

allinea

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

Enabling life-changing scientific discoveries
with Allinea's tools



Agenda for the day

- 08:30 – 09:00 : Arrival and Welcome
- 09:00 – 10:20 : Adding vectorization to linear algebra applications
- 10:20 – 10:40 : Coffee break
- 10:40 – 13:00 : Preparing applications for large scale
- 13:00 – 14:00 : Lunch Break
- 14:00 – 15:00 : Commercial presentation

HPC Ultimate target



Example: Weather and Forecasting models



Building blocks for Exascale



Scalability

- Enable multi-physics simulations
- Run larger, more accurate models
- Resolve ground-breaking scientific problems

Efficiency

- Reduce wasted resources (energy...)
- Maximize science output per \$
- Minimize time to result

Simplicity

- Pro-actively and automatically detect faults
- Provide applications on various hardware
- Facilitate technical dialogue with scientists

Allinea's vision

- **Helping maximize HPC efficiency**

- Reduce HPC systems operating costs
- Resolve cutting-edge challenges
- Promote Efficiency (as opposed to Utilization)
- Transfer knowledge to HPC communities



- **Helping the HPC community design the best applications**



- Reach highest levels of performance and scalability
- Improve scientific code quality and accuracy

Where to find Alinea's tools

Over 65% of Top 100 HPC systems

- From small to very large tools provision

8 of the Top 10 HPC systems

- From 1,000 to 700,000 core tools usage


Future leadership systems

- Millions of cores usage

“Learn” with Alinea Performance Reports




Executable: MADbench2
Resources: 16 processes, 1 node
Machine: sandybridge2
Start time: Mon Nov 4 12:27:50 2013
Total time: 109 seconds (2 minutes)
Full path: /tmp/MADbench2
Notes: 12-core server / HDD / 16 readers + writers



Summary: MADbench2 is I/O-bound in this configuration

The total wallclock time was spent as follows:



Time spent running application code. High values are usually good. This is **low**; it may be worth improving I/O performance first.

Time spent in MPI calls. High values are usually bad. This is **average**; check the MPI breakdown for advice on reducing it.

Time spent in filesystem I/O. High values are usually bad. This is **high**; check the I/O breakdown section for optimization advice.

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

CPU

A breakdown of how the 4.8% 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. No time was spent in vectorized instructions. Check the compiler's vectorization advice to see why key loops could not be vectorized.

I/O


A breakdown of how the 53.9% total I/O time was spent:



Most of the time is spent in write operations, which have a very low transfer rate. This may be caused by contention for the filesystem or inefficient access patterns. Use an I/O profiler to investigate which write calls are affected.

MPI

Of the 41.3% total time spent in MPI calls:



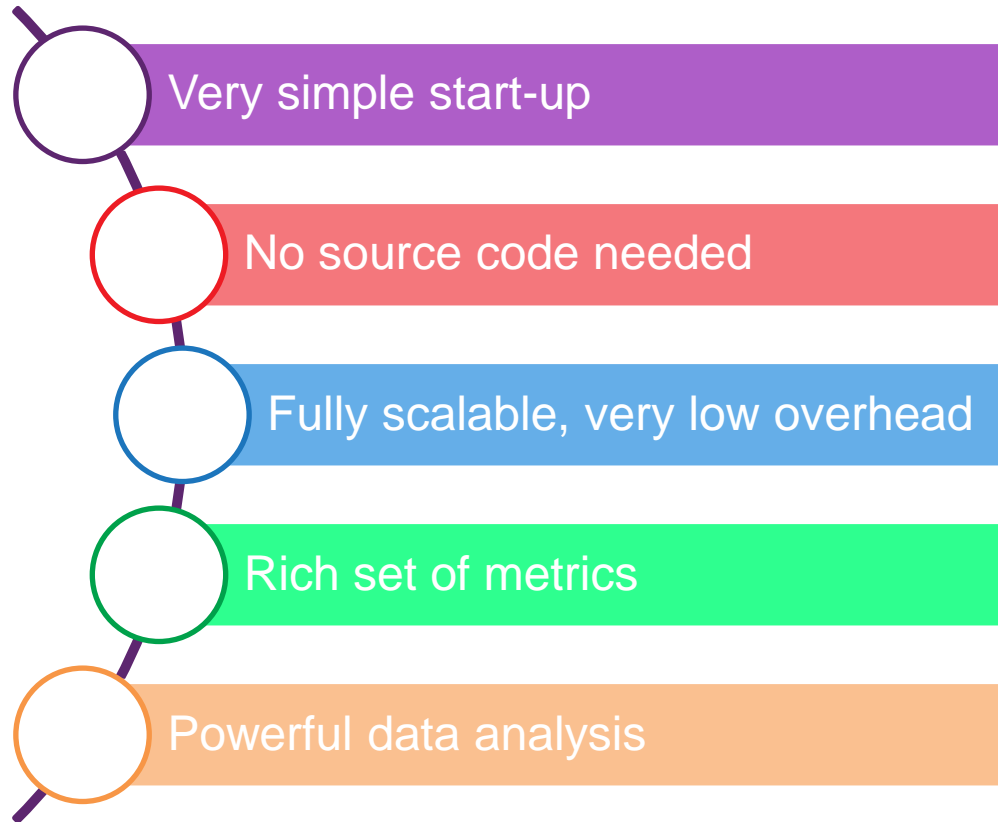
All of the time is spent in collective calls with a very low transfer rate. This suggests a significant load imbalance is causing synchronization overhead. You can investigate this further with an MPI profiler.

Memory

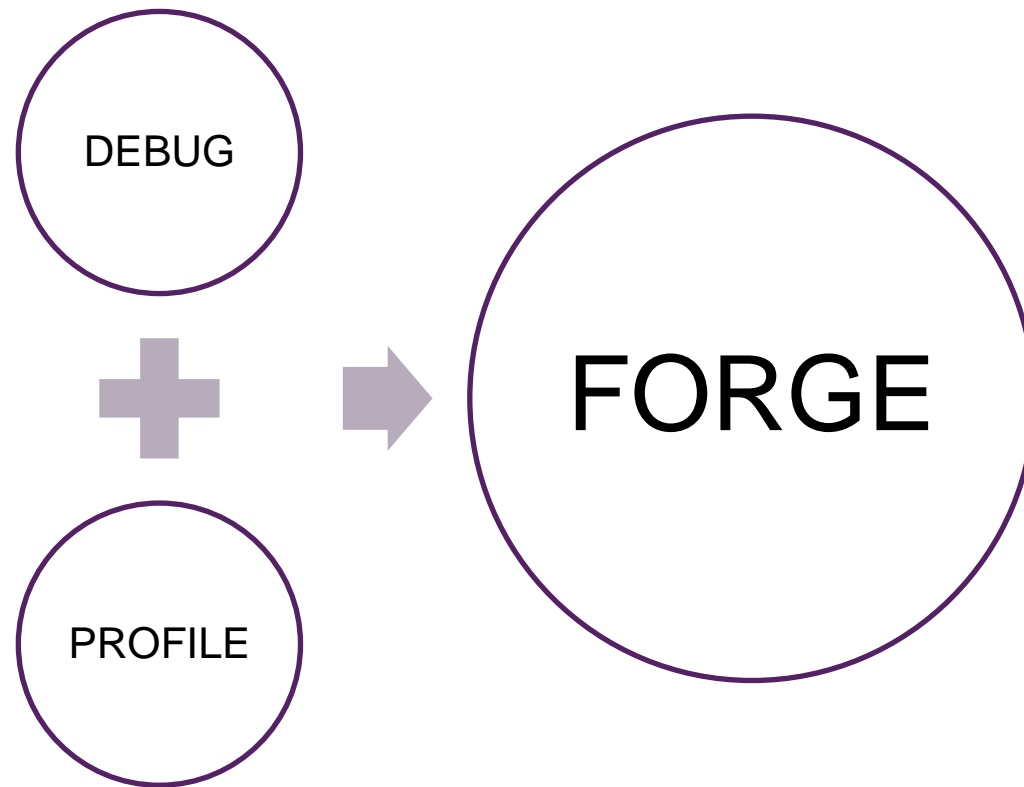
Per-process memory usage may also affect scaling:



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.



