



# ECMWF Workshop on High Performance Computing in Meteorology

3<sup>rd</sup> November 2010

Dean Stewart



# Agenda

---

- Company Overview
- Rogue Wave Product Overview
  - IMSL Fortran
  - TotalView Debugger
  - Acumem ThreadSpotter

# Rogue Wave

---



- **Customers**

- 3,000+ in 36 countries
- Financial services, telecoms, oil and gas, government and aerospace, research and academic

- **History**

- Founded: 1989
- Named after the founder's boat
- Acquisitions:
  - Visual Numerics: 2009
  - TotalView Technologies: 2010
  - Acumem: 2010

- **Global Locations**

- HQ: Boulder, CO
- NA: Houston, TX; Corvallis, OR; Natick MA
- EMEA: France, Germany, Sweden, UK
- APAC: Japan

# Rogue Wave Today

---

The largest independent provider of cross-platform software development tools and embedded components for the next generation of HPC applications



Leading provider of enterprise class C++ components and infrastructure for high performance applications.



Leader in embeddable math and statistics algorithms and visualization software for data-intensive applications.



Industry-leading interactive analysis and debugging tools for the world's most sophisticated software applications.



Leading provider of intelligent software technology which analyzes and optimizes the computing performance in single- and multi-core environments.

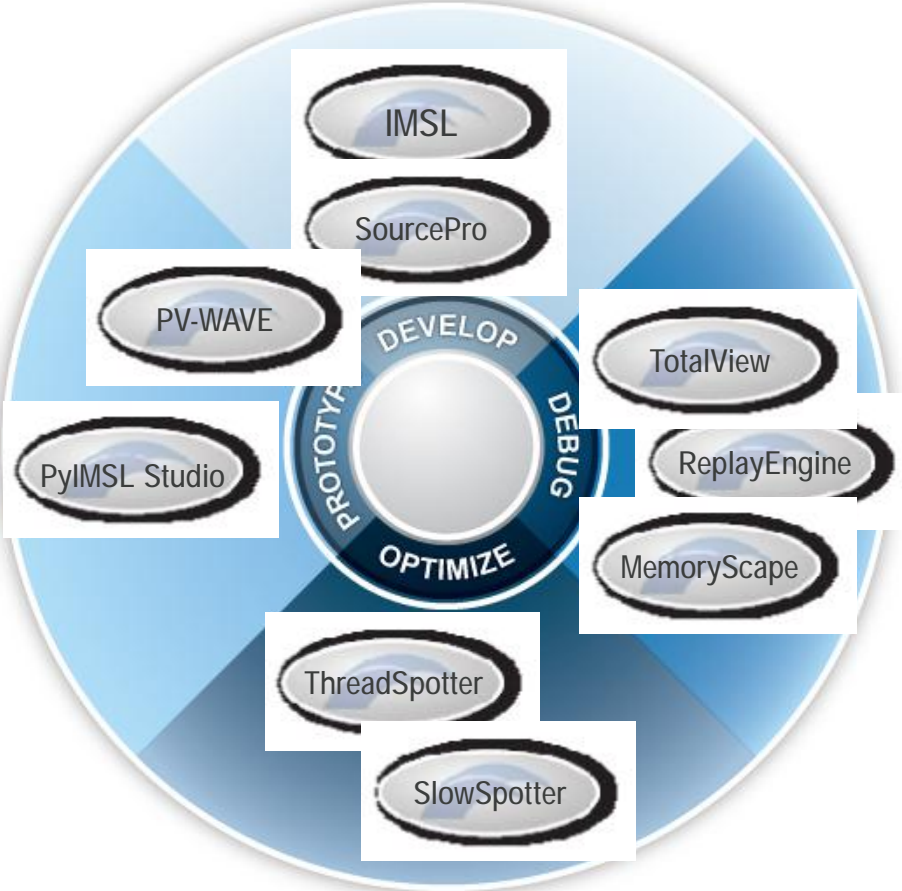


# Agenda

---

- Company Overview
- Rogue Wave Product Overview
  - IMSL Fortran
    - CUDA support
  - TotalView Debugger
    - Remote Display Client
    - TVScript for Batch Debugging
    - MemoryScape and ReplayEngine
  - Acumem ThreadSpotter

