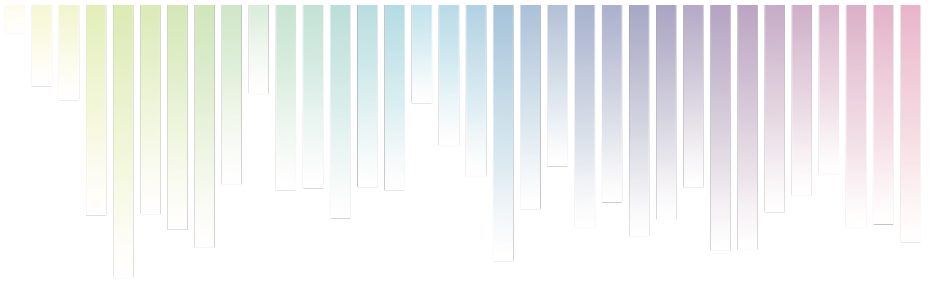

allinea

High performance tools to debug, profile, and analyze your applications

Debugging and Profiling your HPC Applications

Ryan Hulguin, Application and Support Analyst
rhulguin@allinea.com

About this talk

allinea FORGE DDT + MAP

- Learn how to debug and profile your code
  - Techniques to take home
- Tools we will use: Allinea Forge
  - Debugging with Allinea DDT
  - Profiling with Allinea MAP
  - NB. Allinea MAP is not supported on BG/Q
- Where to find Allinea's tools
  - > 70% of Top 500 have at least one Allinea tool

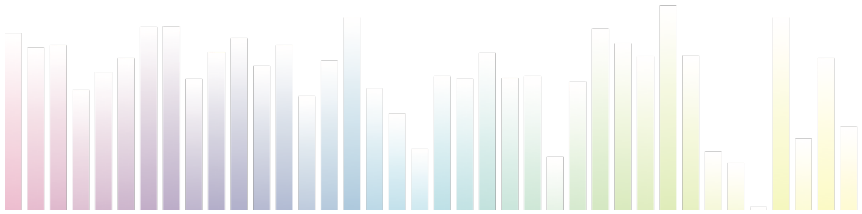

allinea

Motivation

- "Without capable highly parallel software, large supercomputers are less useful"
  - Council on Competitiveness
- "1% of HPC application codes can exploit 10,000 cores"
  - IDC, 2011