Allinea Forge: a toolkit to save developers' time



Allinea Forge: a toolkit to save developers' time

ACCESSIBLE

- Unique single interface
- Easy to start and use

POWERFUL

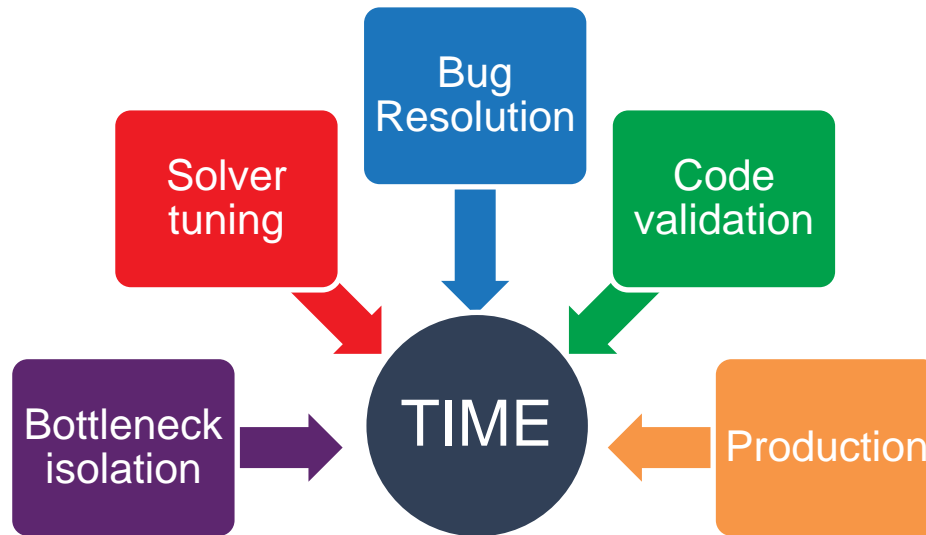
- State of the art features
- Fully scalable

INNOVATIVE

- Tackles new challenges
- For latest HPC systems



We are here today for a single reason.
Help you save your time.




Convergent Science: developing faster software

- **Challenges**

Developing faster and more capable software to meet the growing demands of Convergent Science clients for precision and performance

- **Huge speed up on CONVERGE: from 2h down to 4 sec**

- Now possible to run jobs efficiently on hundreds of cores
- Now possible to scale up from 2 million to 20 million nodes



Full case study:

<http://www.allinea.com/news/201509/convergent-science-ignites-combustion-simulation-performance-allinea-forge>

allinea

allinea

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

Questions?

Feel free to ask anything...!



allinea

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

Technical session


Adding vectorization to linear algebra applications



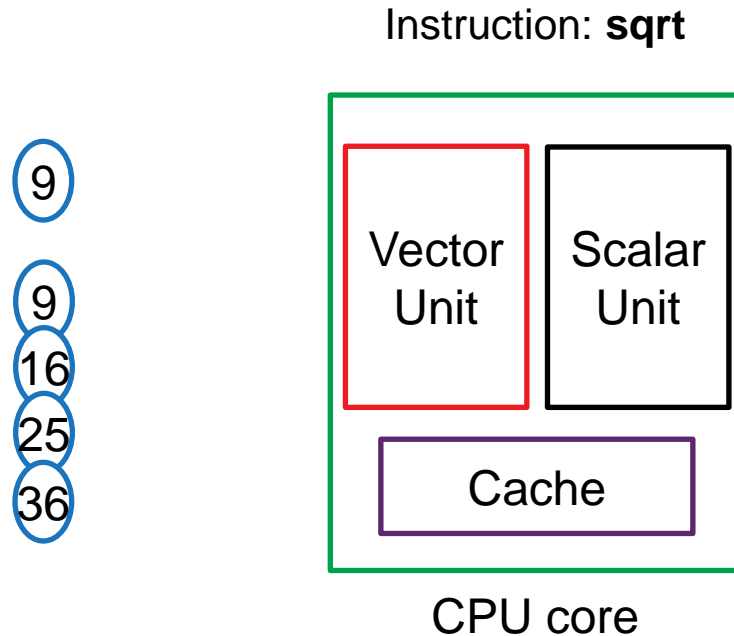
Agenda: Adding vectorization to an application

- What is vectorization?
- Why is vectorization important?
- When can my code use vectorization?
- How to actually enable vectorization?

What is vectorization?



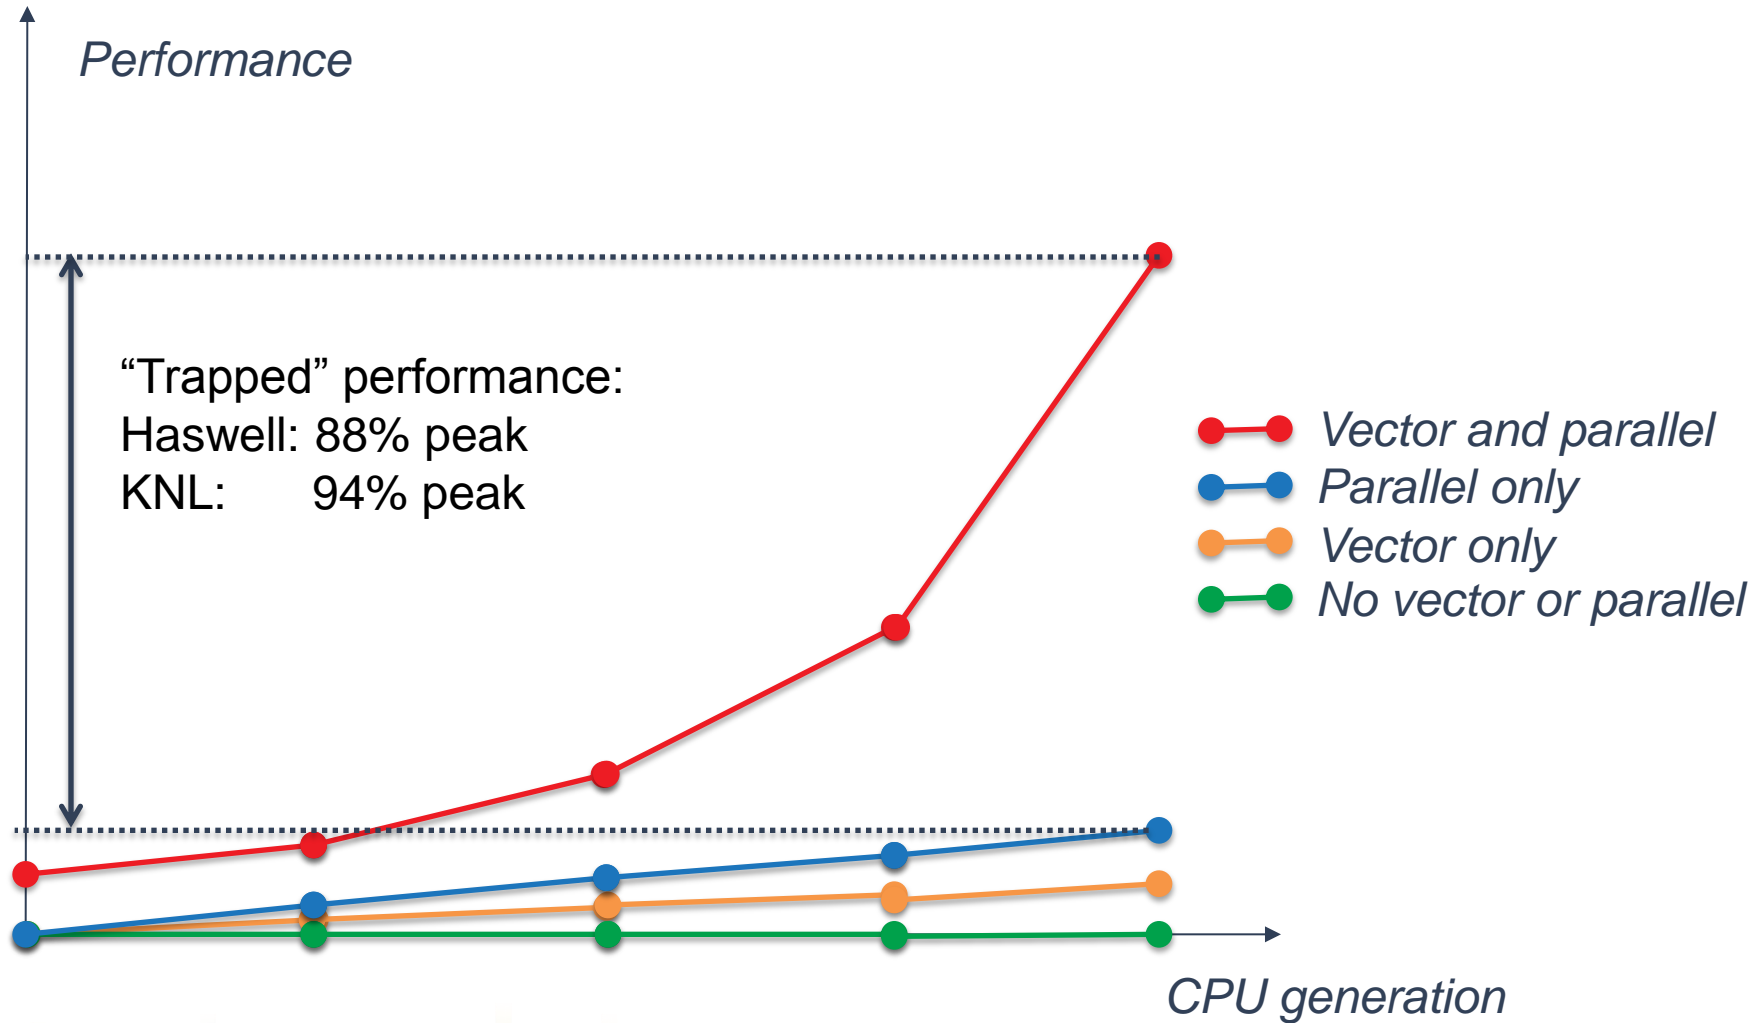
What is vectorization?