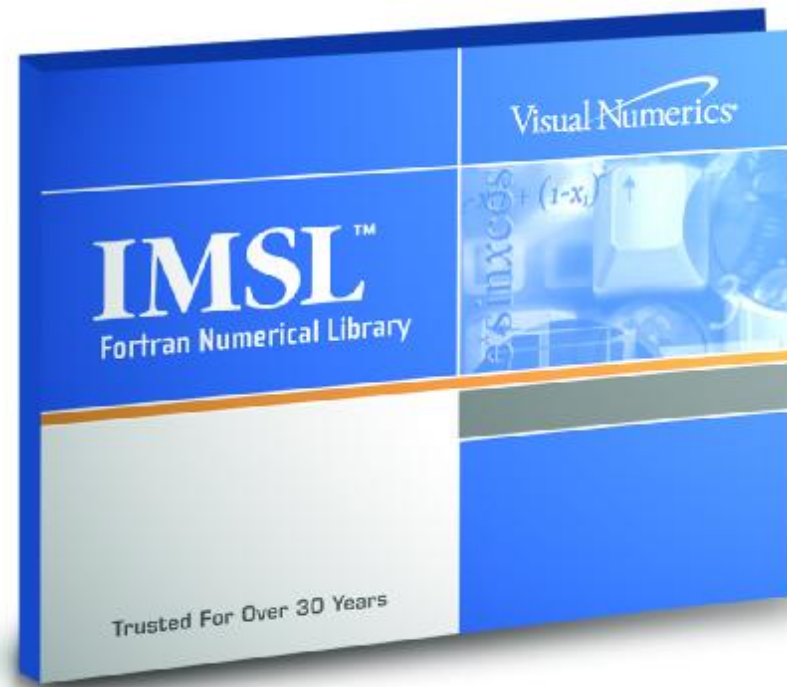
# Rogue Wave Product Family



PROTOTYPE è DEVELOP è DEBUG è OPTIMIZE

# IMSL Fortran Library

---



*Now supports NVIDIA GPUs*

# Parallelism Using the IMSL Fortran Library

---

- OpenMP
  - Linear systems, Eigensystem Analysis, Transforms, etc.
  - The corresponding routines differ by environment
  - IMSL uses the OpenMP directive to parallelize internally; users do not need to write parallel program source
  
- MPI
  - Linear Algebra Operator and Generic Function, Linear systems, etc.
  - Corresponding routines are the same in all environments
  - Users need to call the IMSL function `MP_SETUP()`