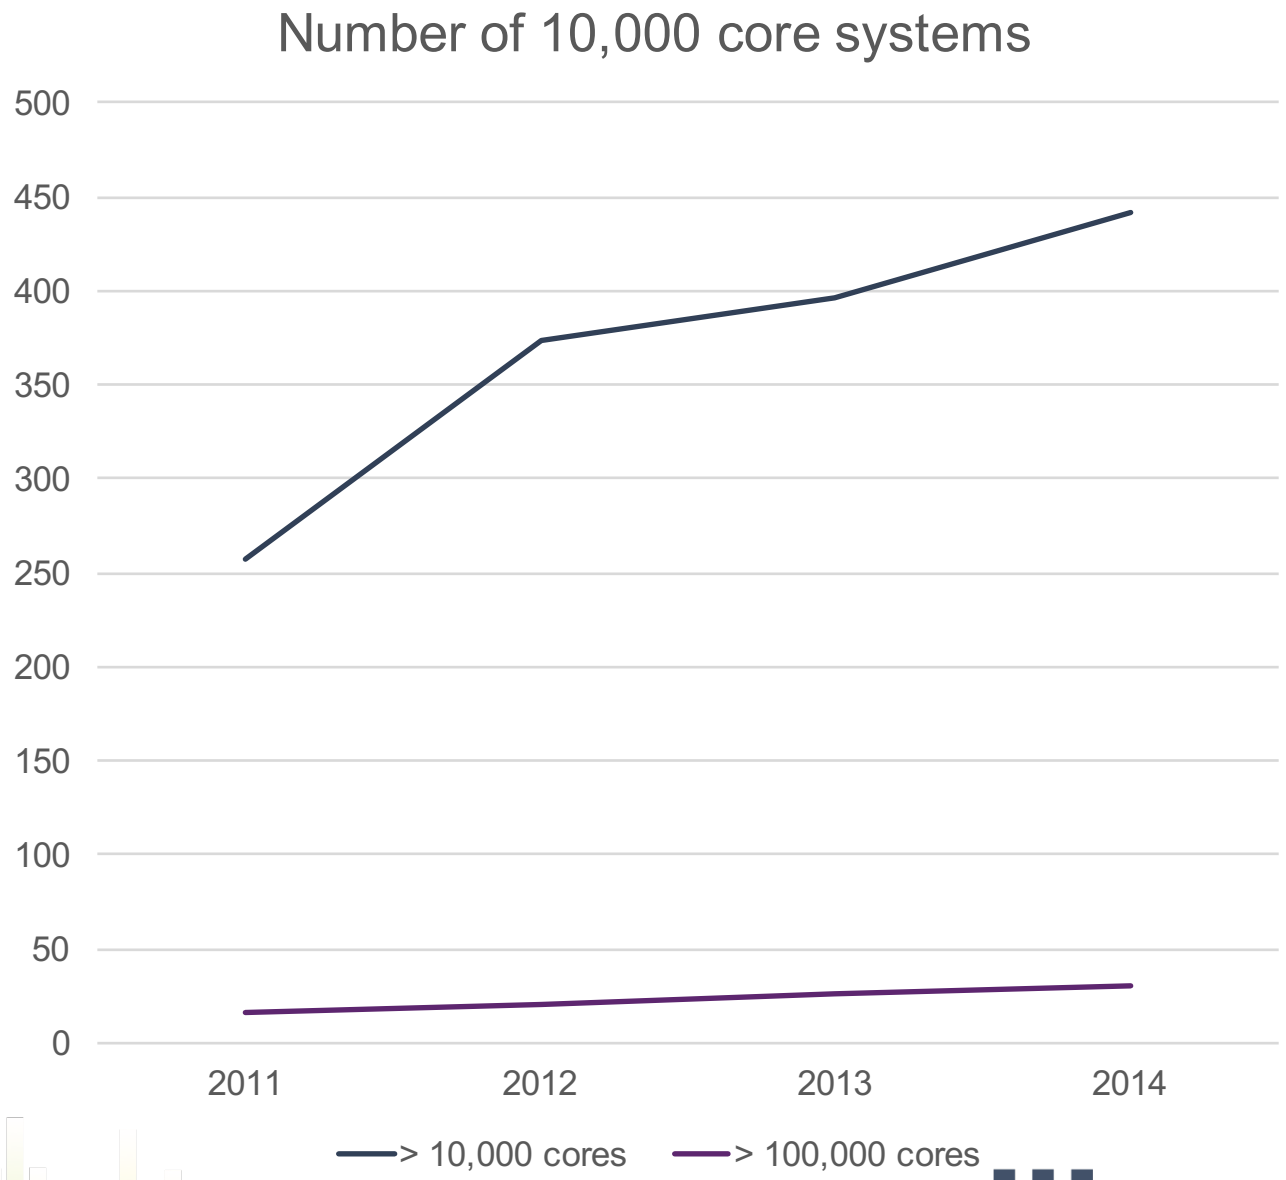

Use the right software tool to be faster

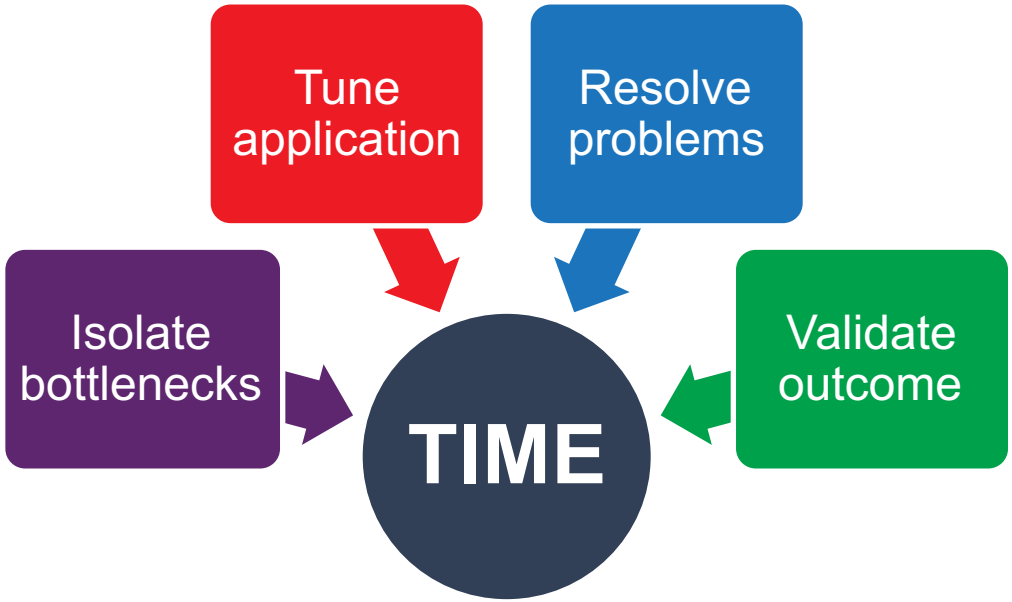

- What parts of the code would benefit most from being rewritten?
- How should I modify a code to make it better (or work at all)?

Application Development Workflow

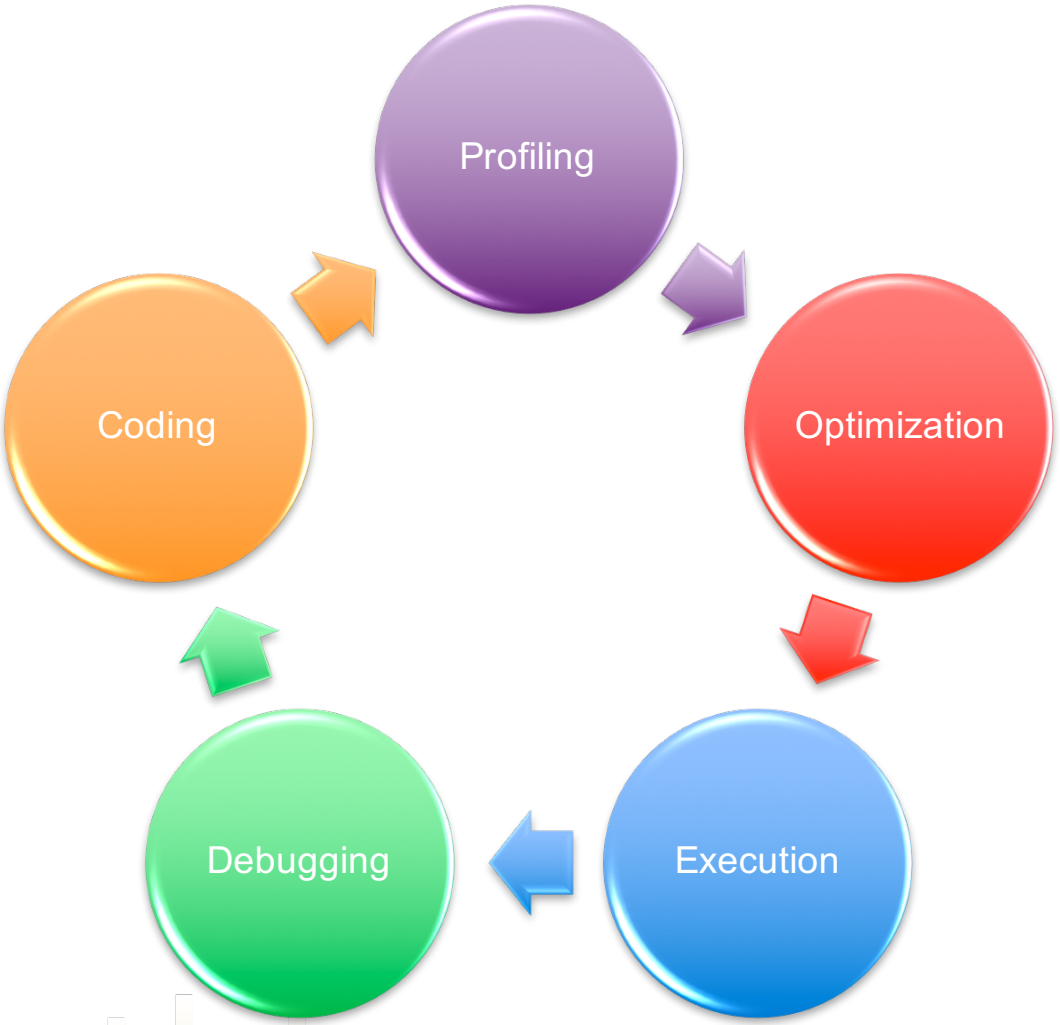

Hello Allinea Forge!

Allinea MAP to find performance bottleneck

Increasing memory usage? *Memory leak!*
Workload imbalance? *Possible partitioner bug!*

