	i--	<unavailable>	<unavailable>
▶ p1.2	1	wait_barrier	0x0	ip-	.omp_outlined..7	[tx_omp_task.c#26]
▶ p1.3	2	work_parallel	0x0	-p-	.omp_task_entry..6	[tx_omp_task.c#34]
▶ p1.4	3	wait_barrier	0x0	ip-	.omp_outlined..7	[tx_omp_task.c#26]

OpenMP > ICVs (TV 2024.2)

- Hierarchical view of OpenMP internal control variables
 - For the focus process
- Organized by OpenMP scope
 - Global / address space scope
 - Thread scope
 - Parallel region scope
 - Task / implicit task scope

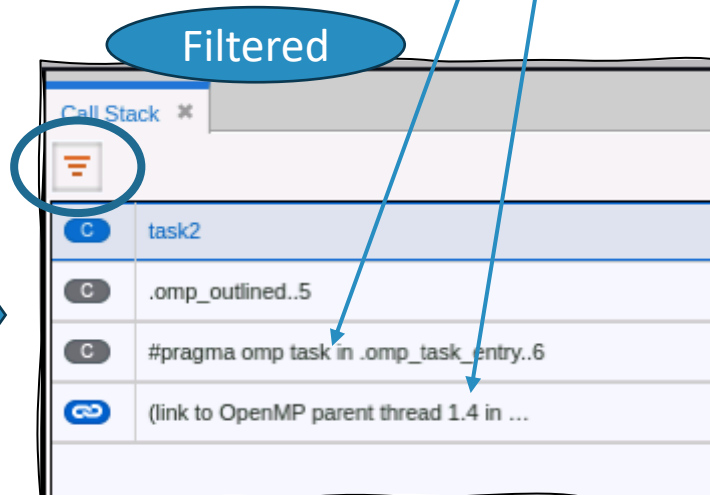


ICV Name	Value
▼ Process 1 (global / address space scope)	
affinity-format-var	"OMP: pid %P tid %i thread %n b...
cancel-var	0
debug-var	1
display-affinity-var	0
max-task-priority-var	0
num-procs-var	32
stacksize-var	8388608
tool-libraries-var	""
tool-var	1
tool-verbose-init-var	""
▼ Thread 1.1 (thread scope)	
default-device-var	0
dyn-var	0
nthreads-var	"32"
thread-num-var	0
▼ Region 0 (task / implicit task scope)	
bind-var	0
final-task-var	0
implicit-task-var	0
max-active-levels-var	1
run-sched-var	"static,0"
thread-limit-var	2147483647

OpenMP Stack Transformations



- Select the filter icon to filter
- OpenMP thread-stack transformations
 - Filters-out OMP runtime frames
 - Annotates parallel/task regions
 - Inserts parent thread links (click to focus on parent thread)



OpenMP Debugging Caveats

- OpenMP support prior to TotalView 2024.2 is a *prototype*
 - Not fully supported w/ limited compiler support
 - Has bugs and other problems with OpenMP displays
- LLVM/Clang-based compilers do not do a good at generating DWARF debug information
 - Program variables inside regions are all marked artificial, so TotalView does not display them
 - Use TotalView “**-compiler_vars**” option to display program variables, but compiler-generated variables are also displayed
 - Parallel “for” loop variables do not have correct values
 - Many other DWARF debug information problems exist
- GNU compilers seem to do a much better job in general
- Linking applications with OMPD support varies by compiler
 - Check the documentation

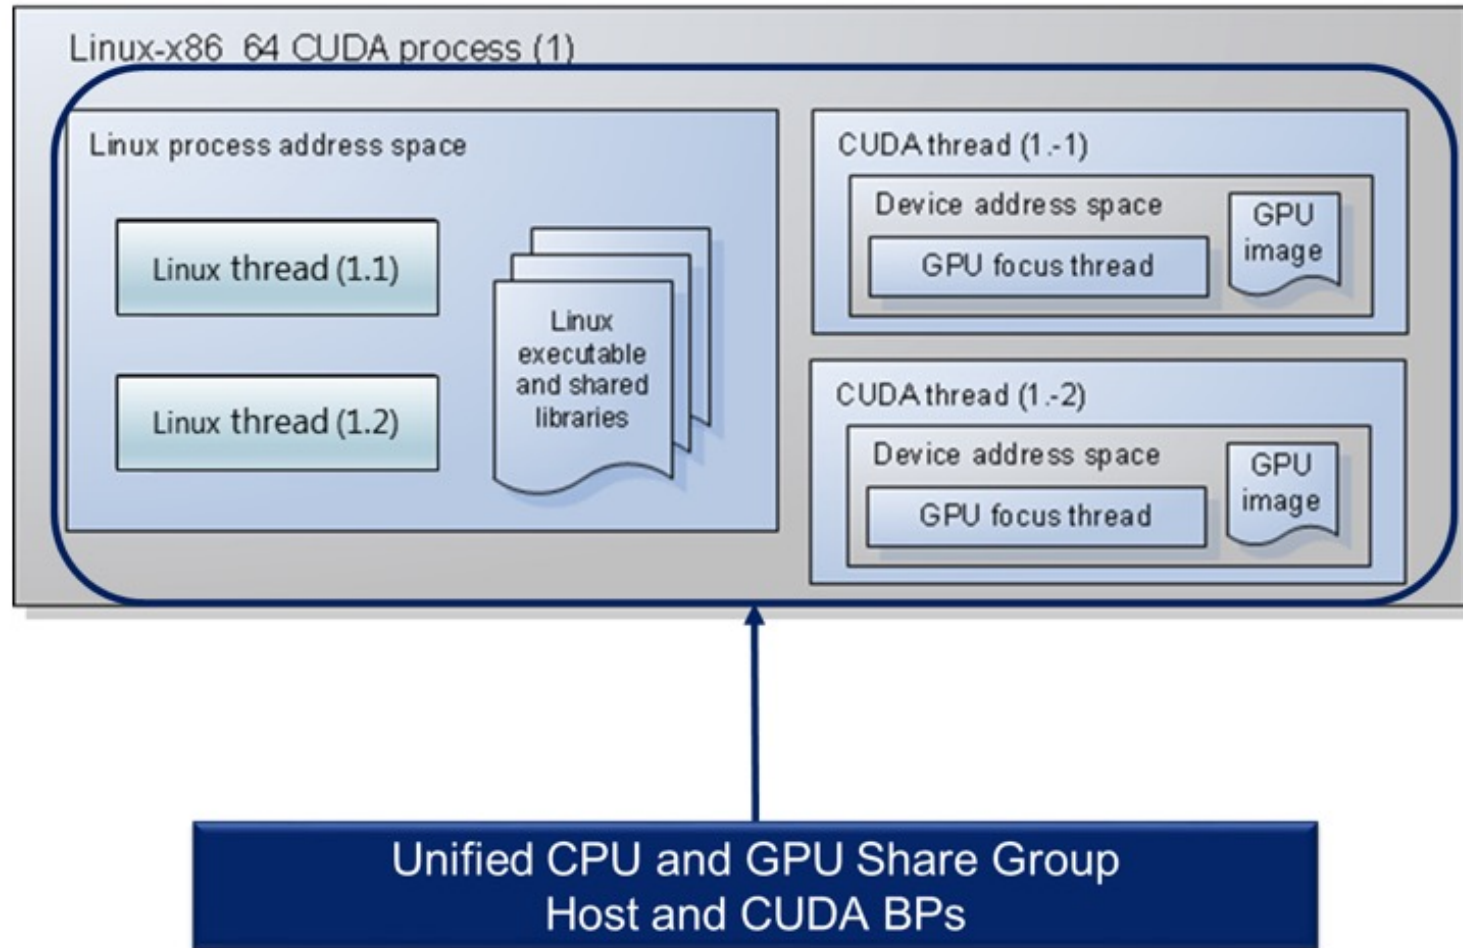
Debugging NVIDIA GPUs and CUDA with TotalView

TotalView for the NVIDIA[®] GPU Accelerator

- NVIDIA Tesla, Fermi, Kepler, Pascal, Volta, Turing, Ampere, Hopper
- NVIDIA CUDA 9.2, 10, 11 and 12
 - With support for Unified Memory
- NVIDIA and Cray OpenACC support
- 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
 - Support for types and separate memory address spaces
 - GPU Status view reveals what is running where

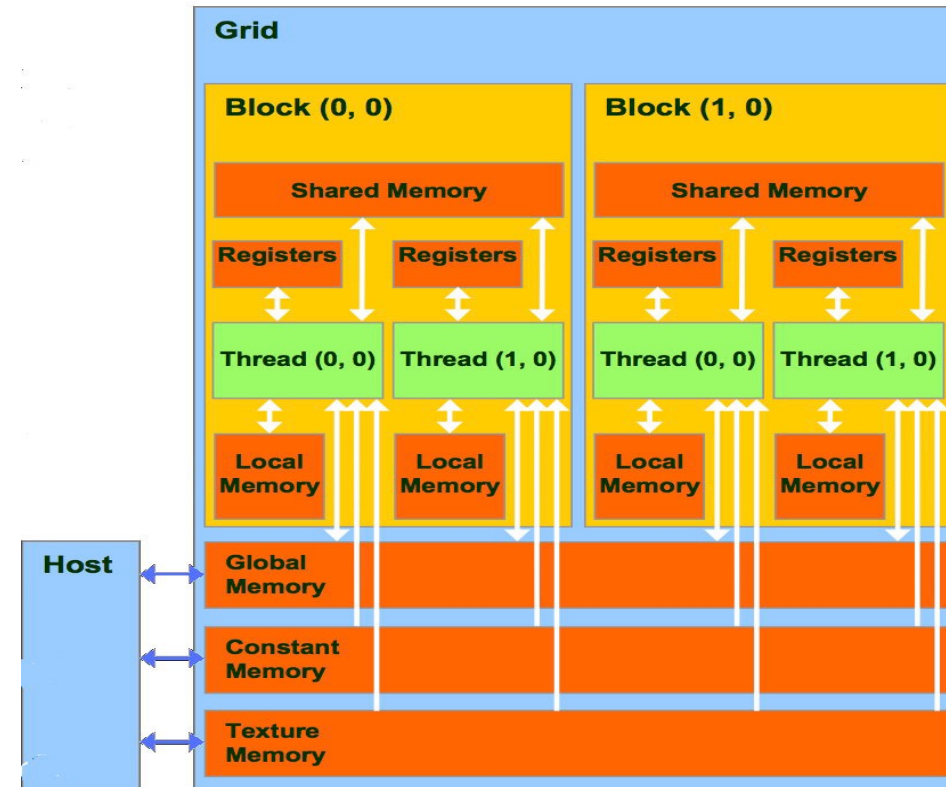


TotalView CUDA Debugging Model



GPU Memory Hierarchy

- Hierarchical memory
 - Local (thread)
 - Local
 - Register
 - Shared (block)
 - Shared Memory
 - Global (GPU)
 - Global
 - Constant
 - Texture
 - System (host)



Supported Type Storage (aka, Address Space) Qualifiers

@generic	An offset within generic storage
@frame	An offset within frame storage
@global	An offset within global storage
@local	An offset within local storage
@parameter	An offset within parameter storage
@iparam	Input parameter
@oparam	Output parameter
@shared	An offset within shared storage
@surface	An offset within surface storage
@texsampler	An offset within texture sampler storage
@texture	An offset within texture storage
@rtvar	Built-in runtime variables
@register	A PTX register name
@sregister	A PTX special register name

Control of Threads and Warps

- Warps advance synchronously
 - They share a PC
- Single step operation advances all GPU threads in the same warp
- Stepping over a `__syncthreads()` call will advance all relevant threads
- To advance more than one warp
 - Continue, possibly after setting a new breakpoint
 - Select a line and “Run To”

NVIDIA GPU and CUDA Parallelization

- CUDA uses the single instruction multiple thread (SIMT) model of parallelization.
- CUDA GPUs made up of many computing units called cores
 - Cores includes an arithmetic logic unit (ALU) and a floating-point unit (FPU).
- Cores collected into groups called streaming multiprocessors (SMs).
- Computing tasks are parallelized by breaking them into numerous subtasks called threads.
- Threads are organized into blocks.
- Blocks are divided into warps whose size matches the number of cores in an SM.
- Each warp gets assigned to a particular SM for execution. GPUs have one or more SMs.
- SM control unit directs each of its cores to execute the same instructions simultaneously for each thread in the assigned warp.