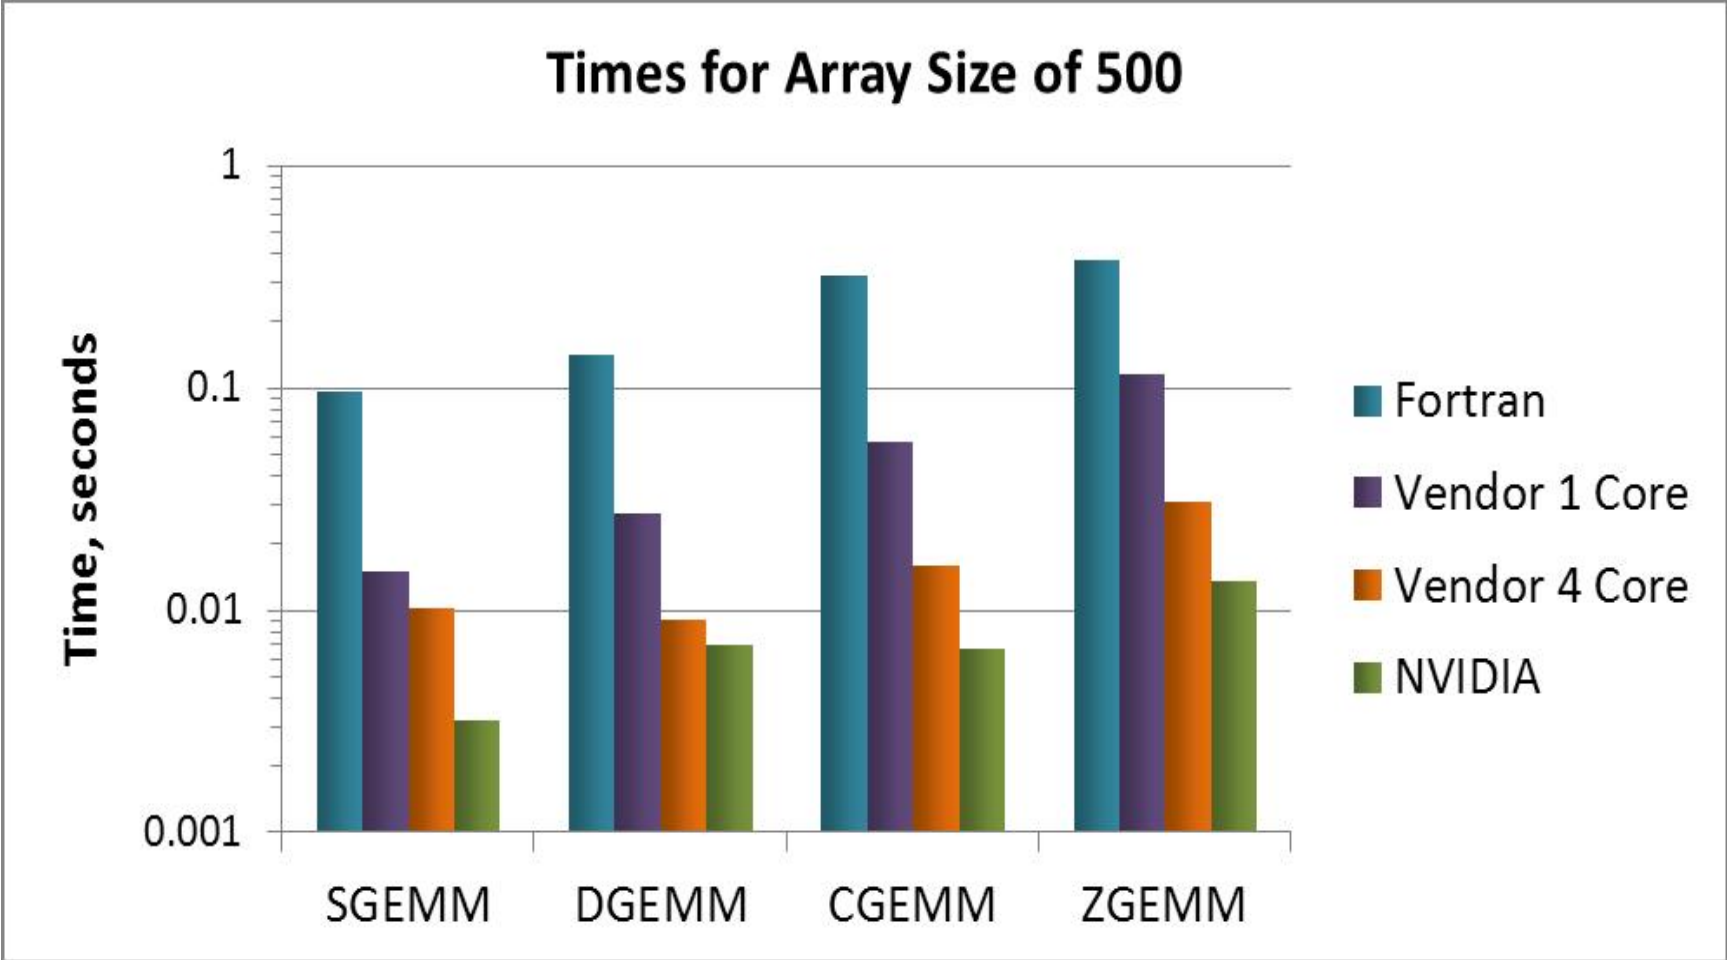


# IMSL Fortran Library 7.0

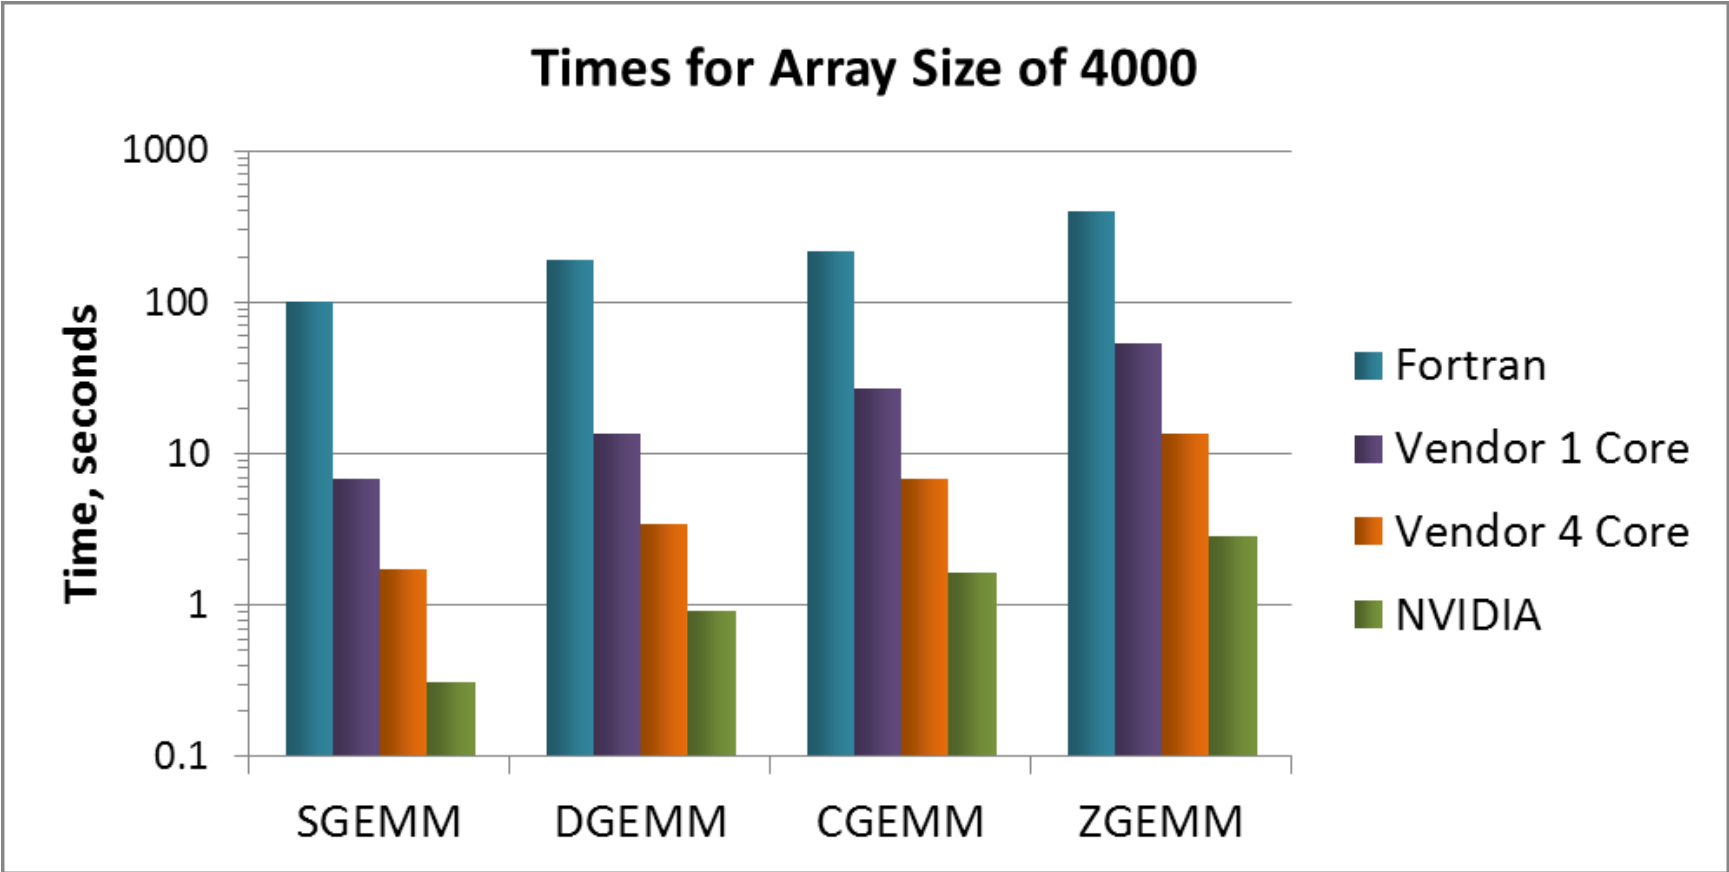
---

- Integration of NVIDIA CUDA BLAS
  - Support NVIDIA GPU hardware (Tesla, Fermi)
  - Call CUDA Level 2 and 3 BLAS functions
    - Problem size must be large enough to see the benefit
    - Use IMSL functions as usual, calls NVIDIA's library behind the scenes
    - If no GPU hardware or smaller problem, execute on CPU
  - Double and Single precision
  - Benchmark white paper available  
<http://www.vni.com/products/imsl/fortran/>

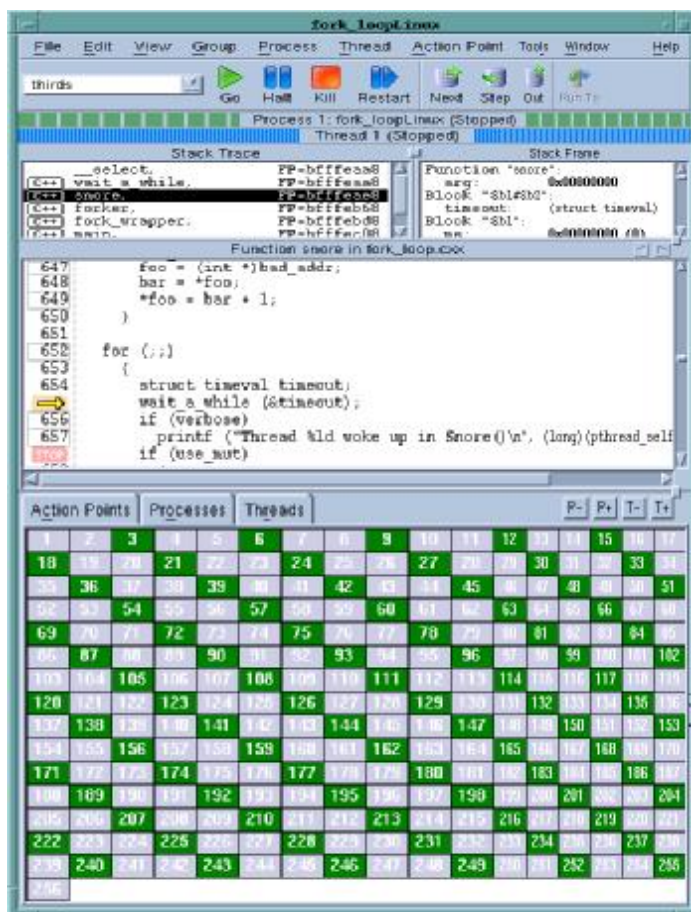
# Throughput Results



# Throughput Results (cont.)



# TotalView Debugging Technology



- What is TotalView?
  - Parallel and Multithreaded Debugging and Analysis Tool
  - For scientists and engineers working with C/C++ and Fortran
  - Makes developing, maintaining and supporting critical and cutting edge applications easier and less risky
- Major Features
  - Supports Linux, Unix and Mac OS X
  - Parallel Debugging
    - MPI, Pthreads, OMP, UPC
  - Includes a Remote Display Client freeing users to work from anywhere
  - Memory Debugging with MemoryScope
  - Optional Reverse Debugging with ReplayEngine
  - Batch Debugging with TVScript and the CLI
- Advantages
  - Easy to learn graphical user interface with data visualization
  - Wide variety of features so users can tackle unexpected bugs
  - Consistent functionality and look and feel across a wide range of platforms
  - Works robustly with open source and vendor compilers
  - Native debugger core is highly scalable to large clusters, large code and massive datasets

# How can TotalView help you ?

---

**Debugging means examining a specific controlled instance of program execution**  
**Provides an answer to the question : “What is my program *really* doing?”**

- **Threads and/or MPI**

- **When you have**
  - Deadlocks and hangs
  - Race conditions
- **It provides**
  - Asynchronous thread control
  - Barrier points
  - Powerful group mechanism