Flick to Allinea DDT
Common interface and settings files

Observe and debug your code step by step

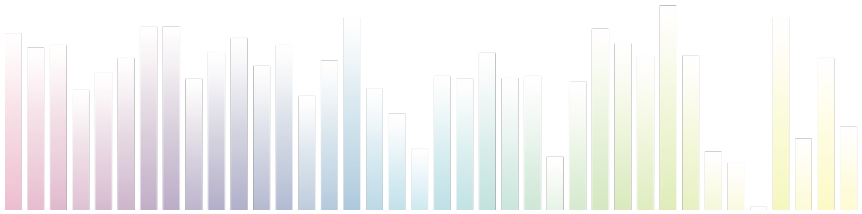

allinea

HPC means being productive on remote machines

- Linux
- OS/X
- Windows
- Multiple hop SSH
- RSA + Cryptocard
- Uses server license

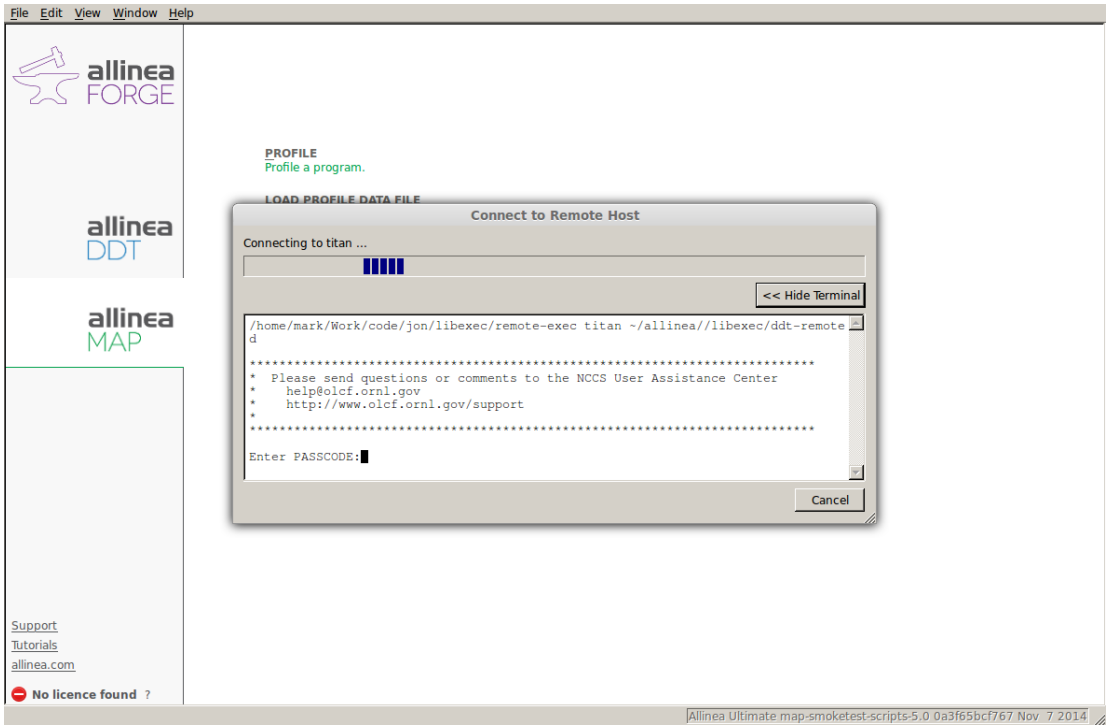

allinea FORGE

Profiling for performance

allinea MAP

- Code optimisation can be time-consuming…

  - (image courtesy of xkcd.com)

