allinea

Leaders in parallel software development tools

From Thousands to Millions:

Visual and System Scalability for Debugging and Profiling

Mark O'Connor

VP Product Management

Introduction: Bandwidth and Complexity

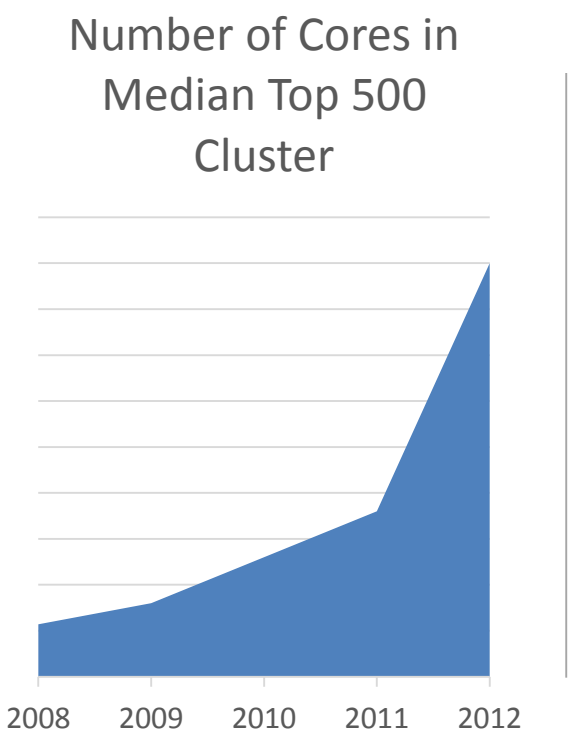

- Parallelism is increasing exponentially in HPC clusters ($R^2 \sim 0.95$)
- Performance data size and bandwidth requirements are increasing exponentially too
- More parallel execution contexts than lines of code
- Storage, networking and human visual acuity can no longer keep up

