allinea

Leaders in parallel software development tools

Allinea DDT
Intelligence Within

14 September 2012

Agenda

- Introduction
- Allinea DDT overview
- Conclusion

What is happening ?
Extreme Machine Size

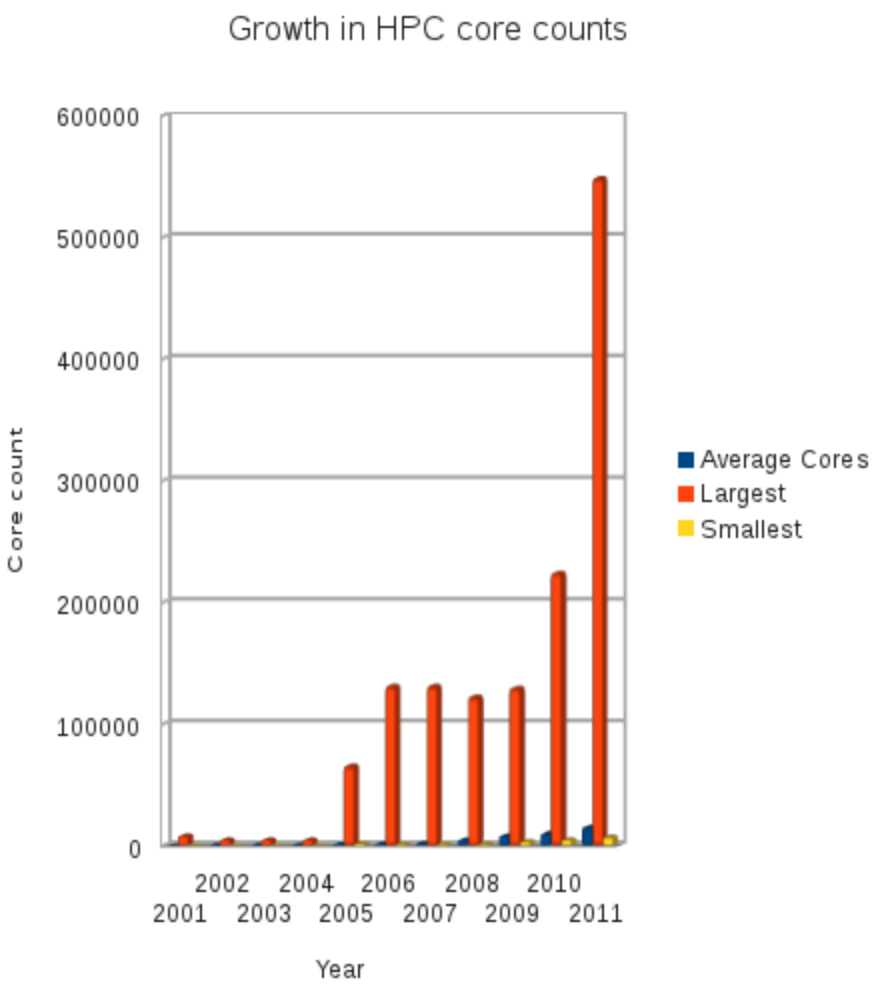

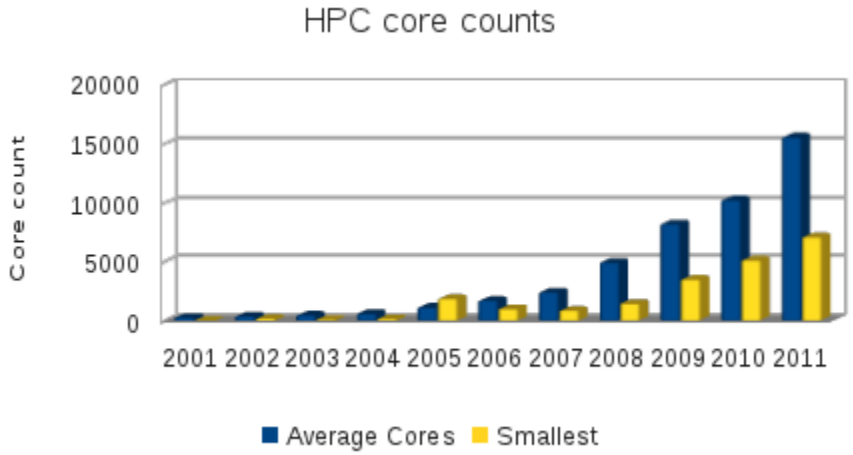

- Scientific progress requires more CPU hours
  - Maximum machine size is exploding
  - Average machine size grows exponentially

The Company

- Development tools company
  - Leading in HPC software tools market worldwide
  - Global customer base
    - Blue-chip engineering, government and academic research
- Allinea DDT
  - The leading debugger in parallel computing
  - Production use from desktop to extreme scale
  - World's only scalable debugger
    - Record holder for debugging software on largest machines
    - First at Petascale – and first for GPUs and ARM support!
- Allinea OPT
  - Profiling tool for parallel applications

allinea
www.allinea.com

Collaboration – National Labs

Partnership to develop Petascale debugger with NVIDIA support, with Cray and Caps Entreprise

Partnership on Full Scale debugging on IBM Blue Gene /P & /Q

Partnership with CEA French Atomic Energy Authority on scalable programming, CUDA and Allinea OPT

European partnership to develop techniques and solutions which address the exascale challenges

Partnership with BSC through its utilization of ARM technology to develop energy efficient HPC systems

allinea DDT

allinea
www.allinea.com

Application Development Life Cycles

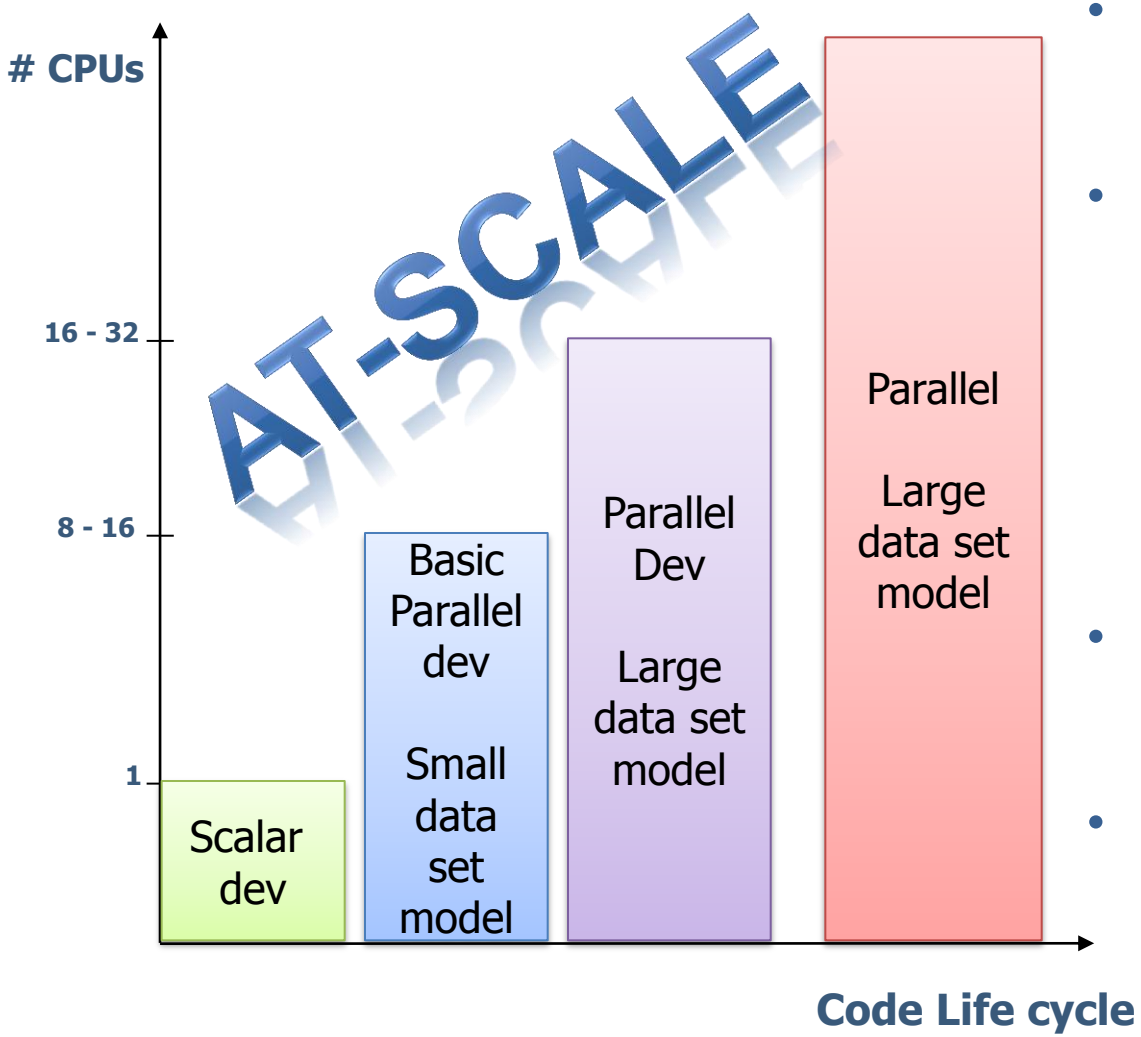

- Bugs are present at any stage
- Bugs get more and more tricky
  - More threads/processes
  - More data
  - Easy bugs already found
- Time to debug increases
- Debugging at scale : this is what you mostly do !