Compiling for CUDA debugging

When compiling an NVIDIA CUDA program for debugging, it is necessary to pass the **-g -G** options to the `nvcc` compiler driver. These options disable most compiler optimization and include symbolic debugging information in the driver executable file, making it possible to debug the application.

```
% /usr/local/bin/nvcc -g -G -c tx_cuda_matmul.cu -o tx_cuda_matmul.o
```

```
% /usr/local/bin/nvcc -g -G -Xlinker=-R/usr/local/cuda/lib64 \  
tx_cuda_matmul.o -o tx_cuda_matmul
```

```
% ./tx_cuda_matmul
```

```
A:
```

```
[ 0][ 0] 0.000000
```

```
...output deleted for brevity...
```

```
[ 1][ 1] 131.000000
```

Compiling for a specific GPU architecture (avoids JIT'ing from PTX)

Compiling for Ampere

`-gencode arch=compute_80,code=sm_80`

Compiling for Volta

`-gencode arch=compute_70,code=sm_70`

Compiling for Pascal

`-gencode arch=compute_60,code=sm_60`

Compiling for Kepler

`-gencode arch=compute_35,code=sm_35`

Compiling for Fermi and Tesla

`-gencode arch=compute_20,code=sm_20 -gencode arch=compute_10,code=sm_10`

Compiling for Fermi

`-gencode arch=compute_20,code=sm_20`

A TotalView Session with CUDA

A standard TotalView installation supports debugging CUDA applications running on both the host and GPU processors.

TotalView dynamically detects a CUDA install on your system. To start the TotalView GUI or CLI, provide the name of your CUDA host executable to the `totalview` or `totalviewcli` command.

For example, to start the TotalView GUI on the sample program, use the following command:

```
% totalview tx_cuda_matmul
```

* This example is just a single node, no MPI application

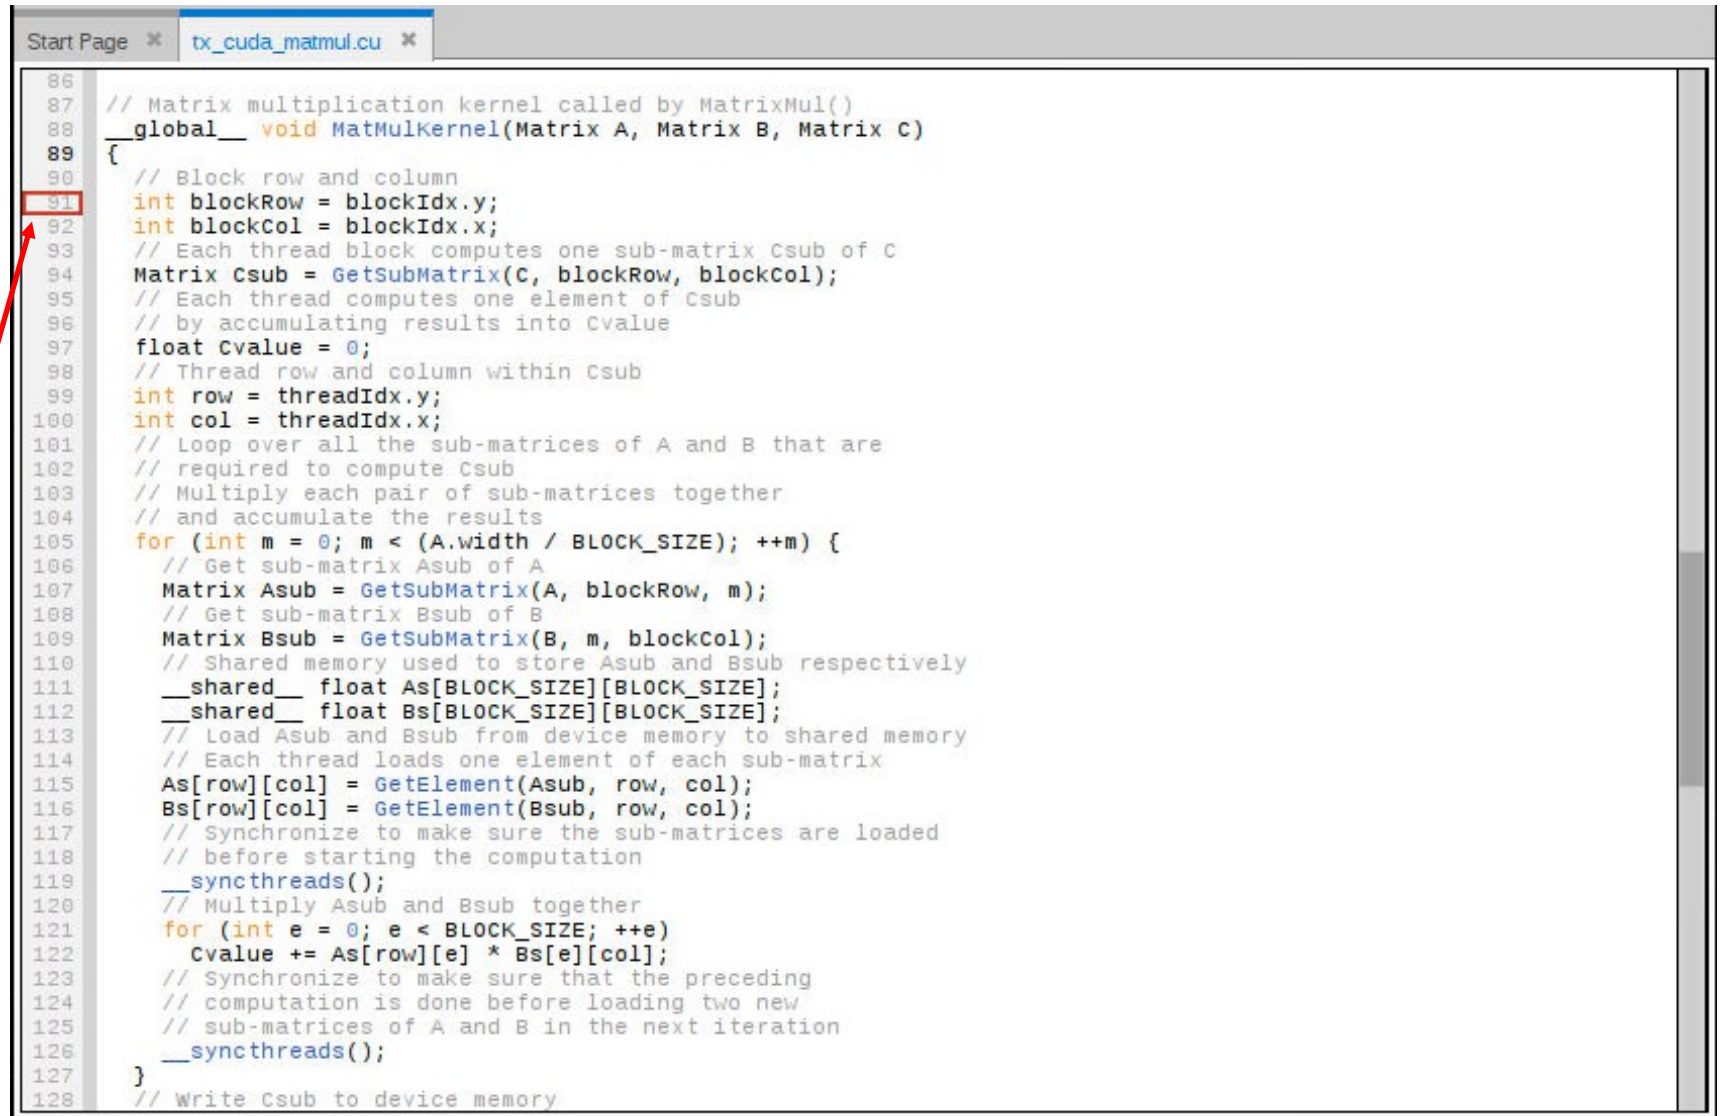
Source View Opened on CUDA host code

```
Start Page * tx_cuda_matmul.cu *
139 Matrix A;
140 A.width = width_;
141 A.height = height_;
142 A.stride = width_;
143 A.elements = (float*) malloc(sizeof(*A.elements) * width_ * height_);
144 for (int row = 0; row < height_; row++)
145     for (int col = 0; col < width_; col++)
146         A.elements[row * width_ + col] = row * 10.0 + col;
147 return A;
148 }
149
150 static void
151 print_Matrix (Matrix A, const char *name)
152 {
153     printf("%s:\n", name);
154     for (int row = 0; row < A.height; row++)
155         for (int col = 0; col < A.width; col++)
156             printf ("%5d[%5d] %f\n", row, col, A.elements[row * A.stride + col]);
157 }
158
159 // Multiply an m*n matrix with an n*p matrix results in an m*p matrix.
160 // Usage: tx_cuda_matmul [ m [ n [ p ] ] ]
161 // m, n, and p default to 1, and are multiplied by BLOCK_SIZE.
162 int main(int argc, char **argv)
163 {
164     cudaSetDevice(0);
165     const int m = BLOCK_SIZE * (argc > 1 ? atoi(argv[1]) : 1);
166     const int n = BLOCK_SIZE * (argc > 2 ? atoi(argv[2]) : 1);
167     const int p = BLOCK_SIZE * (argc > 3 ? atoi(argv[3]) : 1);
168     Matrix A = cons_Matrix(m, n);
169     Matrix B = cons_Matrix(n, p);
170     Matrix C = cons_Matrix(m, p);
171     MatMul(A, B, C);
172     print_Matrix(A, "A");
173     print_Matrix(B, "B");
174     print_Matrix(C, "C");
175     return 0;
176 }
177
178 /*
179  * Update log
180  *
181  * Feb 25 2015 NYP: Removed forceinline , it is making cli too fast
```

Set Breakpoints in CUDA Kernel Code Before Launch

Set breakpoints in the CUDA or OpenMP TARGET region code before you start the process.

