

# Uni.lu HPC School 2019

## PS4b: Monitoring & Profiling II: Advanced Performance engineering

---



Uni.lu High Performance Computing (HPC) Team

V. Plugaru, X. Besson

University of Luxembourg (UL), Luxembourg

<http://hpc.uni.lu>



## Latest versions available on Github:



UL HPC tutorials:

<https://github.com/ULHPC/tutorials>

UL HPC School:

<http://hpc.uni.lu/hpc-school/>

PS4b tutorial sources:

[ulhpc-tutorials.rtf.d.io/en/latest/debugging/advanced/](http://ulhpc-tutorials.rtf.d.io/en/latest/debugging/advanced/)





# Summary

- 1 Introduction**
- 2 Debugging and profiling tools
- 3 Conclusion



# Foreword

So we have some news...

# Foreword



The screenshot shows the homepage of the Luxembourg government website. At the top, there is a navigation bar with the logo of the Government of Luxembourg and the text "gouvernement.lu". Below this, there is a search bar and a menu with categories like "Actualités", "Le gouvernement", "Ministères", "Administrations", "Dossiers", and "Système politique". The main content area features a news article titled "Le superordinateur luxembourgeois 'Meluxina' fera partie du réseau européen EuroHPC". The article is dated 14.06.2019 and is categorized as "Communiqué". The text of the article states that on June 14, 2019, the Vice-Premier ministre, ministre de l'Économie, Étienne Schneider, the directeur général adjoint de la Direction générale des réseaux de communication, du contenu et des technologies à la Commission européenne, Khalil Rouhana, and the CEO of LuxConnect, Roger Lampach, presented the future Luxembourg supercomputer, named "Meluxina", which will join the EuroHPC network of supercomputers. Below the text is a photograph of four men sitting at a table during a press conference. The background of the photo shows the Luxembourg flag and the European Union flag, along with a banner for the "MINISTÈRE DE L'ÉCONOMIE".

# Foreword



NEWS CLUB BUSINESS GUIDE JOBS

PAPERJAM  
BUSINESS ZU LËTZEBUERG

POLITIQUE & INSTITUTIONS — POLITIQUE

## SUPERCALCULATEUR

# Meluxina, l'outil du futur du Luxembourg

Écrit par **Thierry Labro**  
Publié Le 14.06.2019 • Édité Le 17.06.2019

Partager



TOP LUS | RECOMMANDÉS

- 1** **FINTECH**  
Revolut s'installe au Luxembourg
- 2** **POLITIQUE**  
Le Benelux approuve les transports gratuits
- 3** **BANQUES**  
BGL BNP Paribas, un siècle d'histoire bancaire

# Foreword



NEWS CLUB BUSINESS GUIDE JOBS

PAPERJAM  
BUSINESS ZU LETZEBUERG

ENTREPRISES & STRATÉGIES — TECHNOLOGIES

ÉCONOMIE DE LA DONNÉE

## Meluxina, le supercalculateur en 10 questions

Écrit par **Thierry Labre**  
Publié Le 14.06.2019 • Édité Le 14.06.2019

Partager



**TOP**

LUS | RECOMMANDÉS

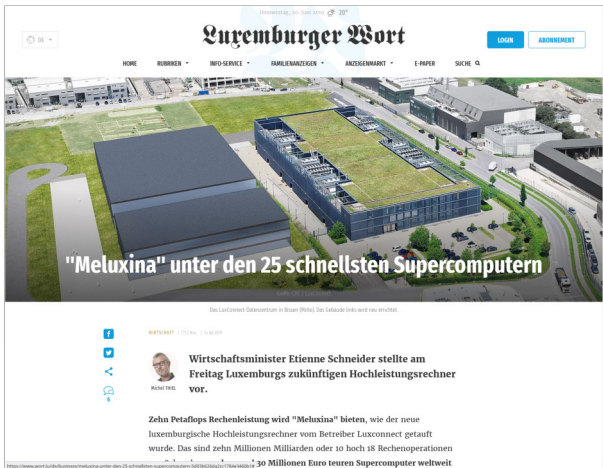
- 1** **FINTECH**  
Revolut s'installe au Luxembourg
- 2** **POLITIQUE**  
Le Benelux approuve les transports gratuits
- 3** **BANQUES**  
BGL, BNP Paribas, un siècle d'histoire bancaire

