

# allinea



Leaders in parallel software development tools

## Allinea Unified Environment

Allinea's unified tools for debugging and profiling HPC  
Codes

Beau Paisley  
Allinea Software  
[bpaisley@allinea.com](mailto:bpaisley@allinea.com)  
720.583.0380

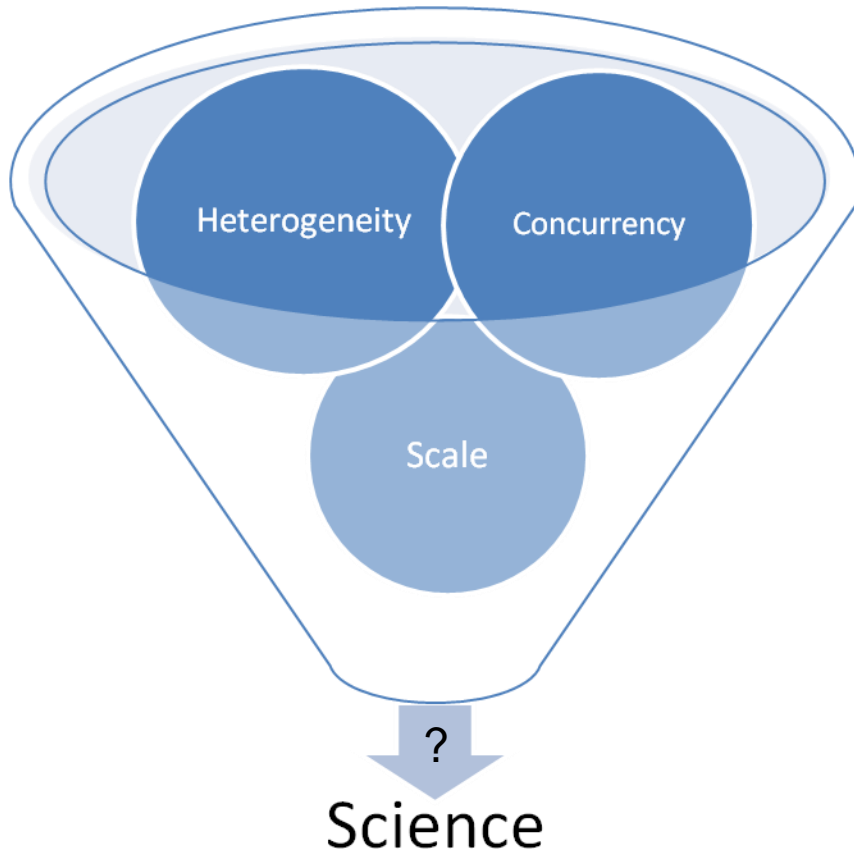
[www.allinea.com](http://www.allinea.com)

# allinea

Leaders in parallel software development tools



# Today's Challenge




Q: What is the impact of current trends in HPC on your application?

Q: How can you make your science run well on the available system?

A: Development.