Introduction: Exploding Bandwidth

Trivial 16,000 process wave equation code

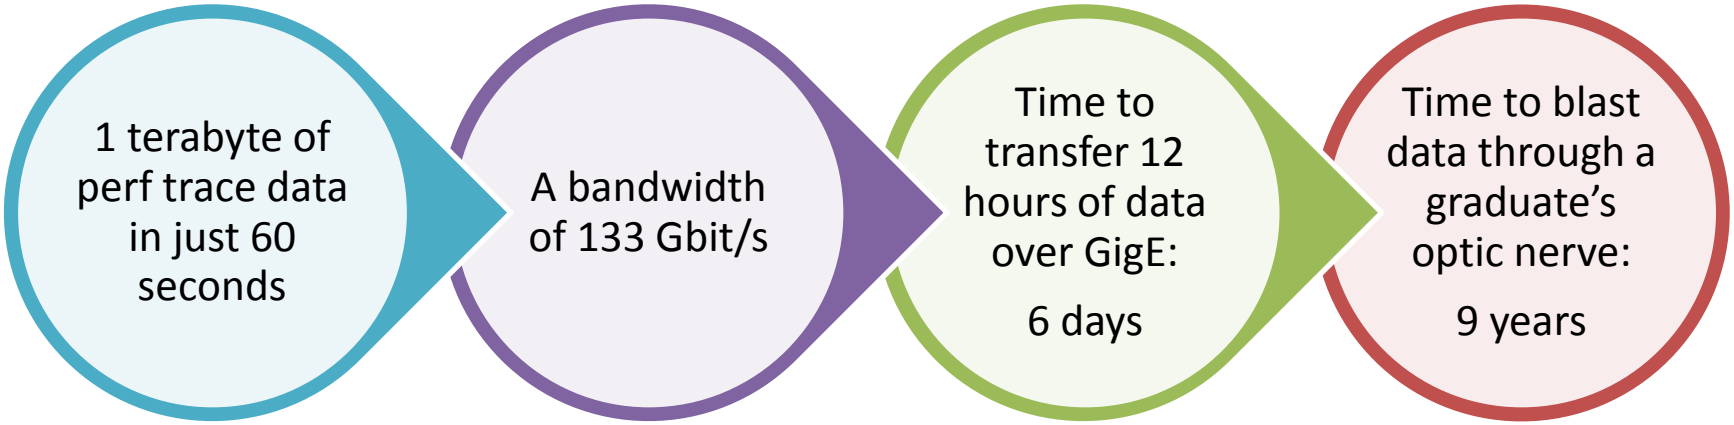

Performance Analysis Approaches

Record Everything

- Collect as much as possible and data mine it afterwards
- Use the cluster to analyse and mine large data files in parallel during analysis
- Implemented by trace-based tools such as Vampir

Statistical Analysis

- Only record data that provides:
  - Actionable information
  - Context for the above
- Example: duration of 16k MPI_Sends
  - Record the distribution shape
  - Record ranks of min / max
- Use the cluster to create small report files

Strengths and Weaknesses

Record Everything

- Can sift through and analyse in extreme detail after one recording
- Shows the inner workings of communication protocols
- Extremely large trace files
- Analysis may require cluster time
- Care must be taken not to accidentally add 1000x overhead

Statistical Analysis

- Reliable performance overview with low (< 5%) overhead
- Small trace files (~20Mb)
- Simple to configure, run and interpret
- Hides the inner workings of communication protocols
- May not contain enough data to explain *why* a line or loop is slow

