

allinea



Leaders in parallel software development tools

Weather and Forecasting

From development to production

With Allinea tools



16th Workshop on High Performance
Computing in Meteorology

Special thanks to DKRZ for their
contribution to this presentation.



Agenda



-
- Weather and forecasting : challenges
 - How to reach the highest quality of service with Allinea?
 - Illustrated with a real-life case study
 - Summary



Weather and Forecasting Challenges



- **Results quality and accuracy**
 - Lack of scalability negatively impacts the size & resolution of meshes
 - Achieving mesh invariance and solution stability is difficult
- **Time to result**
 - Not enough runs per day
 - Pace of research is slowed down by the lack of machine availability
- **New systems : migration and reliability**
 - Code migrations generate software problems
 - Hardware issues need to be identified and solved quickly



Allinea

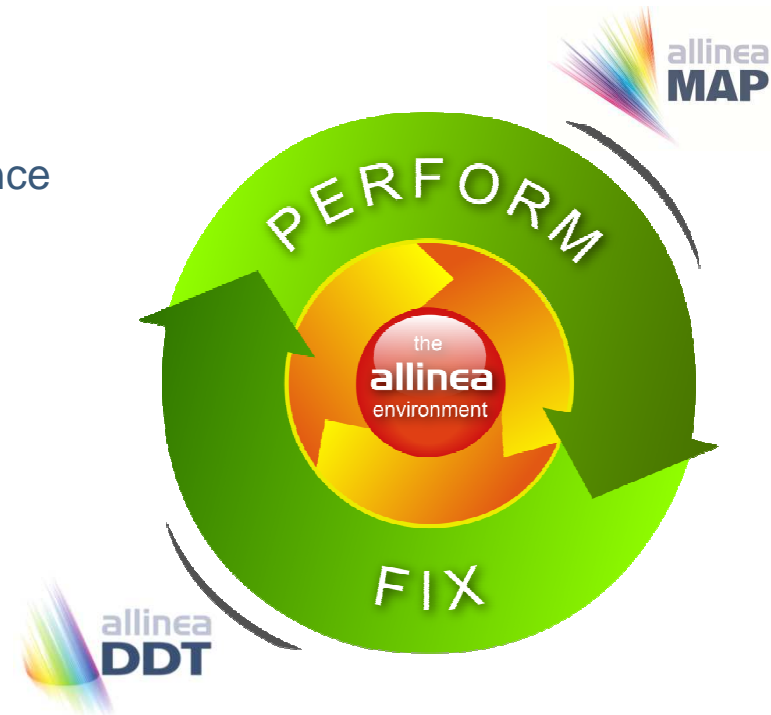


- **HPC tools company since 2002**
 - Leading in HPC software tools market worldwide
 - Large customer base in Weather and Forecasting
 - ECMWF, Met Office, DKRZ, DWD, Meteo France, AEMET, BOM, NOAA...
- **Helping the HPC community design the best applications with Allinea Unified environment**
 - Increase the scale of applications fearlessly
 - Reach the highest level of performance
- **Helping HPC users make the most of their allocations with Allinea Performance Reports**
 - Increase the efficiency of your cluster
 - Identify hardware and software problems in production



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
- Designed for productivity
 - Consistent easy to use tools
 - Enables effective HPC development
- Improve system usage
 - Fewer failed jobs
 - Higher application performance




A real life example

Coupled models: CCLM, CICE, TRIM and Oasis3



- **Once upon a time : new experimental runs in production**
 - New coupling of models in the hope to improve the time to results
 - Scaling up the runs to increase the accuracy of the results
- **And during a profiling workshop at DKRZ...**

