



Debugging at Scale

Lindon Locks

llocks@allinea.com


- At scale debugging - from 100 cores to 250,000
 - Problems faced by developers on real systems
 - Alternative approaches to debugging and how they stack up
 - How Allinea makes debugging at scale work

- HPC tools since 2001
 - Allinea DDT – Scalable parallel debugger
 - Allinea OPT – Optimisation tool for MPI and non-MPI
 - Allinea DDTLite – Parallel debugging plugin for Microsoft Visual Studio
- Large customer base
 - Ease of use – means tools get used
 - Users debugging regularly at all scales: 1 to 100,000 cores
 - World's only Petascale debugger



- Academic
 - Over 200 universities
- Major research centres
 - ANL, CEA, EPCC, GENCI, IDRIS, Juelich, NERSC, ORNL
- Aviation and Defence
 - Airbus, AWE, BAE, Dassault, DLR, EADS
- Energy
 - CGG Veritas, IFP, Total
- EDA
 - Cadence, Intel, Synopsys
- Climate and Weather
 - UK Met Office, Meteo France, NOAA

- Processor counts growing rapidly
- GPUs entering HPC
- Large hybrid systems imminent
- But what happens when software doesn't work?



- Increasing job sizes leads to unanticipated errors
 - Regular bugs
 - Data issues from larger data sets – eg. garbage in..., overflow
 - Logic issues and control flow
 - Increasing probability of independent random error
 - Memory errors/exhaustion – “random” bugs!
 - System problems – MPI and operating system
 - Pushing coded boundaries
 - Algorithmic (performance)
 - Hard-wired limits (“magic numbers”)
 - Unknown unknowns
 -

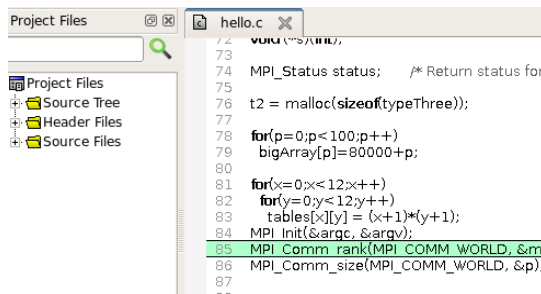

- Improved coding standards
 - Unit tests, assertions and consistency checks
 - Good practice – but tend to be single-process checks
 - Parallel checks also valid and good practice
 - Only checks for things you predict when developed
 - Coverage is rarely perfect
 - Unexpected problems – particularly random/system issues – often missed
 - Debugger still required
 - Combines well with debuggers
 - Find why a failure occurs not just a pass/fail

- Logging – printf and write
 - The oldest debugger still in active use
 - Tried and tested - as easy as “hello world”
 - If you have good intuition into the problem
 - Edit code, insert print, recompile and re-run
 - **Slow and iterative**
 - Use to log exceptions, progress or state
 - Post-mortem analysis only
 - Hard to establish real causal order of output of multiple processes
 - Output can be lost by process termination
 - Rapid growth in log output size
 - **Unscalable**

- Reproduce at a smaller scale
 - Attempt to make problem happen on fewer nodes
 - Often requires reduced data set – the large one may not fit
 - Didn't you already try the code at small scale?
 - Smaller data set may not trigger the problem
 - Does the bug even exist on smaller problems?
 - Is it a system issue – eg. an MPI problem?
 - Is probability stacking up against you?
 - **Example:** 1 in 10,000 independent probability of error?
 - Unlikely to spot on smaller runs – without many many runs
 - But near guaranteed to see it on a 10,000 core run
 - What can a parallel debugger do to help?
 - Debug at the scale of the problem. **Now.**

- Many benefits to graphical parallel debuggers
 - Large feature sets for common bugs
 - Richness of user interface and real control of processes
- Historically **all** parallel debuggers hit scale problems
 - Bottleneck at the frontend: Direct GUI → nodes architectures
 - Linear performance in number of processes
 - Human factors limit – mouse fatigue and brain overload
- Are tools ready for the task?
 - Allinea DDT has changed the game