Attacking Visual Scalability

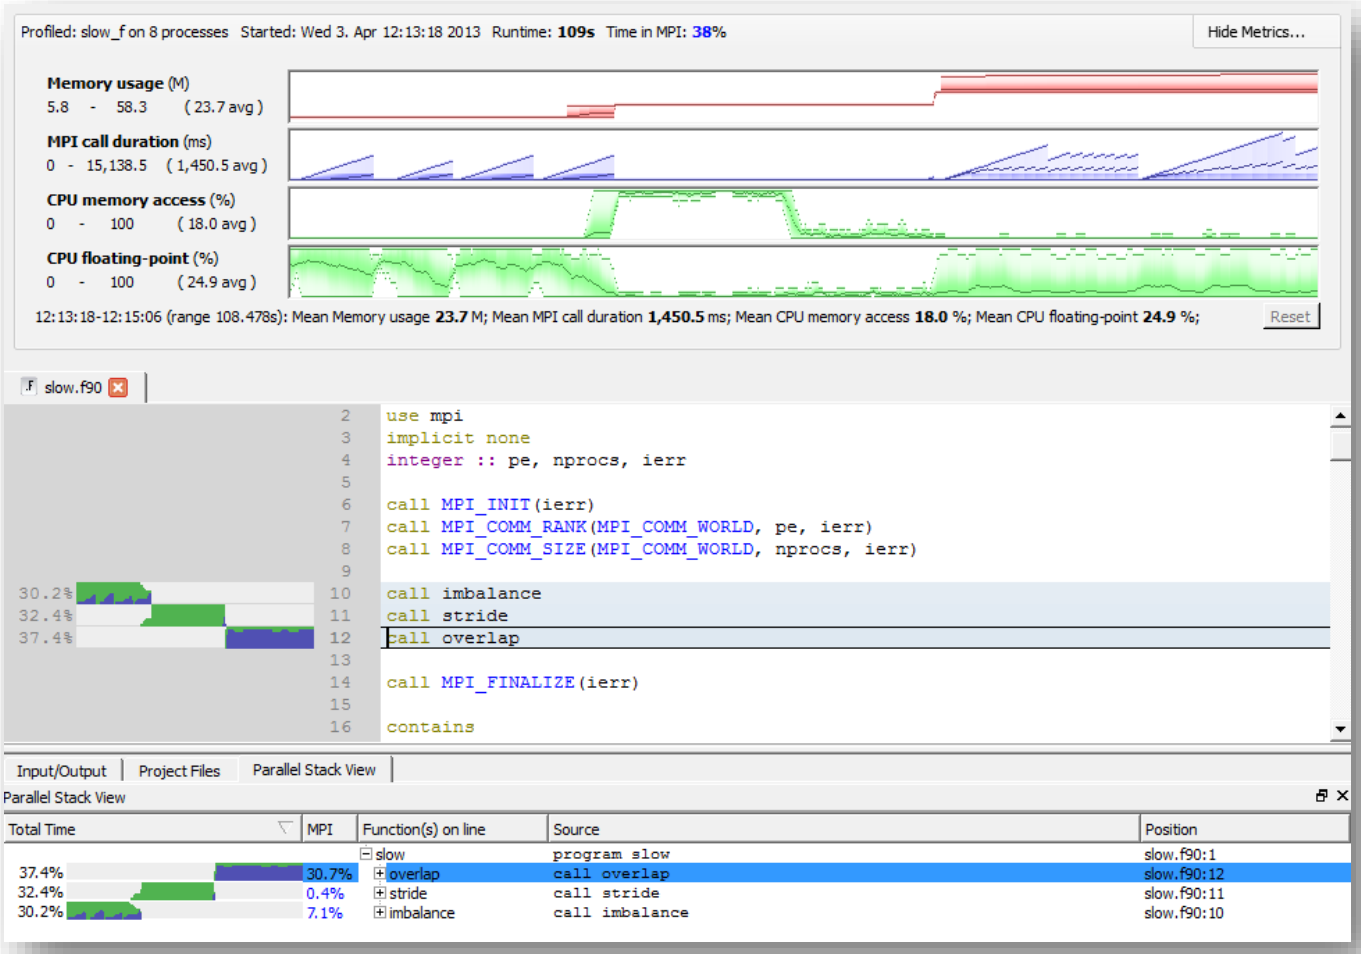

Common horizontal axis

Aggregate across all processes

Highlight imbalance visually

Always refer to source code

allinea

www.allinea.com