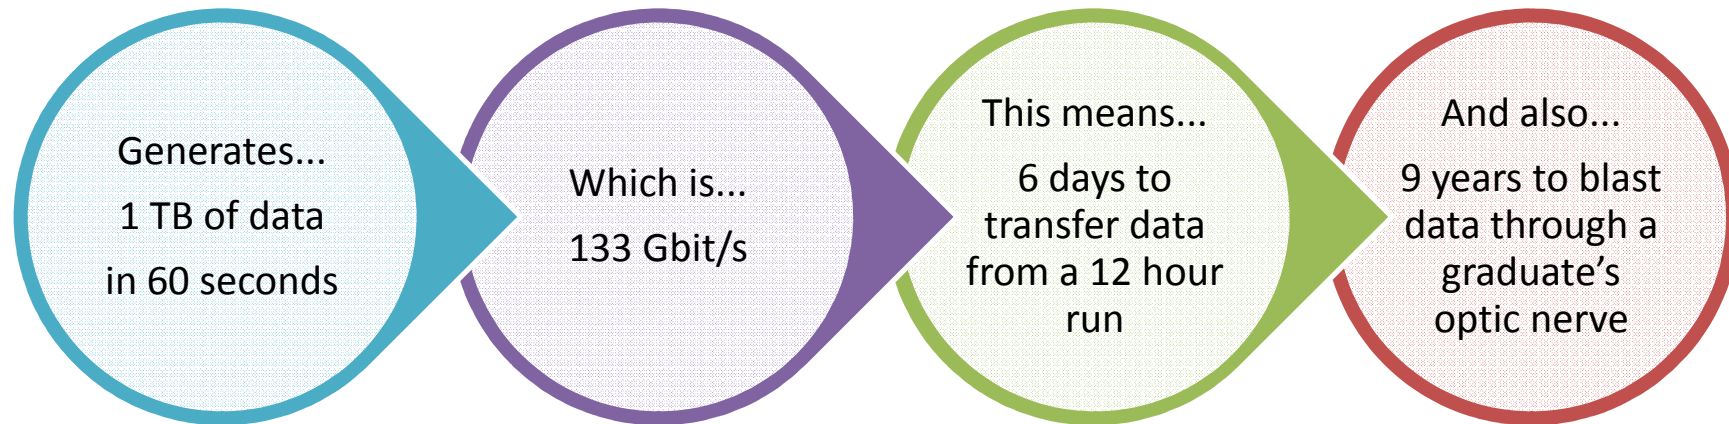


- **Follow-up plan**
 - Creation of a support ticket : major performance problem of the model