Intel Haswell: 256-bit vector unit → 8 SP / 4 DP

Intel Knights Landing: 512-bit vector unit → 16 SP / 8 DP

Why is vectorization important?



Question



Who is in charge of
vectorizing a code?

When can my code use vectorization?

- **A compiler can vectorize a loop if:**
 - The loop is countable at runtime
 - There is a single control flow within the loop
 - The loop does not contain function calls

Is this enough? No!

- Data dependencies
can stop vectorization

Consider the loop:

`a = {0,1,2,3,4}`

`b = {5,6,7,8,9}`

```
for( i=1; i<N; i++)  
    a[i] = a[i-1] + b[i];
```

Applying each operation sequentially:

`a[1] = a[0] + b[1] → a[1] = 0 + 6 → a[1] = 6`

`a[2] = a[1] + b[2] → a[2] = 6 + 7 → a[2] = 13`

`a[3] = a[2] + b[3] → a[3] = 13 + 8 → a[3] = 21`

`a[4] = a[3] + b[4] → a[4] = 21 + 9 → a[4] = 30`

`a = {0, 6, 13, 21, 30}`

When can my code use vectorization?

- **A compiler can vectorize a loop if:**
 - The loop is countable at runtime
 - There is a single control flow within the loop
 - The loop

Now let's try vector operations:

a = {0,1,2,3,4}
b = {5,6,7,8,9}

```
for( i=1; i<N; i++)  
    a[i] = a[i-1] + b[i];
```

Applying vector operations, i={1,2,3,4}:

a[i-1] = {0,1,2,3} (load)
b[i] = {6,7,8,9} (load)
{0,1,2,3} + {6,7,8,9} = {6, 8, 10, 12} (operate)
a[i] = {6, 8, 10, 12} (store)

Is this enough?

- Data dependencies can stop vectorization

```
for( i=1; i<N; i++)  
    a[i] = a[i-1] + b[i];
```

Applying each operation sequentially:

a[1] = a[0] + b[1] → a[1] = 0 + 6 → a[1] = 6
a[2] = a[1] + b[2] → a[2] = 6 + 7 → a[2] = 13
a[3] = a[2] + b[3] → a[3] = 13 + 8 → a[3] = 21
a[4] = a[3] + b[4] → a[4] = 21 + 9 → a[4] = 30

a = {0, 6, 13, 21, 30}

How to actually produce vectorized binaries?

- **Step 1: adopt a “vector-aware” coding methodology**
 - Make sure the code can be vectorized (see above)
 - Enable vectorization during compilation (e.g. `-xhost, ...`)
 - Tell the compiler he can vectorize (`#pragma ivdep`)
 - Use optimized mathematical libraries as much as possible
- **Step 2: use tools to help you...**

... not just any tool. Use the Right Tools.

```
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
```

... maybe?

allinea

allinea

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

Interactive demonstration

Adding vectorization to a linear algebra application



allinea

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

Coffee break



allinea

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

Technical session


Preparing applications for large scale



Agenda: Preparing applications for large scale

- Application design life-cycle
- Key concepts to address scalability issues
- Identifying the limiting factors

Building a scientific application




In your opinion, what is the most critical step?

Example: a sheet of paper (0.1 mm thick)

How many times do you need to fold a sheet of paper to make it as thick as the Eiffel Tower (300m)?

Answer: 22 times

Because: $0.1\text{mm} * 2^{22} = 419\text{m}$



Building a scientific application

MODEL

- Science



ALGORITHM(S)

- Complexity
- Parallelism
- Scalability

COMPILATION



APPLICATION PROFILE

- Profile
- Tune

In your opinion

... is the most critical step?



Example: How many times (0.1 mm thick)

How many times do you need to fold a sheet of paper to reach the height of the Eiffel Tower (300m)?

Answer: 22 times

Because: $0.1\text{mm} * 2^{22} = 419\text{m}$

ALGORITHM(S)

- Complexity
- Parallelism
- Scalability

Algorithms Complexity Analysis

Definition:

In computer science, the time complexity of an algorithm quantifies the amount of time taken by the algorithm to run, as a function of the length of the string representing the input.

Example: “Schoolbook Matrix Multiplication”

```
Input:  matrices A(n,m) and B(m,p)
Output: matrice C(n,p)
For i from 1 to n:
  For j from 1 to p:
    Let sum = 0
    For k from 1 to m:
      Set sum ← sum + Aik × Bkj
    Set Cij ← sum
Return C
```

The calculation time is a function of
“n x p x m”

The time complexity of this algorithm is written $O(npm)$ or $O(n^3)$.


This is a **cubic complexity**...

... which is terrible!

Amdahl's Law

Definition:

Amdahl's law gives the theoretical speedup of the execution of a task *at fixed workload* that can be expected of a system whose resources are improved.



Example:

A sequential program runs in 20h.

A 1h-long part of this program cannot be parallelized (parallel portion = 95%)

Amdahl's law says:

This program will never run in less than 1h (speedup of 20).

Scalability in practice

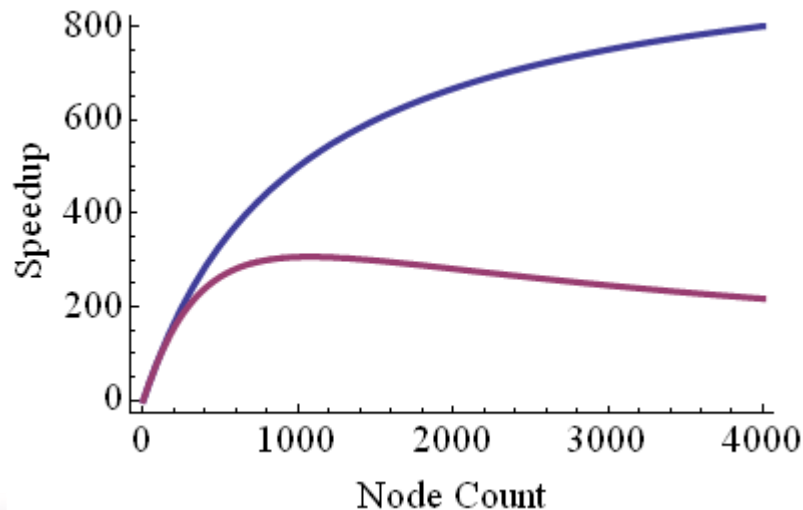
Total Time = Computation + Message Latency + Message Transfer + Noise

Computation = Serial time + (Parallel Time / N)

Message Latency = number of messages * latency

Message Transfer = size of messages / bandwidth

Noise = pre-empting processes, swapping, faulting pages...



Blue : Amdahl's law

Purple : Amdahl's law with messages

Tips for better scalability

Communication between neighbors as opposed to global

- All tasks communicate to all tasks: $O(n^2)$!
- Tasks communicate with their neighbors only: $O(1)$!

Try and avoid global communications!