**Development implies both fixing problems and optimizing the computation.**

# Machine Size Growth



Machine sizes are exploding

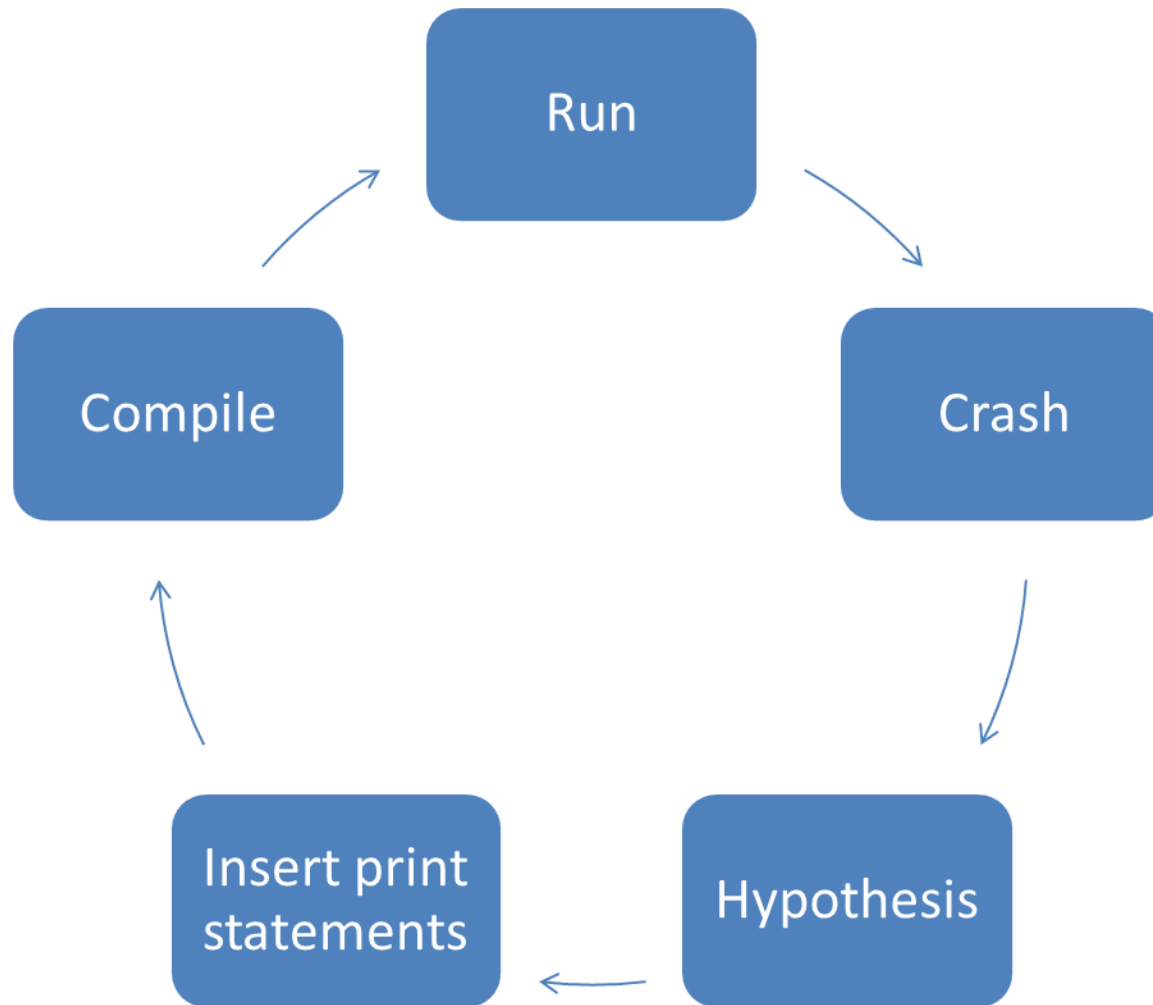
Software scale grows as machines grow

# Compilers Can't do it All, ...

```
mg.f(2432): (col. 10) remark: loop was not vectorized: not inner loop.
mg.f(2431): (col. 7) remark: loop was not vectorized: not inner loop.
mg.f(993): (col. 13) remark: LOOP WAS VECTORIZED.
mg.f(992): (col. 10) remark: loop was not vectorized: not inner loop.
mg.f(991): (col. 7) remark: loop was not vectorized: not inner loop.
mg.f(243): (col. 7) remark: loop was not vectorized: existence of vector dependence.
mg.f(993): (col. 13) remark: LOOP WAS VECTORIZED.
mg.f(992): (col. 10) remark: loop was not vectorized: not inner loop.
mg.f(991): (col. 7) remark: loop was not vectorized: not inner loop.
mg.f(753): (col. 13) remark: loop was not vectorized: vectorization possible but seems inefficient.
mg.f(762): (col. 13) remark: loop was not vectorized: vectorization possible but seems inefficient.
mg.f(749): (col. 10) remark: loop was not vectorized: not inner loop.
mg.f(746): (col. 7) remark: loop was not vectorized: not inner loop.
mg.f(993): (col. 13) remark: LOOP WAS VECTORIZED.
mg.f(992): (col. 10) remark: loop was not vectorized: not inner loop.
mg.f(991): (col. 7) remark: loop was not vectorized: not inner loop.
mg.f(2255): (col. 16) remark: loop was not vectorized: existence of vector dependence.
mg.f(2254): (col. 13) remark: loop was not vectorized: not inner loop.
mg.f(2251): (col. 7) remark: loop was not vectorized: not inner loop.
mg.f(2433): (col. 13) remark: LOOP WAS VECTORIZED.
mg.f(2433): (col. 13) remark: loop was not vectorized: not inner loop.
mg.f(2432): (col. 10) remark: loop was not vectorized: not inner loop.
mg.f(2431): (col. 7) remark: loop was not vectorized: not inner loop.
mg.f(2433): (col. 13) remark: LOOP WAS VECTORIZED.
mg.f(2433): (col. 13) remark: loop was not vectorized: not inner loop.
mg.f(2432): (col. 10) remark: loop was not vectorized: not inner loop.
mg.f(2431): (col. 7) remark: loop was not vectorized: not inner loop.
mg.f(527): (col. 7) remark: loop was not vectorized: nonstandard loop is not a vectorization candidate.
mg.f(552): (col. 7) remark: loop was not vectorized: nonstandard loop is not a vectorization candidate.
mg.f(1150): (col. 7) remark: loop was not vectorized: loop was transformed to memset or memcpy.
mg.f(1150): (col. 7) remark: loop was not vectorized: loop was transformed to memset or memcpy.
mg.f(1645): (col. 7) remark: loop was not vectorized: loop was transformed to memset or memcpy.
mg.f(1655): (col. 7) remark: loop was not vectorized: loop was transformed to me
```