Profiling large scale codes

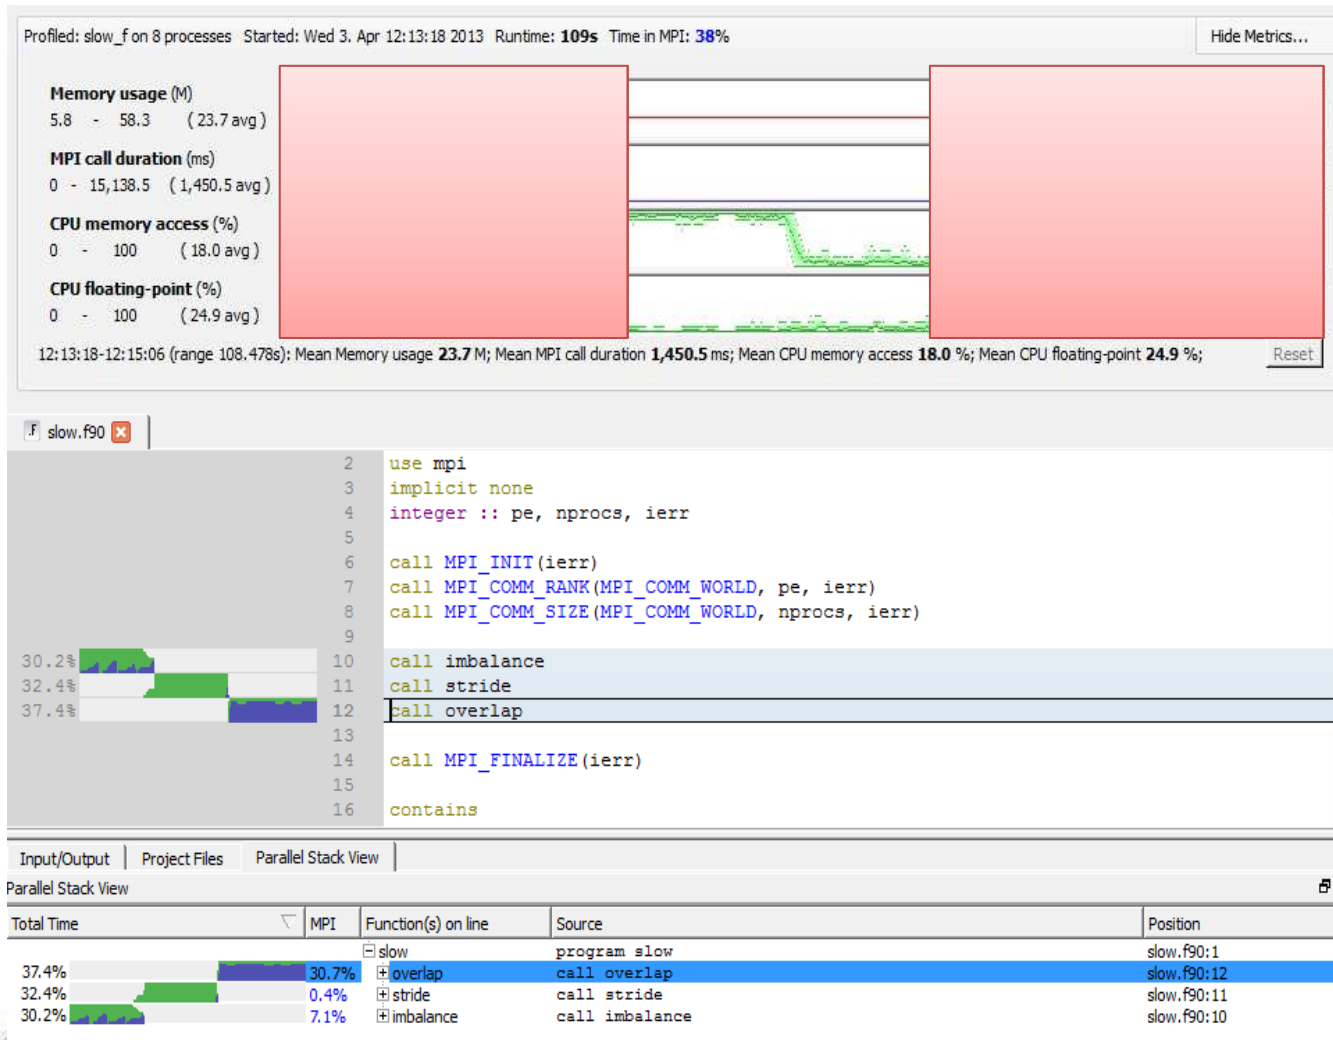
Application running at 16k procs with a tracer



```
void do_math(int i)
{
    const double dtime = 0.3;
    const double c = 1.0;
    const double dx = 1.0;
    double tau, sqtau;

    tau = (c * dtime / dx);
    sqtau = tau * tau;
    newval[i] = (2.0 * values[i]) - oldval[i]
        + (sqtau * (values[i-1] - (2.0 * values[i]) + values[i+1]));
}
```

High level profiling with Alinea MAP



Common horizontal axis



Aggregate across all processes and threads



Highlight imbalance visually



Always refer to source code



Optional: Target specific areas for instrumentation

Statistic sampling or tracing ?