Example

```
$> mpirun –np 4 ./cpi.exe
      Process 0 on dopey
      Process 2 on dopey
      Process 1 on dopey
      Process 3 on dopey
      pi is approximately 3.1416009869231249
      Error is 0.0000083333333318
      wall clock time = 0.012800
```

Happy !!

```
$> mpirun –np 8 ./cpi.exe
      Process 5 on dopey
      Process 1 on dopey
      Process 4 on dopey
      Process 6 on dopey
      Process 3 on dopey
      Process 0 on dopey
      Process 2 on dopey
      Process 7 on dopey

      ^Cmpirun: killing job...
```

Deadlock !!

How to debug At-Scale ?

- Natural reflex : printf !
- No fully integrated solution : time consuming
  - Problems do not appear at smaller scale
  - Takes time to move the problems to a smaller size
  - Need to use multiple tools for multiple needs
  - More cores means more bugs...
  - ... And more debugging information available

What can we do... ?

HPC's current challenge

- GPUs – a rival to traditional processors
  - AMD, Intel, ARM
  - NVIDIA, OpenCL, CUDA
- A big challenge for HPC developers
  - Data transfer
  - Several memory levels
  - Grid/block layout and thread scheduling
  - Synchronization
- New languages, compilers, potential standards

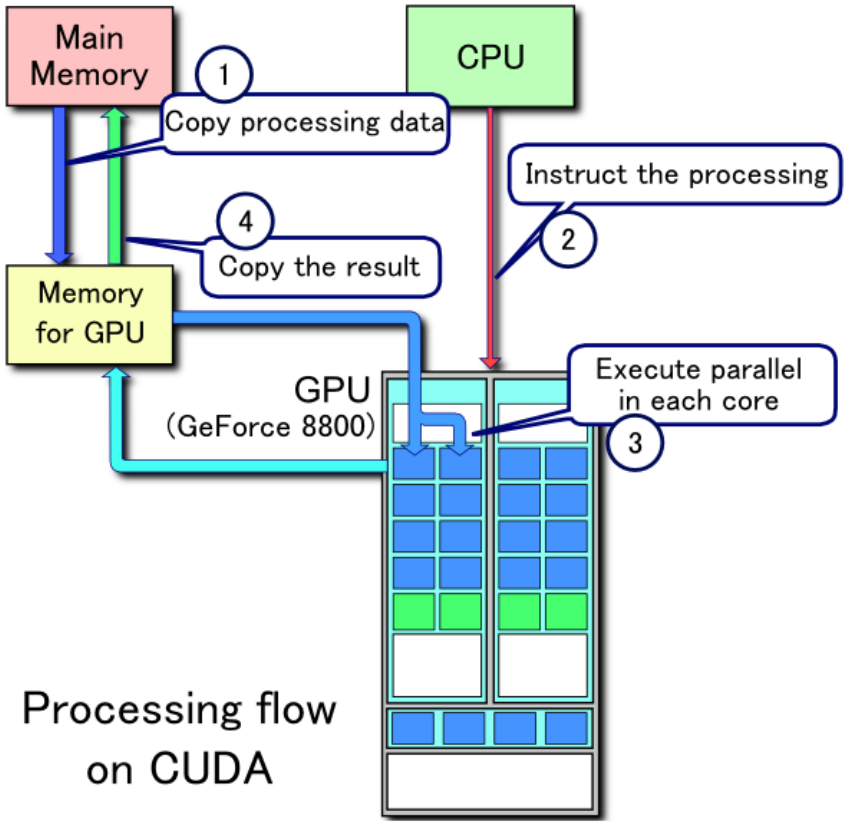

Allinea DDT
In a Nutshell

- Graphical debugger designed for:
  - C/C++, Fortran, UPC, CUDA
  - Multithreaded code
    - Single address space
  - Multiprocess code
    - Interdependent or independent processes
  - GPU codes
    - Hybrid software
  - Any mix of the above
- Managing concurrency
  - Emphasizing differences
  - Collective control
- Strong feature set
  - Memory debugging
  - Data analysis

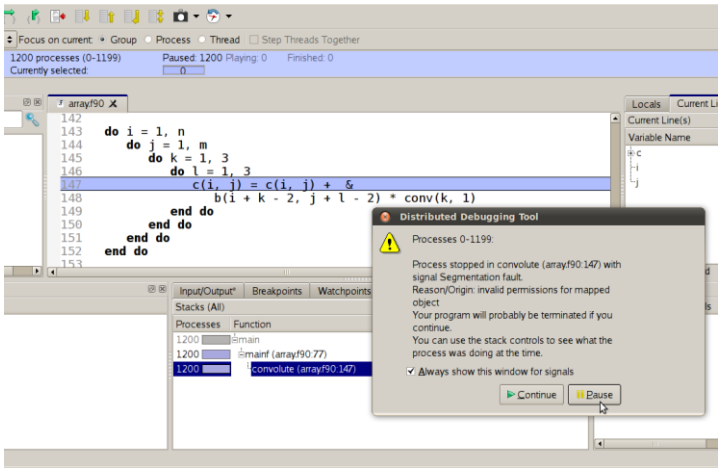

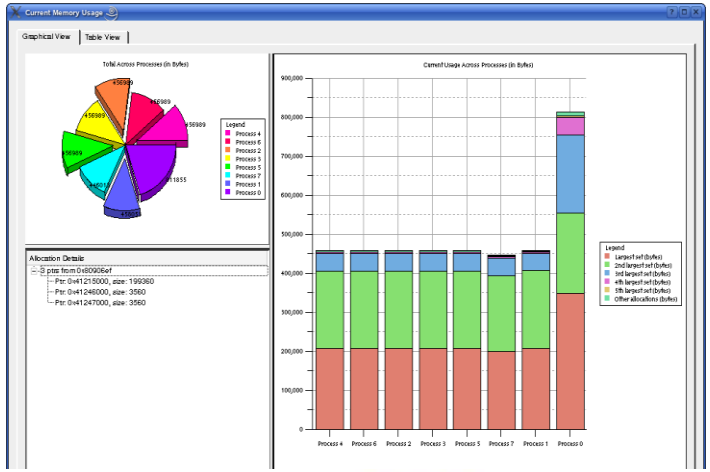

allinea
www.allinea.com

Allinea DDT

Gather, Sort and Display information

- User and administrator friendly
- Flow control
- Data monitoring
- Many environments

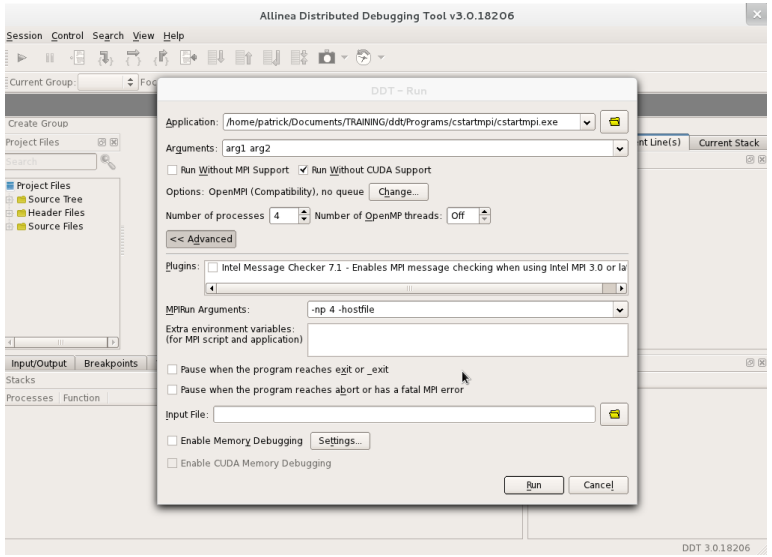

allinea
www.allinea.com

Allinea DDT

Gather, Sort and Display information

- **User and administrator friendly**
  - Get started easily
  - Fast, reliable, simple and intuitive GUI interface
  - Offline debugging
- Flow control
- Data monitoring
- Many environments