# Foreword





# Foreword




Donnerstag, 06. Juli 2018 20°  
**Luxemburger Wort** LOGIN ABONNEMENT  
 HOME RUBRIKEN INFO-SERVICE FAMILIENANZEIGEN ANZEIGENMARKT E-PAPER SUCHE

**"Meluxina" unter den 25 schnellsten Supercomputern**

Das LuConnect Datenzentrum in Bissen (Recht) Das Gebäude links wird neu errichtet.

WIRTSCHAFT / 17.06.18 / 17.06.2018


**Wirtschaftsminister Etienne Schneider stellte am Freitag Luxemburgs zukünftigen Hochleistungsrechner vor.**

RECHNEN TEIL

**Zehn Petaflops Rechenleistung wird "Meluxina" bieten, wie der neue luxemburgische Hochleistungsrechner vom Betreiber Luxconnect getauft wurde. Das sind zehn Millionen Milliarden oder 10 hoch 18 Rechenoperationen**

**30 Millionen Euro teuren Supercomputer weltweit**

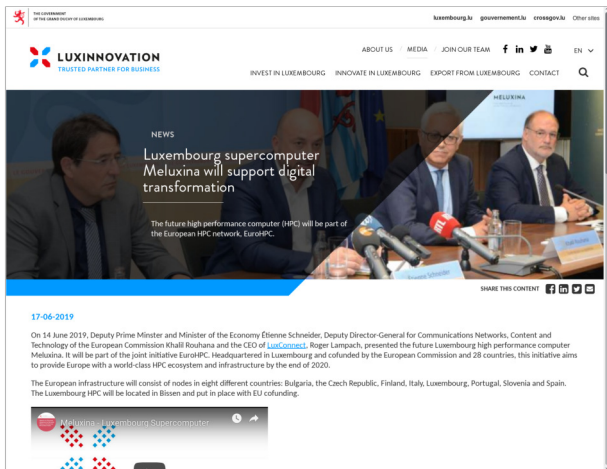
https://www.wort.lu/de/aktuelles/meluxina-unter-den-25-schnellsten-supercomputern-30-millionen-euro-teuren-supercomputer-weltweit

# Foreword



The screenshot shows a news article on the website 'Luxemburger Wort'. The main headline reads 'Meluxina, le nouveau superordinateur du Luxembourg'. Below the headline is a photograph of a server room with rows of server racks. The article text states: 'L'engin, d'une puissance de dix pétaflops, sera hébergé dans le datacenter de LuxConnect à Bissen. Dédié à la recherche, la médecine personnalisée et aux projets eHealth, il devrait à terme employer 50 personnes.' It also mentions that the supercomputer was unveiled by the Minister of Economy, Etienne Schneider (LSAP), on Tuesday.

## Foreword



THE GOVERNMENT OF THE GRAND-DUCHY OF LUXEMBOURG

luxembourg.lu gouvernement.lu crossgov.lu Other sites

LUXINNOVATION  
TRUSTED PARTNER FOR BUSINESS

ABOUT US / MEDIA / JOIN OUR TEAM f in t EN

INVEST IN LUXEMBOURG INNOVATE IN LUXEMBOURG EXPORT FROM LUXEMBOURG CONTACT

NEWS

### Luxembourg supercomputer Meluxina will support digital transformation

The future high performance computer (HPC) will be part of the European HPC network, EuroHPC.