- Scalar features
 - Advanced C++ and STL
 - Fortran 90, 95 and 2003: modules, allocatable data, pointers, derived types
 - Memory debugging
- Multithreading & OpenMP features
 - Step, breakpoint etc. one or all threads
- MPI features
 - Easy to manage groups
 - Control processes by groups
 - Compare data
 - Visualise message queues



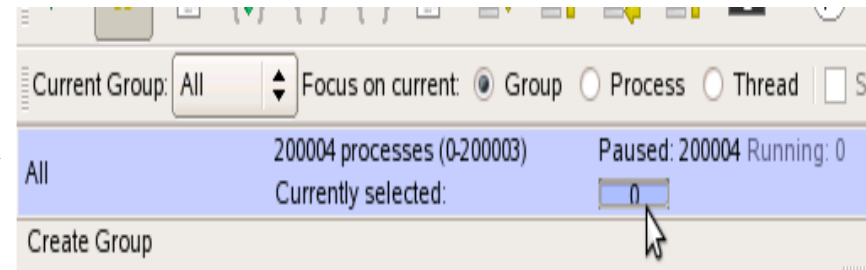

Stacks (All)	
Processes	Function
150120	_start
150120	__libc_start_main
150120	main
150120	pop (POP.f90:81)
150120	initialize_pop (initial.f90:119)
150120	init_communicate (communicate.f90:87)
150119	create_ocn_communicator (communicate.f90:300)
1	create_ocn_communicator (communicate.f90:303)

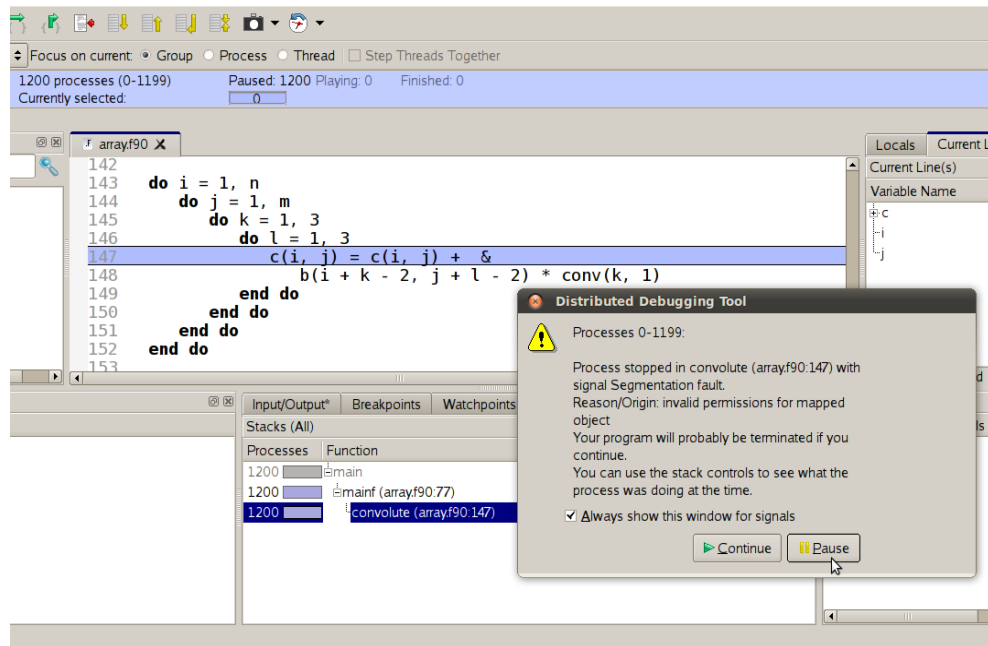
- Parallel Stack View

- Find rogue processes quickly
- Identify classes of process behaviour
- Rapid grouping of processes

- Control Processes by Groups


- Set breakpoints, step, play, stop for groups
- Scalable groups view: compact group display