Improve load balance to reduce synchronization

- When processes synchronize, they go as quickly as the slowest
- Check the time spent in synchronization tasks (e.g. MPI_Barrier, etc.)*

Monitor the supercomputer noise

- Lost microseconds and milliseconds can turn into seconds


Pin your tasks, etc...



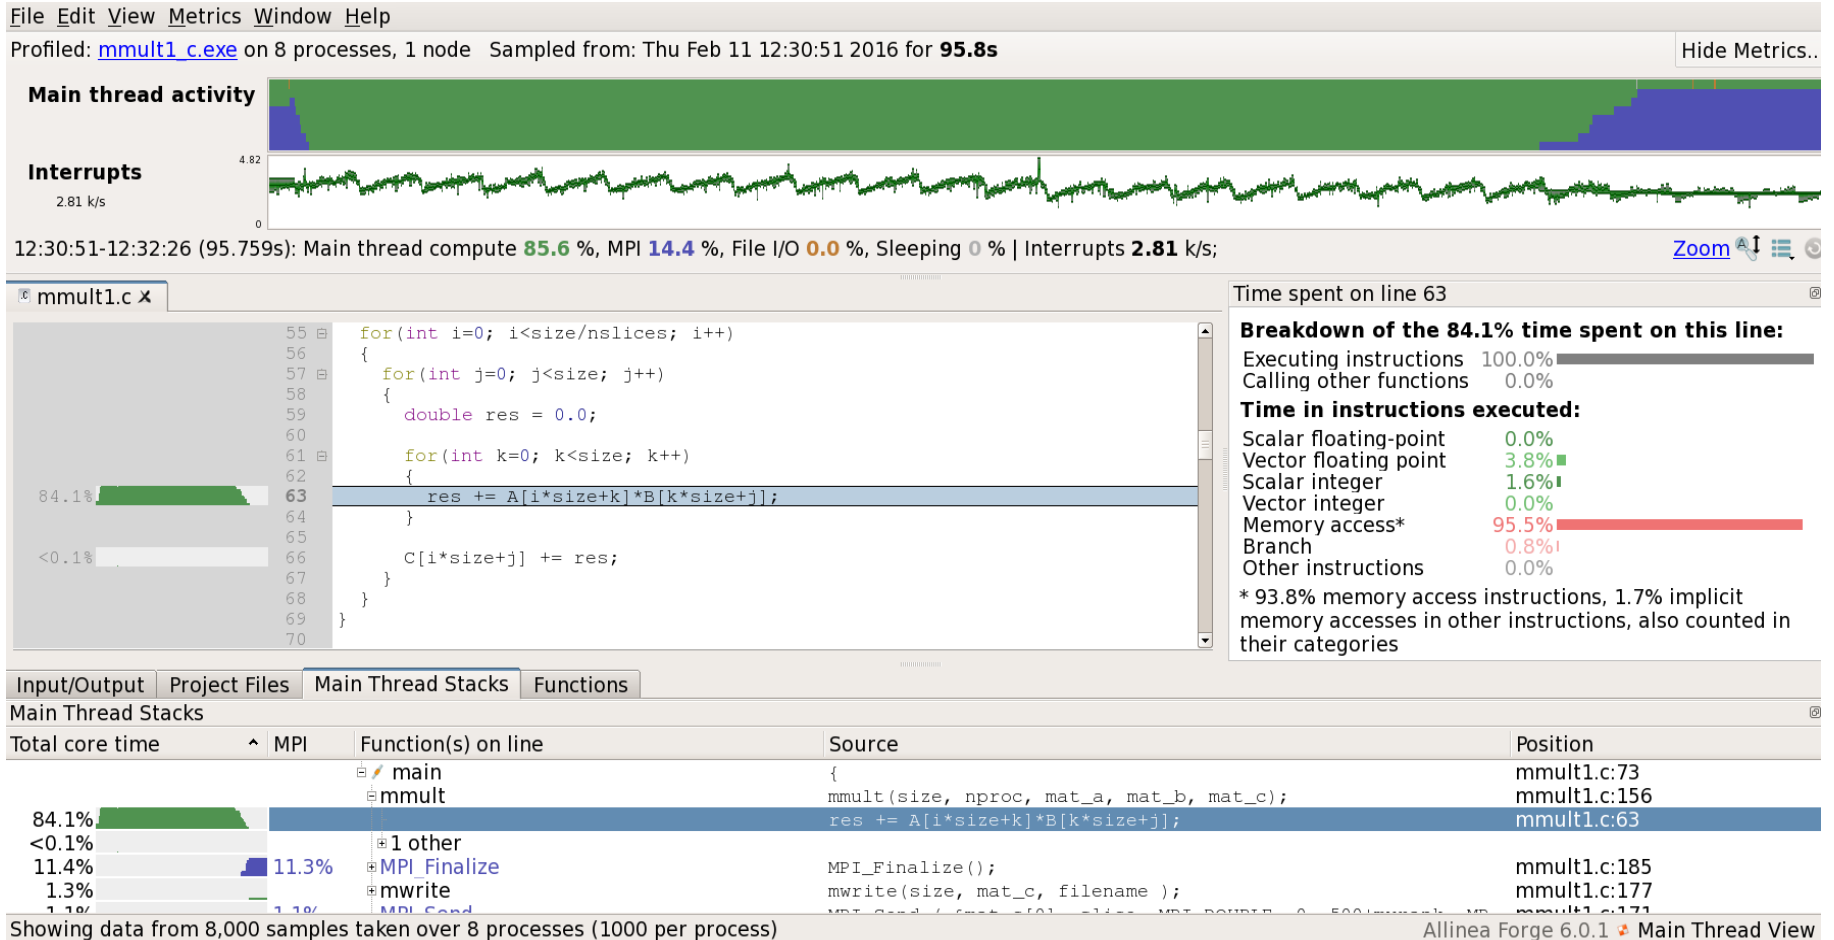
So... what is a scalable algorithm?

As a rule of thumb, an application will scale if:

- Its algorithms have a good complexity
- Its workload can be split into independent tasks
- Communications are infrequent or unnecessary
- Lots of calculations take place before messaging or I/O occurs
- The system is under control
- All the above remains true as the number of tasks grows.



Talking about system under control...



allinea

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

Interactive demonstration

Preparing an application for large scale



allinea

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

A real life example

Bringing HemelB to Petascale with Allinea




UCL: HPC to fight aneurysms and save lives

- **Challenges**

Make surgical decisions within minutes (instead of hours) following patients MRI scans

- **UCL got HemelB to perform and scale to above 50k cores**
 - Fixed limiting performance bottlenecks and crashes at scale
 - Study hemodynamics within the Circle of Willis for the 1st time