- **Fortran and/or C++**

- **Complex data structures**
  - Diving and recursive dive
- **STL Collection Classes**
  - STLView
- **Rich class hierarchies**
  - Powerful type-casting features

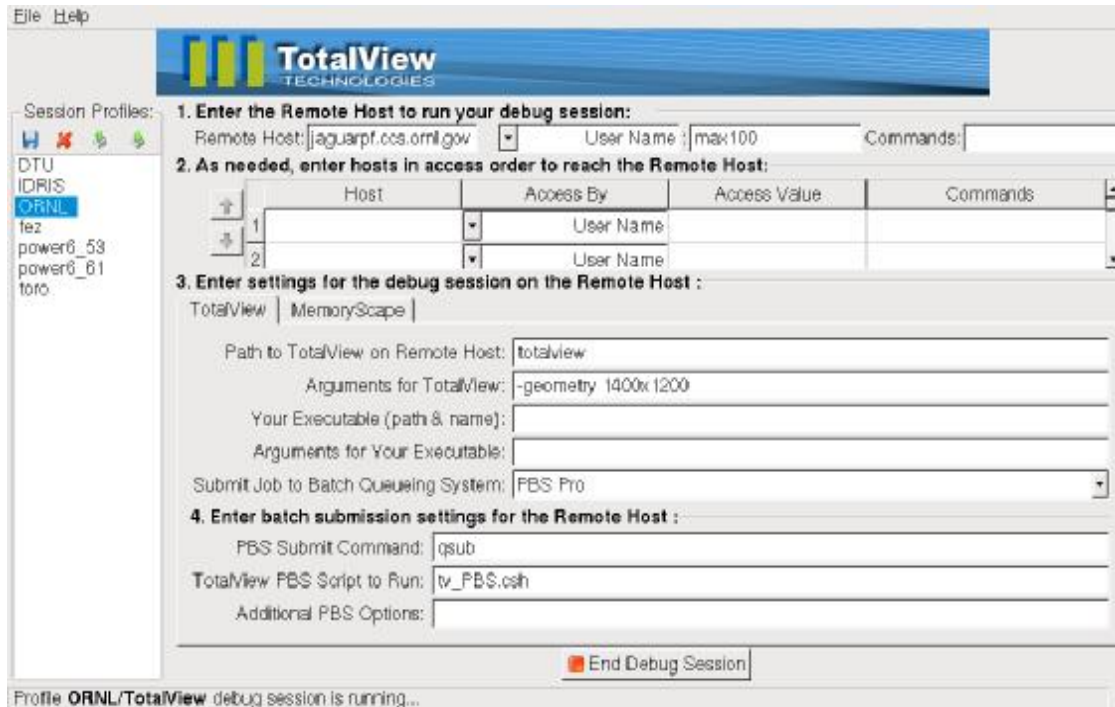
- **Memory Analysis**

- **Leaks and Bounds Errors**
  - Automatic error detection tools
- **Out of Memory Errors**
  - Analysis of heap memory usage by file function and line

- **Data Analysis**

- **Numerical errors**
  - Extensible data visualization
  - Slicing and filtering of arrays
  - Powerful expression system
  - Conditional watchpoints

# TotalView Remote Display Client



- The Remote Display Client offers users the ability to easily set up and operate a TotalView debug session that is running on another system.
- Provides for a connection that is
  - Easy
  - Fast
  - Secure
- The Remote Display Client is available for:
  - Linux x86
  - Linux x86-64
  - Windows XP
  - Windows Vista
  - Mac OS X Leopard and Snow Leopard
- The Client also provides for submission of jobs to batch queuing systems PBS Pro and LoadLeveler

# Batch Debugging with TVScript

---

- TVScript
  - Defines events
    - Breakpoints, memory errors, etc..
  - Actions to take in response to these events
    - Print variables or create memory reports
  - Runs a serial or MPI program towards completion
    - With no user interaction
- More powerful and flexible than Printf-style debugging
  - Use to prepare and guide interactive debugging
  - Use whenever jobs need to be submitted into a managed environment
  - Can be used to automate test/verify environments

Here, for example, is how tvscript is invoked on a program:

```
tvscript \  
-create_actionpoint "method1=>display_backtrace" \  
-show_arguments \  
-create_actionpoint "method2#37=>display_backtrace" \  
-show_locals -level 1 \  
-display_specifiers "nowshow_pid,noshow_tid" \  
-maxruntime "00:00:30" \  
filterapp -a 20
```