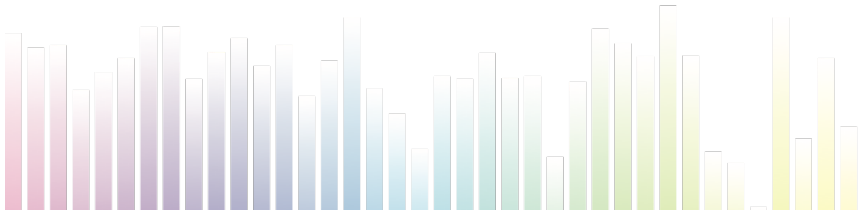

allinea

MAP in a nutshell

- Small data files
- <5% slowdown
- No instrumentation
- No recompilation

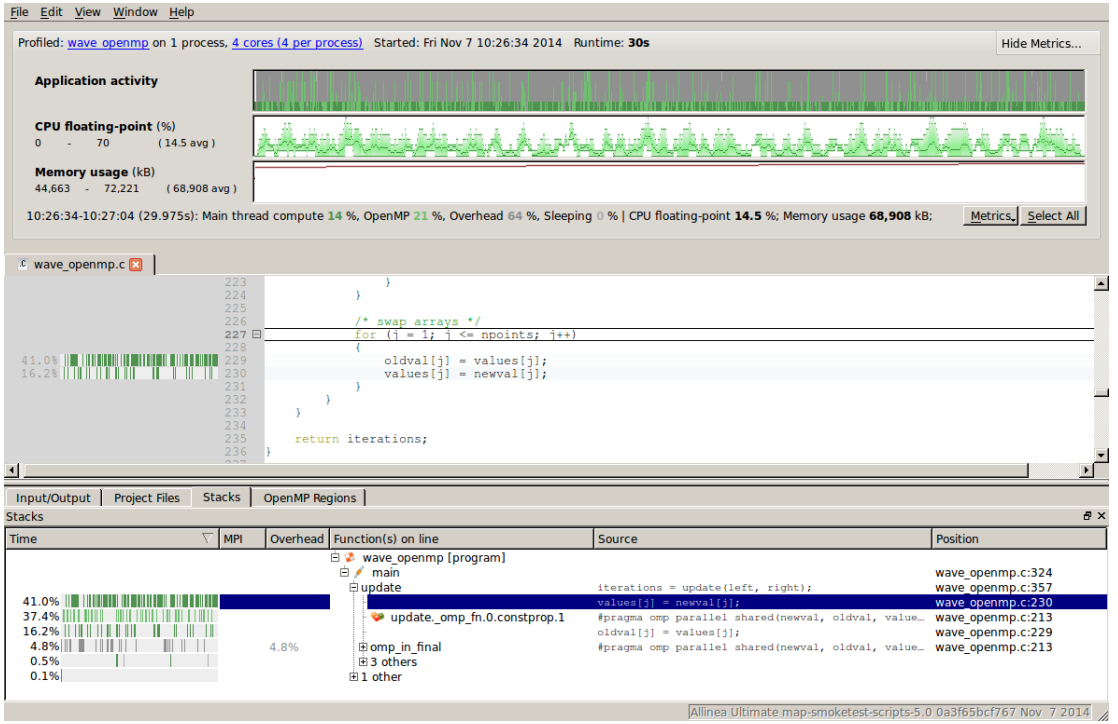

Scaling issue – 512 processes

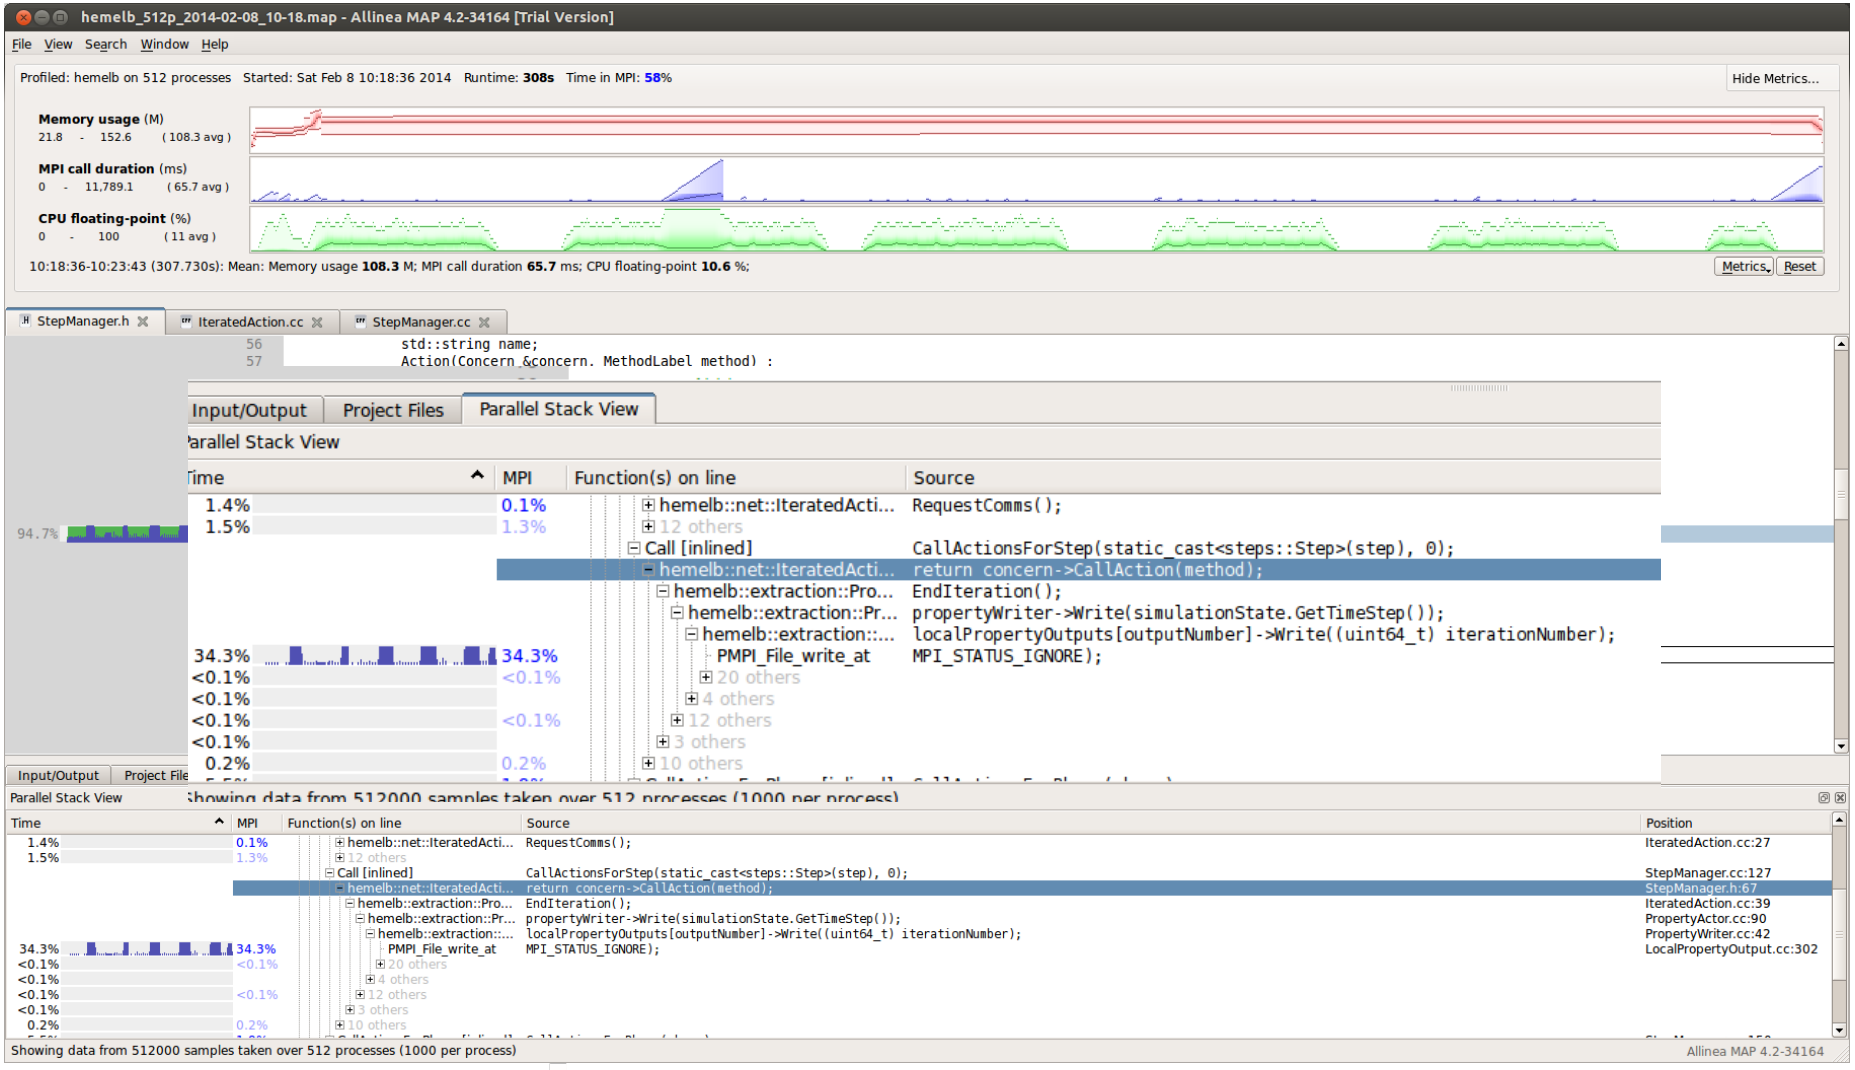

Simple fix… reduce periodicity of output

allinea FORGE

Some types of bug

allinea DDT

Bohrbug	Steady, dependable bug
Heisenbug	Vanishes when you try to debug (observe)
Mandelbug	Complexity and obscurity of the cause is so great that it appears chaotic
Schroedinbug	First occurs after someone reads the source file and deduces that it never worked, after which the program ceases to work

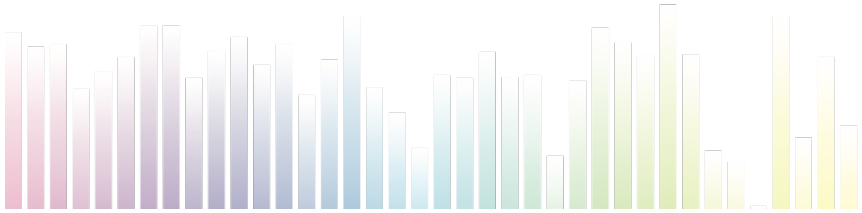

allinea

allinea FORGE

Debugging

allinea DDT

- Transforming a broken program to a working one
- How? TRAFFIC!
  - Track the problem
  - Reproduce
  - Automate - (and simplify) the test case
  - Find origins – where could the “infection” be from?
  - Focus – examine the origins
  - Isolate – narrow down the origins
  - Correct – fix and verify the testcase is successful

- Suggested Reading:
  - Zeller A., “Why Programs Fail”, 2nd Edition, 2009