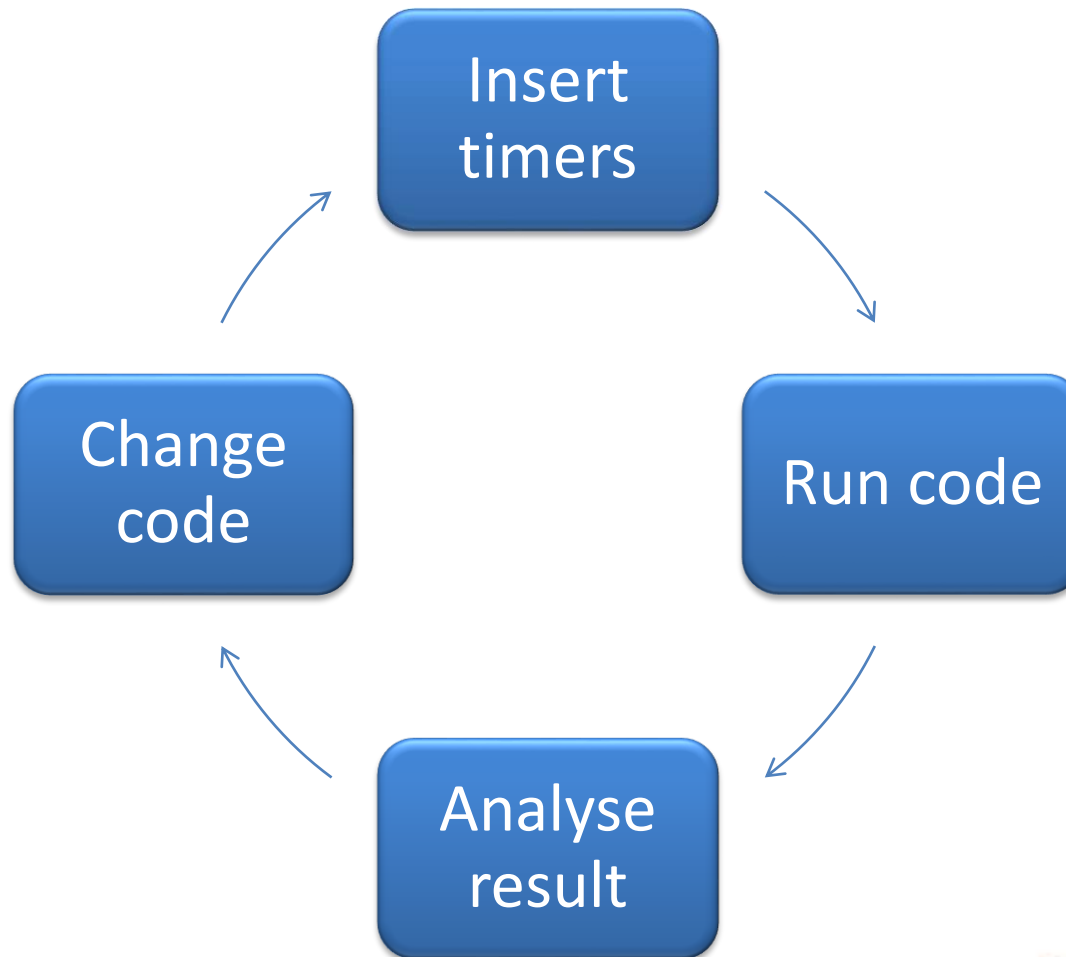
# Debugging in practice...