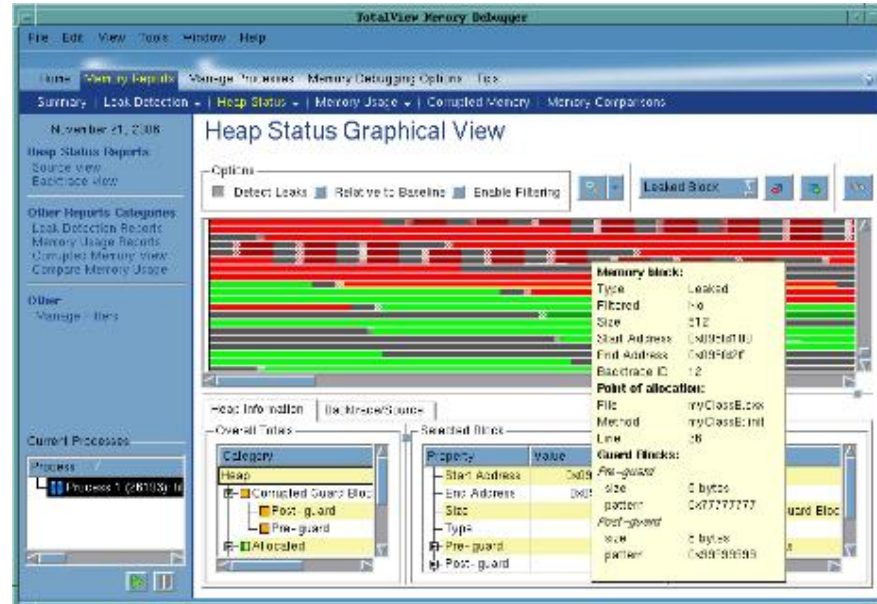
You can also execute MPI programs using tvscript. Here is a small example:

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

# MemoryScape

## Simple to use, intuitive memory debugging

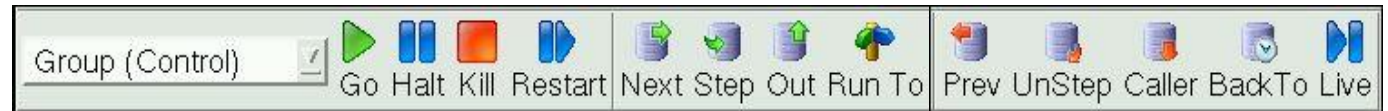
- What is MemoryScape?
  - Streamlined
  - Lightweight
  - Intuitive
  - Collaborative
  - Memory Debugging
- Features
  - Shows
    - Memory errors
    - Memory status
    - Memory leaks
    - Buffer overflows
  - MPI memory debugging
  - Remote memory debugging



- Technical Advantages
  - Low overhead
  - No Instrumentation
- Interface
  - Inductive
  - Collaboration
  - Multi-process



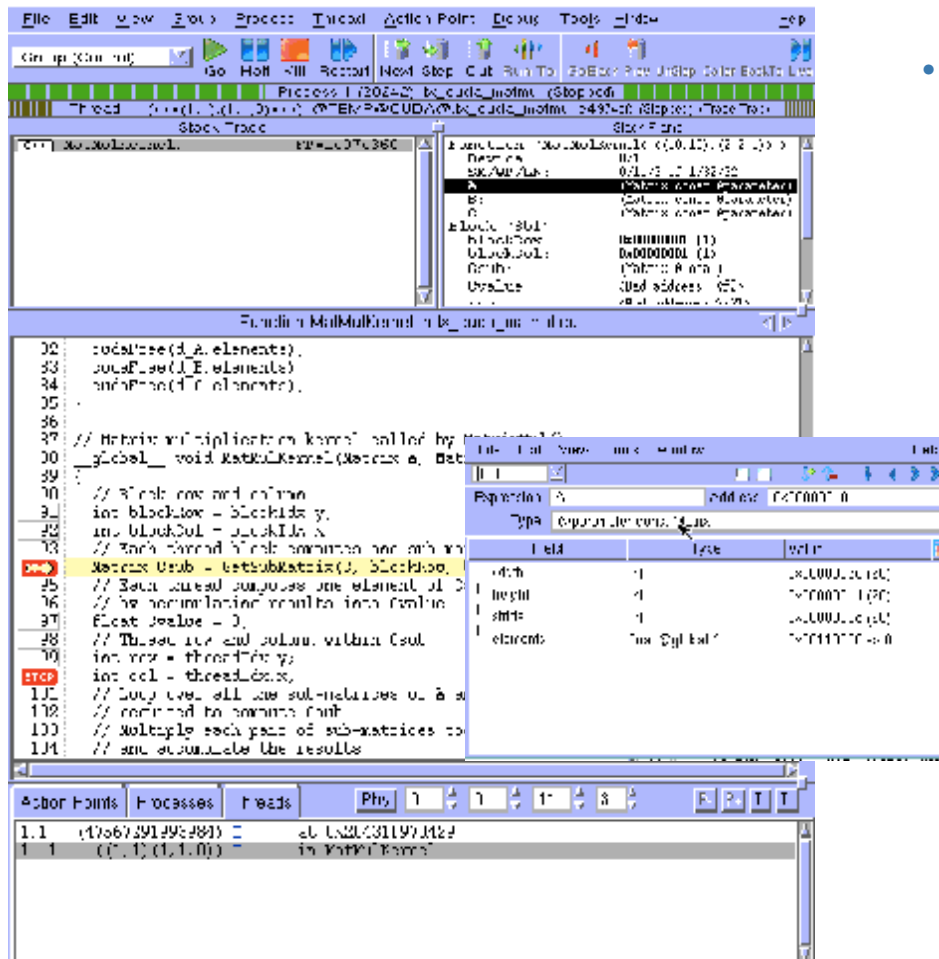
# ReplayEngine



- Enhances debugging experience
  - Add-on to TotalView
- Captures execution history
  - Record all external input to program
  - Records internal sources of non-determinism
- Replays execution history
  - Examine any part of the execution history
  - Step as easily back through code as you do forwards
  - Jump to points of interest
- Simple extension to TotalView
  - No recompilation or instrumentation
  - The user just says where they want to go
  - Explore data and state in the past just like a live process
- Supported on Linux x86 and x86-64
- Supports MPI, Pthreads, and OpenMP

```
40
41
42 int funcB(int
43 int c;
44 int i;
45 int v[MAXDEPT
46 int *p;
47 c=b+2;
48 p=&c;
49 if(c<MAXDEPTH
50 c=funcA(c);
51 for (i=arrayl
52 v[i]=*p;
```