Record Everything Example

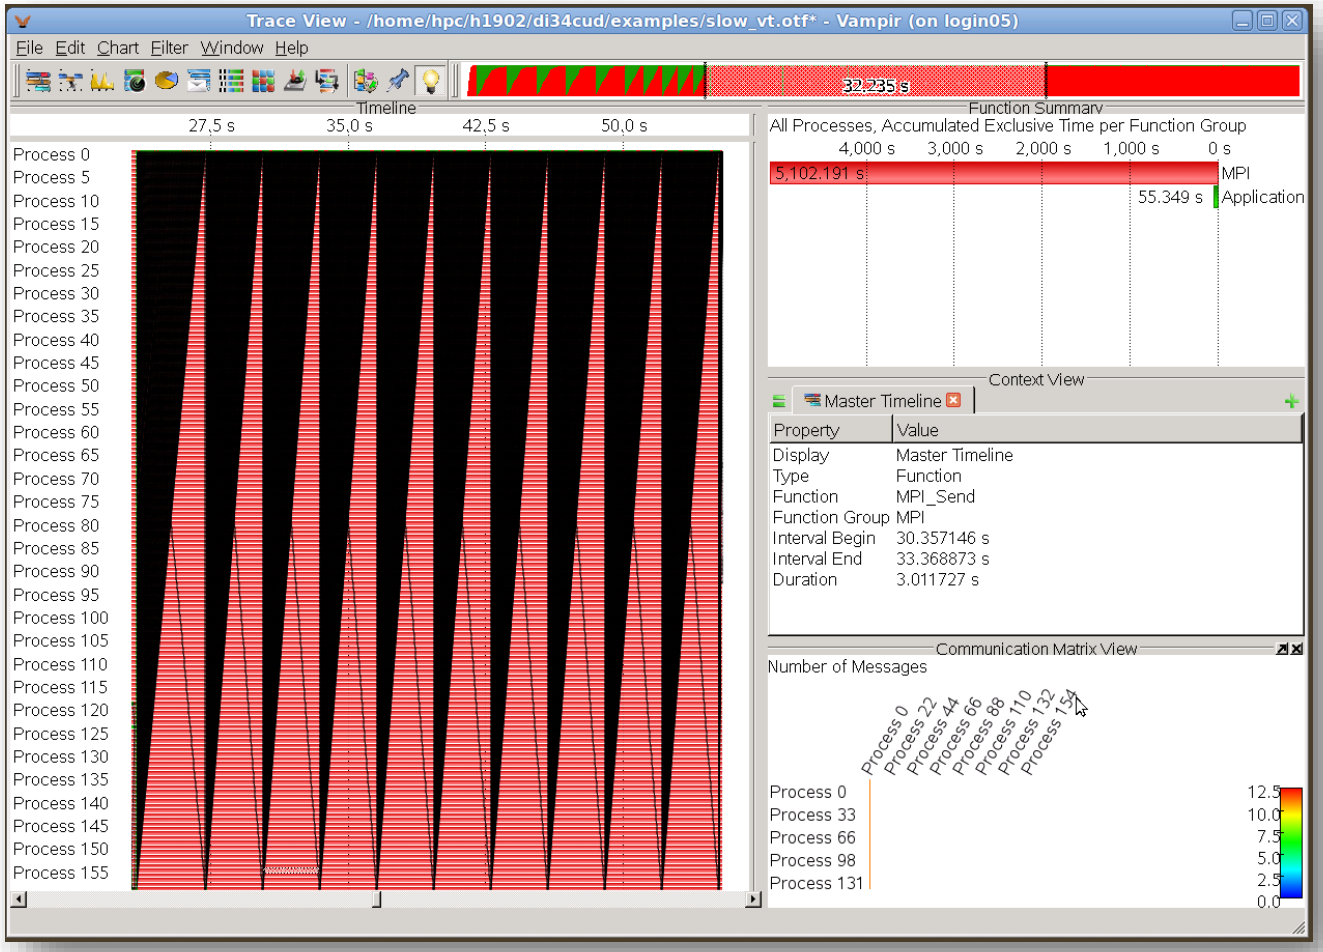

Pick one MPI call and view its specific data

Recorded data first, source code second

Can see underlying Send + Barrier traffic

All waiting for process 0?