Hollow breakpoint indicates a breakpoint will be set when the code is loaded onto the GPU.



```
Start Page * tx_cuda_matmul.cu *
86
87 // Matrix multiplication kernel called by MatrixMul()
88 __global__ void MatMulKernel(Matrix A, Matrix B, Matrix C)
89 {
90     // Block row and column
91     int blockRow = blockIdx.y;
92     int blockCol = blockIdx.x;
93     // Each thread block computes one sub-matrix Csub of C
94     Matrix Csub = GetSubMatrix(C, blockRow, blockCol);
95     // Each thread computes one element of Csub
96     // by accumulating results into Cvalue
97     float Cvalue = 0;
98     // Thread row and column within Csub
99     int row = threadIdx.y;
100    int col = threadIdx.x;
101    // Loop over all the sub-matrices of A and B that are
102    // required to compute Csub
103    // Multiply each pair of sub-matrices together
104    // and accumulate the results
105    for (int m = 0; m < (A.width / BLOCK_SIZE); ++m) {
106        // Get sub-matrix Asub of A
107        Matrix Asub = GetSubMatrix(A, blockRow, m);
108        // Get sub-matrix Bsub of B
109        Matrix Bsub = GetSubMatrix(B, m, blockCol);
110        // Shared memory used to store Asub and Bsub respectively
111        __shared__ float As[BLOCK_SIZE][BLOCK_SIZE];
112        __shared__ float Bs[BLOCK_SIZE][BLOCK_SIZE];
113        // Load Asub and Bsub from device memory to shared memory
114        // Each thread loads one element of each sub-matrix
115        As[row][col] = GetElement(Asub, row, col);
116        Bs[row][col] = GetElement(Bsub, row, col);
117        // Synchronize to make sure the sub-matrices are loaded
118        // before starting the computation
119        __syncthreads();
120        // Multiply Asub and Bsub together
121        for (int e = 0; e < BLOCK_SIZE; ++e)
122            Cvalue += As[row][e] * Bs[e][col];
123        // Synchronize to make sure that the preceding
124        // computation is done before loading two new
125        // sub-matrices of A and B in the next iteration
126        __syncthreads();
127    }
128    // Write Csub to device memory
```

Stopped at a Breakpoint in CUDA Kernel Code

- Bold line numbers indicate source code lines where the compiler generated code, which are good places to set breakpoints

The screenshot shows the TotalView GPU debugger interface. The main window displays CUDA kernel code with line numbers 78 to 108. Line 91, `int blockRow = blockIdx.y;`, is highlighted in yellow and has a red breakpoint icon next to it. The code includes comments and function definitions for memory management and matrix multiplication. A GPU toolbar at the top shows 'GPU (Logical)', 'Block', and 'Thread' selectors. A GPU focus thread selector on the right allows changing block and thread indices. A variable window at the bottom shows parameters A, B, and C, all of type Matrix @local. A PC arrow on the left points to the current instruction address.

VAR	Type	Value
Arguments		
A	Matrix @local	(Matrix @local)
B	Matrix @local	(Matrix @local)
C	Matrix @local	(Matrix @local)

CUDA thread IDs and Coordinate Spaces

Description	#P	#T	Members
tx_cuda_matmul (S3)	1	1	p1
▼ Breakpoint	1	1	p1.-1
▼ MatMulKernel	1	1	p1.-1
▼ tx_cuda_matmul.c...	1	1	p1.-1
1-1	1	1	p1.-1
▼ Stopped	1	3	p1.1-3
▶ __poll_nocancel	1	1	p1.3
▶ accept4	1	1	p1.2
▼ cuVDPAUCtxCreate	1	1	p1.1
▼ <unknown line>	1	1	p1.1
1.1	1	1	p1.1

Host thread IDs have a positive thread ID (p1.1)

CUDA thread IDs have a negative thread ID (p1.-1)

GPU Physical and Logical Focus Toolbars



Logical toolbar displays the Block and Thread coordinates.

Physical toolbar displays the Device number, Streaming Multiprocessor, Warp and Lane.

To view a CUDA host thread, select a thread with a positive thread ID in the Process and Threads view.

To view a CUDA GPU thread, select a thread with a negative thread ID, then use the GPU focus controls in the logical or physical toolbar to focus on a specific GPU thread or lane.

Displaying CUDA Program Variables

Name	Type	Thread ID	Value
▼ A	Matrix @local	1-1	(Matrix @local)
width	int	1-1	0x00000002 (2)
height	int	1-1	0x00000002 (2)
stride	int	1-1	0x00000002 (2)
▼ elements	float @generic *	1-1	0x7f724e800000 -> 0
*(elements)	@generic float	1-1	0
[Add New Expression]			

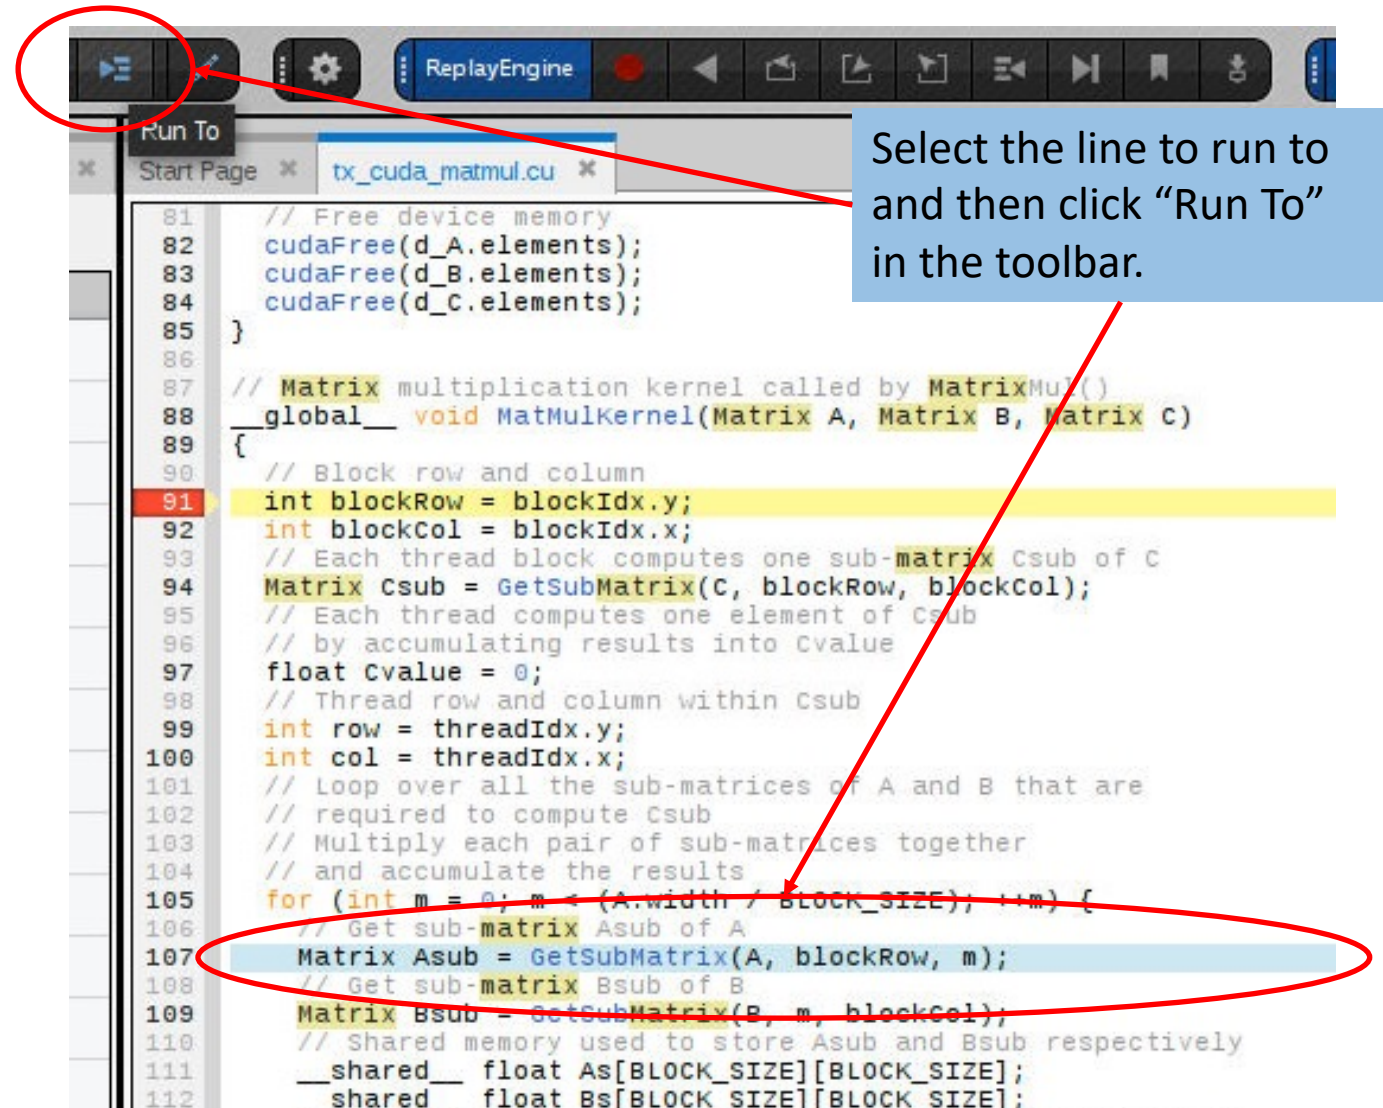
@local type qualifier indicates that variable A is in local storage.

"elements" is a pointer to a float in @generic storage.

- The identifier @local is a TotalView built-in type storage qualifier that tells the debugger the storage kind of "A" is local storage.
- The debugger uses the storage qualifier to determine how to locate A in device memory

Stepping GPU Code

- Single-step operations advance all the GPU hardware lanes in the same warp
- Note that stepping operations Step and Next are slow in GPU code; the following is faster...
- To advance the execution of more than one warp, you may either:
 - Set a breakpoint and continue the process, or
 - Select a line number in the source pane and select “Run To”.



GPU Status View

Displays the state of all the GPUs being debugged.

Fully configurable to allow aggregating, sorting and filtering based on physical or logical attributes.

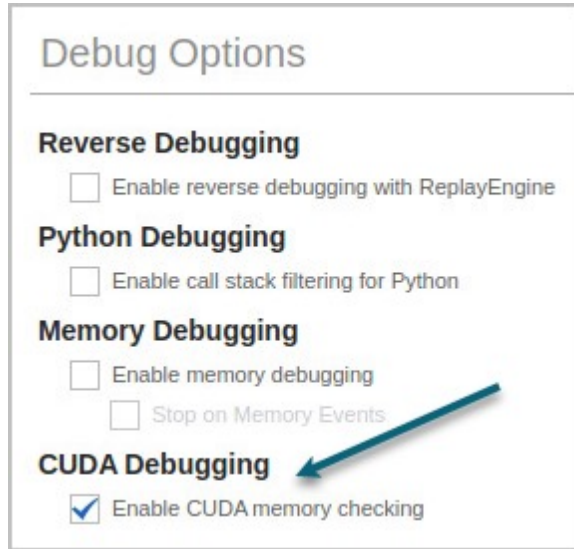
The screenshot shows a debugger window with the 'GPU Status' tab selected. The top pane displays C++ code for a matrix multiplication kernel. The bottom pane shows a tree view of the GPU hierarchy and a table of thread status.

```
122     Cvalue += As[row][e] * Bs[e][col];
123     // Synchronize to make sure that the preceding
124     // computation is done before loading two new
125     // sub-matrices of A and B in the next iteration
126     __syncthreads();
127 }
128 // Write Csub to device memory
129 // Each thread writes one element
```

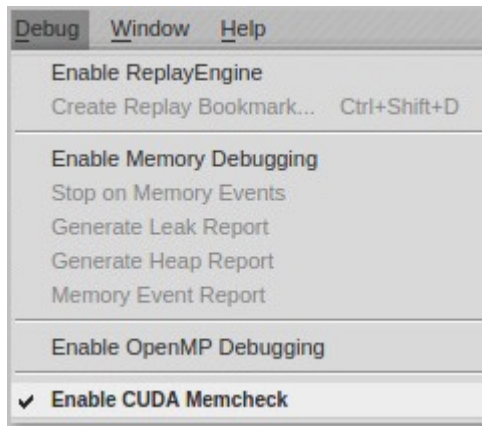
Variables	Status
▼ Devices	
▼ Device(0)	
▼ SM(0)	
▼ Warp(0)	
▼ Lane(\$lane)	for lane in {0 1 2 3}
	Function = MatMulKernel
	State = stopped

Frame: MatMulKernel File: /home/dstewart/cuda/tx_cuda_matmul.cu Line: 115

Enabling CUDA Memory Checker Feature



From the Program Session Dialog



From the Debug Menu

Demo

The screenshot displays the TotalView 2022 IDE interface during a debugging session. The main window shows the source code of a CUDA kernel in `vecAddWrapperCXX.cu`. A breakpoint is set at line 10, which is highlighted in yellow. The code defines a `vecAdd` kernel function that calculates a thread ID and performs a self-addition on a float array element.