---



# Optimization in practice...

---



# Some Bug Types

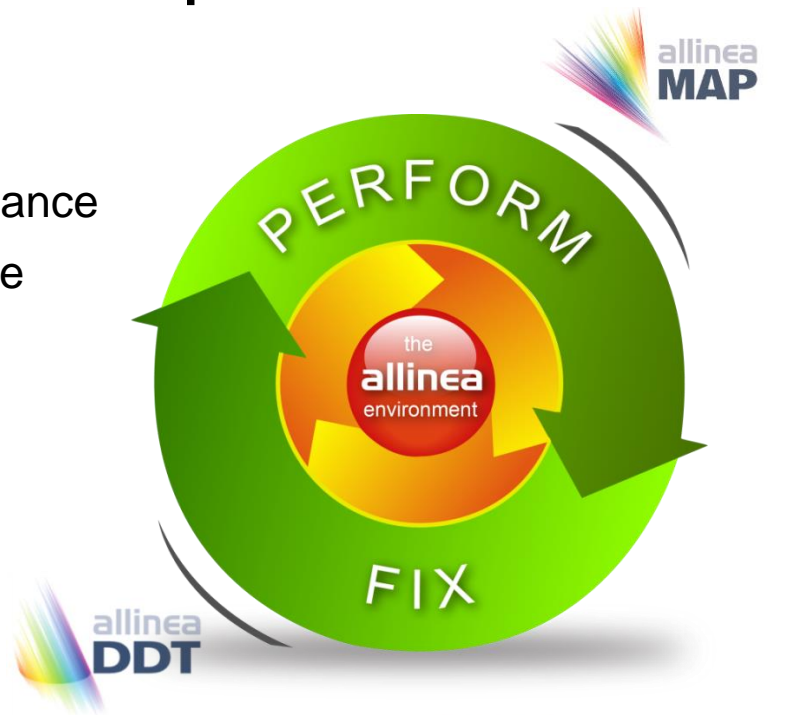
---

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

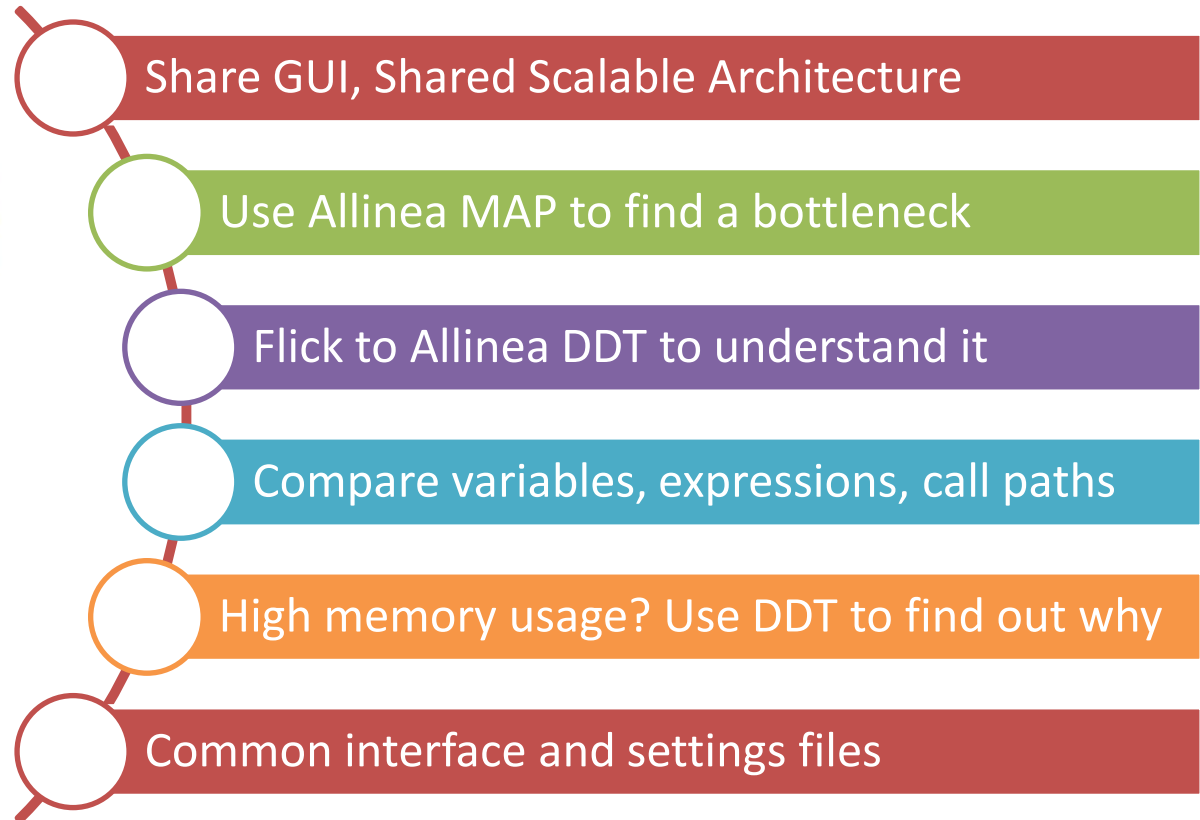


# Allinea Unified Environment

- A modern integrated environment for HPC developers
- Supporting the lifecycle of application development and improvement
  - Allinea DDT : Productively debug code
  - Allinea MAP : Enhance application performance
  - Allinea Performance Reports: Characterize Application performance
- Designed for productivity
  - Consistent easy to use tools
  - Enables effective HPC development
- Improve system usage
  - Fewer failed jobs
  - Higher application performance



# An Integrated Environment




# Alinea DDT

## Fix software problems - fast

- **Graphical debugger designed for:**
  - C/C++, Fortran, UPC, CUDA, CUDA Fortran, OpenACC
  - Multi-threaded code
  - Multi-process code
  - Accelerated codes