allinea
www.allinea.com

Statistical Analysis Example

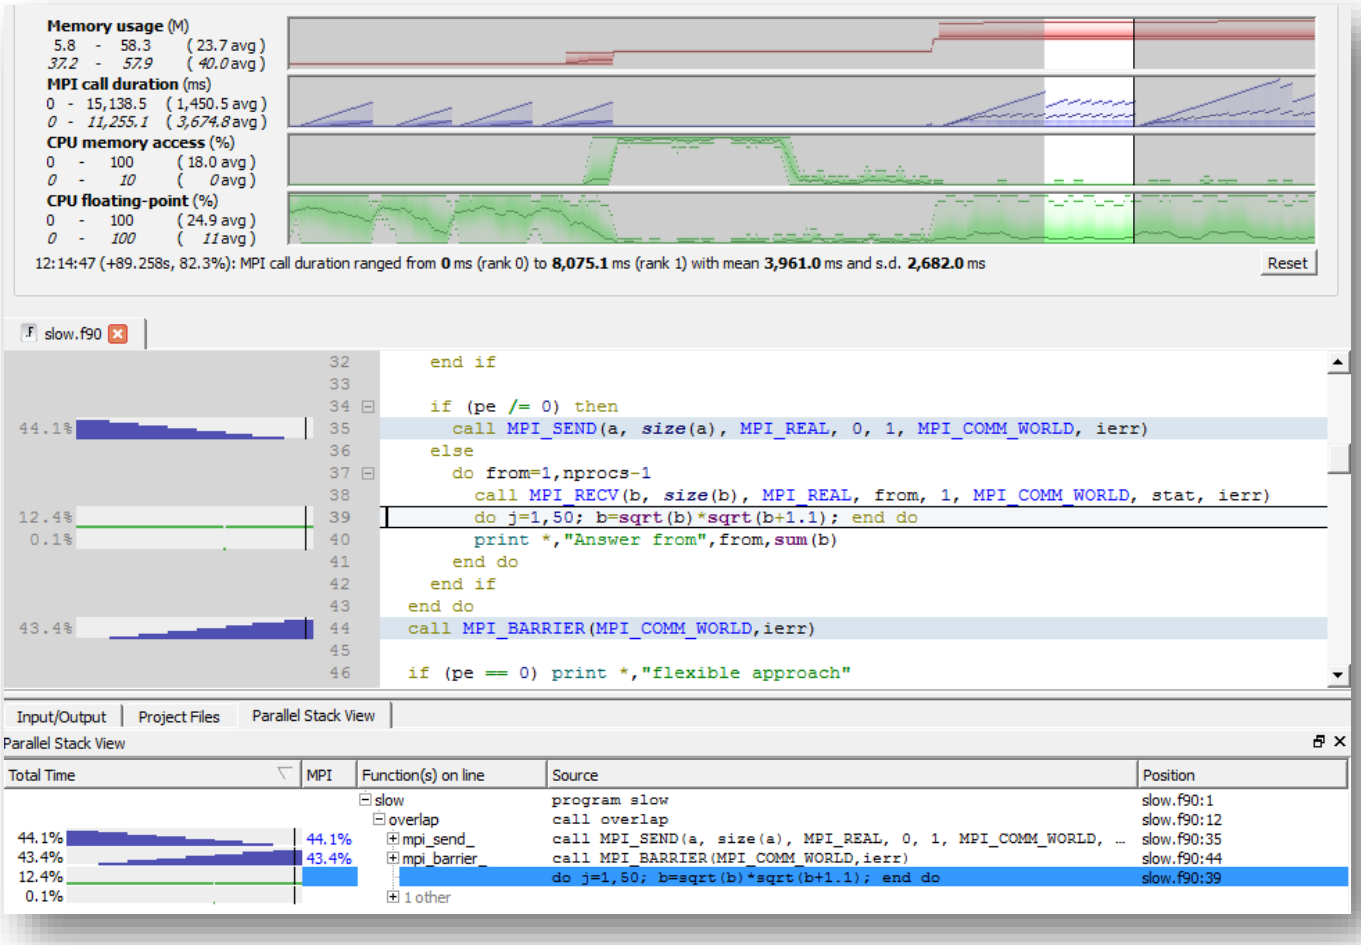

Show distributions and min/max ranks

Show per-line information

Focus on movement through code

Process 0 busy computing on line 39

allinea
www.allinea.com

Statistical CPU Analysis

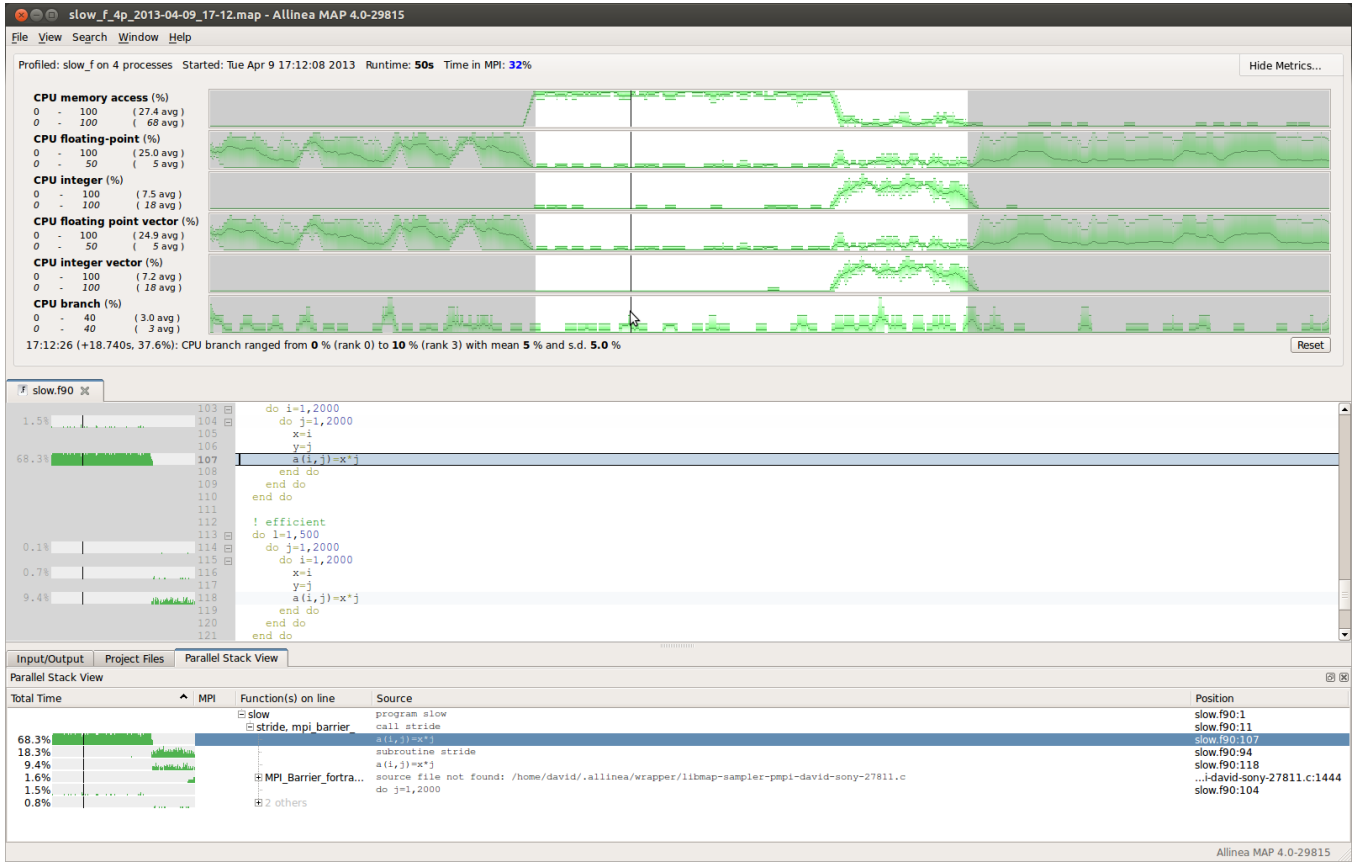

- Able to see cache performance, floating-point or integer operations
- … and other MPI key data

allinea
www.allinea.com

Complimentary Approaches

Allinea MAP

- Quick, low-overhead way to characterize performance
- See which lines of code are hotspots
- Identify common problems at once

Record one hotspot

- Pass more obscure problems to an expert