SHARE THIS CONTENT f in t

17-06-2019

On 14 June 2019, Deputy Prime Minister and Minister of the Economy Étienne Schneider, Deputy Director-General for Communications Networks, Content and Technology of the European Commission Khalil Rouhana and the CEO of [LuxConnect](#), Roger Lampach, presented the future Luxembourg high performance computer Meluxina. It will be part of the joint initiative EuroHPC. Headquartered in Luxembourg and cofunded by the European Commission and 28 countries, this initiative aims to provide Europe with a world-class HPC ecosystem and infrastructure by the end of 2020.

The European infrastructure will consist of nodes in eight different countries: Bulgaria, the Czech Republic, Finland, Italy, Luxembourg, Portugal, Slovenia and Spain. The Luxembourg HPC will be located in Bissen and put in place with EU cofunding.

Meluxina - Luxembourg Supercomputer

# Foreword

[Anzeigen-Preise](#)
[Jobs](#)
[Immobilien](#)
[Verkehr](#)

[Box Finder](#)
[E-Paper](#)
[Newsletter](#)

---

fr: de Luxembourg 22° 

[Luxemburg](#)
[Ausland](#)
[Panorama](#)
[Wirtschaft](#)
[Sport](#)
[Community](#)
[People](#)
[Lifestyle](#)
[Digital](#)
[Buzz](#)
[Entertainment](#)
[Mehr](#)

 L'essentiel Radio

---

IN LUXEMBURG 14. Juni 2019 17:07, Akt. 14.06.2019 20:36 

## «Meluxina» soll Lücke zu China und USA schließen

*BISSEN – Der Supercomputer, der bis Ende 2020 installiert wird, soll einer der 20 schnellsten Rechner der Welt werden. Das teilte Wirtschaftsminister Schneider am Freitag mit.*



Mario Grotz, Étienne Schneider, Khalif Rouhana und Roger Lampech (v.l.) haben am Freitag den Supercomputer vorgestellt. (Bild: L'essentiel)

**TIERISCHE FRACHT**  
**Cargolux liefert die Beluga-Wale wohlbehalten ab**



LUXEMBURG – In einer Bucht bei Island entsteht ein Freiwasserreservat für Wale und Delfine - als Alternative zu Freizeitparks. Zwei Beluga-Wale wurden nun aus China eingeflogen.

**CHAMBER-PRÄSIDENT**  
**Etgen zieht nach dem erstem Halbjahr Bilanz**



LUXEMBURG – Vor einem halben Jahr hat Fernand Etgen das Amt des Parlamentspräsidenten übernommen. Zeit für eine Zwischenbilanz.

**RADARE IN LUXEMBURG**  
**Am Donnerstag heißt es wieder aufpassen!**



LUXEMBURG – Die Police Grand-Ducale zögert keine Gnade für Raser. Am Donnerstag hat sie sich dieserorts im Großherzogtum auf die Lauer gelegt.

IN LUXEMBURG



Thursday, 30 Jun 2019

## Chronicle.lu

HOME NEWS FEATURES EVENTS OPINION CLASSIFIEDS PROMOTIONS ABOUT LUXEMBOURG EMERGENCY

HOME / NEWS / ECONOMICS / LUXEMBOURG MELUXINA SUPERCOMPUTER TO JOIN EUROPEAN EUROHPC NETWORK

### Luxembourg Meluxina Supercomputer to Join European EuroHPC Network

Published on Friday, 14 Jun 2019 17:11 by RD

SHARE THIS ARTICLE: [f](#) [t](#) [g+](#) [s](#) [w](#) [e](#) [p](#)

RATE THIS ITEM: ★★★★★

On Friday 14 June 2019, Luxembourg's Deputy Prime Minister and Minister of the Economy, Etienne Schneider, together with the Deputy Director General of the Directorate General of Communication Networks, Content and Technologies (DG CNECT) at the European Commission, Khalil Raufhana, as well as LucConnect CEO, Roger Lampach, presented the future Luxembourg supercomputer, named 'Meluxina', which will join the European network of EuroHPC supercomputers.