# TotalView for CUDA



- Characteristics
  - Debugging of application running on the GPU device (not in an emulator)
  - Full visibility of both Linux threads and GPU device threads
    - Device threads shown as part of the parent Unix process
    - Correctly handle all the differences between the CPU and GPU
  - Fully represent the hierarchical memory
    - Display data at any level (registers, local, block, global or host memory)
    - Making it clear where data resides with type qualification
  - Thread and Block Coordinates
    - Built in runtime variables display threads in a warp, block and thread dimensions and indexes
    - Displayed on the interface in the status bar, thread tab and stack frame
  - Device thread control
    - Warps advance synchronously
  - Handles CUDA function inlining
    - Step into or over inlined functions
    - Functions show on stack trace
  - Reports memory access errors
    - CUDA memcheck
  - Multi-Device Support
  - Can be used with MPI

# Acumem ThreadSpotter

---

- Optimization tool that provides real performance analysis of your running code
- Focus attention to the code and the problems that matter and can be solved
- What-if analysis from a single sampling
- The tool is your mentor – learn as you go, become proficient
- Performance regression test tool:
  - because resource contention is a global concern
  - Problems don't always show up at intuitive places



# Cache Optimization Technology

---

- Analyze binary behavior
  - Collect sparse runtime data
  - Model off-line using math
- ThreadSpotter
  - Application performance analysis for programmers
  - We tell what the problem is, where it is and how to fix it

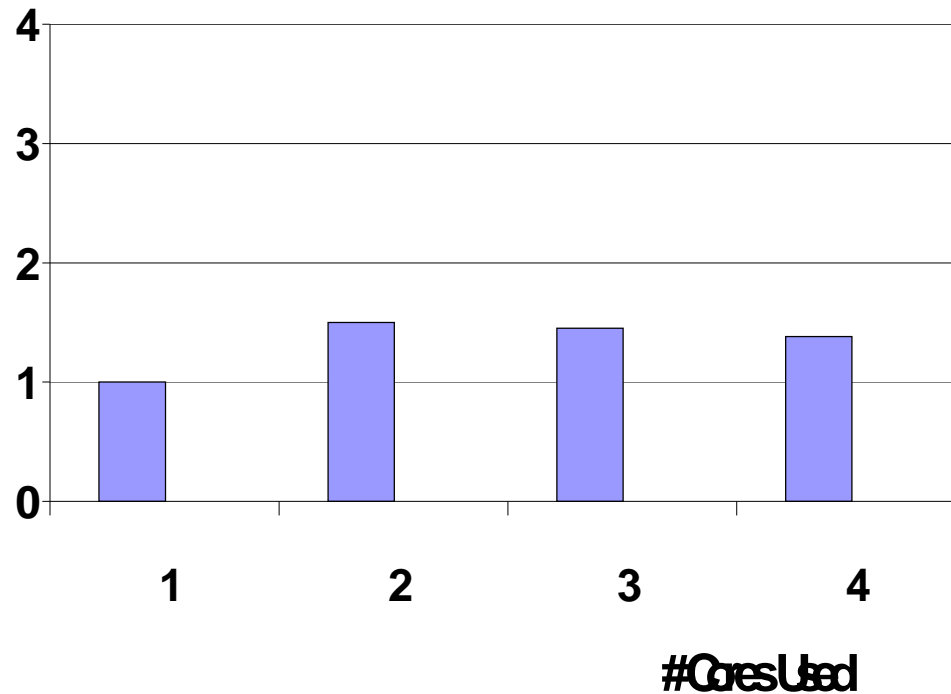
Makes the performance experts more productive