- Now know which loop to instrument and which performance counters should be recorded

Surprising uses for Statistical Analysis

- <5% runtime overhead
- 20Mb output files
- No instrumentation needed
- Run regularly – or in regression tests
- Keep XML output files in source control

Success with Allinea MAP

"Experience with other profilers had left us more confused than informed. MAP is the opposite."

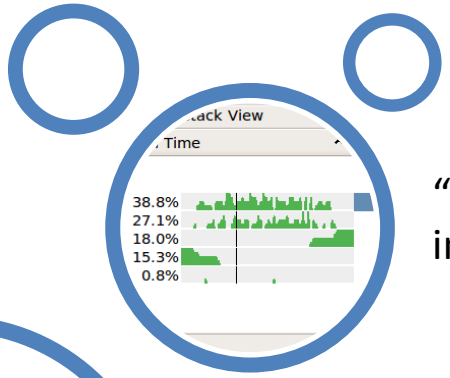

"We got a 20% speed improvement in just 3 days"

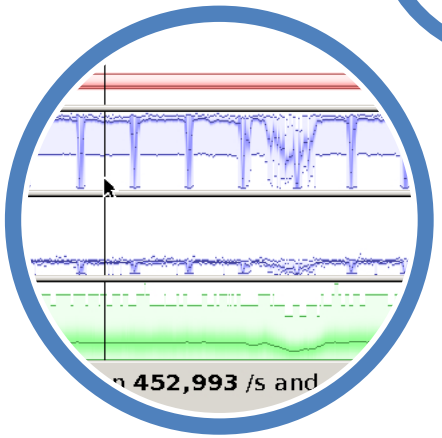

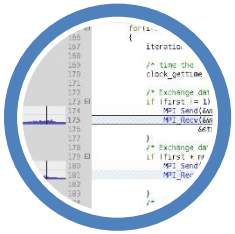

"I found a performance problem that I'd been chasing for 3 weeks on my very first run"