  - Zen and the Art of Motorcycle Maintenance, Robert M. Pirsig

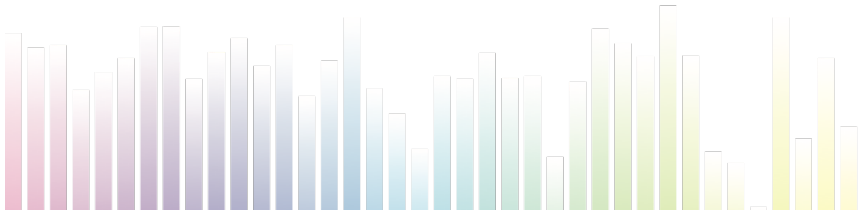

allinea

Solving Software Defects

- Who had a rogue behavior ?
  - Merges stacks from processes and threads
- Where did it happen?
  - leaps to source
- How did it happen?
  - Diagnostic messages
  - Some faults evident instantly from source
- Why did it happen?
  - Unique “Smart Highlighting”
  - Sparklines comparing data across processes

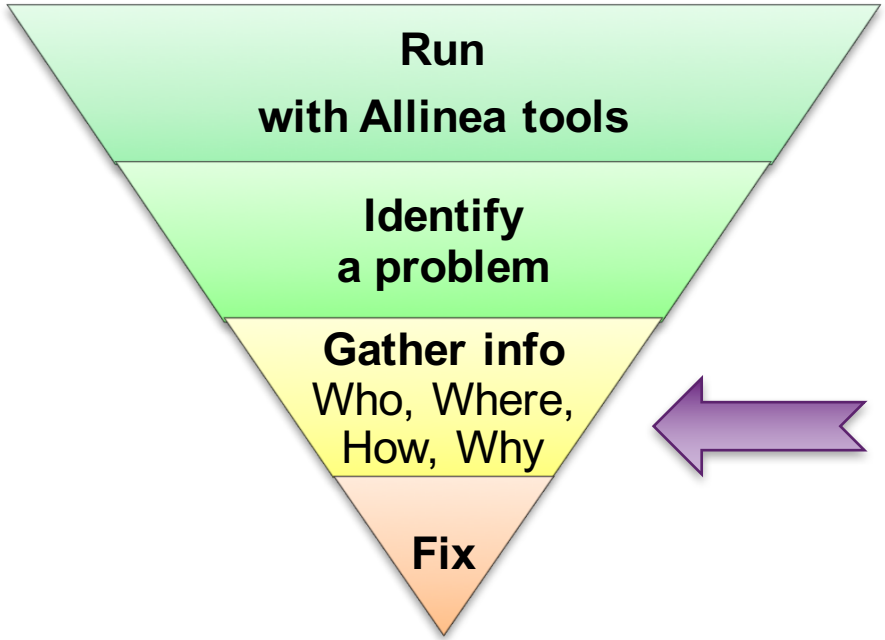

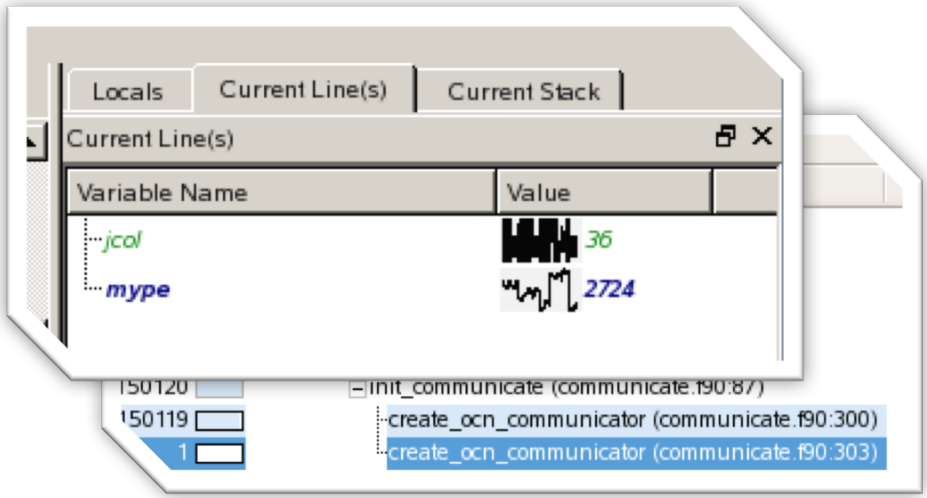

HPC could be brain surgery

- Brain aneurysms
  - 2-5% of population – most are undiagnosed
  - 30,000 rupture in US each year – 40% fatal
  - Early discovery and treatment increases survival rates
- Neurosurgery as HPC
  - MRI provides the blood vessel structure
  - Intra-cranial blood flow and pressures is just complex CFD
  - Full brain 3D model is 2-10GB geometry
- Individualized HPC
  - Patient's MRI scan enables surgical decision: whether to operate, how to operate, …
  - Circle of Willis requires super-Petascale ~~machine~~ software
  - Need answer in minutes or hours
- … but it crashes at 49152 cores