    - GPUs, Intel Xeon Phi
  - Debugging 1 to 700k processes
- **Slash your time to debug :**
  - Reproduces and triggers your bugs instantly
  - Helps you to fix them as swiftly as possible
  - Helps you easily understand where issues come from quickly



# Allinea DDT: Debugging that Scales




## Where?

- Leaps to source automatically
- Powerful instantaneous memory debugging



## How?

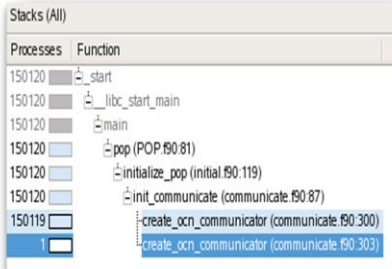
- Real-time data comparison and consolidation
- Identify outliers and unusual threads



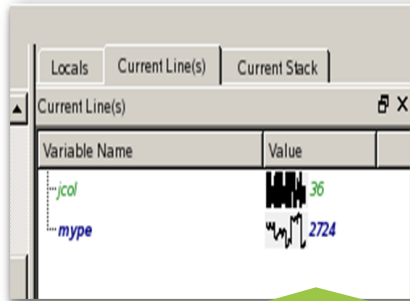
## Why?

- “Smart Highlighting” of differences and changes
- Sparklines comparing data across processes

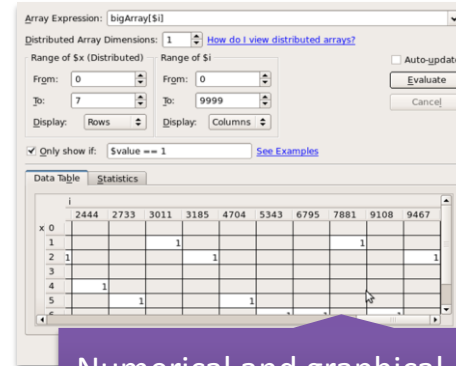
# Top Features for HPC Debugging




Parallel stack view



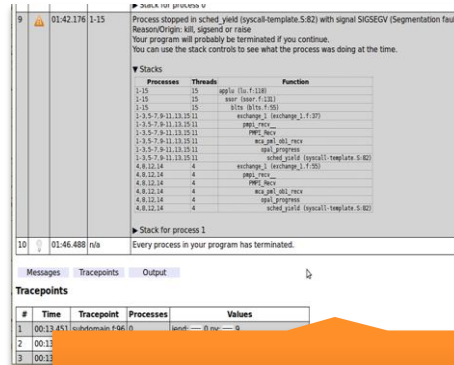
Automated data comparison: sparklines



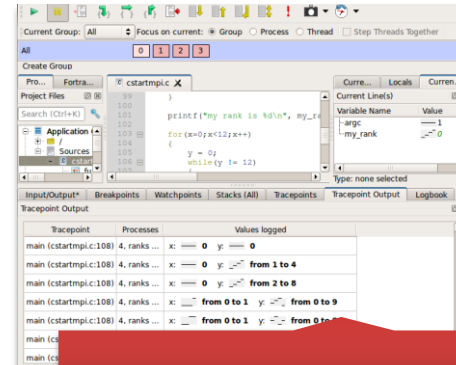
Numerical and graphical data visualization