- Immediate stop on crash
 - Segmentation fault, or other memory problems
 - Abort, exit, error handlers
 - CUDA errors
- Scalable handling of error messages
- Leaps to the problem
 - Source code highlighted
 - Affected processes shown
 - Process stacks displayed clearly in parallel

- Full class/structure browsing
 - Local variables and current line(s)
 - Show variables relevant to current position
 - Drag in the source code to see more
 - C, C++, F90: object members, static members, derived types
- Automatic comparison and change detection
 - Scalable and fast



The screenshot shows the Allinea ierr tool interface. At the top, the 'Expression' is set to 'ierr'. Below it, the text reads 'Processes in current group (All, 8196 procs)'. There are checkboxes for 'Limit comparison to 1 s.f.' and 'Only show if:'. A 'Compare' button is visible. The main area is divided into a table and a 'Statistics' panel.

Values	Process(es)
-1	2195
0	0-2194,2196-8195

Statistics:

- Count: 8196
- Not shown: 0
- Errors: 0
- Aggregate: 0
- Numerical: 8196
- Sum: -1
- Minimum: -1
- Maximum: 0
- Range: 1
- Mean: -0.000122011
- Variance: 0.000122011
- nan: 0
- nan: 0
- inf: 0
- inf: 0

Buttons at the bottom: Use as MPI Rank, Create Groups, Export to Spreadsheet..., Close.

- Easy to find where the differences are...
 - Cross process comparison of data
 - Fetches values from every process, compares and then groups by value
 - Summary of NaN, Inf and statistics
 - Easy to spot rogues
- Use to group processes
 - Define a process group and control en-masse

Array Expression:

Distributed Array Dimensions: [How do I view distributed arrays?](#)

Range of \$x (Distributed) Range of \$i

From: From:

To: To:

Display: Display:

Auto-update

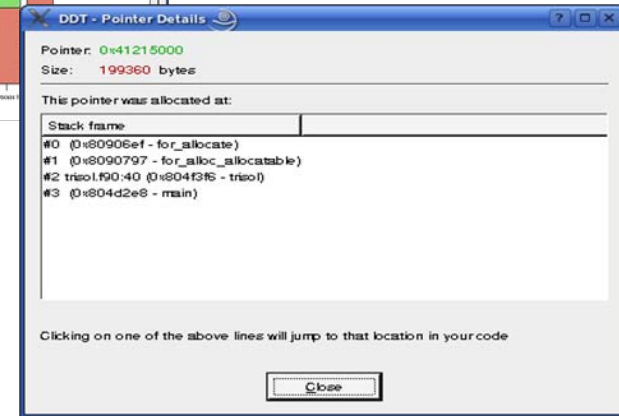
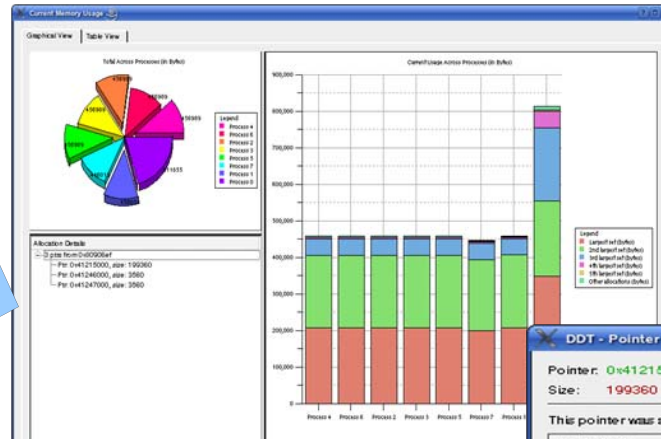
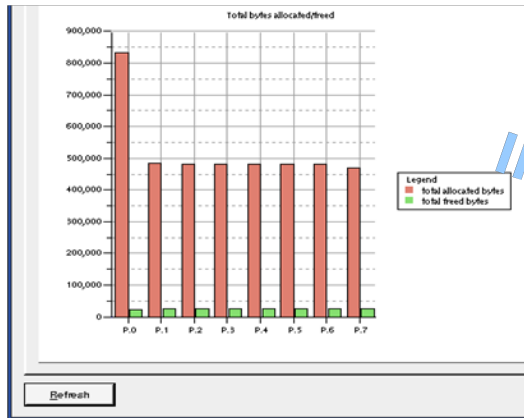
Only show if: [See Examples](#)