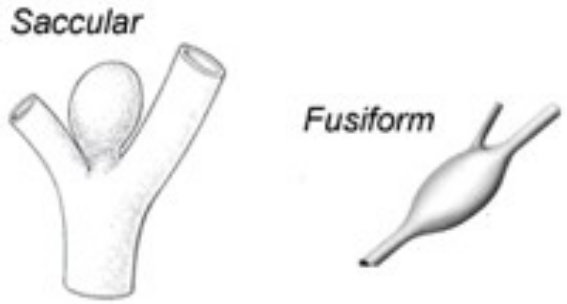

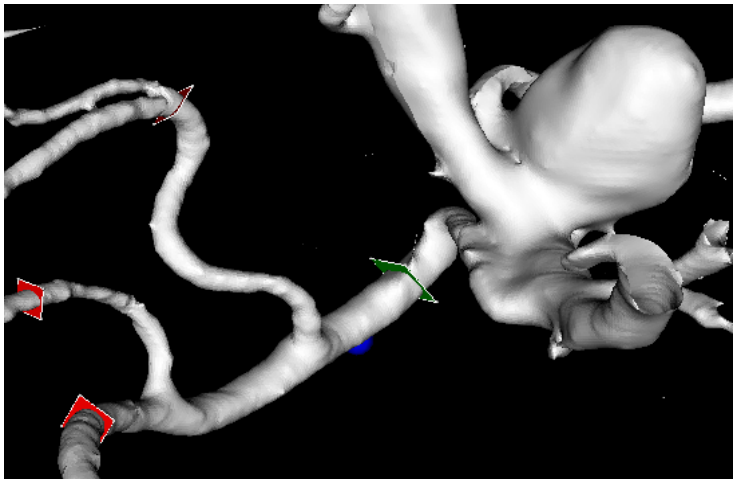

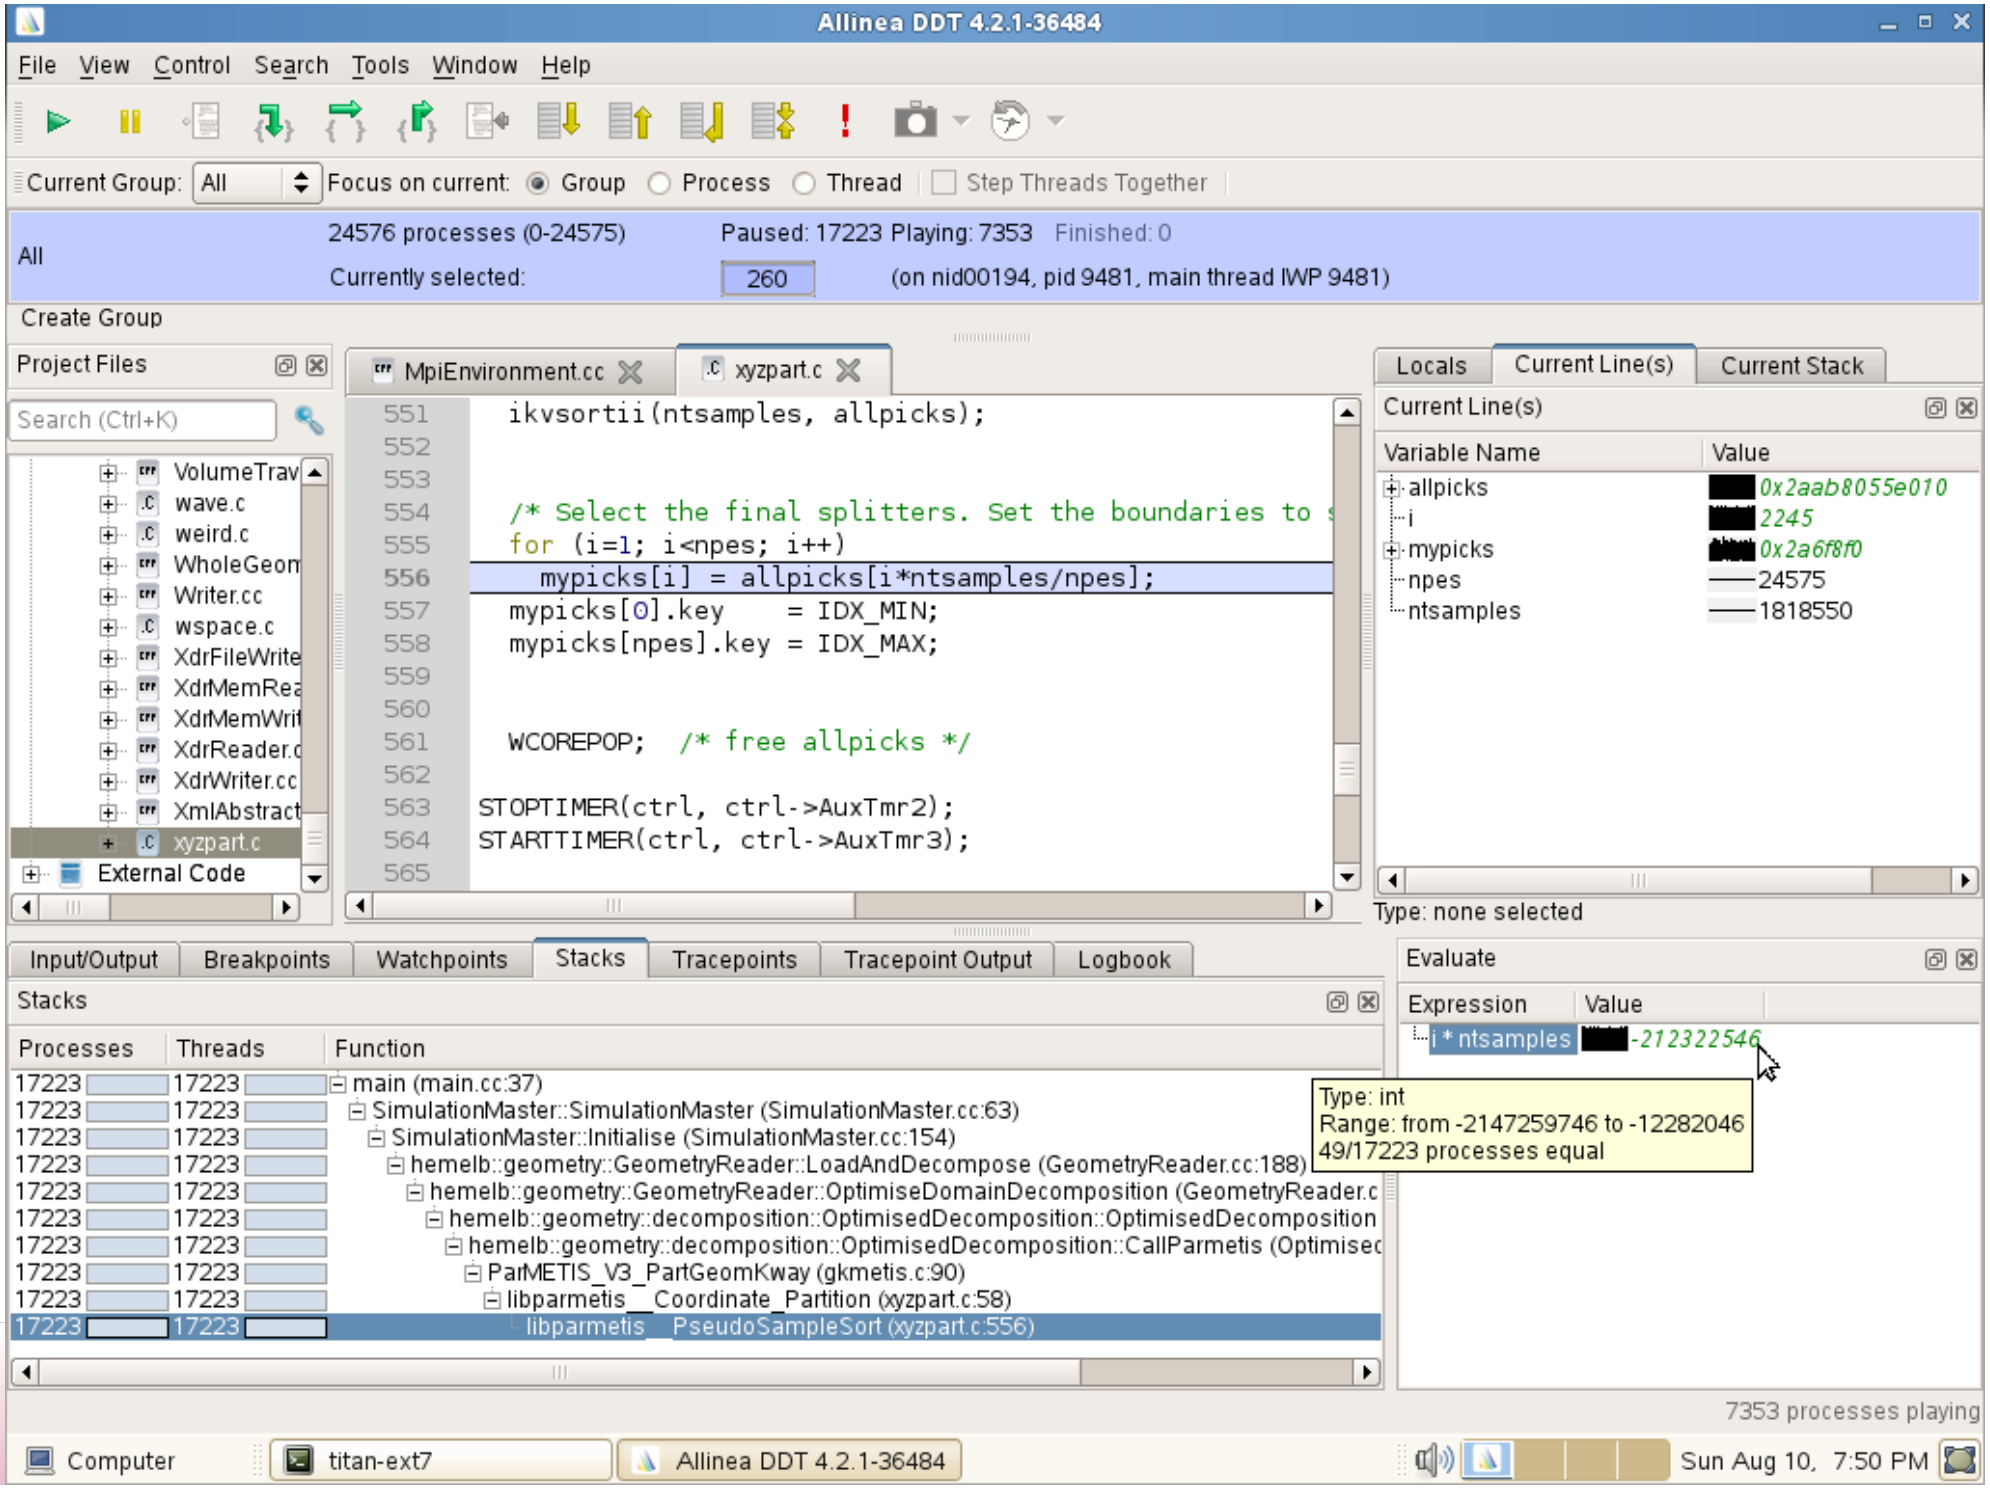
Allinea DDT 4.2.1-36484
File View Control Search Tools Window Help
Current Group: All
Focus on current: Group Process Thread
Step Threads Together
All
24576 processes (0-24575)
Paused: 17223 Playing: 7353 Finished: 0
Currently selected: 260
(on nid00194, pid 9481, main thread IWP 9481)
Create Group
Project Files
Search (Ctrl+K)
VolumeTrav
wave.c
weird.c
WholeGeom
Writer.cc
wspace.c
XdrFileWrite
XdrMemRea
XdrMemWrit
XdrReader.c
XdrWriter.cc
XmlAbstract
xyzpart.c
External Code
MpiEnvironment.cc
xyzpart.c
551 ikvsortii(ntsamples, allpicks);
552
553
554 /* Select the final splitters. Set the boundaries to
555 for (i=1; i<npes; i++)
556 mypicks[i] = allpicks[i*ntsamples/npes];
557 mypicks[0].key = IDX_MIN;
558 mypicks[npes].key = IDX_MAX;
559
560
561 WCOREPOP; /* free allpicks */
562
563 STOPTIMER(ctrl, ctrl->AuxTmr2);
564 STARTTIMER(ctrl, ctrl->AuxTmr3);
565
Locals
Current Line(s)
Current Stack
Current Line(s)
Variable Name
Value
allpicks 0x2aab8055e010
i 2245
mypicks 0x2a6f8f0
npes 24575
ntsamples 1818550
Type: none selected
Input/Output
Breakpoints
Watchpoints
Stacks
Tracepoints
Tracepoint Output
Logbook
Stacks
Processes
Threads
Function
17223 17223 main (main.cc:37)
17223 17223 SimulationMaster::SimulationMaster (SimulationMaster.cc:63)
17223 17223 SimulationMaster::Initialise (SimulationMaster.cc:154)
17223 17223 hemelb::geometry::GeometryReader::LoadAndDecompose (GeometryReader.cc:188)
17223 17223 hemelb::geometry::GeometryReader::OptimiseDomainDecomposition (GeometryReader.c
17223 17223 hemelb::geometry::decomposition::OptimisedDecomposition::OptimisedDecomposition
17223 17223 hemelb::geometry::decomposition::OptimisedDecomposition::CallParmetis (Optimised
17223 17223 ParMETIS_V3_PartGeomKway (gkmetis.c:90)
17223 17223 libparmetis__Coordinate_Partition (xyzpart.c:58)
17223 17223 libparmetis__PseudoSampleSort (xyzpart.c:556)
Evaluate
Expression
Value
i * ntsamples -212322546
Type: int
Range: from -2147259746 to -12282046
49/17223 processes equal
7353 processes playing
Computer
titan-ext7
Allinea DDT 4.2.1-36484
Sun Aug 10, 7:50 PM