The left sidebar contains several panels:

- Processes & Th...:** A table showing the execution context. The 'Running' process is expanded to show 4 threads (1.1, 1.2, 1.3, 1.4).
- Action Points & Replay Bookmarks:** A table with two entries, both marked as 'Break' at the location `...perCXX.cu#10`.

The right sidebar contains:

- Call Stack:** Shows the current function `vecAdd`.
- Local Variables:** A table listing arguments `a`, `b`, `c`, and `n`, along with a register variable `id`.
- GPU Status:** A tree view showing the GPU configuration, including 3 SMs, 31 warps, and 31 lanes. The current state is `Function = vecAdd` and `State = breakpoint`.

The bottom status bar indicates the current execution state: `Rank: 1 (3351@128.55.64.195.mnet) @ TEMP@CUDA@.b.out Thread: 1.-2 (<<<(0,0,0),(0,0,0)>>>) - Breakpoint Frame: vecAdd File: ...al/u1/fjdesign/support-50919/vecAddWrapperCXX.cu Line: 10`

Debugging on Perlmutter

Debugging on Perlmutter (Things to Know)

- If you bind processes to GPUs using srun, the debugger cannot determine which GPUs the processes are using
 - SLURM's use of Linux control groups make it impossible
 - Workaround – Do not use the “--gpu-bind” option when debugging
- Watchpoints in GPU memory are not supported on NVIDIA GPUs, but CPU watchpoints are supported
- On Perlmutter, the environment variable “TVD_DISABLE_CRAY=1” must be set to disable using Cray CTI
 - “module load totalview” sets TVD_DISABLE_CRAY=1 on Perlmutter
 - SSH is used to instantiate the TV/MRNet tree
 - Requires passwordless SSH between nodes

Debugging on Perlmutter (Things to Know)

- Using SSH to between NERSC nodes can generate a lot of terminal output
 - Each SSH generates a long “NOTICE TO USERS” message
- The messages can be suppressed by adding the following lines to your “\$HOST/.ssh/config” file:

```
# The "LogLevel quiet" option stops the "NOTICE TO USERS" messages
Host *
    LogLevel quiet
```

- The above is not necessary, but it does reduce terminal output

Debugging on Perlmutter (Supported Start-ups)

- TotalView supports interactive and batch debugging sessions
- Interactive debugging sessions
 - Use **salloc** to allocate interactive nodes
 - Start TotalView on **srun** within the allocation
 - Allows restarting **srun** multiple times within the same allocation
- Batch debugging sessions
 - Use **sbatch** to submit a batch job
 - Batch script uses **tvconnect srun ...** to request a “reverse connect” to TotalView
 - Start TotalView on a login node and accept the “reverse connect” request
 - To restart srun multiple times, invoke **tvconnect srun** in a loop in the script

Debugging on Perlmutter (Interactive Start-up)

- Load the “totalview” module
 - `module load totalview`
- Allocate some nodes, for example
 - `salloc -A ntrain7 -C gpu -N 2 -G 8 -t 60 -q interactive_ss11`
- An interactive shell (bash, csh, etc.) will start inside the allocation
- Start **totalview** on **srun**, for example
 - `totalview -args srun -n 8 -G 8 -c 32 --cpu-bind=cores ./b.out`
 - Remember, “`--gpu-bind`” does not work, so do not use it while debugging

Debugging on Perlmutter (Batch Start-up)

- Example batch script using **tvconnect**

```
#!/bin/bash -x
#SBATCH -A nvendor
#SBATCH -C gpu
#SBATCH -N 2
#SBATCH -G 8
#SBATCH -t 30
#SBATCH --qos=debug

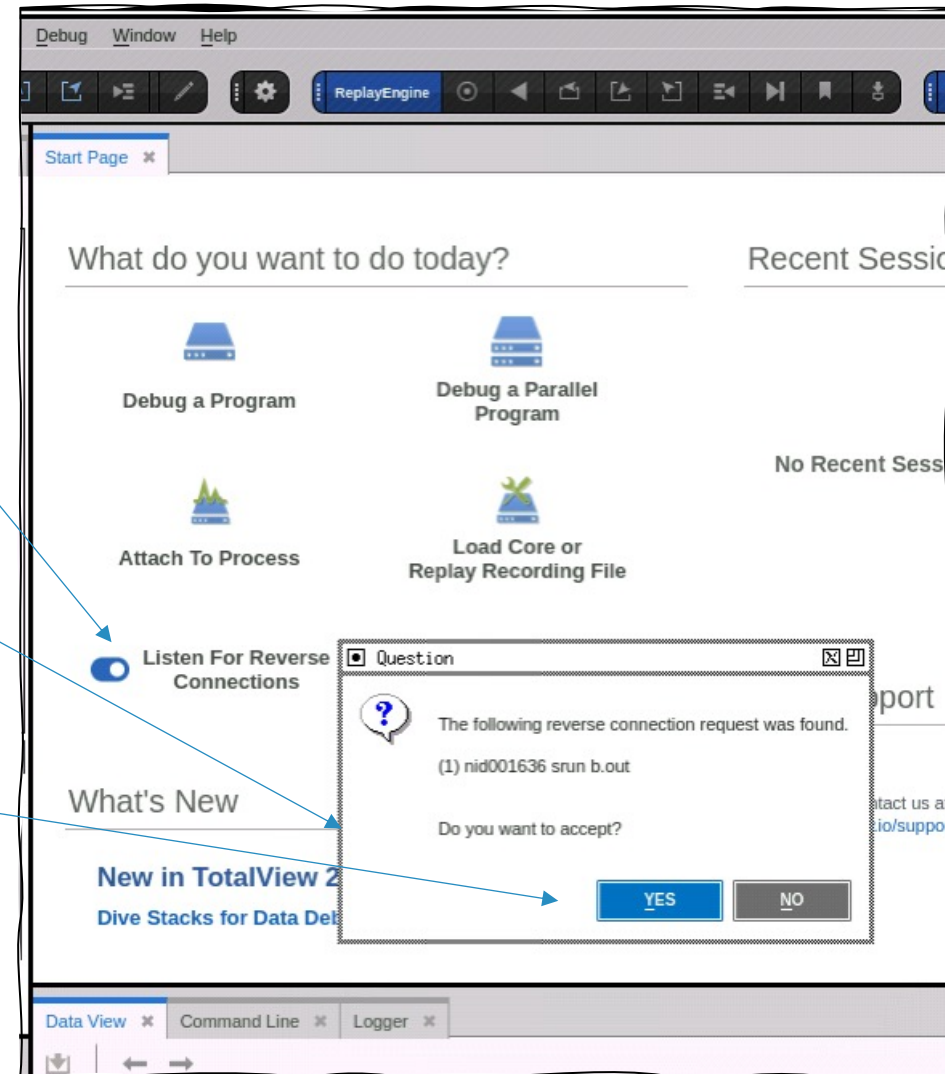
module load totalview
tvconnect srun b.out
```

- When the batch script starts, **tvconnect** blocks until a **totalview** accepts the reverse connect request
- On the login node, load the “totalview” module and start **totalview**

```
module load totalview
totalview
```

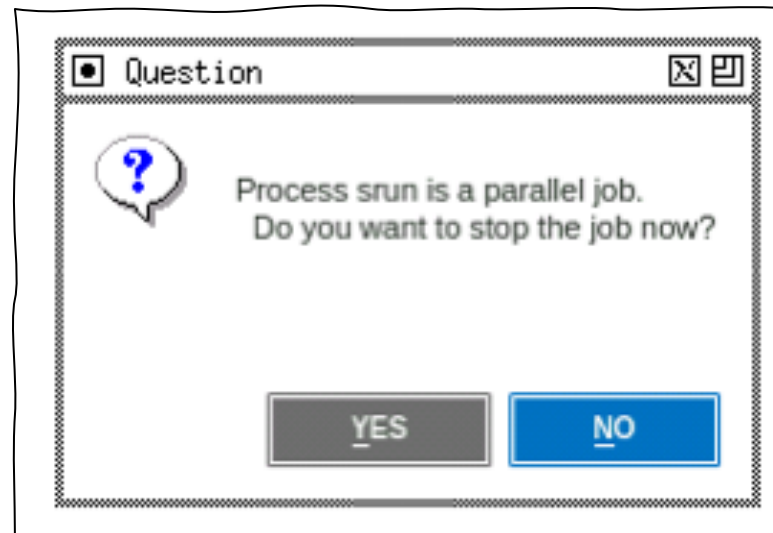
Debugging on Perlmutter (Batch Start-up)

- TotalView will “Listen For Reverse Connections” by default, but make sure the option is enabled
- When the batch script executes the **tvconnect** command, TotalView will post a dialog
- Select “Yes” to connect TotalView to the batch job