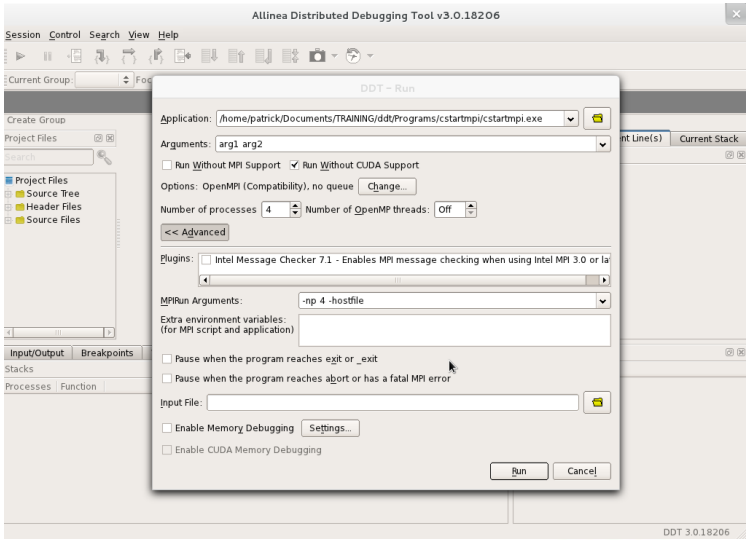

allinea
www.allinea.com

User and administrator friendly
Getting started

- Quick creation of runs
- Well integrated in workload schedulers
  - SLURM, LSF, SGE, PBS,....

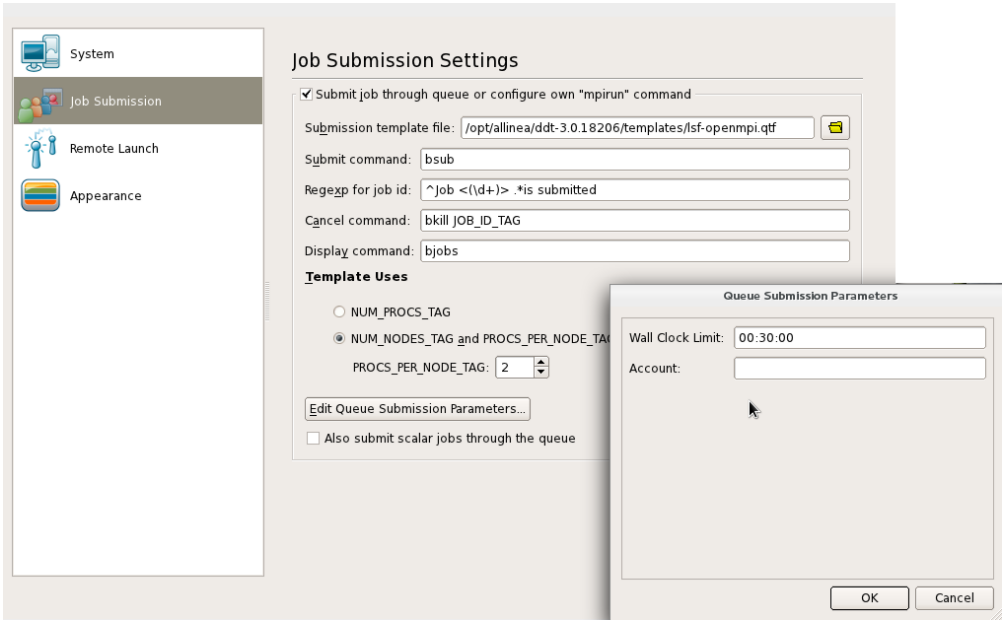

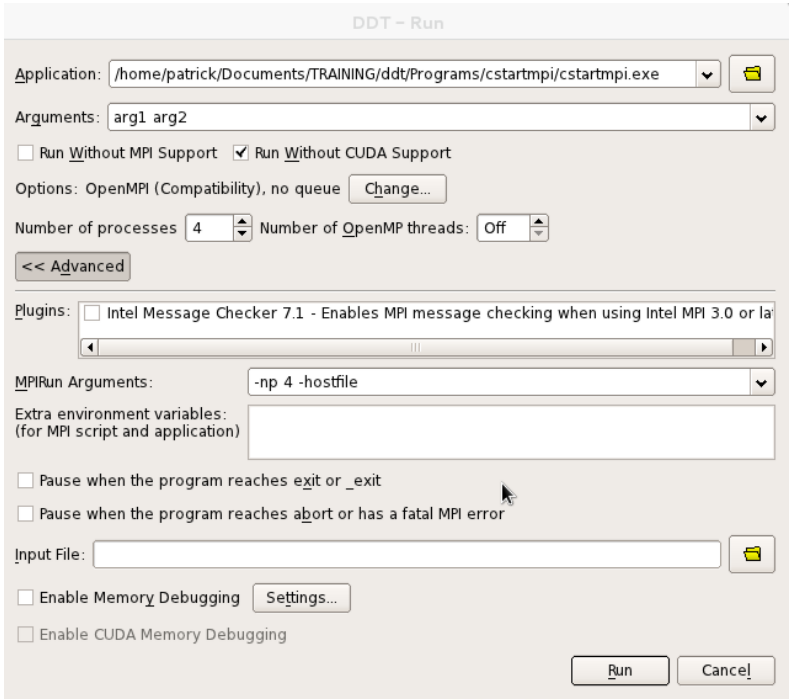

allinea
www.allinea.com

User and administrator friendly
Intuitive GUI interface

- Intelligent GUI that adapts to the environment
  - From workstation to large scale clusters