Favorite Allinea DDT Features for Scale

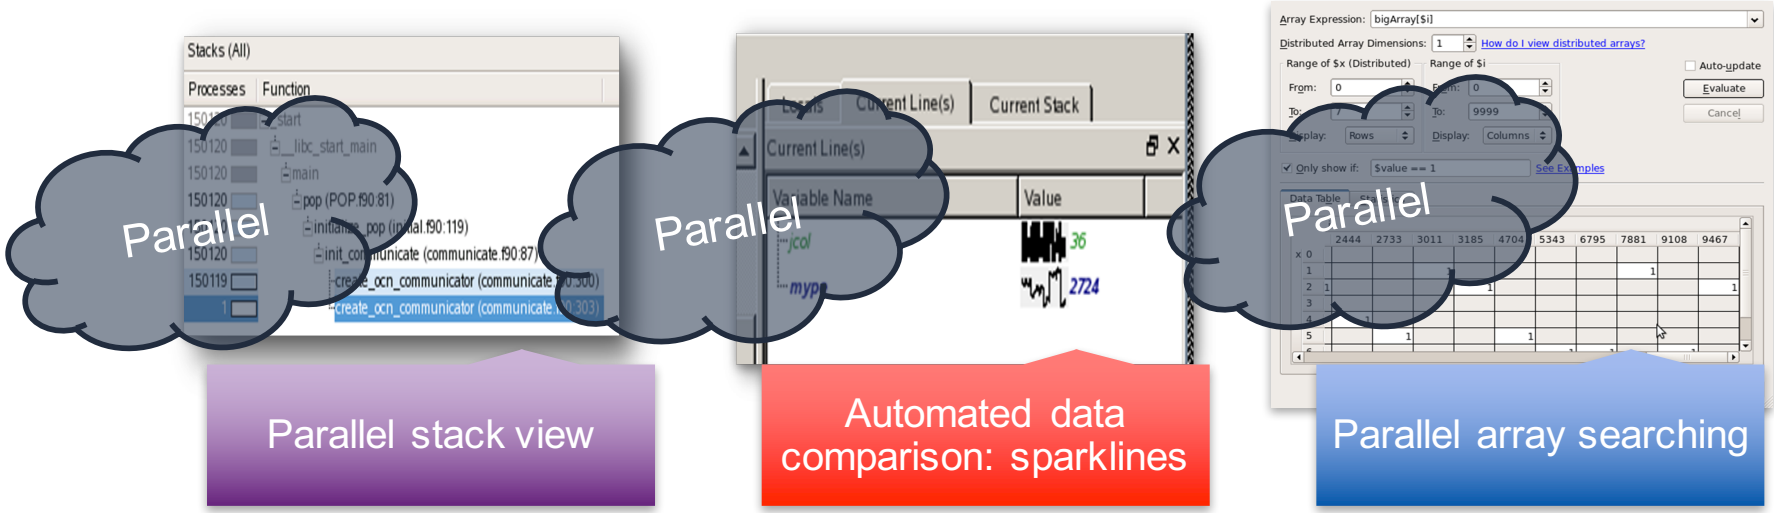

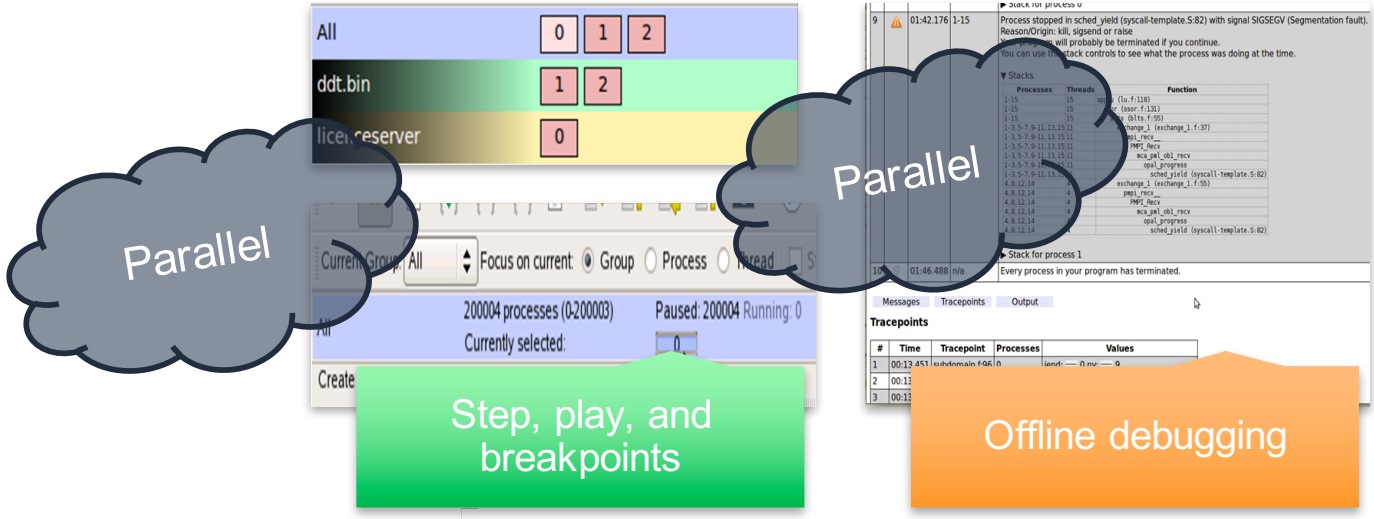

Today's Status on Scalability

- Debugging and profiling
  - Active users at 100,000+ cores debugging
  - 50,000 cores is largest profiling tried to date (and was Very successful)
  - … and active users with just 1 process too
- Deployed on
  - ORNL's Titan, NCSA Blue Waters, ANL Mira etc.
  - Hundreds of much smaller systems – academic, research, oil and gas, genomics, etc.
- Tools help the full range of programmer ambition
  - Very small slow down with either tool (< 5%)

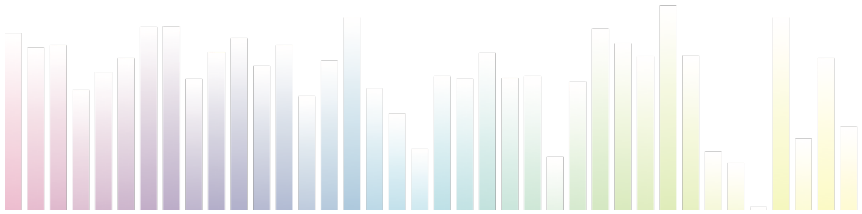

allinea

Getting started on Mira/Cooley

- Install local client on your laptop
  - www.allinea.com/products/forge/downloads
    - Linux – installs full set of tools
    - Windows, Mac – just a remote client to the remote system
  - Run the installation and software
  - “Connect to remote host”
  - Hostname:
    - username@cetus.alcf.anl.gov
    - username@cooley.alcf.anl.gov
  - Remote installation directory: /soft/debuggers/ddt
  - Click Test
- Congratulations you are now ready to debug on Mira/Vesta/Cetus – or debug and profile on Cooley.