The EuroHPC Joint Undertaking, headquartered in the Grand Duchy, is an initiative co-financed by the European Commission and 28 countries, including Luxembourg, which aims to provide Europe with an ecosystem and a computing infrastructure. By June 2019, following a call for projects, EuroHPC selected eight sites in different Member States to host supercomputers. The Luxembourg project to install the petascale supercomputer Meluxina at LucConnect in Bissen has

Meluxina Supercomputer (3-4), Mario Grotz, Ministry of Economy, Minister Etienne Schneider, Khalil Raufhana, European Commission, Roger Lampach, CEO LucConnect, Credit: MESCO

Commission and 28 countries, including Luxembourg, which aims to provide Europe with an ecosystem and a computing infrastructure. By June 2019, following a call for projects, EuroHPC selected eight sites in different Member States to host supercomputers. The Luxembourg project to install the petascale supercomputer Meluxina at LucConnect in Bissen has

<https://subject-chronicle.lu/body-https://chronicle.lu/category/economies/2019/Luxembourg-meluxina-supercomputer-to-join-european-eurohpc-network-high-performance>

**TRENDING NEWS**

**Jean Asselborn Advocates Defence of Human Rights, Support for Jordan at EU Foreign Affairs Council**  
Abroad 16 Jun, 2019 09:59

**Minister Gramigna Reveals 60 Companies Established in Grand Duchy re Brexit**  
17 Jun, 2019 07:27

**Nordic Chamber (Nobelus & Sverbehus)**  
17 Jun, 2019 07:27

**Nicolas Grass Formally Appointed Honorary Consul General of India in Luxembourg**  
Embassies & Consulates 17 Jun, 2019 20:05

**Luxembourg-UK Agreement Guarantees Citizens' Voting Rights Post-Brexit**  
Politics 16 Jun, 2019 15:04

**Grosbusch Introduces 100% Recyclable, Biodegradable Packaging**  
Environment 17 Jun, 2019 10:30

Luxembourg, LU

# Foreword

D
NEWS
AGENDA
MAGAZINE
EXPAT GUIDE
JOBS
REAL ESTATE
INSIGHTS
Q

## DELANO

LUXEMBOURG IN ENGLISH

### MEET MELUXINA, LUXEMBOURG'S NEW SUPERCOMPUTER

NEWS • BUSINESS • 17.06.2019 • DELANO STAFF

[f](#)
[t](#)
[in](#)
[g+](#)
[v](#)

**Luxembourg's petascale supercomputer, planned to open in Bissen before the end of 2020, was presented to the public on Friday.**

Industry and startups will be able to access Meluxina, as the computer has been dubbed, at the offices of LuxConnect in Bissen, where project will eventually generate up to 50 jobs. The Luxembourg supercomputer is one of eight to join the European network of supercomputers.

According to a government statement published on Friday, it will have a computing power of 10 petaflops/second, which corresponds to 10,000,000,000,000,000 computing operations per second.

Meluxina will be dedicated to applications in research, personalised medicine and eHealth projects, as well as the needs of companies, in



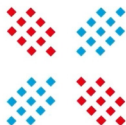
Meluxina will join the European EuroHPC network of supercomputers. Photo: Shutterstock

LATEST NEWS





# Foreword



## MELUXINA

HIGH PERFORMANCE  
COMPUTING IN LUXEMBOURG



# MeluXina National Supercomputer

## MeluXina - coming in 2020

- 10 PetaFlop supercomputer
- Modular architecture covering a wide variety of needs
- High performance network & storage for HPC, BigData & AI





# MeluXina National Supercomputer

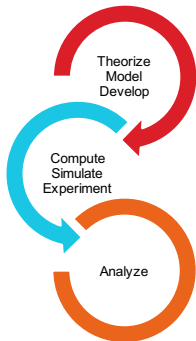
## MeluXina - coming in 2020

- 10 PetaFlop supercomputer
- Modular architecture covering a wide variety of needs
- High performance network & storage for HPC, BigData & AI