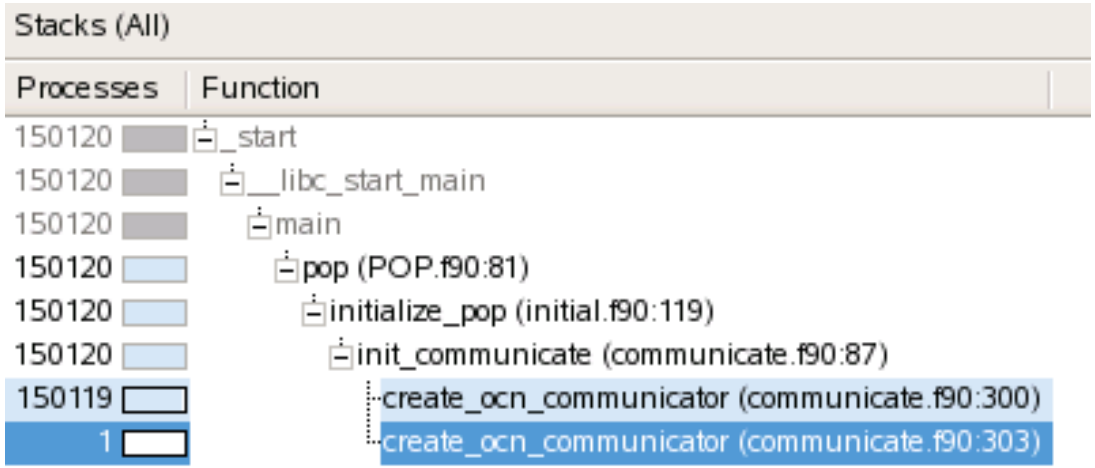

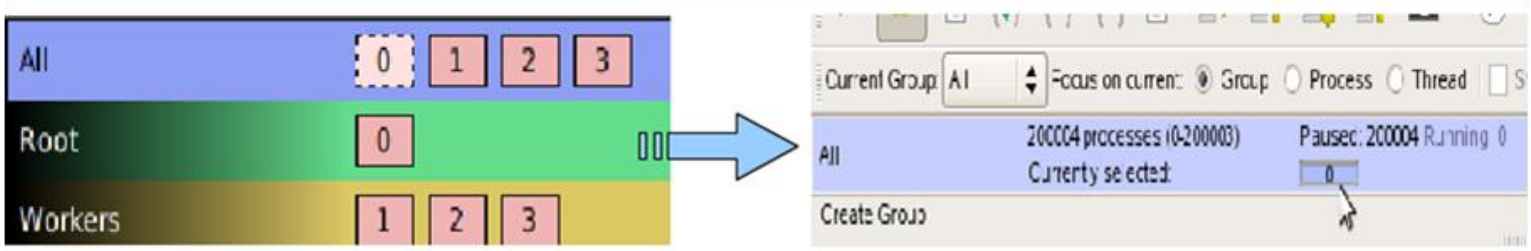

allinea
www.allinea.com

User and administrator friendly
Intuitive GUI interface

- Intelligent GUI that adapts to the environment
  - Even in CUDA environments

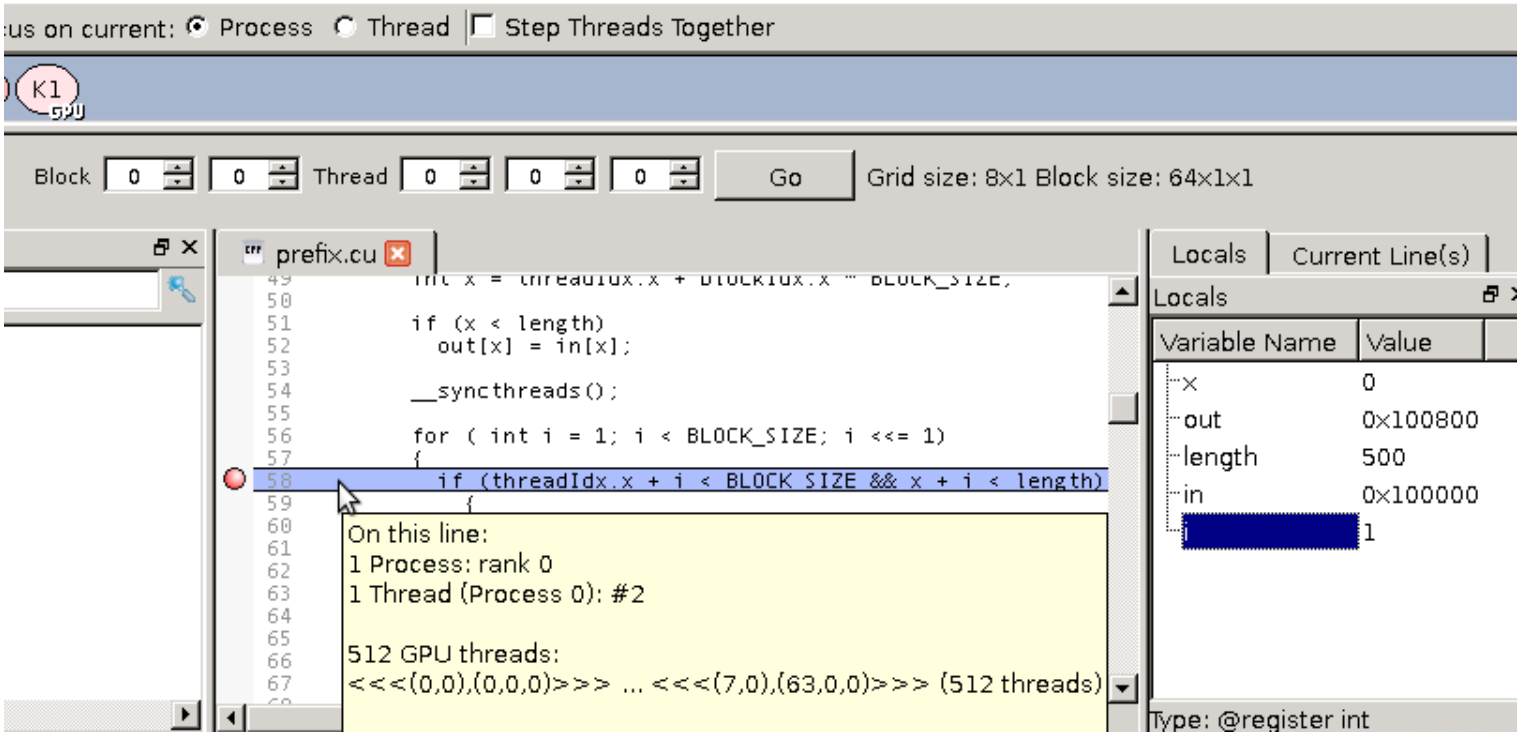

allinea
www.allinea.com

User and administrator friendly Offline debugging

- Using a workload scheduler
  - Machines are available when the scheduler decides (by night ?)
  - Can be tricky to get a big cluster exactly when the developer wants it
- Offline debugging : printf replacement
  - Tracepoints and offline debugging
  - Job runs without debugger interface and record variables
- Worlds first scalable batch debugger
  - Set tracepoints, breakpoints, and run !
  - Memory debugging errors, crashes
  - Reports in HTML or plain text

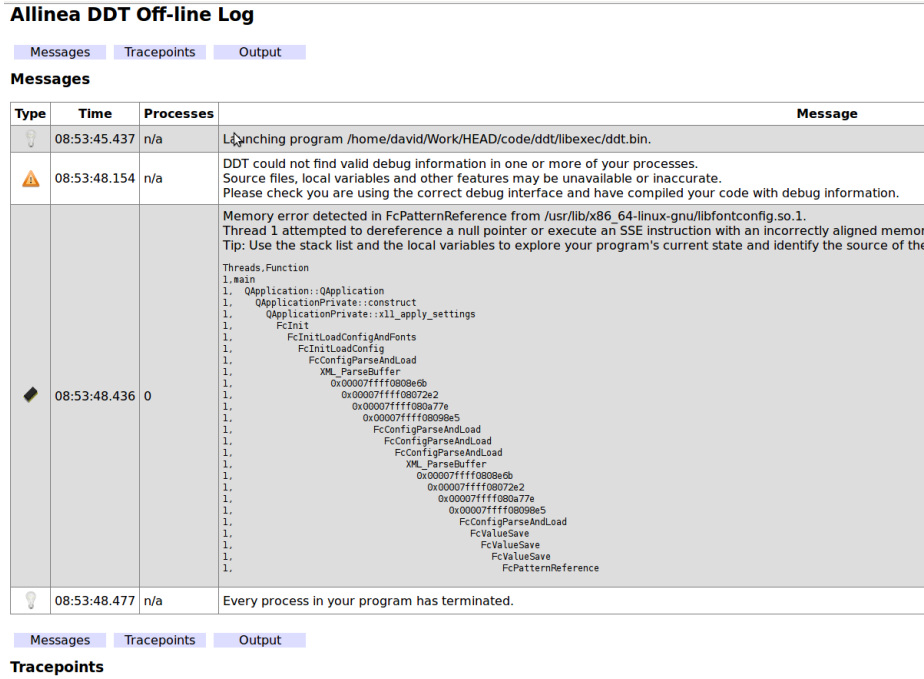

Allinea DDT Off-line Log

Messages | Tracepoints | Output

Messages

Type	Time	Processes	Message
	08:53:45.437	n/a	Launching program /home/david/Work/HEAD/code/ddt/libexec/ddt.bin.
	08:53:48.154	n/a	DDT could not find valid debug information in one or more of your processes. Source files, local variables and other features may be unavailable or inaccurate. Please check you are using the correct debug interface and have compiled your code with debug information.
	08:53:48.436	0	Memory error detected in FcPatternReference from /usr/lib/x86_64-linux-gnu/libfontconfig.so.1. Thread 1 attempted to dereference a null pointer or execute an SSE instruction with an incorrectly aligned memor Tip: Use the stack list and the local variables to explore your program's current state and identify the source of the Threads,Function 1,main 1, QApplication::QApplication 1, QApplicationPrivate::construct 1, QApplicationPrivate::x11_apply_settings 1, FcInit 1, FcInitLoadConfigAndFonts 1, FcInitLoadConfig 1, FcConfigParseAndLoad 1, XML_ParseBuffer 1, 0x00007ffff0808e6b 1, 0x00007ffff08072e2 1, 0x00007ffff080a77e 1, 0x00007ffff08098e5 1, FcConfigParseAndLoad 1, FcConfigParseAndLoad 1, FcConfigParseAndLoad 1, XML_ParseBuffer 1, 0x00007ffff0808e6b 1, 0x00007ffff08072e2 1, 0x00007ffff080a77e 1, 0x00007ffff08098e5 1, FcConfigParseAndLoad 1, FcValueSave 1, FcValueSave 1, FcValueSave 1, FcPatternReference
	08:53:48.477	n/a	Every process in your program has terminated.