allinea
www.allinea.com

Integrated with Allinea DDT

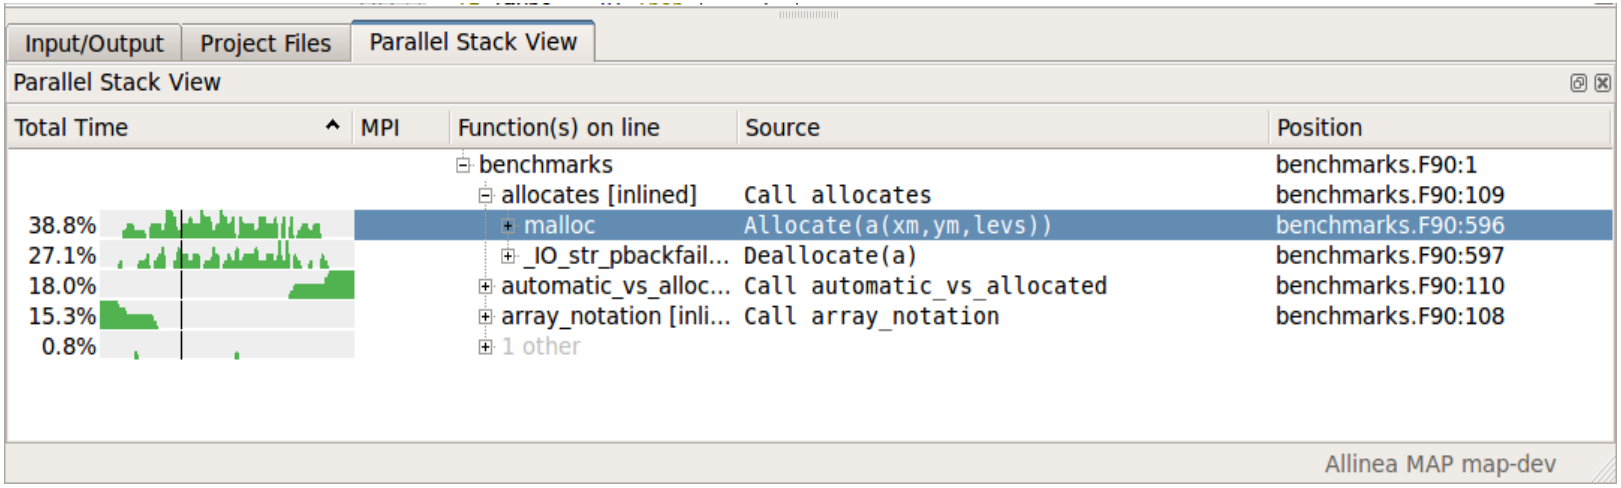

- World-class scalability
  - Shares Allinea DDT tree architecture – proven beyond Petascale
  - Data is merged on the cluster: no huge files.

allinea
www.allinea.com

Allinea DDT at Scale

Full scale on Blue Waters

- Full interactive GUI at 700,000+ processes, 30x faster than required
- "We can ramp up and down and not only pay for the largest possible case"

Full Cray Support

- "Allinea has proven a great partner on multiple installations"
- "Known for its scalable performance and interface"

Full scale on Titan

- "The transition has been smoother than previously thought possible"
- "DDT is tightly-integrated into the Cray programming environment"

allinea
www.allinea.com

Unified Products: DDT + MAP

- Use Allinea MAP to find a bottleneck
- Flick to Allinea DDT to understand it
- Compare variables, expressions, call paths
- High memory usage? Use DDT to find out why
- Common interface and settings files

allinea
www.allinea.com

Thank-you! Any Questions?

Try Allinea MAP and Allinea DDT for yourself:
www.allinea.com

allinea MAP

allinea DDT

allinea
www.allinea.com