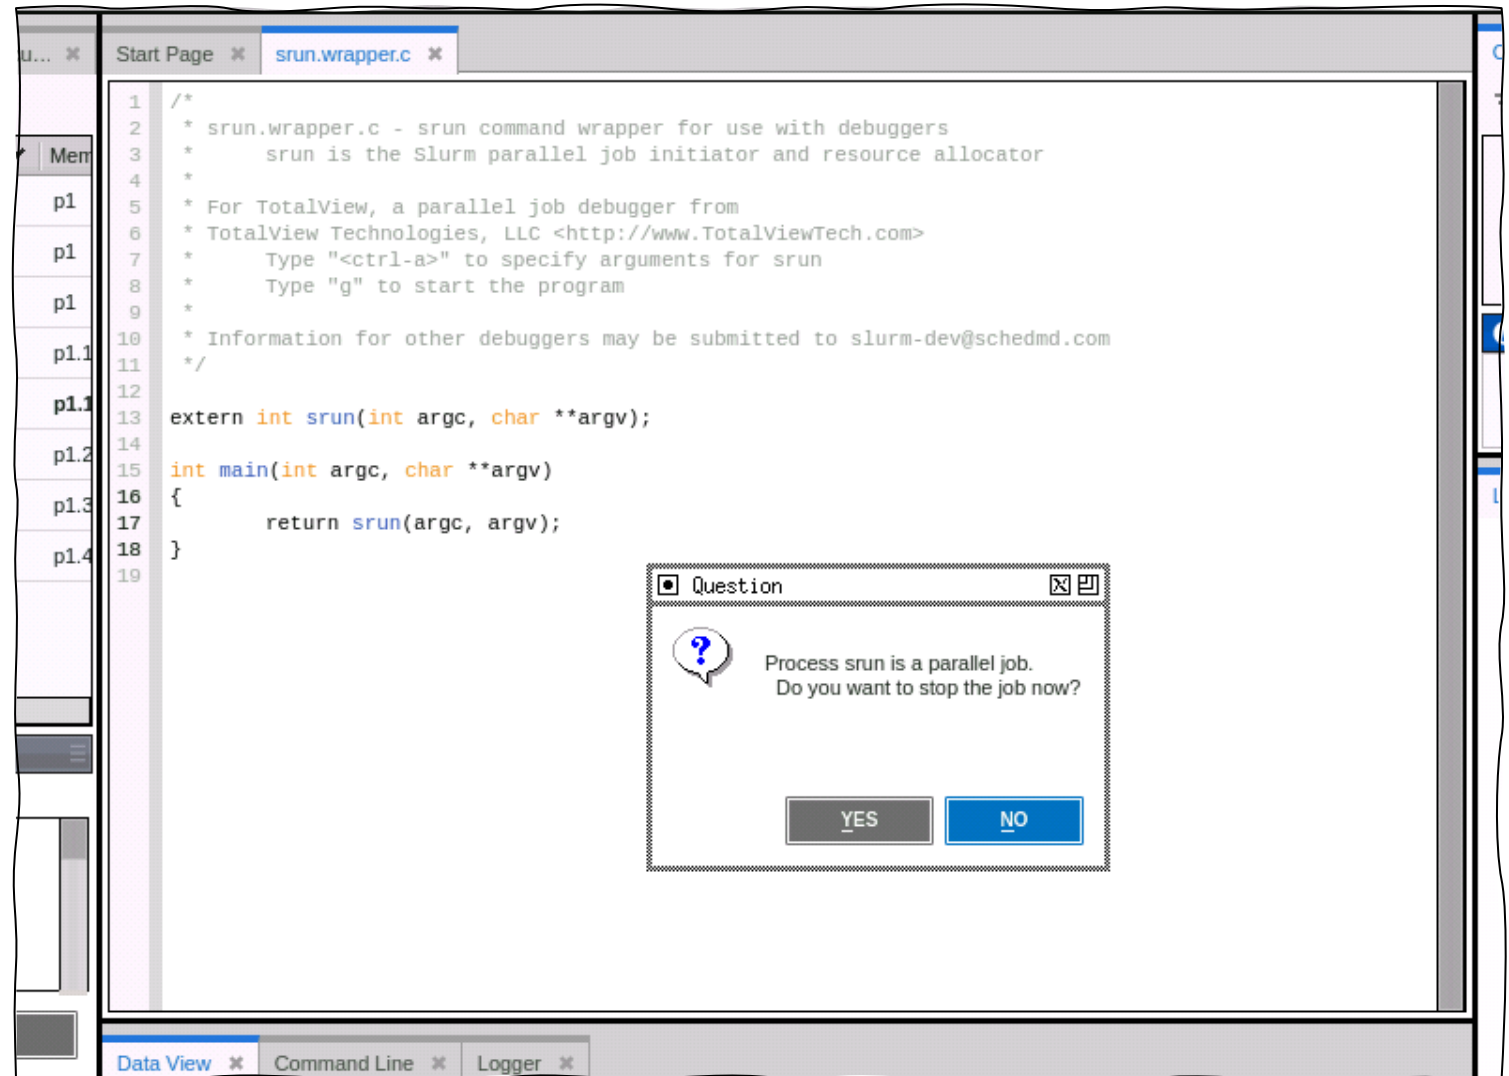
Debugging on Perlmutter (Common to Interactive/Batch)

- Once TotalView starts-up on **srun**, the following steps are common to interactive / batch debugging
- Typically
 - Select “**Go**” to start **srun**
 - **srun** will launch the parallel program
 - TotalView detects the parallel program launch and attaches to the MPI processes
- When the jobs goes parallel, TotalView will post a dialog

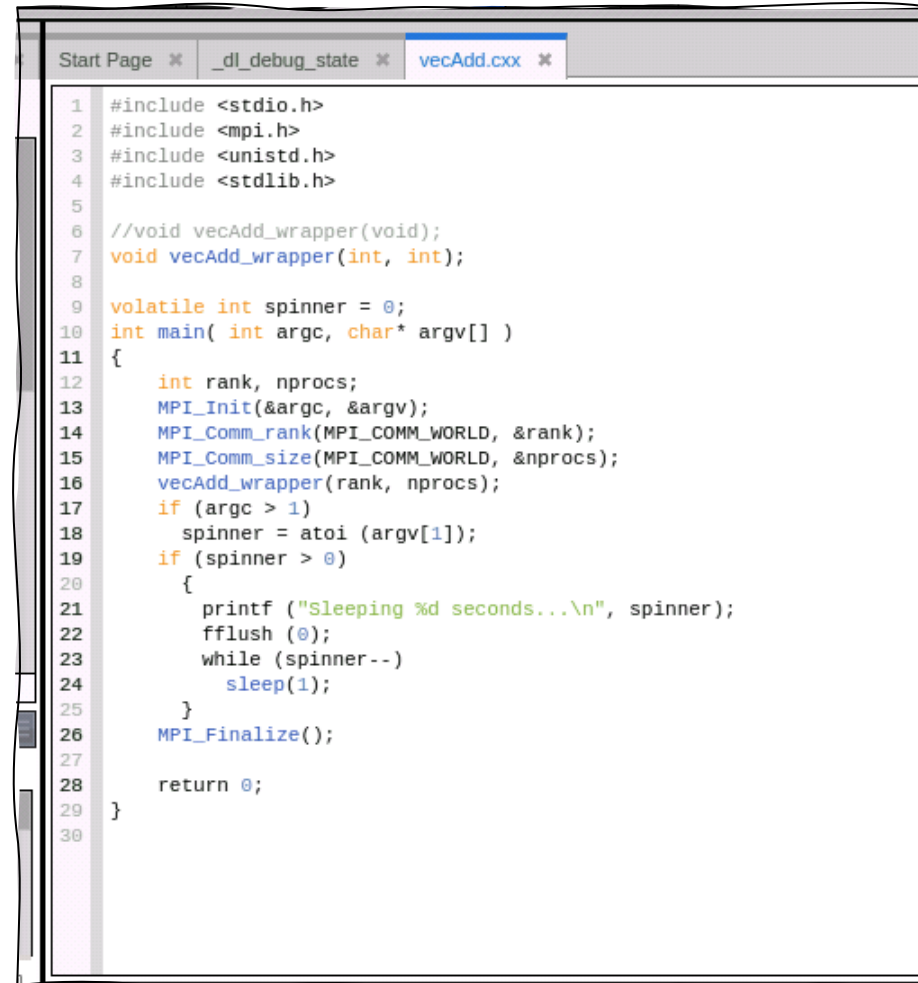


Stop the job when it goes parallel?

- Click “Yes” to stop the parallel job, which is useful if you want to
 - Navigate to source files / functions
 - Set breakpoints
- Click “No” to allow the job to run, which is useful if you
 - Have saved breakpoints from a previous session
 - Know the program is going to crash (SEGV, etc.)



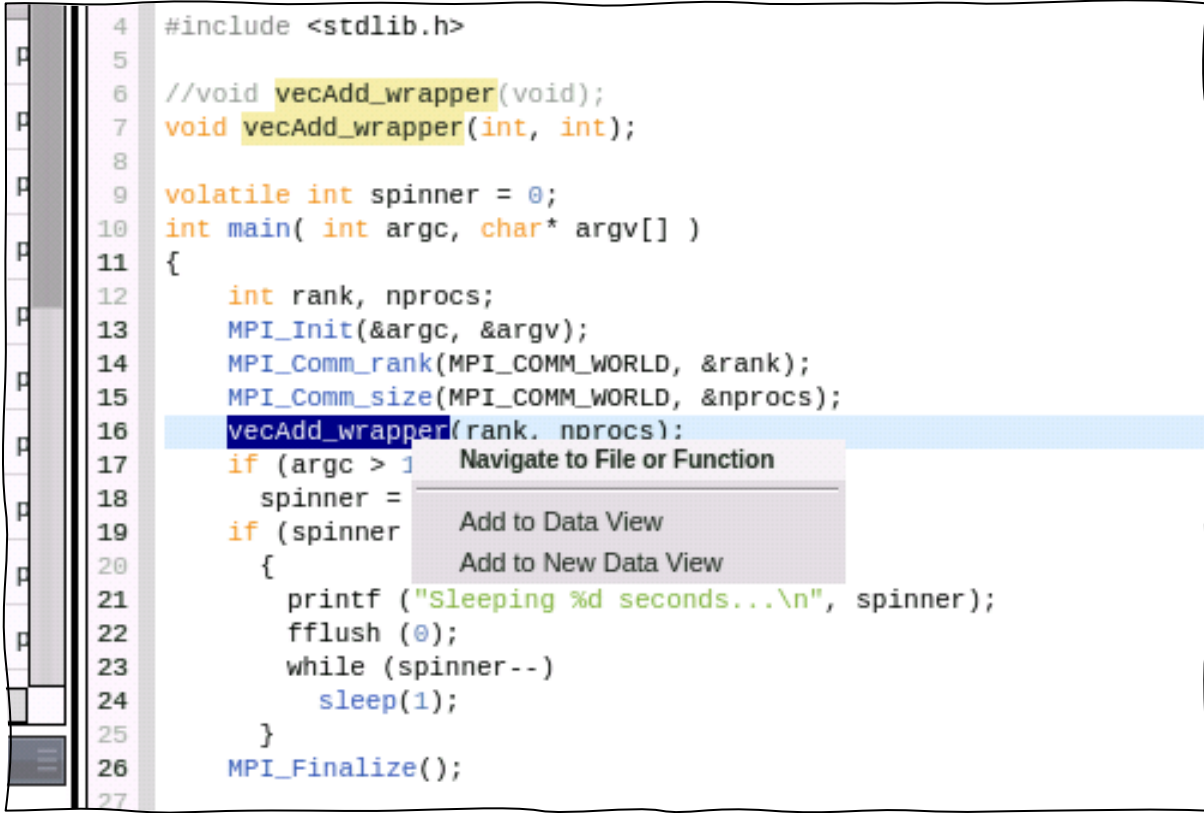
TotalView will focus on main() in rank 0



```
1 #include <stdio.h>
2 #include <mpi.h>
3 #include <unistd.h>
4 #include <stdlib.h>
5
6 //void vecAdd_wrapper(void);
7 void vecAdd_wrapper(int, int);
8
9 volatile int spinner = 0;
10 int main( int argc, char* argv[] )
11 {
12     int rank, nprocs;
13     MPI_Init(&argc, &argv);
14     MPI_Comm_rank(MPI_COMM_WORLD, &rank);
15     MPI_Comm_size(MPI_COMM_WORLD, &nprocs);
16     vecAdd_wrapper(rank, nprocs);
17     if (argc > 1)
18         spinner = atoi (argv[1]);
19     if (spinner > 0)
20     {
21         printf ("Sleeping %d seconds...\n", spinner);
22         fflush (0);
23         while (spinner--)
24             sleep(1);
25     }
26     MPI_Finalize();
27
28     return 0;
29 }
30
```

Navigate to a file or function you want to debug

```
4 #include <stdlib.h>
5
6 //void vecAdd_wrapper(void);
7 void vecAdd_wrapper(int, int);
8
9 volatile int spinner = 0;
10 int main( int argc, char* argv[] )
11 {
12     int rank, nprocs;
13     MPI_Init(&argc, &argv);
14     MPI_Comm_rank(MPI_COMM_WORLD, &rank);
15     MPI_Comm_size(MPI_COMM_WORLD, &nprocs);
16     vecAdd_wrapper(rank, nprocs);
17     if (argc > 1)
18         spinner =
19         if (spinner
20             {
21                 printf ("Sleeping %d seconds...\n", spinner);
22                 fflush (0);
23                 while (spinner--)
24                     sleep(1);
25             }
26     MPI_Finalize();
27
```

A screenshot of a code editor window showing a C program. The code is as follows:

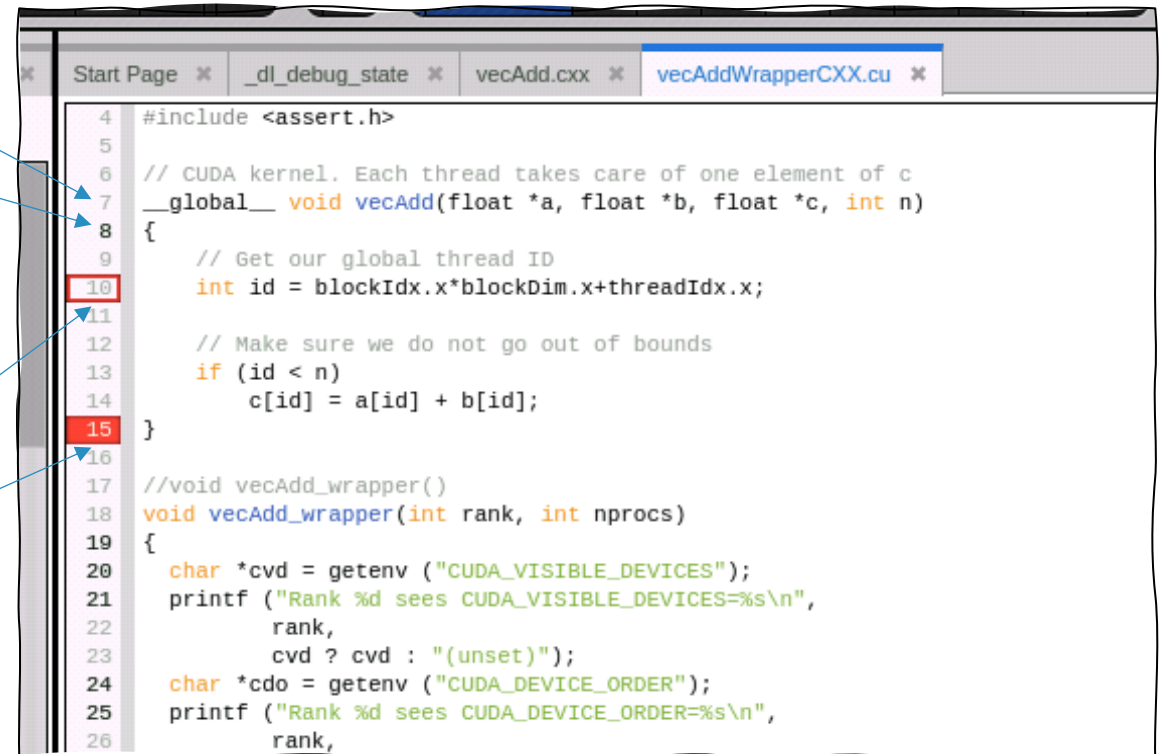
```
4 #include <stdlib.h>
5
6 //void vecAdd_wrapper(void);
7 void vecAdd_wrapper(int, int);
8
9 volatile int spinner = 0;
10 int main( int argc, char* argv[] )
11 {
12     int rank, nprocs;
13     MPI_Init(&argc, &argv);
14     MPI_Comm_rank(MPI_COMM_WORLD, &rank);
15     MPI_Comm_size(MPI_COMM_WORLD, &nprocs);
16     vecAdd_wrapper(rank, nprocs);
17     if (argc > 1)
18         spinner =
19     if (spinner
20         {
21             printf ("Sleeping %d seconds...\n", spinner);
22             fflush (0);
23             while (spinner--)
24                 sleep(1);
25         }
26     MPI_Finalize();
27
```

The line number 16 is highlighted in blue. A context menu is open over this line, with the following options:

- Navigate to File or Function
- Add to Data View
- Add to New Data View

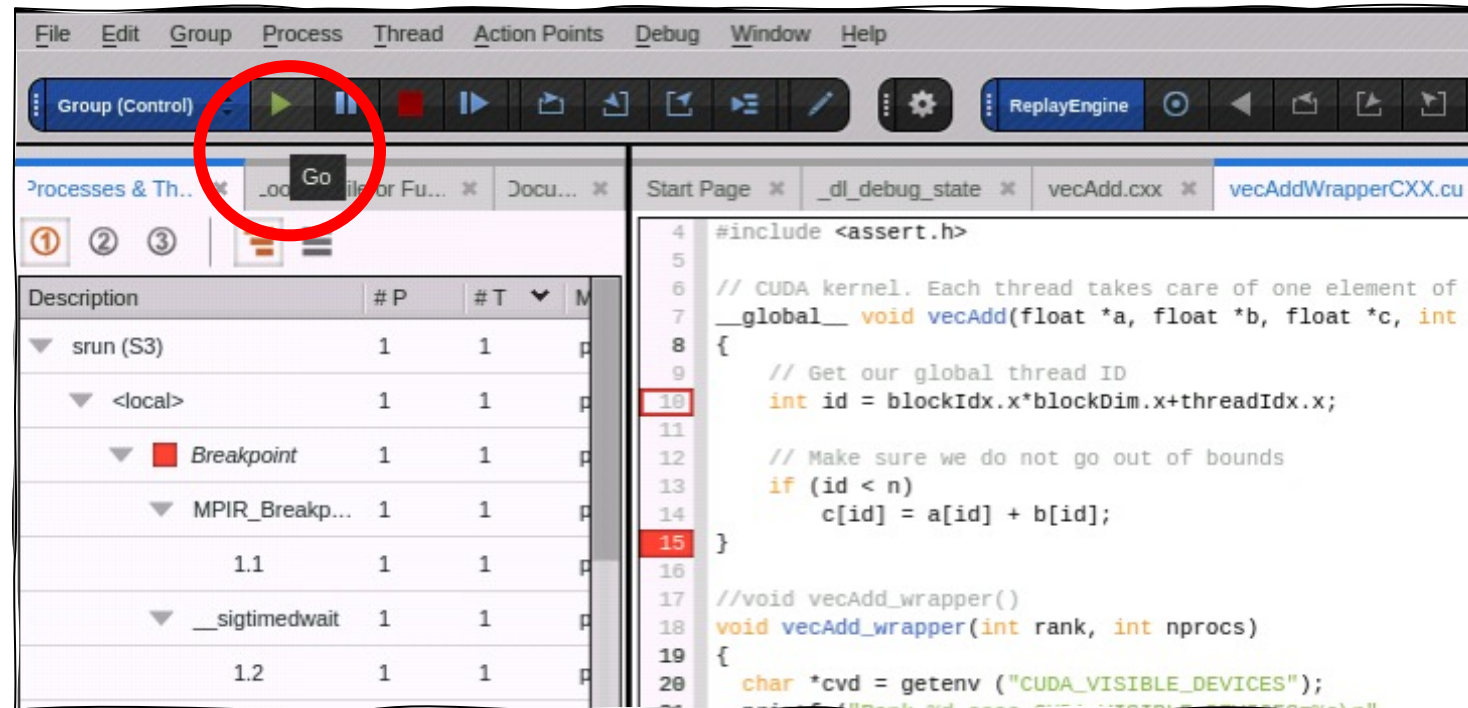
Find the CUDA kernel and select a line number to plant a breakpoint

- Line numbers indicate if there's code at that line
 - Pale line numbers indicate no code (yet)
 - Bold line numbers indicate code
- CUDA code is *dynamically* loaded at runtime, so TotalView does not have any debug information *until* the CUDA kernel is launched
- Select a line number in the CUDA kernel that will have CUDA code loaded
 - Hollow breakpoint markers indicate no code *yet*
 - Solid breakpoint markers indicate code
- Source line information for a source file is *unified* for both GPU and CPU code



```
4 #include <assert.h>
5
6 // CUDA kernel Each thread takes care of one element of c
7 __global__ void vecAdd(float *a, float *b, float *c, int n)
8 {
9     // Get our global thread ID
10    int id = blockIdx.x*blockDim.x+threadIdx.x;
11
12    // Make sure we do not go out of bounds
13    if (id < n)
14        c[id] = a[id] + b[id];
15 }
16
17 //void vecAdd_wrapper()
18 void vecAdd_wrapper(int rank, int nprocs)
19 {
20     char *cvd = getenv ("CUDA_VISIBLE_DEVICES");
21     printf ("Rank %d sees CUDA_VISIBLE_DEVICES=%s\n",
22            rank,
23            cvd ? cvd : "(unset)");
24     char *cdo = getenv ("CUDA_DEVICE_ORDER");
25     printf ("Rank %d sees CUDA_DEVICE_ORDER=%s\n",
26            rank,
```

Click the “Go” button to run the application and launch the kernel



Stopped at a breakpoint in the CUDA kernel

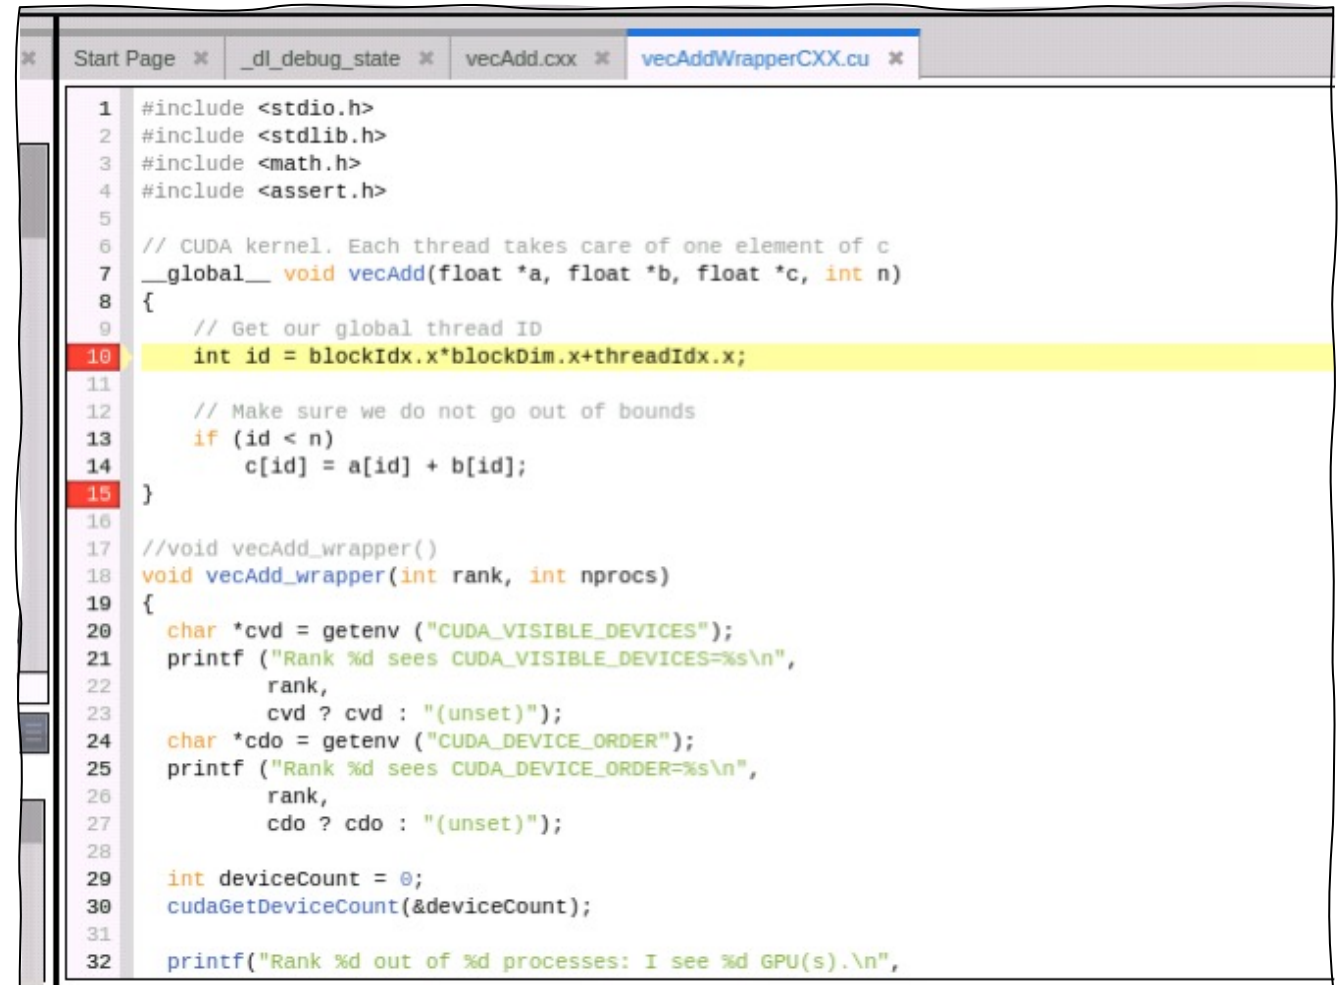
The screenshot displays the TotalView 2024.1 IDE interface. The main window shows a C++ source file named `vecAddWrapperCXX.cu` with a breakpoint set at line 10. The code is a CUDA kernel function `vecAdd` that takes three float pointers and an integer `n`. The kernel iterates over `n` elements, calculating `c[id] = a[id] + b[id]`. The breakpoint is triggered at the assignment statement.