Data Table

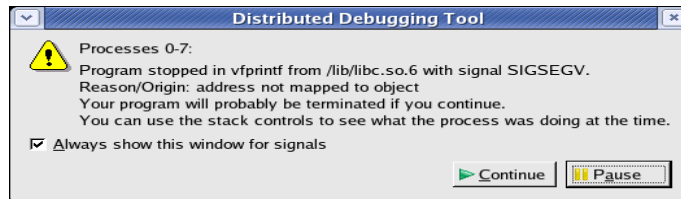
	i	2444	2733	3011	3185	4704	5343	6795	7881	9108	9467
x 0											
1				1					1		
2	1				1						1
3											
4		1									
5			1			1					
6											

- Browse arrays
 - 1, 2, 3, ... dimensions
 - Table view
- Filtering
 - Look for an outlying value
- Export
 - Save to a spreadsheet
- View arrays from multiple processes
 - Search terabytes for rogue data: in parallel with [v3.0]

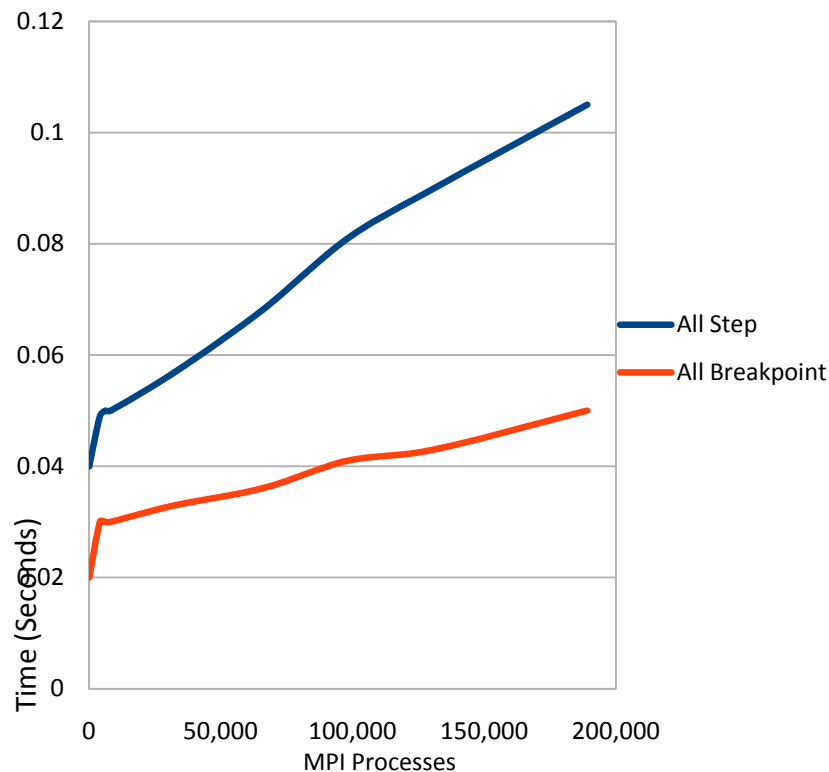
- Find memory leaks



- Or stop on read/write beyond end of array:



DDT 3.0 Performance Figures



- DDT is delivering Petascale debugging **today**
 - Collaborations with ORNL on Jaguar Cray XT and CEA
 - Logarithmic performance
 - Many operations now faster at 220,000 than previously at 1,000 cores
 - **~1/10th of a second** to step and gather all stacks at 220,000 cores

- Debuggers are recognised as the right tools to fix bugs quickly: other methods have limited success, and major issues at scale
- Debugging interfaces must scale to help the user understand what is happening
- Allinea DDT scales in performance and interface – breaking all records and making problems manageable
- See Allinea at Supercomputing 2010: Booth 2305

Questions?