Enables non-experts to optimize code

# Example: Poor parallelism? (LBM)

Lattice Boltzmann methods – Fluid Dynamics

Performance

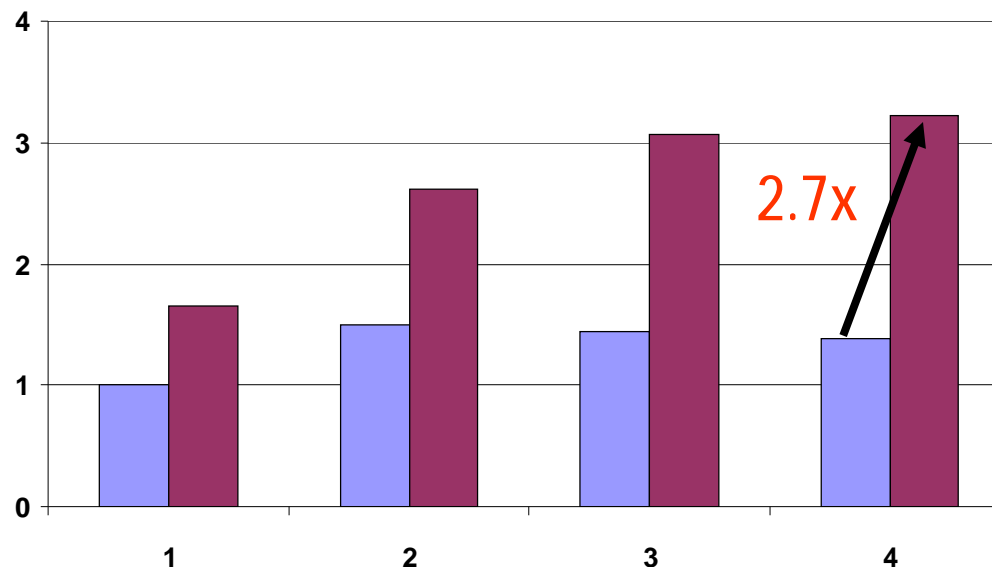


- Actually, it is “embarrassingly parallel”
- Poor memory usage  $\Rightarrow$  super-linear slowdown

# The same application optimized (LBM)

---

Performance



- Optimization can be rewarding, but costly...
  - Requires expert knowledge
  - Weeks of wading through performance data
- ThreadSpotter one-click advice: Change one line

# ThreadSpotter Screenshot

What?

Where?

The screenshot displays the ThreadSpotter interface. On the left, a table lists various bandwidth issues. The 'Issue #1: Fetch utilization' panel provides detailed statistics and a stack trace. The code editor on the right shows the source code for a quantum gate function, with lines 89-93 highlighted in yellow to indicate the location of the issue.

#	Issue type	% of bandwidth	% of fetches	% of write-backs	Fetch utilization	Write-back utilization
1	Fetch utilization	48.5%	64.5%	0.0%	50.9%	100.0%
2	Spot temp. Nacking	48.5%	64.5%	0.0%	50.9%	100.0%
4	Fetch utilization	33.9%	22.6%	68.3%	51.4%	51.4%
5	Write back utilization	33.9%	22.6%	68.3%	51.4%	51.4%
6	Spot temp. Nacking	33.9%	22.6%	68.3%	51.4%	51.4%
8	Fetch utilization	9.1%	12.3%	0.0%	52.4%	100.0%
9	Spot temp. Nacking	9.1%	12.3%	0.0%	52.4%	100.0%
11	Write back utilization	4.5%	0.0%	19.7%	0.0%	21.5%
13	Write back utilization	2.6%	0.0%	10.6%	0.0%	31.2%
15	Write back hot-spot	0.6%	0.4%	1.4%	46.9%	34.7%

**Issue #1: Fetch utilization**

This instruction group also show symptoms of Fetch hot-spot.

**Statistics for instructions of this issue**

**Instructions involved in this issue**

**Instructions previously writing to related data**

Stack	Instruction
+	quantum_sigma_x1(0x04c1c9) [RW] gates.c:174
	quantum_sigma_x1(0x804c1cc) [RW] gates.c:174

**Loop statistics**

**Loop instructions**

Copyright © 2007-2010 Acumen AB. All Rights Reserved.  
Patents pending.

```

74
73
74
75 void
76 quantum_toffoli(int control1, int control2, int target,
77 {
78     int i;
79     int qec;
80
81     quantum_qec_get_status(&qec, NULL);
82
83     if(qec)
84         quantum_toffoli_ft(control1, control2, target, reg);
85     else
86     {
87         if(quantum_objcode_put(TOFFOLI, control1, control2,
88             return;
89
90         for(i=0; i<reg->size; i++)
91         {
92             /* Flip the target bit of a basis state if bot
93             set */
94             if(reg->node[i].state & ((MAX_UNSIGNED) 1 <<
95             {
96                 if(reg->node[i].state & ((MAX_UNSIGNED) 1
97                 {
98                     reg->node[i].state ^= ((MAX_UNSIGNED)
99
100             }
101         }
102         quantum_decohere(reg);
103     }
104 }
105
106 /* Apply an unbounded toffoli gate. This gate is not cor
107 elementary and is not available on all physical realizat
108 quantum computer. Be sure to pass the function the corre
109 controlling qubits. The target is given in the last argu
110
111 void
112 quantum_unbounded_toffoli(int controlling, quantum_reg *
113 {

```

How?



# Resources

---

## Rogue Wave Software website

- [www.roguewave.com](http://www.roguewave.com)

## Demystifying debugging whitepaper

- <http://www-931.ibm.com/bin/newsletter/tool/landingPage.cgi?lpld=3022>

## TotalView video tutorials

- <http://www.totalviewtech.rsvp1.com/support/videos.html#0>

## Acumem ThreadSpotter On-line Demo

- <https://acumem.webex.com/acumem/ldr.php?AT=pb&SP=MC&rID=19140117&rKey=CD1662E760EA59FE>



# Visiting SC10?

---



Visit Rogue Wave at Booths 2431 and 2432 and Acumem at 3749

Contact me if you would like us to schedule a meeting with our senior product and management staff at the event .

Dean Stewart [dean.stewart@roguewave.com](mailto:dean.stewart@roguewave.com)  
Tel: +44 (0)6450 549957, Mob: +44 (0) 7713 150529



Thank You  
I U9UUK I on