Step, play, and breakpoints



Offline debugging



Tracepoints

# Alinea MAP

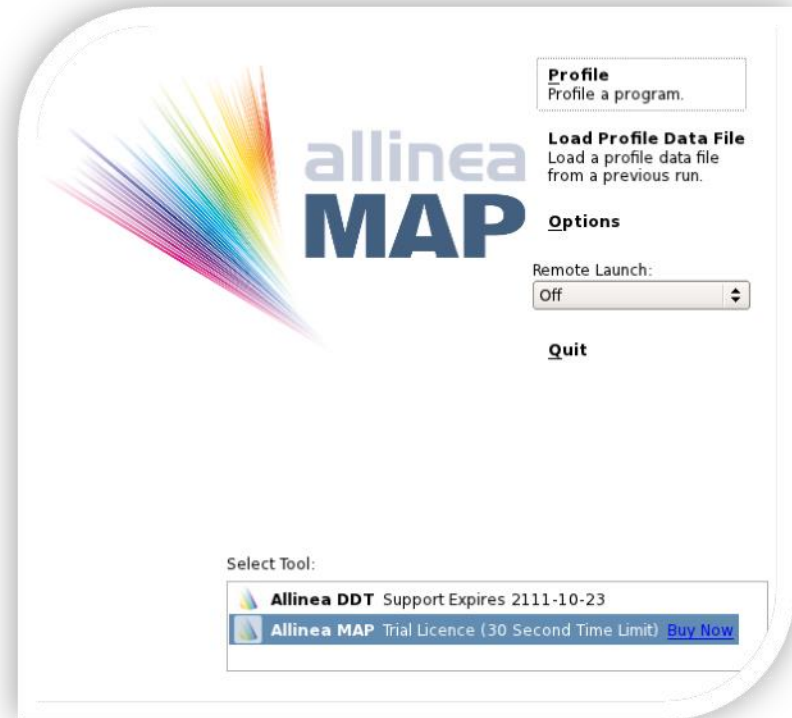
## Increase Application Performance

- **Parallel profiler designed for:**

- C/C++, Fortran
- Multi-process code
  - Interdependent or independent processes
- Multi-threaded code
  - Monitor the main threads for each process
- Accelerated codes
  - GPUs, Intel Xeon Phi

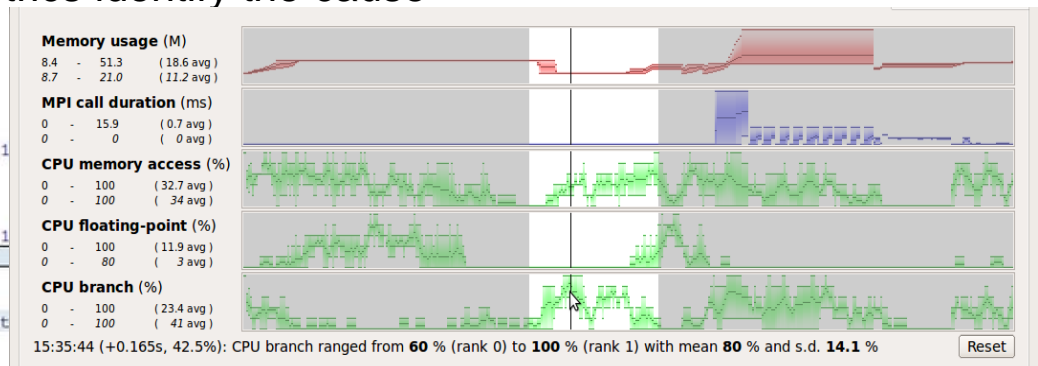
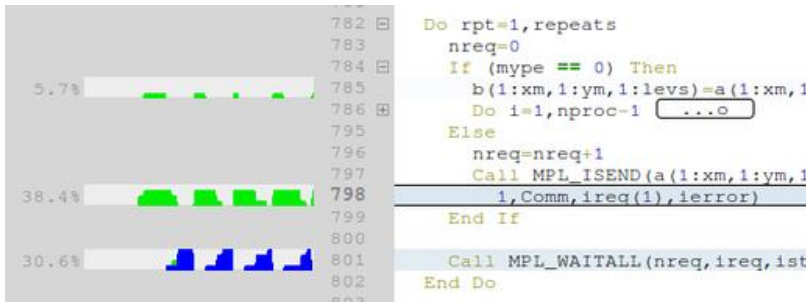
- **Improve productivity :**

- Helps you detect performance issues quickly and easily
- Tells you immediately where your time is spent in your source code
- Helps you to optimize your application efficiently