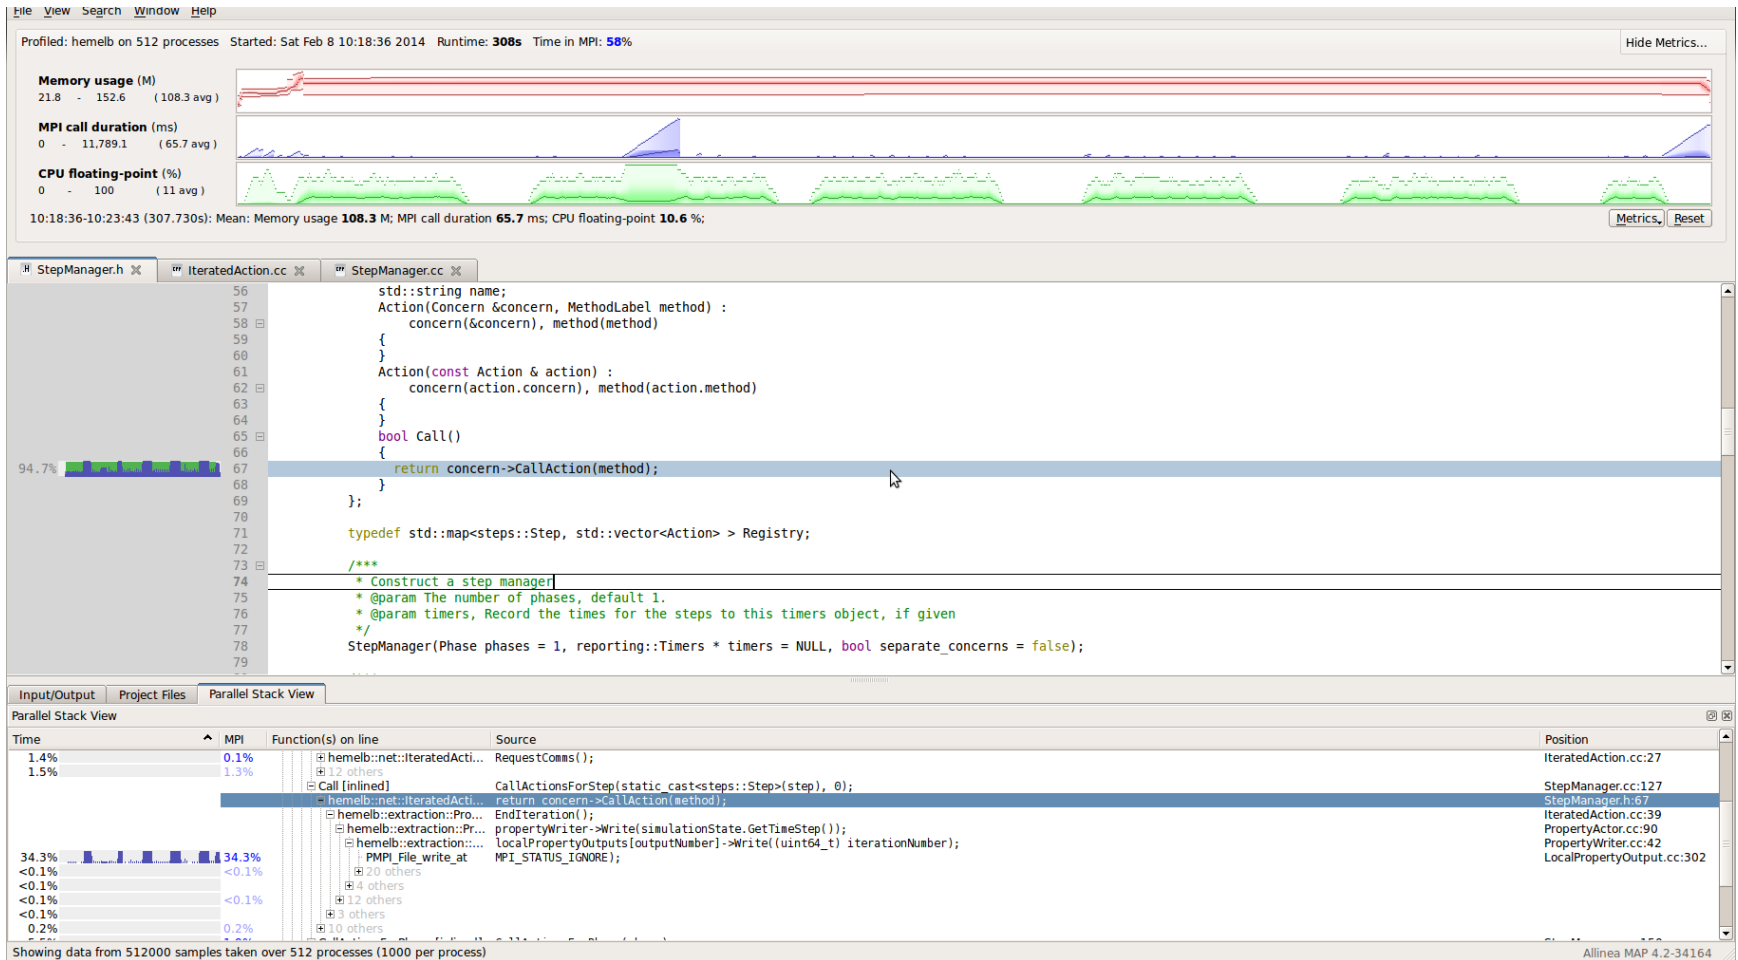


Full case study:

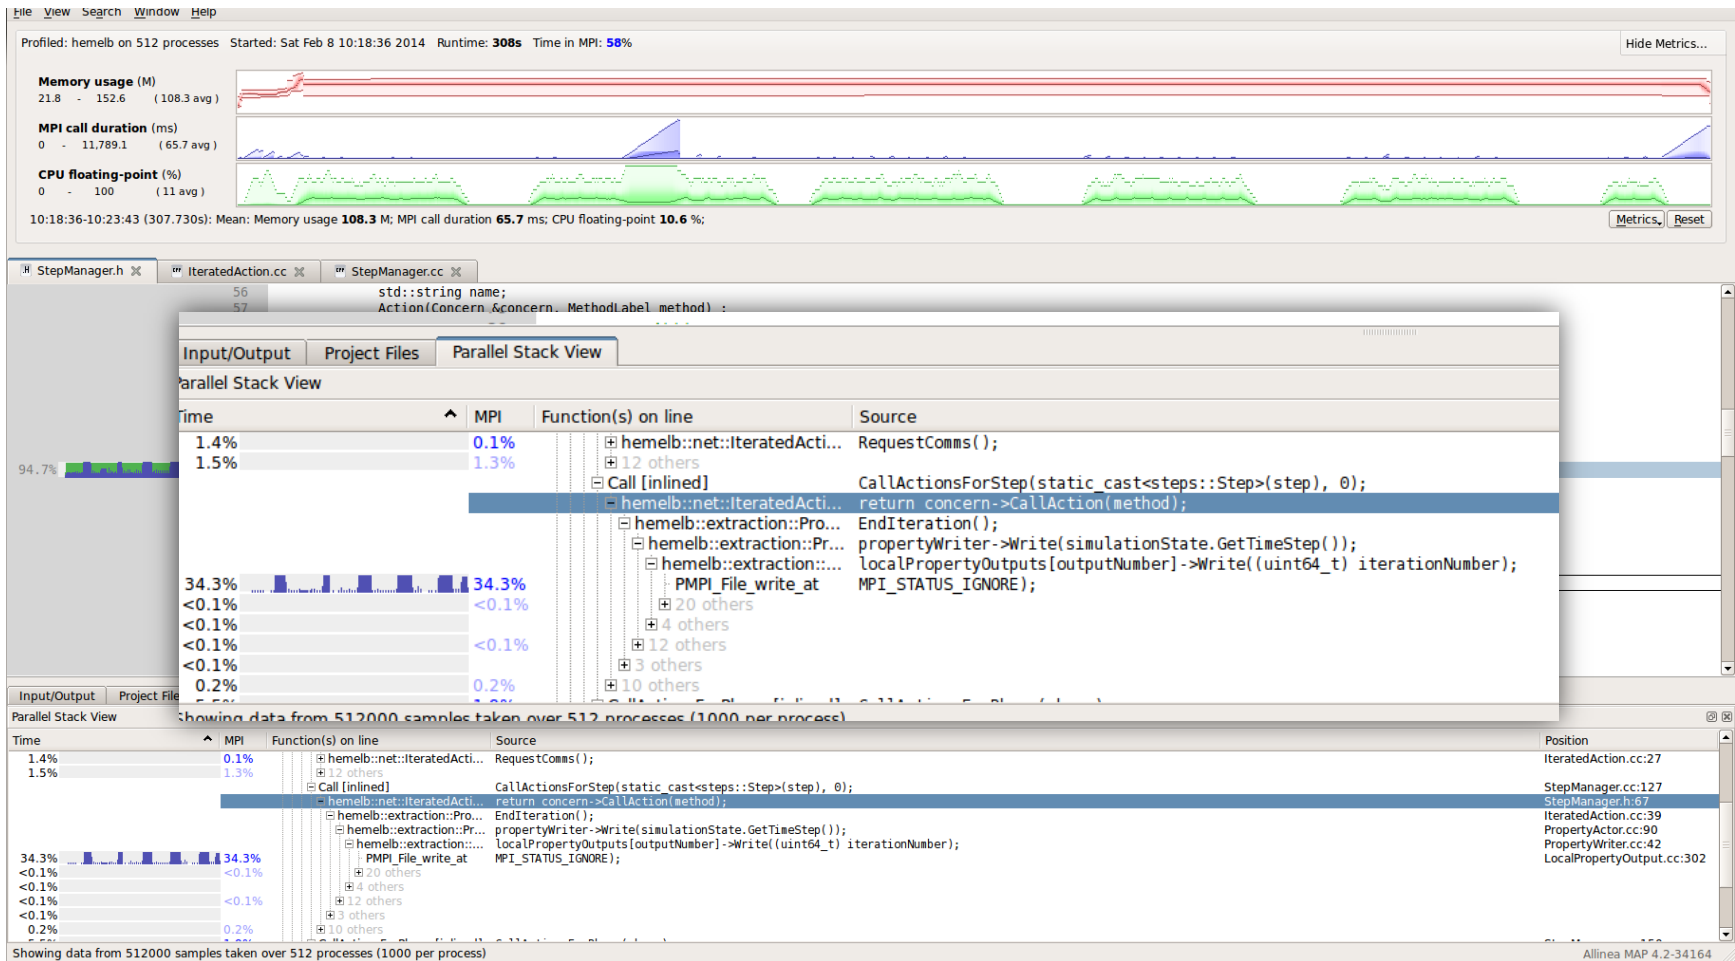
<http://www.allinea.com/case-studies/education/application-soars-above-petascale-after-tools-collaboration>

allinea

Scaling issue – at 512 processes



Scaling issue – at 512 processes



Simple fix... reduce periodicity of output

... leading to a bug!