The **Processes & Threads** panel shows a tree view of the execution environment. The **Configure** panel is open, showing options for stopping the process or thread, with **Thread ID** selected. The **Variables** panel shows the current execution context, including the process ID, device ID, and warp/lane information. The **Local Variables** panel shows the values of the arguments `a`, `b`, `c`, and `n`, as well as the `id` variable, which is currently `<Bad address: (Optimize...`.

```
1 #include <stdio.h>
2 #include <stdlib.h>
3 #include <math.h>
4 #include <assert.h>
5
6 // CUDA kernel. Each thread takes care of one element of c
7 __global__ void vecAdd(float *a, float *b, float *c, int n)
8 {
9     // Get our global thread ID
10    int id = blockIdx.x*blockDim.x+threadIdx.x;
11
12    // Make sure we do not go out of bounds
13    if (id < n)
14        c[id] = a[id] + b[id];
15 }
16
17 //void vecAdd_wrapper()
18 void vecAdd_wrapper(int rank, int nprocs)
19 {
20     char *cud = getenv ("CUDA_VISIBLE_DEVICES");
21     printf ("Rank %d sees CUDA_VISIBLE_DEVICES=%s\n",
22            rank,
23            cud ? cud : "(unset)");
24     char *cdo = getenv ("CUDA_DEVICE_ORDER");
25     printf ("Rank %d sees CUDA_DEVICE_ORDER=%s\n",
26            rank,
27            cdo ? cdo : "(unset)");
28
29     int deviceCount = 0;
30     cudaGetDeviceCount(&deviceCount);
31 }
```

Source view stopped in a CUDA kernel

- Line number information for the GPU code is *unified* with the CPU code
- The hollow breakpoint marker turns solid, indicating that there is now code at that line
- The PC arrow and highlighted source line indicates where the warp is stopped



```
1 #include <stdio.h>
2 #include <stdlib.h>
3 #include <math.h>
4 #include <assert.h>
5
6 // CUDA kernel. Each thread takes care of one element of c
7 __global__ void vecAdd(float *a, float *b, float *c, int n)
8 {
9     // Get our global thread ID
10    int id = blockIdx.x*blockDim.x+threadIdx.x;
11
12    // Make sure we do not go out of bounds
13    if (id < n)
14        c[id] = a[id] + b[id];
15 }
16
17 //void vecAdd_wrapper()
18 void vecAdd_wrapper(int rank, int nprocs)
19 {
20     char *cvd = getenv ("CUDA_VISIBLE_DEVICES");
21     printf ("Rank %d sees CUDA_VISIBLE_DEVICES=%s\n",
22            rank,
23            cvd ? cvd : "(unset)");
24     char *cdo = getenv ("CUDA_DEVICE_ORDER");
25     printf ("Rank %d sees CUDA_DEVICE_ORDER=%s\n",
26            rank,
27            cdo ? cdo : "(unset)");
28
29     int deviceCount = 0;
30     cudaGetDeviceCount(&deviceCount);
31
32     printf("Rank %d out of %d processes: I see %d GPU(s).\n",
```

GPU thread focus and navigation controls

- “GPU (Logical)” control displays and allows focusing on a specific Block and Thread



- “GPU (Physical)” control displays and allows focusing on a specific Device, SM, Warp, and Lane



CUDA stack backtrace and local variables

- Call Stack

- Open the drawer for details

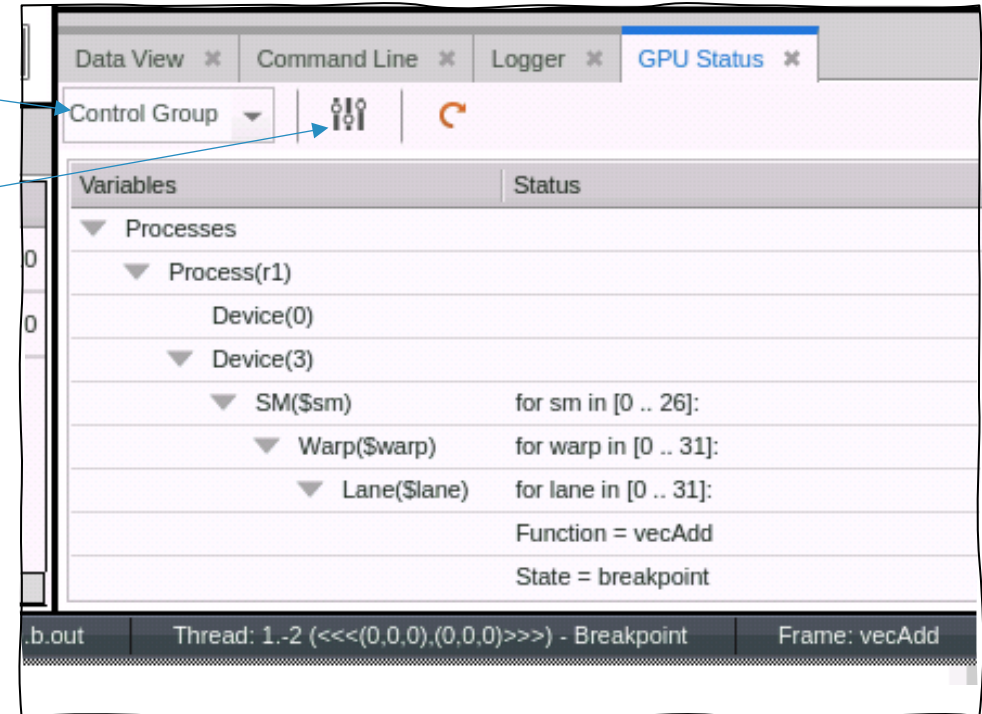
- Local Variables

The screenshot shows a debugger window with two main panes. The top pane is titled 'Call Stack' and shows a single entry for the function 'vecAdd'. The bottom pane is titled 'Local Variables' and displays a table of variables. The 'Info' pane between them shows the function name 'vecAdd' and the source file path.

Name	Type	Value
Arguments		
a	float @generic ...	0x15318b200000 -> 0
b	float @generic ...	0x15318b261c00 -> 0.382683
c	float @generic ...	0x15318b2c3800 -> 0
n	int @parameter	0x000186a0 (100000)
Block at Lin...		
id	int @register	<Bad address: (Optimized Out)>

GPU Status view

- The “GPU Status” view displays an aggregated overview of one or more of the GPUs in the whole job, in a single process, or on a single GPU
- The “GPU Status” view controls allow
 - Selecting the set of properties to display
 - Aggregation by the selected properties
 - Sorting by the selected properties
 - Creating compound filters to include/exclude properties that are equal, not equal, greater, etc.
- Allows you to get a “big picture” of the state of one or more of the GPUs in your job



Demo

The screenshot displays the TotalView 2024.1 debugger interface. The main window title is "@TEMP@CUDA@.b.out - Rank 0, Thread 0.-4 (<<<(0,0,0),(0,0,0)>>>) (Breakpoint) - TotalView 2024.1 (on nid001681)".

Process & Thread List:

Description	# P	# T	Members
srtn (S3)	1	1	p1
b.out (S4)	8	8	0-7
Breakpoint	8	8	0-7-4
2.-4	1	1	0.-4
3.-4	1	1	1.-4
4.-4	1	1	2.-4
5.-4	1	1	3.-4
6.-4	1	1	4.-4
7.-4	1	1	5.-4
8.-4	1	1	6.-4

Code Editor: The central pane shows the source code for `vecAddWrapperCXX.cu`. A breakpoint is set at line 10: `int id = blockIdx.x*blockDim.x+threadIdx.x;`.

```
1 #include <stdio.h>
2 #include <stdlib.h>
3 #include <math.h>
4 #include <assert.h>
5
6 // CUDA kernel. Each thread takes care of one element of c
7 __global__ void vecAdd(float *a, float *b, float *c, int n)
8 {
9     // Get our global thread ID
10    int id = blockIdx.x*blockDim.x+threadIdx.x;
11
12    // Make sure we do not go out of bounds
13    if (id < n)
14        c[id] = a[id] + b[id];
15 }
16
17 //void vecAdd_wrapper()
18 void vecAdd_wrapper(int rank, int nprocs)
19 {
20     char *cvd = getenv ("CUDA_VISIBLE_DEVICES");
21     printf ("Rank %d sees CUDA_VISIBLE_DEVICES=%s\n",
22            rank,
23            cvd ? cvd : "(unset)");
24     char *cdo = getenv ("CUDA_DEVICE_ORDER");
25     printf ("Rank %d sees CUDA_DEVICE_ORDER=%s\n",
26            rank,
27            cdo ? cdo : "(unset)");
28
29     int deviceCount = 0;
30     cudaGetDeviceCount(&deviceCount);
31 }
```

Call Stack: Shows the current function `vecAdd`.

Info Panel:

Field	Value
Function	vecAdd
Source	...es/f/jdelsign/support-50919/vecAddWrapperCXX.cu
Line	10
FP	0x ffdc0

Variables Panel:

Name	Type	Value
a	float @generic ...	0x7f1eb9600000 -> 0
b	float @generic ...	0x7f1eb9661c00 -> 0
c	float @generic ...	0x7f1eb96c3800 -> 0
n	@parameter int	0x000186a0 (100000)
id	@register int	<Bad address: (Optimize...

GPU Status Panel:

Process	Status
Process(\$dpid)	for dpid in [r0 .. r7]:
Device(\$dev)	for dev in {0 1 2}
Device(3)	
SM(\$sm)	for sm in [0 .. 26]:
Warp(\$warp)	for warp in [0 .. 31]:
Lane(\$lane)	for lane in [0 .. 31]:
	Function = vecAdd
	State = breakpoint
SM(27)	
Warp(\$warp)	for warp in [0 .. 23]:
Lane(\$lane)	for lane in [0 .. 31]:
	Function = vecAdd

Action Points Panel:

ID	Type	Stop	Location
1	Break	Process	...vecAddWrapperCXX.cu#10
1	Break	Process	...vecAddWrapperCXX.cu#10

Status Bar: Rank: 0 (833128@nid001681) @TEMP@CUDA@.b.out Thread: 0.-4 (<<<(0,0,0),(0,0,0)>>>) - Breakpoint Frame: vecAdd File: ...homes/f/jdelsign/support-50919/vecAddWrapperCXX.cu Line: 10

Batch Debugging with TVScript

tvscript

- A straightforward language for unattended and/or batch debugging with TotalView and/or MemoryScape
- Usable whenever jobs need to be submitted or batched
- Can be used for automation
- A more powerful version of printf, no recompilation necessary between runs
- Schedule automated debug runs with *cron* jobs
- Expand its capabilities using TCL

tvscript

```
tvscript [options] [filename] [ -a program_args]
```

options

TotalView and tvscript command-line options.

filename

The program being debugged.