Complementary approaches

Optimize with
Allinea MAP

- Characterize performance at-scale with a lightweight tool
- See which lines of code are hotspots
- Identify common problems at once

Prepare strategy
with Allinea MAP

- Pass more obscure problems to an expert
- Identify loop(s) to instrument
- Identify performance counter(s) to record

Record traces

- Retrieve low level details
 - without generating huge traces
 - without huge overheads

Back to our “performance” issue

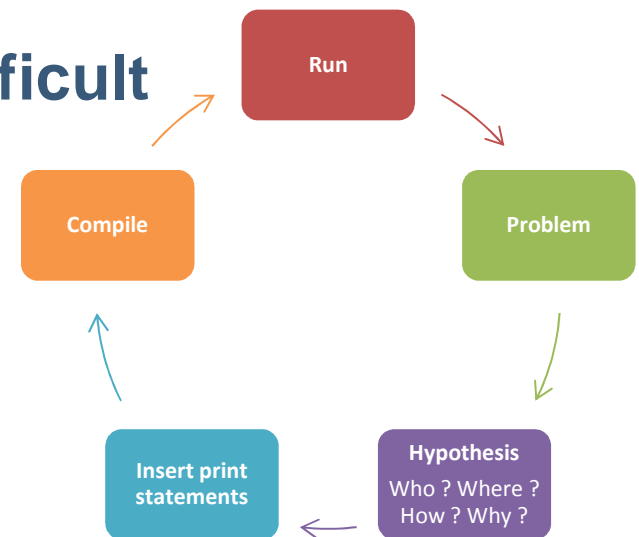
Coupled models: CCLM, CICE, TRIM and Oasis3



- **The performance issue turned out to be a destructive bug**
 - Scaling up to 192 procs : crash !
 - Only happened in this context, on one particular task

- **Bug analysis at large scale can be difficult**

- Core files were useless
 - Crash not related to the actual bug location
- Print statements : unusable
 - Scale of the program was too high
 - Location of the bug was unknown



Allinea MAP & Allinea DDT

One Unified Solution

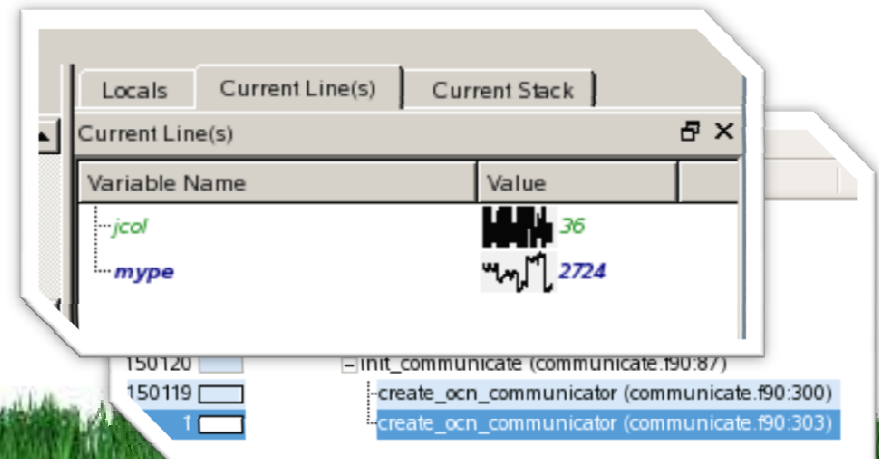


- 1 Use Allinea MAP to find a bottleneck
- 2 Flick to Allinea DDT to understand it
- 3 Compare variables, expressions, call paths
- 4 High memory usage? Use Allinea DDT !
- 5 Common interface and settings files



Allinea DDT helps to understand

- **Who had a rogue behavior ?**
 - Merges stacks from processes and threads
- **Where did it happen?**
 - Allinea DDT leaps to source automatically
- **How did it happen?**
 - Detailed error message given to the user
 - Some faults evident instantly from source
- **Why did it happen?**
 - Unique “Smart Highlighting”
 - Sparklines comparing data across processes



Resolving the TRIM issue

Looking at ocean domain decomposition

The screenshot displays a software interface with several windows:

- Code Editor:** Shows Fortran code for domain decomposition. Key lines include:

```
DO jj=1,nrow
  DO ji=1,ncol
    !b(ji,ii) = p
    !e(ji,ii) = p
    !r(ji,ii) = p
    WRITE(0,*) ii,
  ENDDO
ENDIF ! (old_trim_dec
! if ( m%mynum == 0 )
! print *, "wet po
! print *, "Total v
! print *, "Sub-do
endif
c1_num = (ig-1) * m%
!p = 0
! Setup_sub_domain_bou
```
- DDT - Multi-Dimensional Array Viewer (on p249):** Shows array dimensions and a data table. The array expression is $jk(i, j)$. The data table is:

	1	2	3	4	5	6
1	59	89	134	163	192	230
2	98	136	165	187	209	230
3	128	144	162	187	212	230
- DDT - Visualization (on p249):** Shows a 3D visualization of the domain decomposition. The domain is a rectangular prism divided into several smaller rectangular blocks. The vertical axis is labeled with values 0, 30, 60, 90, and 10. The horizontal axes are labeled with values 1, 2, 3, 4, 5, and 6. A checkbox for "Process 14" is checked.



Resolving the TRIM issue

Looking at ocean domain decomposition

The screenshot shows a software interface for ocean domain decomposition. The main window, titled "DDT - Visualization (on p249)", displays a 3D grid with axes labeled "column", "row", and "depth". A "Process 77" is highlighted in the visualization. The interface includes a "Project Files" pane, a "Code Editor" with Fortran code, and a "Data Table" with numerical data.

```
DO jj=1,nrow
DO ii=1,ncol
  ib(ii,jj) = pbounds(ii,jj)%vs
  is(ii,jj) = pbounds(ii,jj)%vs
  WRITE(0,*) ii, jj, ib(ii,jj), is(ii,jj)
```

j	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23					
1	1	40	48	56	64	72	80	88	96	104	121	133	141	149	157	165	173	181	189	197	205	213	538976288				
2		1	35	42	49	56	64	76	83	92	102	115	131	143	154	162	169	176	183	190	198	205	212	219			
3			1	78	86	97	117	126	135	142	151	157	163	168	173	178	183	188	193	198	203	208	213	218	223		
4				1	106	113	118	123	128	133	138	143	148	154	162	170	178	186	192	197	202	207	212	217	222	227	
5					1	121	125	128	132	135	139	143	147	151	154	157	161	167	173	177	181	186	195	204	211	217	222



With Allinea Unified environment...

Helping developers and support teams



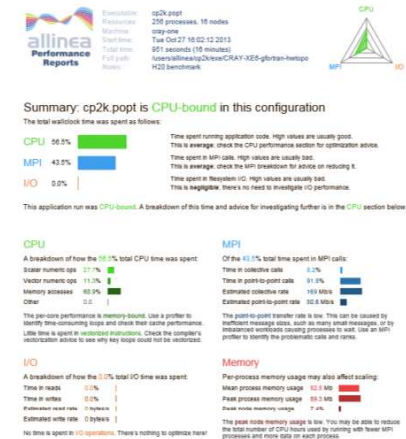
- **... Results quality and accuracy**
 - Resolve scalability issues to increase the size and resolution of meshes
 - Reach the highest level of performance at scale with minimum efforts
- **... Time to result**
 - Increase the amounts of runs per day and the productivity of the system
 - Free machine time to increase the research pace
- **... Migrations and experiments**
 - Resolve bugs and performance issues quickly with Allinea tools

In this case : resolution in a few minutes instead of 1 day.