The screenshot shows a debugger window with the following components:

- Code Editor:** Shows the source code for `xyzpart.c`. The current line is 556: `mypicks[i] = allpicks[i*ntsamples/npes];`. The code includes a sorting function `ikvsortii` and a loop for selecting final splitters.
- Locals:** A table showing local variables and their values:

Variable Name	Value
allpicks	0x2aab8035e010
ctrl	0x13d94f0
elmnts	0x25ad950
firstvtx	<value optimized out>
graph	0xbca010
i	<value optimized out>
j	<value optimized out>
k	<value optimized out>
lastvtx	1143373824
mype	19
mypicks	0x2692160
nlsamples	<value optimized out>
npes	<value optimized out>
nrcvc	1065353216
ntsamples	<value optimized out>
ntvtx	9224
- Stacks:** A table showing the call stack:

Processes	Threads	Function
17220	17220	main (main.cc:37)
17220	17220	SimulationMaster::SimulationMaster (SimulationMaster.cc:63)
17220	17220	SimulationMaster::Initialise (SimulationMaster.cc:154)
17220	17220	hemelb::geometry::GeometryReader::LoadAndDecompose (GeometryReader.cc:188)
17220	17220	hemelb::geometry::GeometryReader::OptimisedDomainDecomposition (GeometryReader.cc:809)
17220	17220	hemelb::geometry::decomposition::OptimisedDecomposition::OptimisedDecomposition (OptimisedDecomposition.cc:65)
17220	17220	hemelb::geometry::decomposition::OptimisedDecomposition::CallParmetis (OptimisedDecomposition.cc:181)
17220	17220	ParMETIS_V3_PartGeomKway (gkmetis.c:90)
17220	17220	libparmetis_CoordinatePartition (xyzpart.c:58)
17220	17220	libparmetis_PseudoSimpleSort (xyzpart.c:556)
- Input/Output, Breakpoints, Watchpoints, Stacks, Tracepoints, Tracepoint Output, Logbook:** These panels are visible at the bottom of the debugger.

7356 processes playing Connected to: lecomber@titan.ccs.ornl.gov

Looking at the bug location in the code

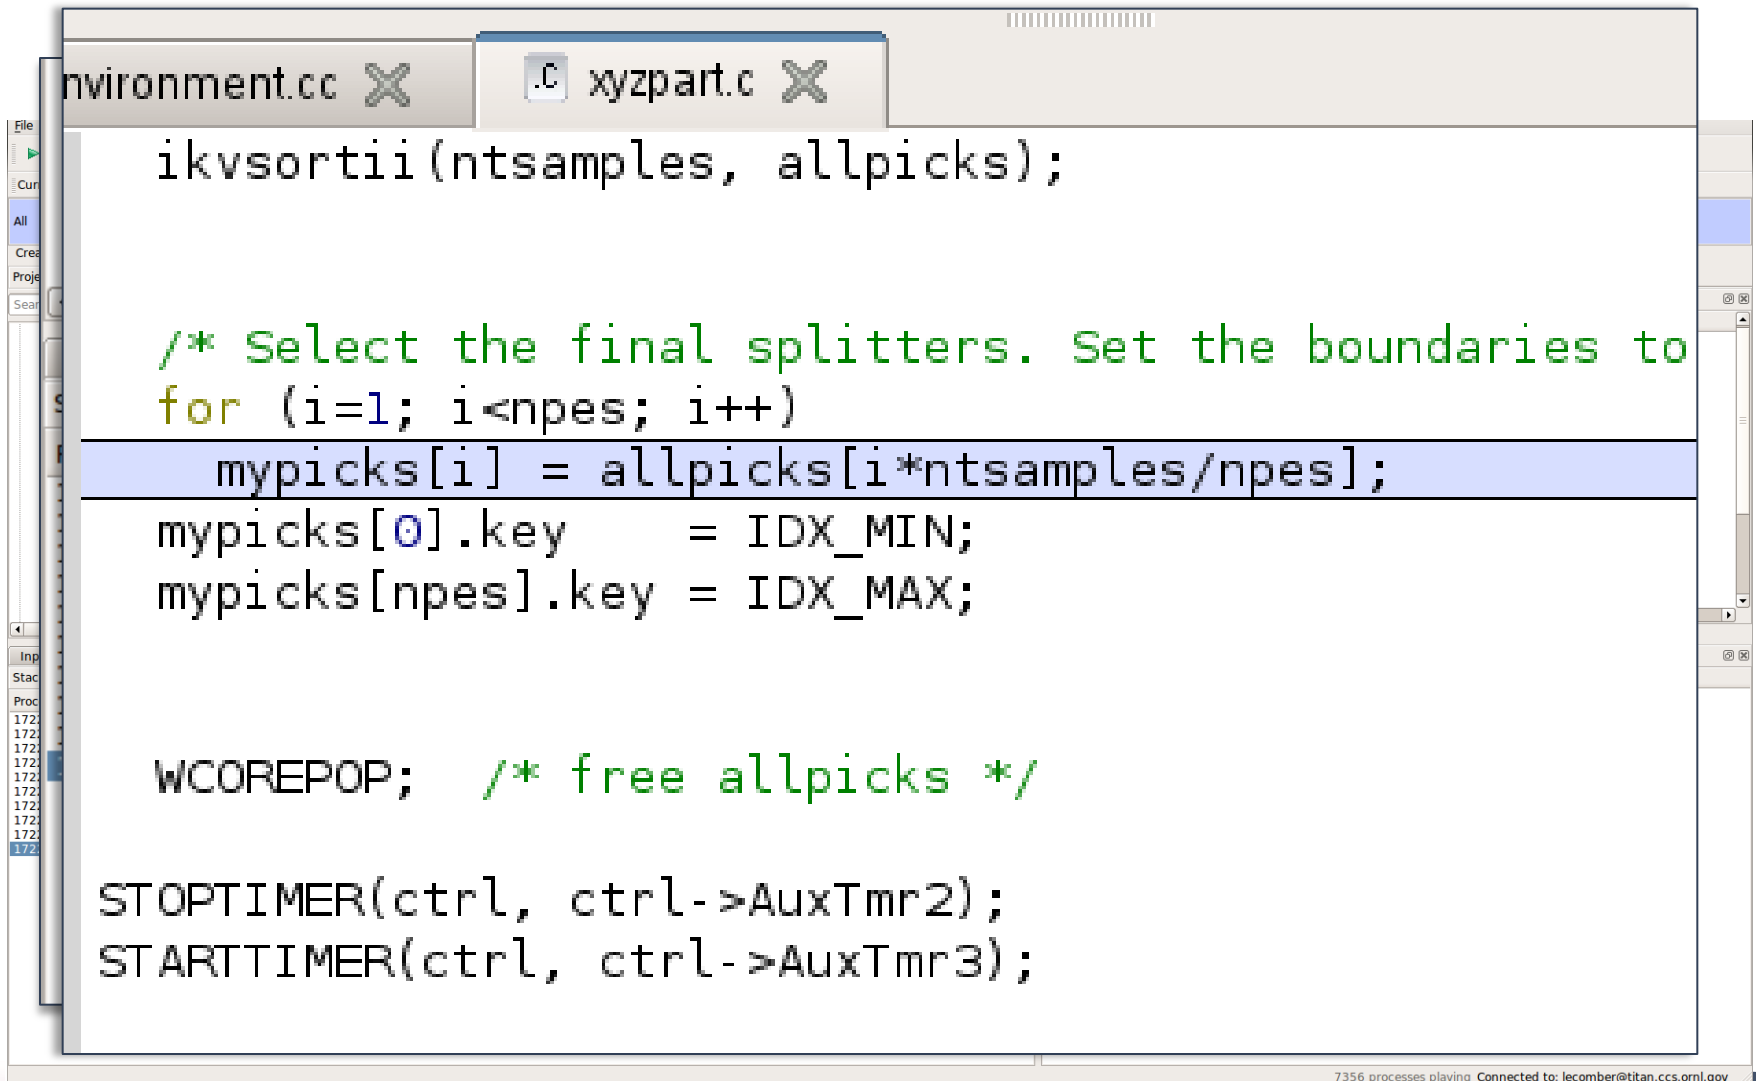
The screenshot displays a debugger interface with two main windows. The top window shows a code editor with the following code:

```
560  
561 WCOREPOP; /* free allpicks */  
562  
563 STOPTIMER(ctrl, ctrl->AuxTmr2);  
564 STARTTIMER(ctrl, ctrl->AuxTmr3);  
565  
566 /* Compute the number of elements that be
```