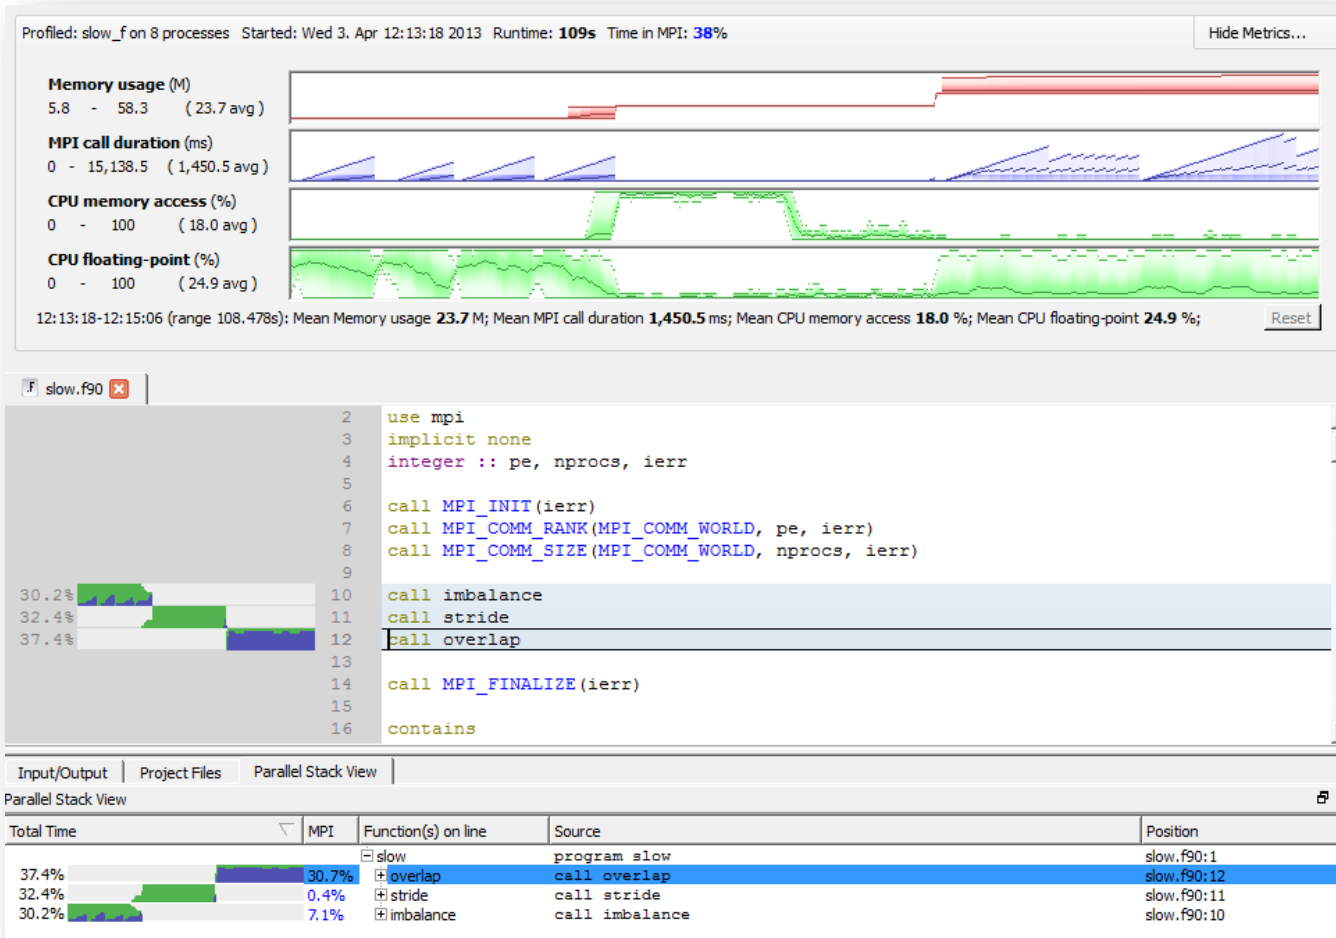


# Find Performance Issues Quickly

- **Look at the entire application on real data sets**
  - Visualize the entire run at full scale, not just reduced sets
  - Zoom in to explore iterations, functions and loops
- **Non-Destructive Profiling**
  - Less than 5% overhead
  - No need to instrument your code
  - Small output files (10-20Mb is typical)
- **Understand the nature of bottlenecks**
  - Source code viewer pinpoints bottleneck locations
  - CPU, MPI, I/Os and memory metrics identify the cause