Messages | Tracepoints | Output

Tracepoints

allinea
www.allinea.com

Allinea DDT and Offline Debugging Using the CLI

- General Options :

```
-offline <YYY>   : activate offline mode
-n <XXX>         : the number of MPI processes
-memory          : enable memory debugging
-ddtsession      : the session file to use (if exists)
```

- Create breakpoints : `-break-at LOCATION[,N:M:P]`

Example:

```
$> ddt –offline myjob.html –break-at main.c:22,-:2:-,x myapp
```

- Create tracepoints : `-trace-at LOCATION[,N:M:P],VAR1,VAR2...`

Example:

```
$> ddt –offline myjob.html –trace-at trisol.F90:14, -:-:3,x myapp
```

Allinea DDT

Gather, Sort and Display information

- User and administrator friendly
- Flow control
  - Static analysis
  - Control progress at scale
  - Understand deadlocks
  - Start investigation
- Data monitoring
- Many environments

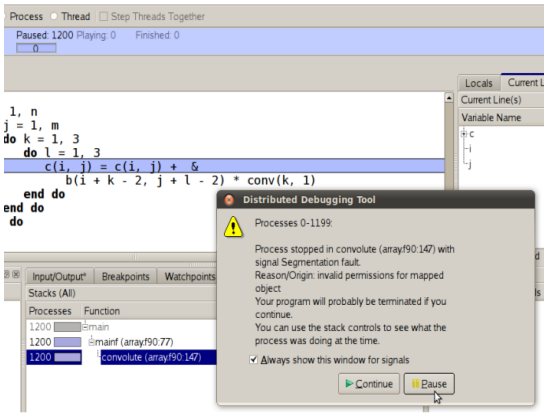

allinea
www.allinea.com

Flow control
Static analysis

- Fix those errors before they bite !
- Static analysis
  - Integrated with cppcheck
  - Also includes ftncheck

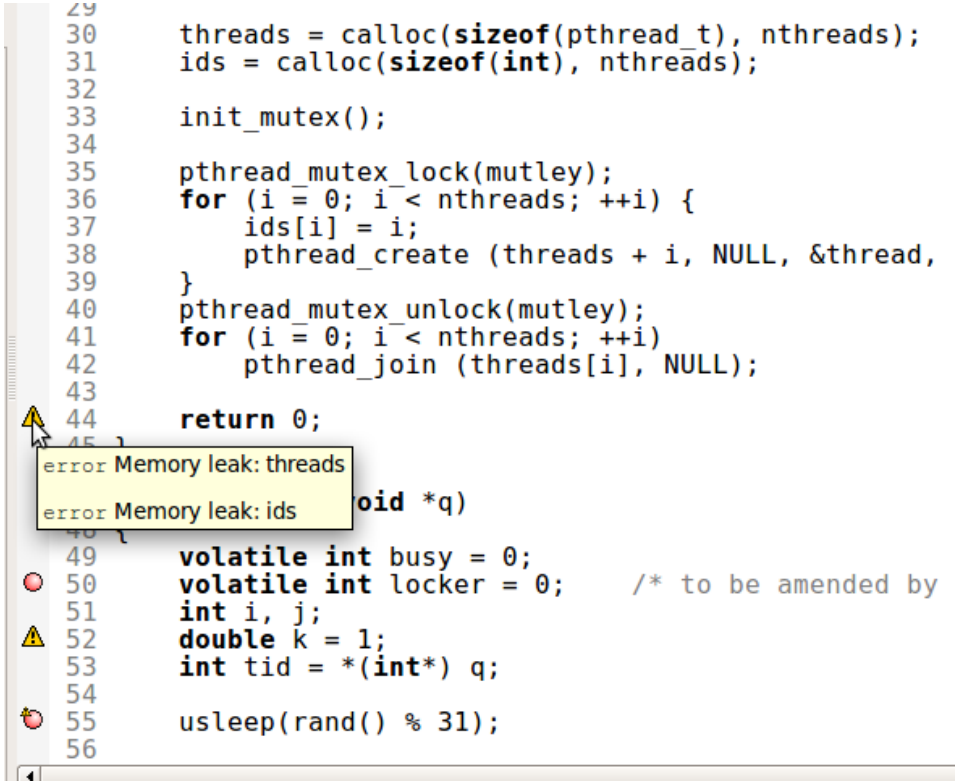

allinea
www.allinea.com

Flow control
Controlling progress at scale

- Bulk control is essential for multicore debugging
  - Group processes together
  - Play, step, reach breakpoints... Based on groups
  - Change interleaving order by stepping/playing selectively

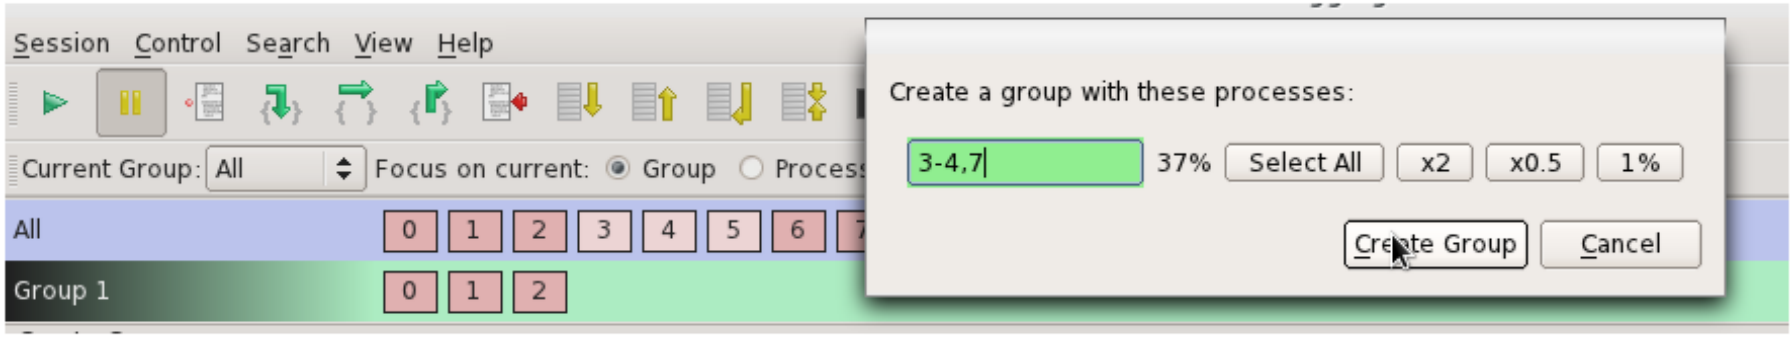

allinea
www.allinea.com

Flow control
Resolving MPI issues

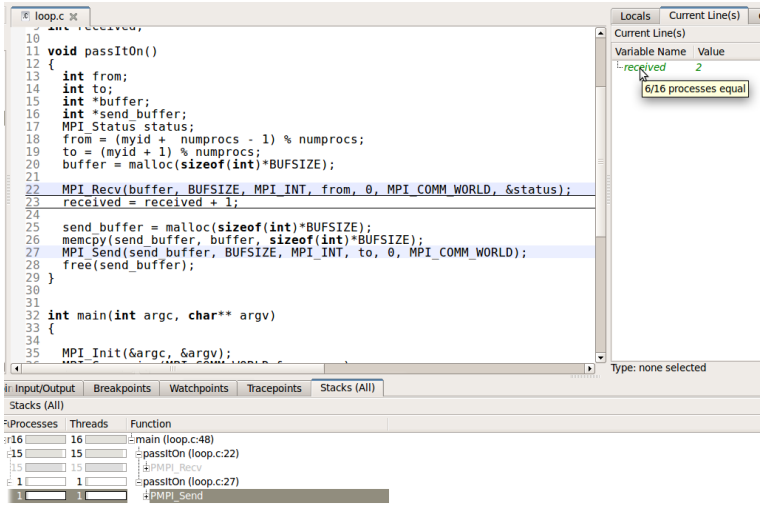


```
void passItOn()
{
  int from;
  int to;
  int *buffer;
  int *send_buffer;
  MPI_Status status;
  from = (myid +  numprocs - 1) % numprocs;
  to = (myid + 1) % numprocs;
  buffer = malloc(sizeof(int)*BUFSIZE);

  MPI_Recv(buffer, BUFSIZE, MPI_INT, from, 0, MPI_COMM_WORLD, &status);
  received = received + 1;

  send_buffer = malloc(sizeof(int)*BUFSIZE);
  memcpy(send_buffer, buffer, sizeof(int)*BUFSIZE);
  MPI_Send(send_buffer, BUFSIZE, MPI_INT, to, 0, MPI_COMM_WORLD);
  free(send_buffer);
}


int main(int argc, char** argv)
{

  MPI_Init(&argc, &argv);
```