The bottom window shows the 'Stacks' panel with the following table:

Processes	Threads	Function
17220	17220	main (main.cc:37)
17220	17220	SimulationMaster::SimulationMaster (SimulationMaster.cc:63)
17220	17220	SimulationMaster::Initialise (SimulationMaster.cc:154)
17220	17220	hemelb::geometry::GeometryReader::LoadAndDecompose (GeometryReader.cc:188)
17220	17220	hemelb::geometry::GeometryReader::OptimiseDomainDecomposition (GeometryReader.cc:218)
17220	17220	hemelb::geometry::decomposition::OptimisedDecomposition::OptimisedDecomposition (OptimisedDecomposition.cc:10)
17220	17220	hemelb::geometry::decomposition::OptimisedDecomposition::CallParmetis (OptimisedDecomposition.cc:10)
17220	17220	ParMETIS_V3_PartGeomKway (gkmetis.c:90)
17220	17220	libpmetis_Coordinate_Partition (xyzpart.c:58)
17220	17220	libpmetis_PseudoSampleSort (xyzpart.c:556)

The stack highlights this particular line



The screenshot shows a code editor window with two tabs: 'environment.cc' and 'xyzpart.c'. The code in 'xyzpart.c' is as follows:

```
ikvsortii(ntsamples, allpicks);

/* Select the final splitters. Set the boundaries to
for (i=1; i<npes; i++)
    mypicks[i] = allpicks[i*ntsamples/npes];
mypicks[0].key    = IDX_MIN;
mypicks[npes].key = IDX_MAX;

WCOREPOP; /* free allpicks */

STOPTIMER(ctrl, ctrl->AuxTmr2);
STARTTIMER(ctrl, ctrl->AuxTmr3);
```

The line `mypicks[i] = allpicks[i*ntsamples/npes];` is highlighted in blue. On the left side of the editor, a stack trace is visible, showing the current function and its callers. The stack trace includes the following entries:

```
Inp
Stac
Proc
1722
1722
1722
1722
1722
1722
1722
1722
1722
1722
```

Here are the related variables... Optimized out by the compiler ! Need to compile with `-O0`!

Environment cc x .C xyzpart.c x

... k	<value optimized out>
... lastvtx	1143373824
... mype	19
+ mypicks	0x2692160
... nlsamples	<value optimized out>
... npes	<value optimized out>
... nrecv	1065353216
ntsamples	<value optimized out>
... nvtxs	9224

Type: idx_t

7356 processes playing Connected to: lecomber@titan.ccs.ornl.gov

Variables are much clearer when compiled with low optimization options. Looking good!

The screenshot shows a debugger's variable window with the following data:

Variable Name	Value
+· allpicks	0x2aab8055e1
··· i	2245
+· mypicks	0x2a6f8f0
··· npes	24575
··· ntsamples	1818550

Below the table, the text "Type: none selected" is visible. At the bottom of the window, there are buttons for "Evaluate" and "Expression", and a "Value" column header. The email address "ecomber@titan.ccs.ornl.gov" is printed in the bottom right corner of the window.

But the actual array index is not looking good at all... this shows...

The screenshot shows a debugger's Evaluate window. At the top, it says "Type: none selected". Below that is an "Evaluate" button. A table with two columns, "Expression" and "Value", is displayed. The first row shows the expression "...i * ntsamples" with a value of "-212322546". A mouse cursor is pointing at the value. A yellow tooltip box is overlaid on the value, containing the text: "Type: int", "Range: from -2147259746 to -12282046", and "49/17223 processes equal". The background shows a stack of processes with addresses 1722.

Expression	Value
...i * ntsamples	-212322546

Type: int
Range: from -2147259746 to -12282046
49/17223 processes equal

Ah... Integer overflow!