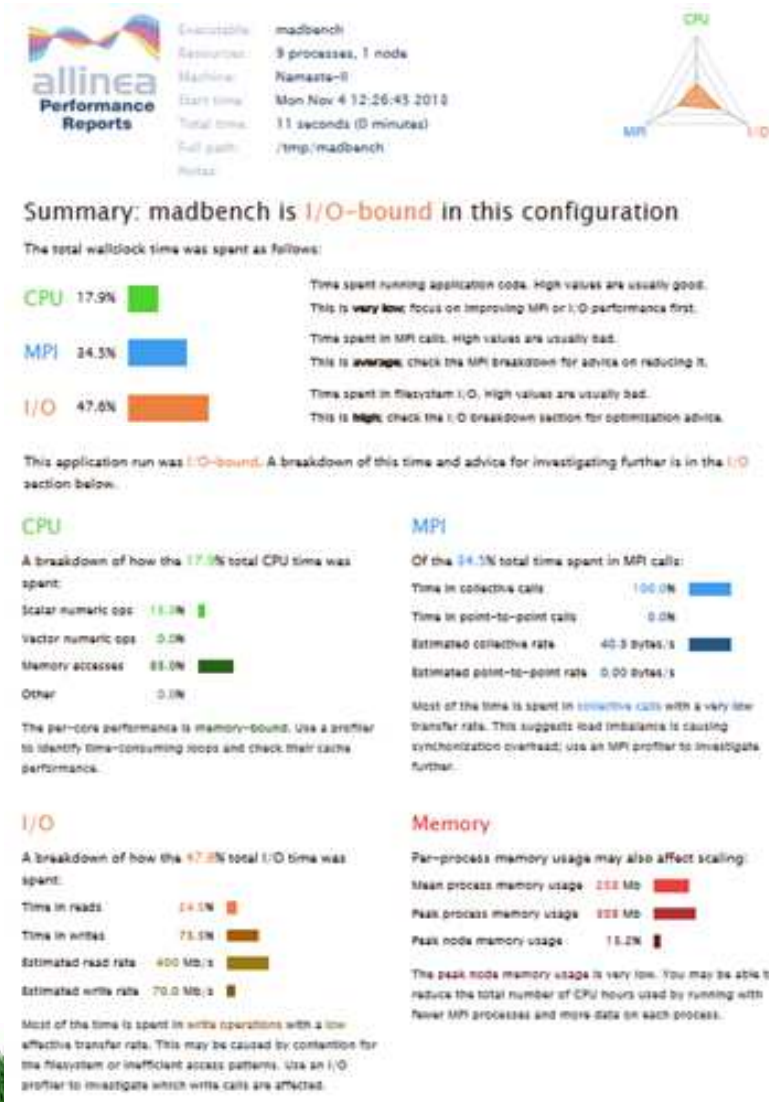


HPC is not just about development


- **Optimization is not always synonym of efficiency**
 - Cluster productivity vs cluster usage
- **Possible efficiency needs during production**
 - Define and enforce best practices (scale, parameters...)
 - Provision and validate cluster upgrades and changes
 - Quickly detect & resolve hardware or software faults impacting performance
- **Effortless one-touch reports with Alinea**
 - Generates explicit and readable reports with metrics and explanations
 - Understand optimized HPC applications effortlessly



Designed for better runs, quickly



- No instrumentation needed
- No source code needed
- No recompilation needed
- Less than 5% runtime overhead
- Fully scalable
- Run regularly – or in regression tests
- Explicit and usable output



With Allinea Performance Reports...

Helping managers, support teams



- **... Retrieve free time on the cluster**
 - Detect heavy demanding users to help them
 - Generate quick feedback on various tweaks and configuration settings
 - Detect software bottlenecks limiting solver performance
- **... Migrations and experiments**
 - Find significant improvements without purchasing extra hardware
 - Determine what type of system or upgrades will benefit research
 - Run and validate codes on new systems in less time



Summary

With Allinea tools



- **Better results quality & accuracy**
 - Improve application scalability to resolve bigger and better scientific problems
 - Achieve mesh invariance and solution stability more easily
- **Better time to result**
 - Squeeze more jobs within a given time frame
 - Improve the research pace by freeing machine time without hardware investment
- **Easier migration, experimentation and systems validation**
 - Resolve quickly code migrations and experimentations problems
 - Identify and resolve hardware issues quickly by looking at the applications



allinea



Leaders in parallel software development tools

Thank you !

Your contacts :

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

Special thanks to DKRZ for their contribution to this presentation.