- We can see messages
  - MPI standard exists for debugging message queues
- Integrates with MPI correctness tools
  - Check correctness of the messages
  - Intel MPI Checker ; Marmot

- Find cycles or blockage in message queue display
  - Parallel stacks
  - Variables
- More details than examining variables and processes alone

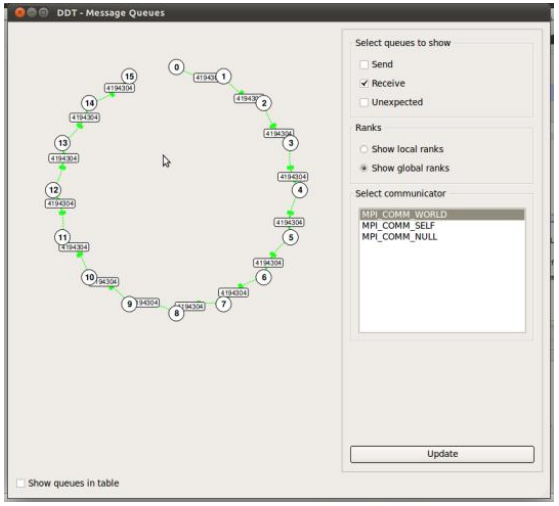

allinea
www.allinea.com

Flow control
Understanding what happens

- Application crashes
  - Threads/processes can be anywhere
  - Impossible to scroll through them individually
- Finding where processes crashed is essential
  - Allinea DDT merges stacks from processors and threads into a tree
  - Common faults patterns instantly evident
    - Divergence, deadlocks...
  - Information scalable without overload

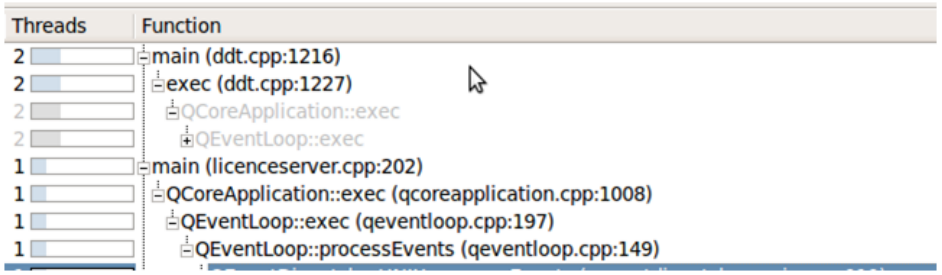

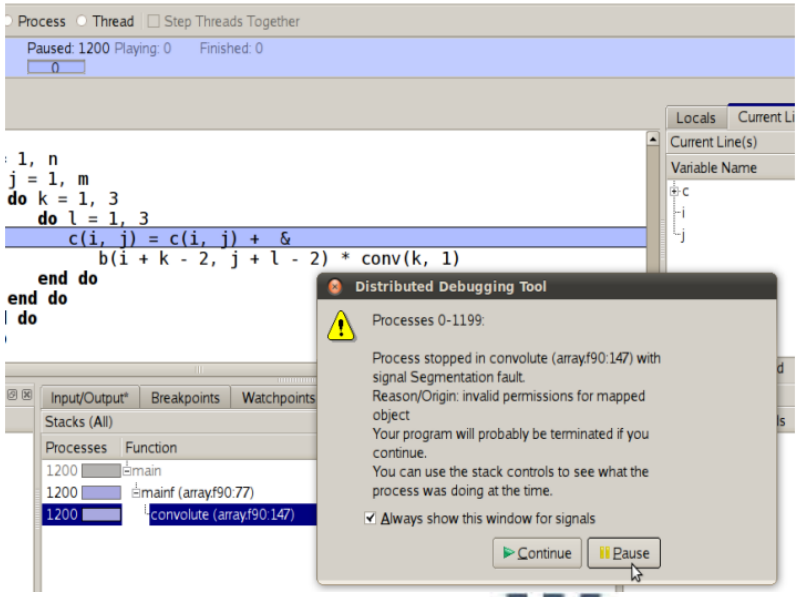

allinea
www.allinea.com

Allinea DDT

Gather, Sort and Display information

- User and administrator friendly
- Flow control
- Data monitoring
  - Monitor variables
  - Detect memory errors
  - Check the calculation data
- Many environments

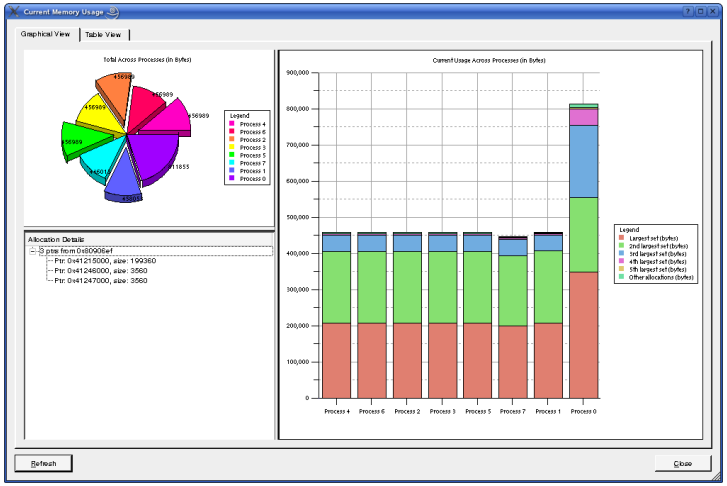

allinea
www.allinea.com

Data monitoring
Monitor variables

- Developers need to see data
  - Too many variables to trawl them manually
  - Too many tasks or thread to display them at the same time

- Intelligent data management inside the debugger
  - Automatic monitoring of the data
  - Subtle highlights differences
  - Sparklines and smart display

- Even more detailed analysis if needed
  - Cross process comparison
  - Historical values of variables

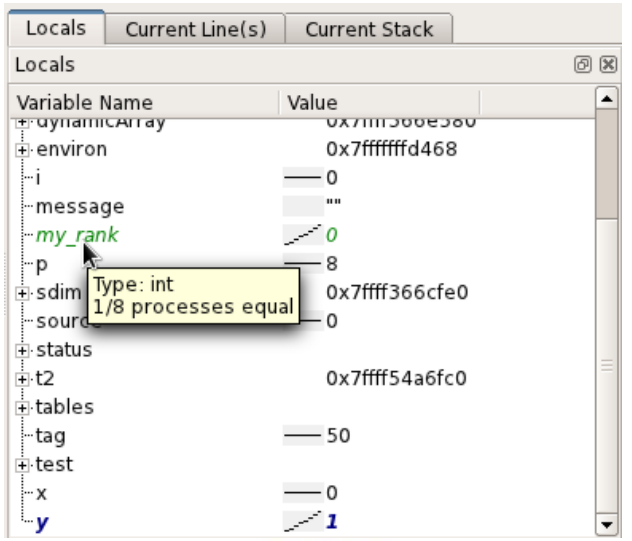

Data monitoring
Smart displays : tracepoints

- “printf” is still mostly used but serious drawbacks:
  - Need to recompile the code
  - Information randomly printed on screen (depending on interleaving)
  - One line per process

- Scalable and advanced printf :
  - First step to reconcile printf and GUI
  - Information sorted by steps
  - Merged by groups of processes
  - No information overload
  - Possibility to filter the printed values
  - Save output for offline analysis