# Providing Visual Scalability



Common horizontal axis



Aggregate across all processes



Highlight imbalance visually




Always refer to source code



# Allinea Performance Reports

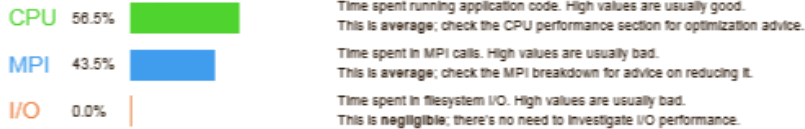


Executable: cp2k.popt  
Resources: 256 processes, 16 nodes  
Machine: cray-one  
Start time: Tue Oct 27 16:02:12 2013  
Total time: 951 seconds (16 minutes)  
Full path: /users/allinea/cp2k/exe/CRAY-XE6-gfortran-hwtopo  
Notes: H20 benchmark



Summary: cp2k.popt is **CPU-bound** in this configuration


The total wallclock time was spent as follows:



This application run was **CPU-bound**. A breakdown of this time and advice for investigating further is in the **CPU** section below.

## CPU


A breakdown of how the **58.5%** total CPU time was spent:



The per-core performance is **memory-bound**. Use a profiler to identify time-consuming loops and check their cache performance. Little time is spent in **vectorized instructions**. Check the compiler's vectorization advice to see why key loops could not be vectorized.

## MPI


Of the **43.5%** total time spent in MPI calls:



The **point-to-point** transfer rate is low. This can be caused by inefficient message sizes, such as many small messages, or by imbalanced workloads causing processes to wait. Use an MPI profiler to identify the problematic calls and ranks.

## I/O

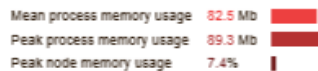
A breakdown of how the **0.0%** total I/O time was spent:



No time is spent in **I/O operations**. There's nothing to optimize here!

## Memory

Per-process memory usage may also affect scaling:





The **peak node memory usage** is low. You may be able to reduce the total number of CPU hours used by running with fewer MPI processes and more data on each process.

- Effortless one-touch reports
  - Add **one command** to your run script
  - A **one-page report** is generated **automatically**
- Characterize and understand application performance
  - With **< 5%** application slowdown

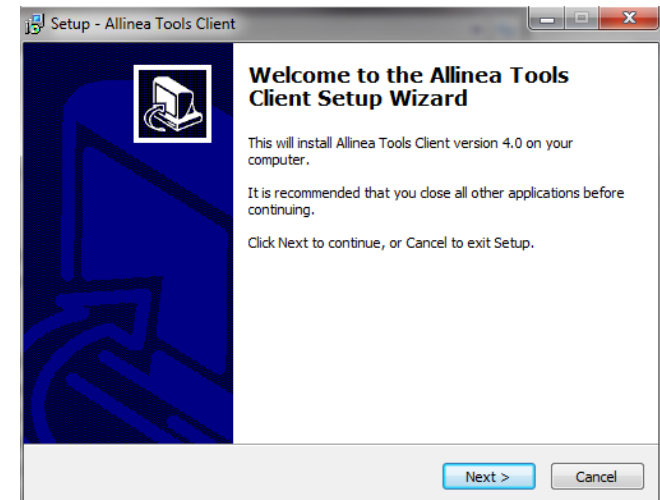
# Top Features for HPC Code Optimization

- Allinea's tools provide extensive performance metrics, with low overhead
- Allinea's tools provide a graphical, easy-to-use presentation that is easily understood by scientists, engineers, and software developers
- Allinea MAP shows exactly which lines of source code are slow and why without modifications or instrumentation
- Allinea Performance Reports offers application level performance characterization and advice



# Remote Access Clients for Mac, Windows and Linux

- **Easier access to distant clusters**
  - Scalable debugging tree already cuts down network traffic
  - Secure low-latency debugging and profiling clients
- **Extends existing remote cluster support to cover**
  - No shared filesystem
  - Remote/local source-viewing
  - Support for multi-hop SSH and OTP systems
- **Allinea DDT feature set available remotely**
  - Linux, Windows and OS/X clients
  - Real native GUI – no ‘VNC’ or ‘X-forwarding’ lag



# What Our Users are Saying



“My group routinely debugs code at over 100,000 processes using Allinea DDT. No other debugger comes close – obviously it’s a hit with users,” Oak Ridge National Laboratory



“Allinea’s experience and tools will make a big impact in the speed at which scientists can complete their research,” NCSA Blue Waters



“Previous experiences with other profilers had left us more confused than informed. Allinea MAP is the opposite.”

# allinea

Leaders in parallel software development tools

## Thank You

Try it out at:

<http://www.allinea.com/products/trials/>