-a *program_args*

Program arguments.

tvscript

- All of the following information is provided by default for each print
 - Process id
 - Thread id
 - Rank
 - Timestamp
 - Event/Action description
- A single output file is written containing all of the information regardless of the number of processes/threads being debugged

Supported tvscript events

Event Type	Event	Definition
General event	any_event	A generated event occurred.
Memory debugging event	addr_not_at_start	Program attempted to free a block using an incorrect address.
	alloc_not_in_heap	The memory allocator returned a block not in the heap; the heap may be corrupt.
	alloc_null	An allocation either failed or returned NULL; this usually means that the system is out of memory.
	alloc_returned_bad_alignment	The memory allocator returned a misaligned block; the heap may be corrupt.
	any_memory_event	A memory event occurred.
	bad_alignment_argument	Program supplied an invalid alignment argument to the heap manager.
	double_alloc	The memory allocator returned a block currently being used; the heap may be corrupt.
	double_dealloc	Program attempted to free an already freed block.
	free_not_allocated	Program attempted to free an address that is not in the heap.
	guard_corruption	Program overwrote the guard areas around a block.

Supported tvscript events

Event Type	Event	Definition
	hoard_low_memory_threshold	Hoard low memory threshold crossed.
	realloc_not_allocated	Program attempted to reallocate an address that is not in the heap.
	rz_ overrun	Program attempted to access memory beyond the end of an allocated block.
	rz_underrun	Program attempted to access memory before the start of an allocated block.
	rz_use_after_free	Program attempted to access a block of memory after it has been deallocated.
	rz_use_after_free_ overrun	Program attempted to access memory beyond the end of a deallocated block.
	rz_use_after_free_underrun	Program attempted to access memory before the start of a deallocated block.
	termination_notification	The target is terminating.
Source code debugging event	actionpoint	A thread hit an action point.
	error	An error occurred.
Reverse debugging	stopped_at_end	The program is stopped at the end of execution and is about to exit.

Supported tvscript actions

Action Type	Action	Definition
Memory debugging actions	check_guard_blocks	Checks all guard blocks and write violations into the log file.
	list_allocations	Writes a list of all memory allocations into the log file.
	list_leaks	Writes a list of all memory leaks into the log file.
	save_html_heap_status_source_view	Generates and saves an HTML version of the Heap Status Source View Report.
	save_memory_debugging_file	Generates and saves a memory debugging file.
	save_text_heap_status_source_view	Generates and saves a text version of the Heap Status Source View Report.
Source code debugging actions	display_backtrace [-level <i>level-num</i>] [<i>num_levels</i>] [<i>options</i>]	<p>Writes the current stack backtrace into the log file.</p> <p>-level <i>level-num</i> sets the level at which information starts being logged.</p> <p><i>num_levels</i> restricts output to this number of levels in the call stack.</p> <p>If you do not set a level, tvscript displays all levels in the call stack.</p> <p><i>options</i> is one or more of the following:</p> <ul style="list-style-type: none">-[no]show_arguments-[no]show_fp-[no]show_fp_registers-[no]show_image-[no]show_locals-[no]show_pc-[no]show_registers

Supported tvscript actions

Action Type	Action	Definition
	print [-slice { <i>slice_exp</i> } { <i>variable</i> <i>exp</i> }	Writes the value of a variable or an expression into the log file. If the variable is an array, the -slice option limits the amount of data defined by <i>slice_exp</i> . A slice expression is a way to define the slice, such as var[100:130] in C and C++. (This displays all values from var[100] to var[130] .) To display every fourth value, add an additional argument; for example, var[100:130:4] . For additional information, see “Examining Arrays” in the <i>TotalView for HPC User Guide</i> .
Reverse debugging actions	enable_reverse_debugging	Turns on ReplayEngine reverse debugging and begins recording the execution of the program.
	save_replay_recording_file	Saves a ReplayEngine recording file. The filename is of the form <ProcessName>-<PID>_<DATE>.<INDEX>.recording .

tvscript examples

Simple example

```
tvscript \  
-create_actionpoint "method1=>display_backtrace -show_arguments" \  
-create_actionpoint "method2#37=>display_backtrace \  
    -show_locals -level 1" \  
-event_action "error=>display_backtrace -show_arguments \  
    -show_locals" \  
-display_specifiers "noshow_pid,noshow_tid" \  
-maxruntime "00:00:30" \  
~/work/filterapp /filterapp -a 20
```

MPI example

```
tvscript -mpi "Open MPI" -tasks 4 \  
-create_actionpoint \  
"hello.c#14=>display_backtrace" \  
~/tests/MPI_hello
```

tvscript examples

Memory Debugging example

```
tvscript -maxruntime "00:00:30" \  
-event_action "any_event=save_memory_debugging_file" \  
-guard_blocks -hoard_freed_memory -detect_leaks \  
~/work/filterapp -a 20
```

ReplayEngine example

```
tvscript \  
-create_actionpoint "main=>enable_reverse_debugging" \  
-event_action "stopped_at_end=>save_replay_recording_file" \  
filterapp
```

Demo

- TVScript demo (`tvscript --script_file file tvscript_example.tvd ex2`)

Common TotalView Usage Hints

Common TotalView Usage Hints

- TotalView can't find the program source
 - Did you compile with -g ?
 - How to adjust the TotalView search paths? Preferences -> Search Path
- Python Debugging
 - Making sure proper system debug packages are installed for Python
- X11 forwarding performance
 - If users are forwarding X11 displays through ssh TotalView UI performance can be bad
- Understanding different ways to stop program execution with TotalView Action Points
 - Using a watchpoint on a local variable
- Focus
 - Diving on a variable that is no longer in scope. Check the Local Variables window for in scope variables
 - TotalView doesn't change focus to the thread hitting a breakpoint. Set Action Point Preferences to "Automatically focus on threads/processes at breakpoint"

Common TotalView Usage Hints (cont.)

- MPI Debugging
 - Differences in launching MPI job from within the TotalView UI vs the command line.
 - TotalView runs an MPI program without stopping. Set the Parallel Preferences to “Ask What To Do” in After Attach Behavior
 - Using wrong attributes in processes and threads view
- Reverse Debugging
 - Running out of memory by not setting the maximum memory allocated to ReplayEngine
 - Defer turning on reverse debugging until later in program execution to avoid slow initialization phases
 - Adjust reverse debugging circular buffer size to reduce resources
- Memory Debugging
 - Starting with All memory debugging options enabled rather than Low
 - Not setting a size restriction for Red Zones
 - Issues with getting memory debugging turned on in an MPI job. May have to set LD_PRELOAD environment variable or worst case, prelink HIA

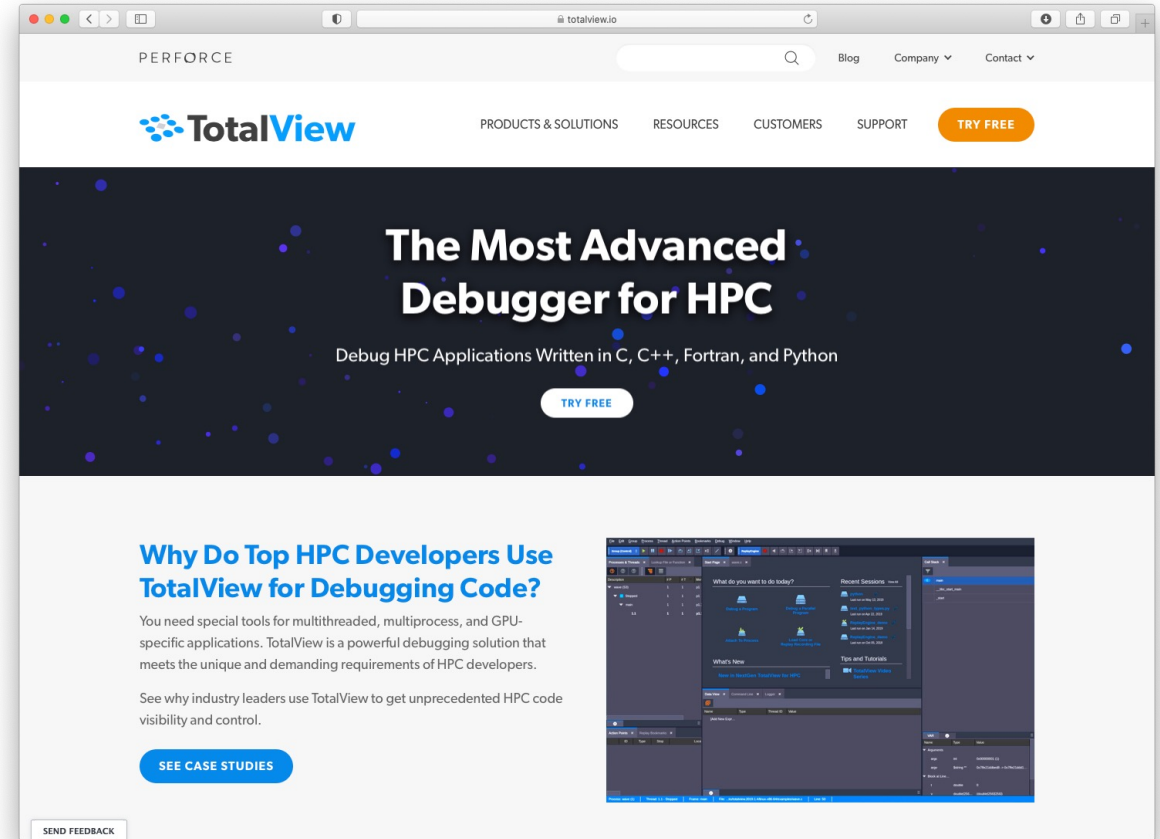
Common TotalView Usage Hints (cont.)

- Differences in functionality between new UI and classic UI
 - How to switch between them. Preferences -> Display or `totalview -newUI` and `totalview -oldUI`
 - Where the gaps still are in functionality
- Reverse Connect with `tvconnect`
 - When I use Reverse Connect I get the following obscure message: *myProgram is an invalid or incompatible executable file format for the target platform*
 - The message indicates an incompatible file format but most often this occurs if the program provided to `tvconnect` for TotalView to debug cannot be found. The easiest way to resolve problem is to provide the full path to the target application, e.g., `tvconnect /home/usr/myProgram`
- How do I get help?
 - How to submit a support ticket? `techsupport@roguewave.com`
 - Where is TV documentation (locally and on the internet). <https://help.totalview.io/>
 - Are there videos I can watch to learn how to use TotalView? <https://totalview.io/support/video-tutorials>

TotalView Resources and Documentation

TotalView Resources and Documentation

- TotalView website:
<https://totalview.io>
- TotalView documentation:
 - <https://help.totalview.io>
 - User Guides: Debugging, Memory Debugging and Reverse Debugging
 - Reference Guides: Using the CLI, Transformations, Running TotalView
- Blog:
<https://totalview.io/blog>
- Video Tutorials:
<https://totalview.io/support/video-tutorials>



Q&A

Contact us

- Bill Burns (Senior Director of Software Engineering and Product Manager)
bburns@perforce.com
- John DelSignore (TotalView Chief Architect)
jdelsignore@perforce.com
- Scot Halverson (NVIDIA Solutions Architect)
shalverson@nvidia.com
- Peter Thompson (Senior Support Engineer)
pthompson@perforce.com
- Bruce Ryan (Senior Account Executive)
bryan@perforce.com
- Ken Hill (Senior Sales Engineer)
khill@perforce.com