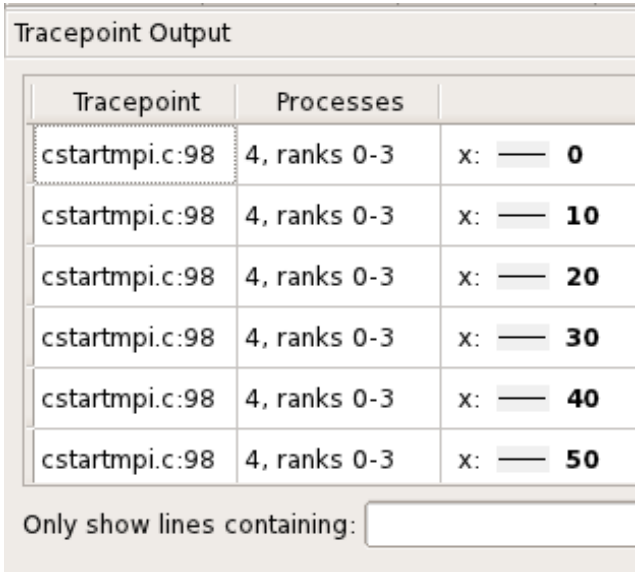

Data monitoring
Memory debugging

- Random errors are sneaky
  - Can’t fix a bug that doesn’t repeat
  - Often caused by memory issues
- Memory debugging can force the bug
  - Better to happen every time, than only during product demos

- Memory debugging :
  - Places agent between memory library and user process
  - Communicates problems to the debuggers
  - Monitors usage : detect memory leaks
  - Automatically protects ends of arrays
  - Trigger instant stop on touching memory
  - Many classes of errors can be checked
  - Also has CUDA support

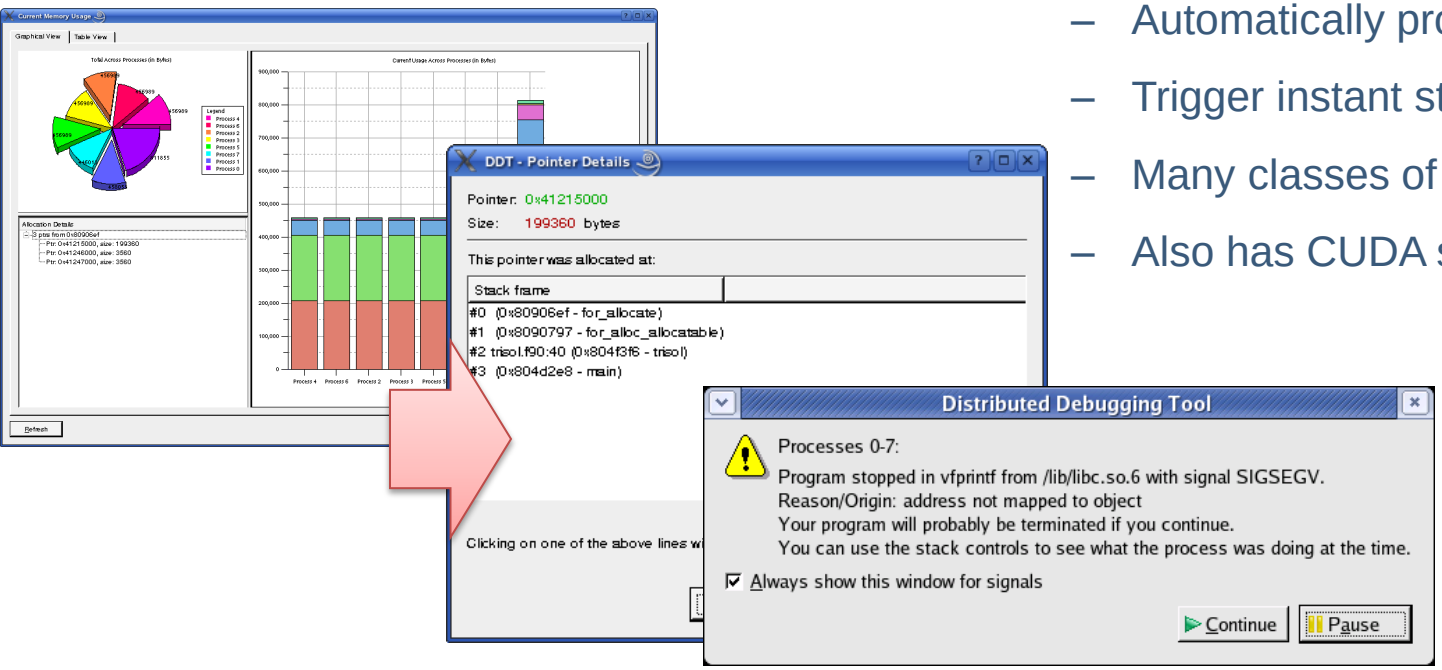

allinea
www.allinea.com

Data monitoring
Searching haystacks

- Arrays are the building blocks of HPC
  - Largest jobs accumulate TB of data
  - Usually 2GB or 4GB per core

- Integrated visualization tool
  - Search data across all tasks or threads
  - Data displayed on a picture as requested by the user
  - Export at runtime

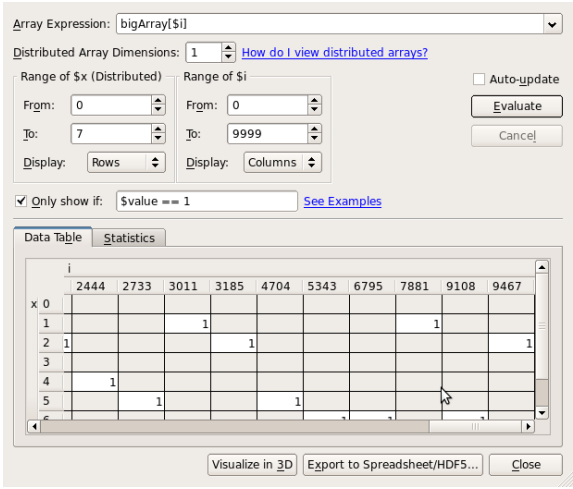

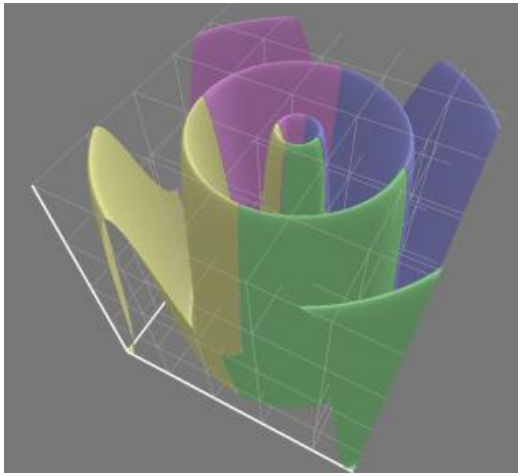

allinea
www.allinea.com

Data Monitoring
Unified Parallel C

- New programming models
  - Support for Cray UPC
  - Support for Cray Co-Array Fortran

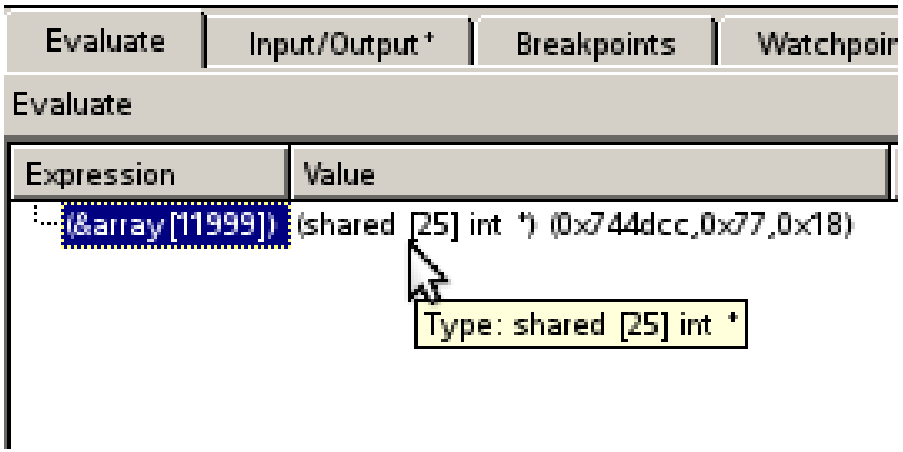

Allinea DDT

Gather, Sort and Display information

- User and administrator friendly
- Flow control
- Data monitoring
- Many environments
  - Debugging at scale
  - GPU Debugging