Type: int
Range: from -2147259200 to -12282046
49/17223 processes equal

Expression Value



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
- XdrMemRead
- XdrMemWrite
- XdrReader.c
- XdrWriter.cc
- XmlAbstract
- xyzpart.c
- External Code

```

551 ikvsortii(ntsamples, allpicks);
552
553
554 /* Select the final splitters. Set the boundaries to s
555 for (i=1; i<nps; i++)
556     mypicks[i] = allpicks[i*ntsamples/nps];
557 mypicks[0].key = IDX_MIN;
558 mypicks[nps].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

Variable Name	Value
allpicks	0x2aab8055e010
i	2245
mypicks	0x2a6f8f0
nps	24575
ntsamples	1818550

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 (Optimise
17223	17223	ParMETIS_V3_PartGeomKway (gkmetis.c:90)
17223	17223	libparmetis_Coordinate_Partition (xyzpart.c:58)
17223	17223	libparmetis_PseudoSampleSort (xyzpart.c:556)

Type: none selected

Evaluate

Expression	Value
i*ntsamples	-212322546

Type: int
Range: from -2147259746 to -12282046
49/17223 processes equal

allinea


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

Allinea new stuff and roadmap

More is available under non-disclosure agreement!



Energy efficiency with Alinea's tools




Energy

A breakdown of how the 3.6 Wh was used:


CPU	62.9%	
System	37.1%	
Mean node power	92.4 W	
Peak node power	94 W	

Significant energy is wasted during MPI communications. It may be more efficient to use fewer nodes with more data on each node.

Significant time is spent waiting for memory accesses. Reducing the CPU clock frequency could reduce overall energy usage.



Quantify gains immediately



CPU

A breakdown of the 94.6% CPU time:

Scalar numeric ops	11.7%	
Vector numeric ops	0.0%	
Memory accesses	88.2%	█
Waiting for accelerators	0.0%	

The per-core performance is **memory-bound**. Use a profiler to identify time-consuming operations.

No time is spent in vectorization advice.

Energy

A breakdown of how the 3.6 Wh was used:

CPU	62.9%	█
System	37.1%	█
Mean node power	92.4 W	█
Peak node power	94 W	█

Significant energy is wasted during MPI communications. It may be more efficient to use fewer nodes with more data on each node.

Significant time is spent waiting for memory accesses. Reducing the CPU clock frequency could reduce overall energy usage.


Accelerators

A breakdown of how accelerators were used:

GPU utilization	0.0%	
Global memory accesses	0.0%	
Mean GPU memory usage	0.0%	
Peak GPU memory usage	0.0%	

GPUs are available but are not used. Identify suitable hot loops with a profiler and try offloading them to the accelerator.

The peak device memory usage is low. It may be more efficient to offload a larger portion of the dataset to each device.



CPU

A breakdown of the 70.2% CPU time:

Scalar numeric ops	2.5%	
Vector numeric ops	0.0%	
Memory accesses	39.7%	█
Waiting for accelerators	61.7%	█

Most of the time is spent waiting for accelerators. Use asynchronous calls to overlap CPU and accelerator workloads.

The per-core performance is **memory-bound**. Use a profiler to identify time-consuming loops and check their cache performance.

Energy

A breakdown of how the 2.84 Wh was used:

CPU	28.4%	█
System	71.6%	█
Mean node power	163 W	█
Peak node power	175.8 W	█

Energy usage appears to be optimal.

Accelerators

A breakdown of how accelerators were used:

GPU utilization	92.5%	█
Global memory accesses	40.4%	█
Mean GPU memory usage	9.6%	
Peak GPU memory usage	15.2%	

Significant time is spent in **global memory accesses**. Try modifying kernels to use shared memory instead and check for bad striding patterns.

The peak device memory usage is low. It may be more efficient to offload a larger portion of the dataset to each device.

Allinea R&D Programs in preparation for Exascale

- **NRE Projects with HPC Centers**

- ORNL: Application-level Trapped Capacity Reports
- CEA : Providing a scalable interface to CEA profiler (MALP)

- **European projects**

- Mont Blanc 2 : R&D on HPC systems using embedded technologies
- Horizon 2020: Towards Exascale computing
 - **ESIWACE** : Excellence in Simulation of Weather and Climate in Europe
 - **NextGenIO** : Next Generation IO for Exascale
 - **SAGE** : Percipient Storage for Exascale Data-Centric Computing
 - **ExaNest** : European Exascale System Interconnect and Storage
 - **ComPat** : Computing Patterns for High Performance Multiscale Computing

- **National Projects**

- TSERO (UK) : Reducing energy consumption of HPC systems

Support for next generation systems

- **Intel KNL Support**

- Support announced at ISC'2016 in Frankfurt (2 phases release)

- **ARMv8 Support (currently supported: Alinea Forge)**

- Adding CPU metrics support for ARMv8 in Alinea MAP
- Alinea Performance Reports for ARMv8 scheduled (H2020 ExaNest)

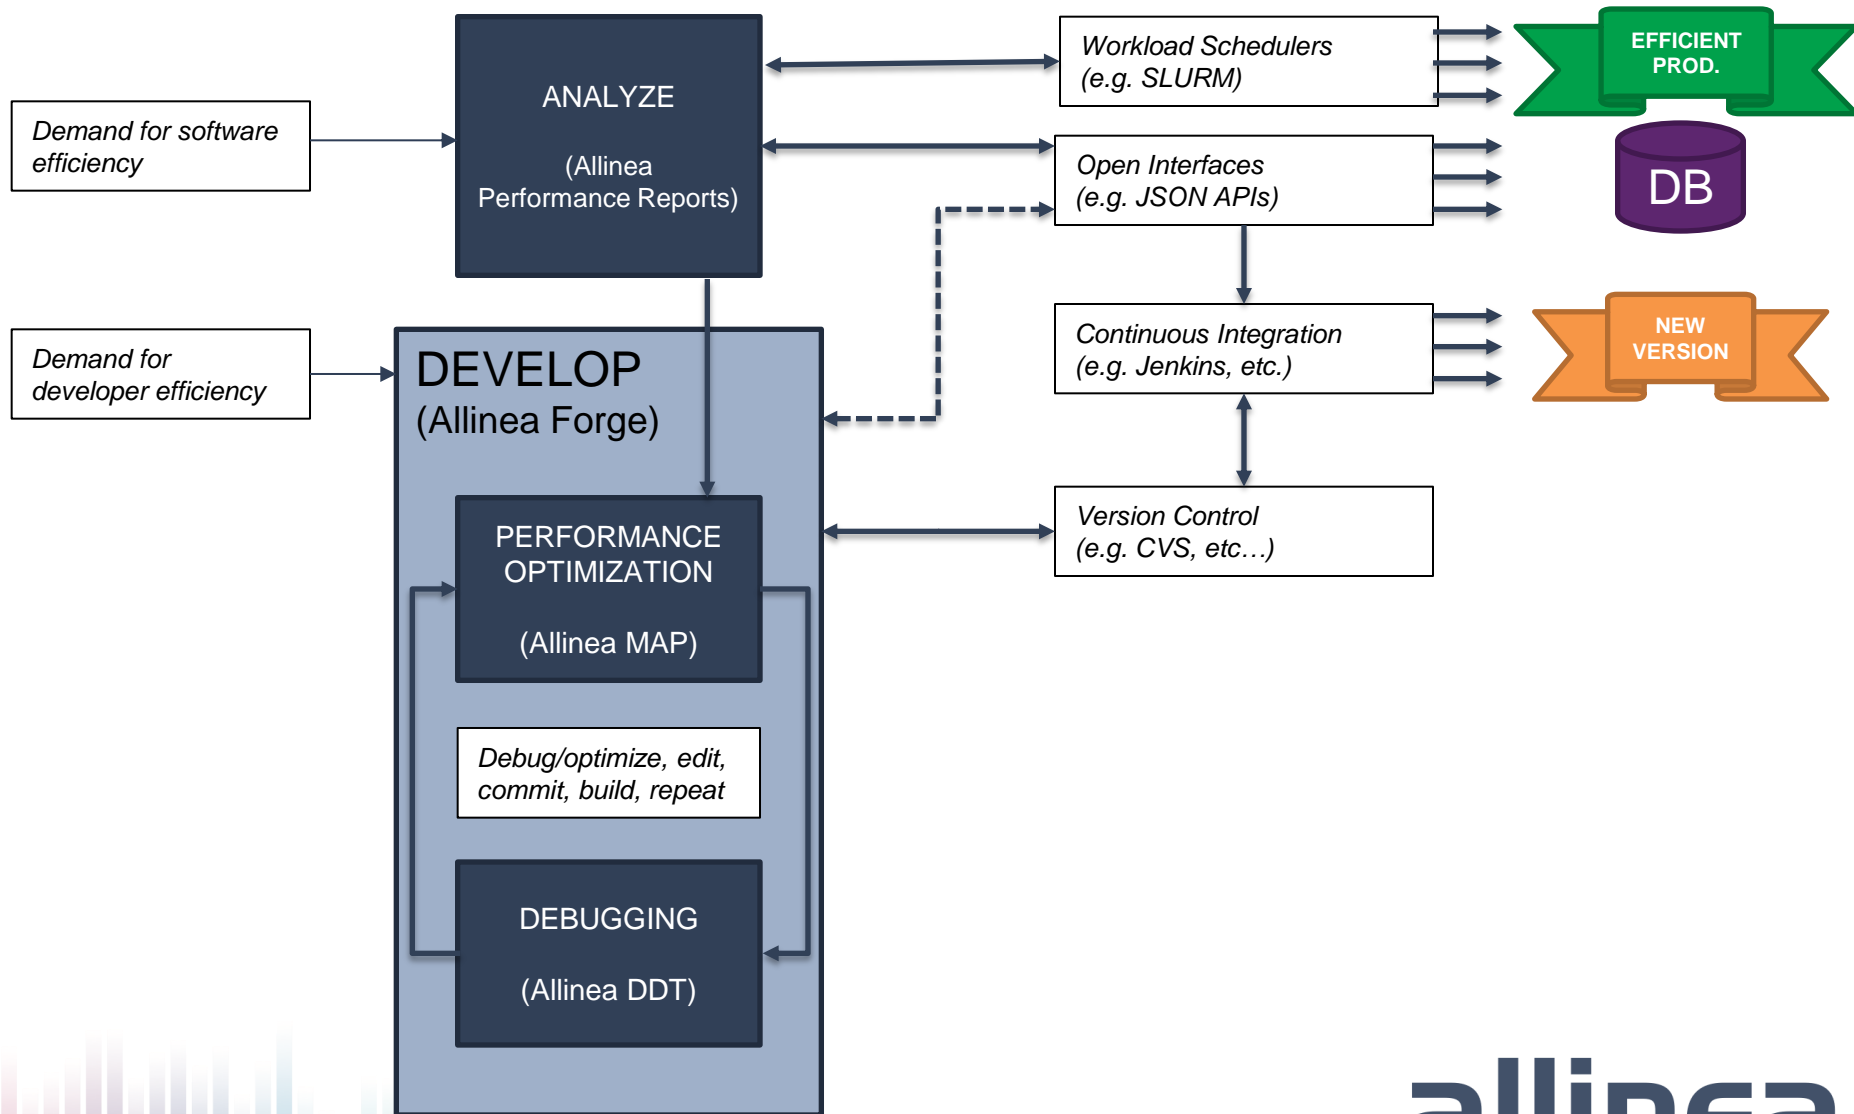
- **Nvidia GPUs Support (currently supported: CUDA 7.5)**

- CUDA 8.0 expected as soon as Nvidia is ready

- **OpenPower Support (currently supported: Alinea Forge)**

- Adding CPU metrics support for Power in Alinea MAP
- Alinea Performance Reports for OpenPower in R&D

Our vision: HPC best practices



Summary



Prepare for application changes and migration

- Change is inevitable

Strengthen professional development workflows

- Reduce your costs

Help the scientific community

- Allinea's mission.

allinea

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

Thank you !

Technical Support team :

support@allinea.com

Sales team :

sales@allinea.com