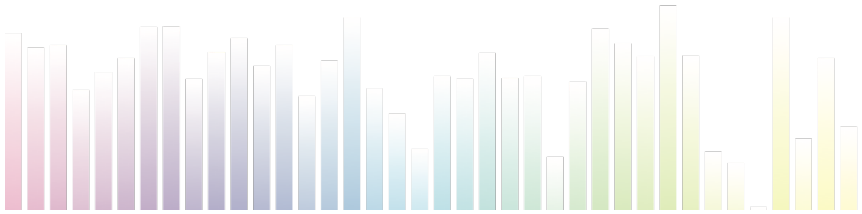

allinea

Five great things to try with Allinea DDT

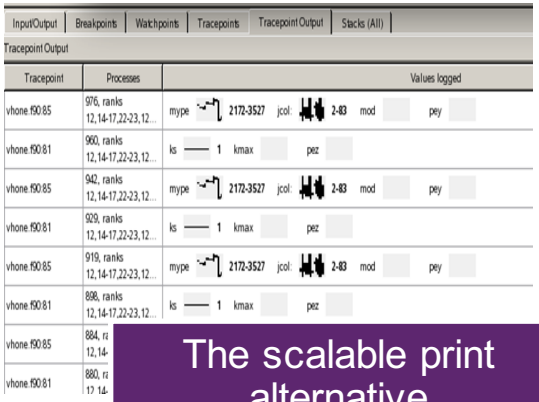

The scalable print alternative

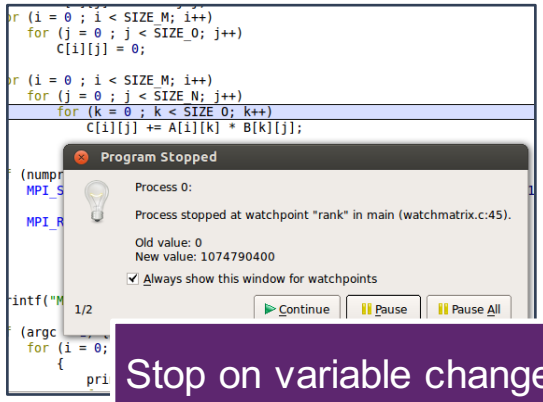

Stop on variable change

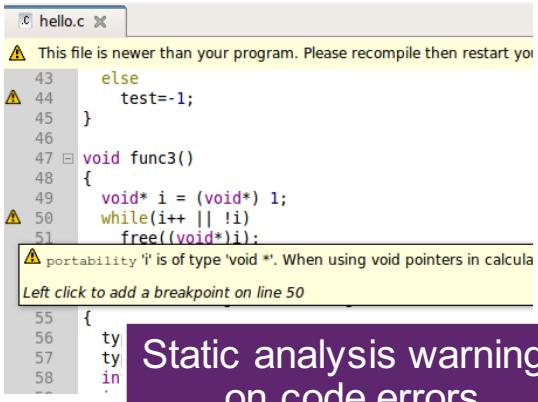

Static analysis warnings on code errors

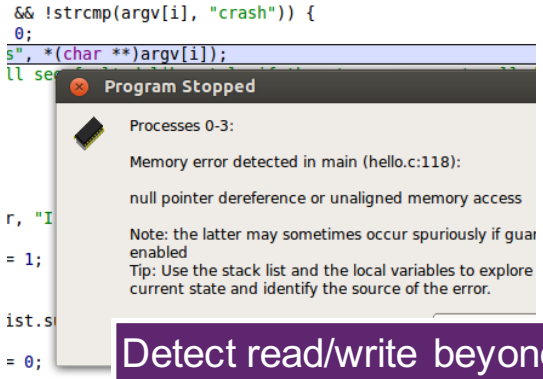

Detect read/write beyond array bounds

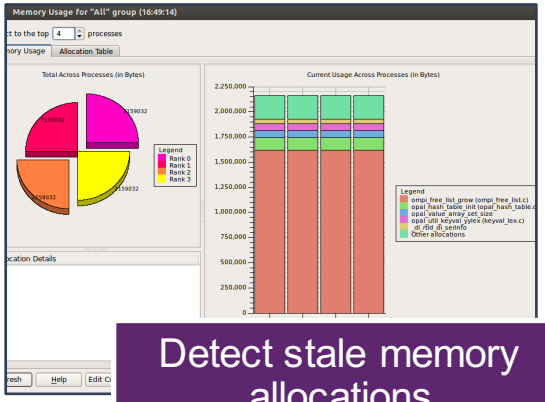

Detect stale memory allocations

allinea MAP

Six Great Things to Try with Allinea MAP

allinea FORGE

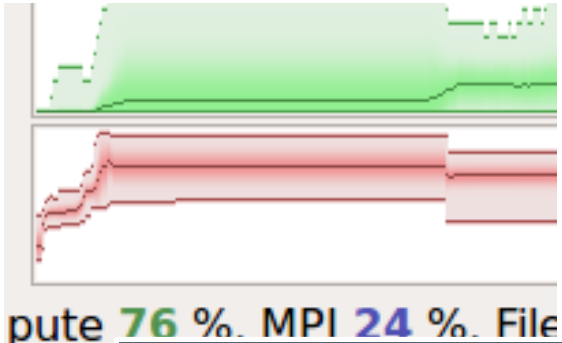

Find the peak memory use

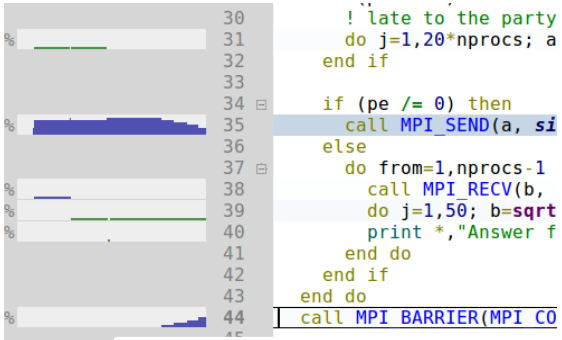

Fix an MPI imbalance

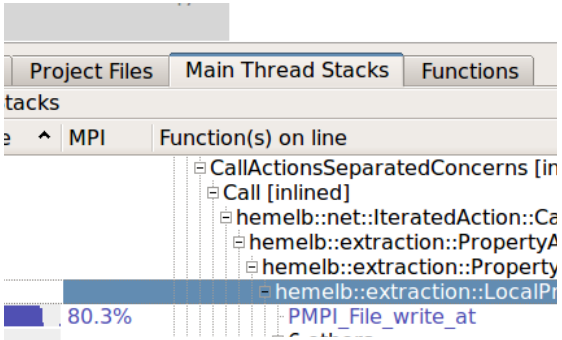

Remove I/O bottleneck

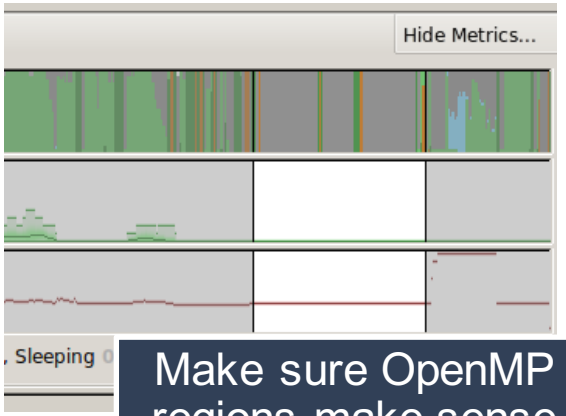

Make sure OpenMP regions make sense

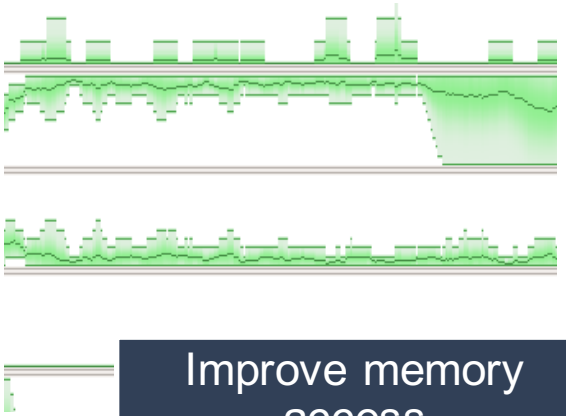

Improve memory access

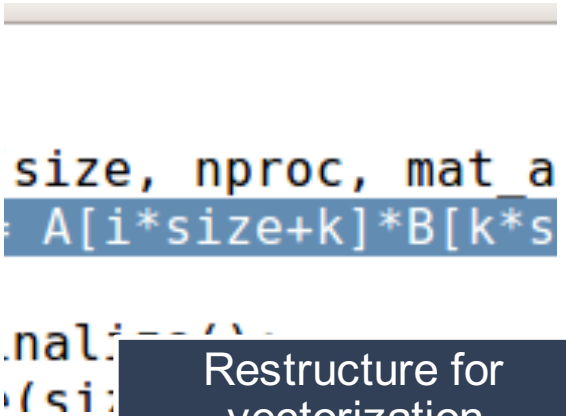

Restructure for vectorization

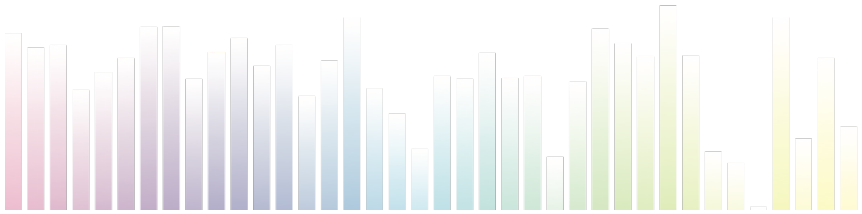

allinea

allinea

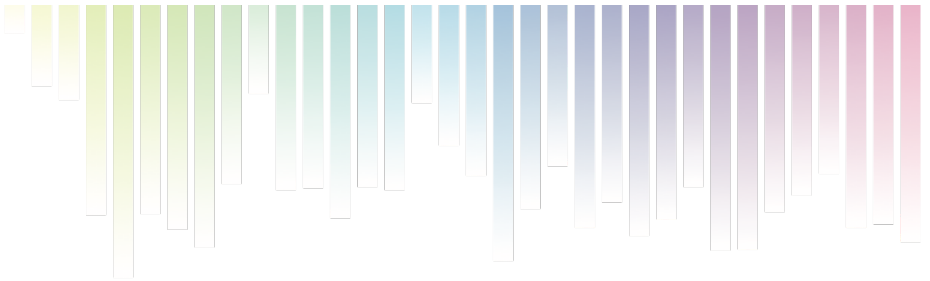

High performance tools to debug, profile, and analyze your applications

Thank you!

Ryan Hulguin, Application and Support Analyst
rhulguin@allinea.com