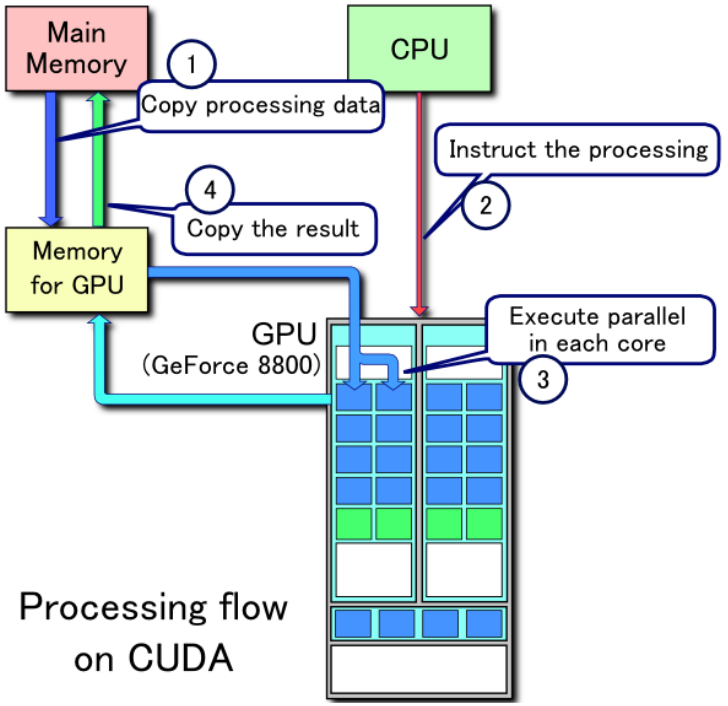

allinea
www.allinea.com

Allinea DDT and Scalability
Simplifies debugging for everyone

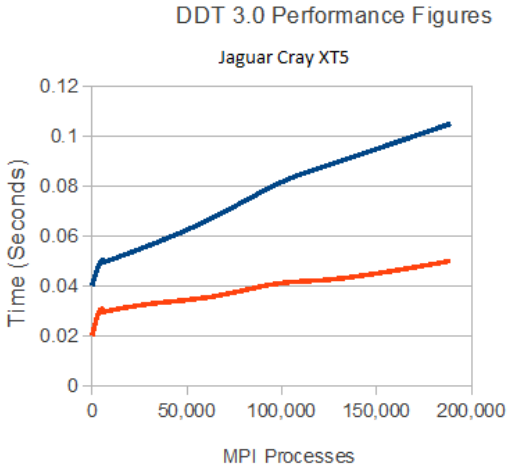

- At scale : new problems appear
  - Small and medium environments : multiple tools
  - More cores means more bugs...
  - ... And more debugging information available
- How can the design of Allinea DDT help your debugging whatever the size of your cluster (workstation or cluster) ?
- High-end architecture (in any Allinea DDT version)
  - Data consolidation (merged and sorted)
  - Highlight differences
  - Very low footprint

allinea
www.allinea.com

Allinea DDT and GPU Successful add-on

- Built on vendors low level efforts
  - Nvidia cuda-gdb, compiler
  - Cray compiler
- Execution model is unusual
  - GUI supports blocks and grids
  - Support 32 thread units (warps)
- Mixed GPU/CPU in one interface
  - Interaction with CPUs
  - Easy to switch between contexts (stacks, threads, data...)
  - Support multiple nodes

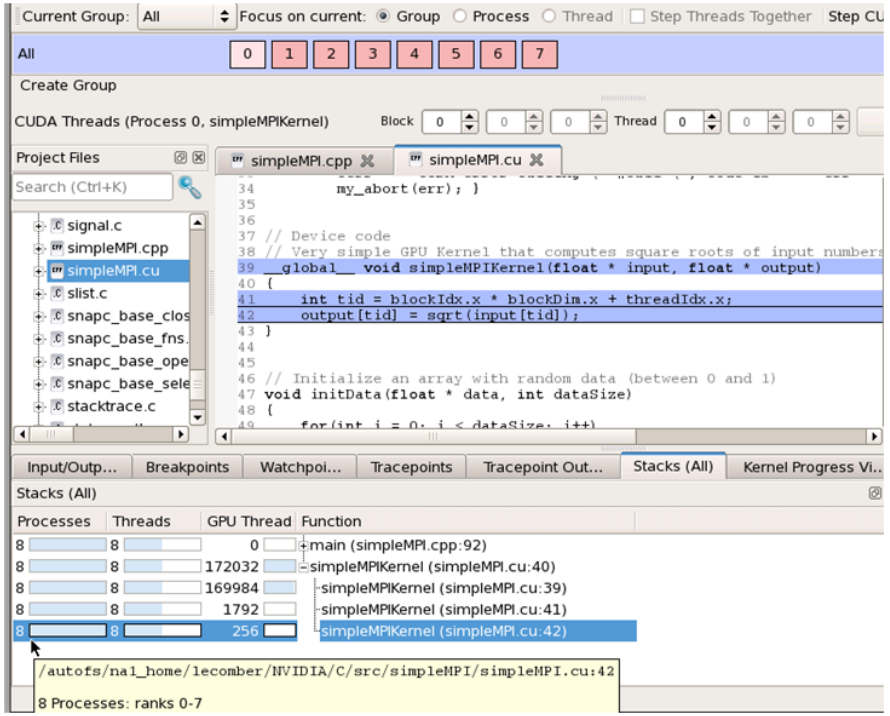

allinea
www.allinea.com

Many environments
Directives support

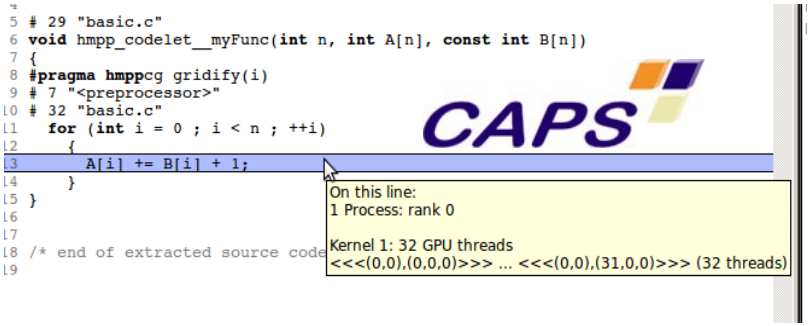

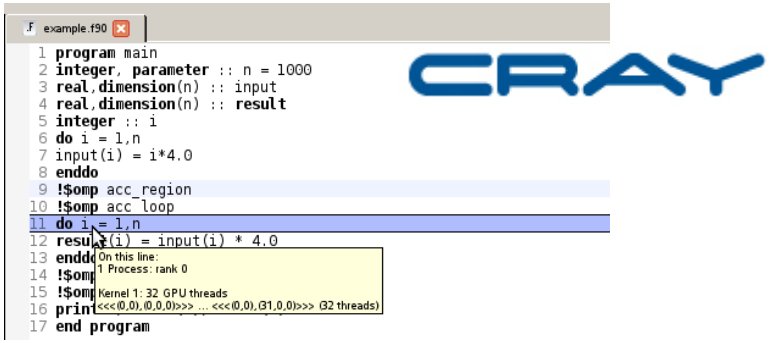

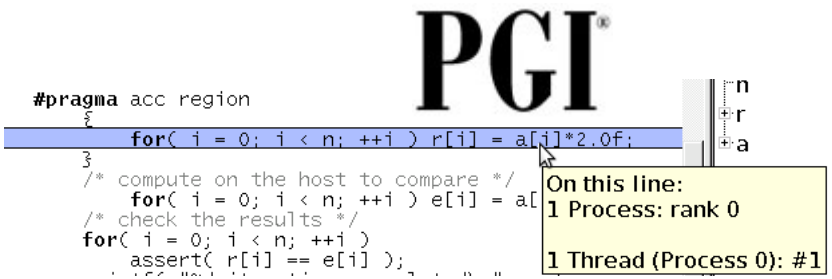

- Wide range of partnerships
  - Support the environments you use for hybrid development
  - You swiftly benefit from the latest updates
  - Read our latest white papers available with PGI and CAPS

allinea
www.allinea.com

Many environments
Directives support - OpenACC

- OpenACC : New parallel programming standard led by
  - CAPS Enterprise, Cray, Nvidia and PGI
- Fully supported within Allinea DDT
  - CAPS and PGI : still adapting their software to the new standard
  - Same mechanisms and features for CUDA programs
  - First implementation to be supported: Cray OpenACC
- You benefit from our partnerships

Summary

- Bugs happen at all stage of the development and during the entire life of the code
- An intelligent debugger is now available to fix bugs quickly
  - Other methods have limited success and issues at scale
  - Intelligence makes debugging easier and faster for the developer
- Allinea DDT scales in both performance and interface
  - Breaking all records and making problems now manageable

allinea

Leaders in parallel software development tools

Thank you

Your contacts :

- Technical Support team : support@allinea.com
- Sales team : sales@allinea.com