## What this means for you

- Algorithms and applications must be run **at scale**
- **Code development** will play a large role
- Need to use different computing elements and memory hierarchy
  - ↔ will play a critical role in your **application performance**

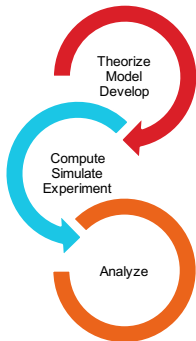
# Main Objectives of this Session



This session is meant to show you some of the various tools you have at your disposal on the UL HPC platform to:

**understand and solve development & runtime problems**  
**understand and improve application performance**

# Main Objectives of this Session



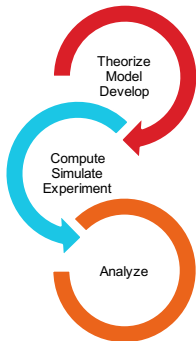
This session is meant to show you some of the various tools you have at your disposal on the UL HPC platform to:

**understand and solve development & runtime problems**  
**understand and improve application performance**

During the session we will (also):

- discuss what happens when an application runs **out of memory** and how to discover how much memory it actually requires.
- see **debugging tools** that help you understand **why your code is crashing**.
- see **profiling tools** that show the **bottlenecks of your code** - and **how to improve it**.

## Main Objectives of this Session



This session is meant to show you some of the various tools you have at your disposal on the UL HPC platform to:

- understand and solve development & runtime problems**
- understand and improve application performance**

During the session we will (also):

- discuss what happens when an application runs **out of memory** and how to discover how much memory it actually requires.
- see **debugging tools** that help you understand **why your code is crashing**.
- see **profiling tools** that show the **bottlenecks of your code** - and **how to improve it**.

**Knowing what to do when you experience a problem is half the battle.**



# Summary

- 1 Introduction
- 2 Debugging and profiling tools**
- 3 Conclusion

## Tools at your disposal (I)

### Common tools used to understand problems

- Common tools at a glance:
  - ↳ **htop**: process/thread live CPU & memory utilization, affinity
  - ↳ **strace**: understand what system calls a process is making & where it gets 'stuck'
  - ↳ **gdb**: general purpose debugger
  - ↳ **valgrind**: memory debugging and profiling
- SLURM scheduler built-in tools and utilities:
  - ↳ **sstat**: running job and jobstep statistics
  - ↳ **sacct**: historic job and jobstep statistics
  - ↳ **seff**: postmortem CPU/memory utilization and efficiency

Some times simple tools help you solve big issues.

## Tools at your disposal (II)

### HPC specific tools - Arm (prev. Allinea)

- Arm DDT (part of Arm Forge)
  - ↳ Visual debugger for C, C++, Fortran & **Python** // code
- Arm MAP (part of Arm Forge)
  - ↳ Visual profiler for C, C++, Fortran & **Python**
- Arm Performance Reports
  - ↳ Application characterization tool

## Tools at your disposal (II)

### HPC specific tools - Arm (prev. Allinea)

- Arm DDT (part of Arm Forge)
  - ↳ Visual debugger for C, C++, Fortran & **Python** // code
- Arm MAP (part of Arm Forge)
  - ↳ Visual profiler for C, C++, Fortran & **Python**
- Arm Performance Reports
  - ↳ Application characterization tool

### Arm tools are licensed

- license check integrated in SLURM: `scontrol show license`
- ask for licenses at job submission with e.g. `srun -L forge:16`

2019 software set being prepared, check new tools with  
`module load swenv/default-env/devel && module avail`



## Tools at your disposal (III)

### HPC specific tools - Intel

- Intel Advisor
  - ↔ Vectorization + threading advisor: check blockers and opport.
- Intel Inspector
  - ↔ Memory and thread debugger: check leaks/corrupt., data races
- Intel Trace Analyzer and Collector
  - ↔ MPI communications profiler and analyzer: evaluate patterns
- Intel VTune Amplifier
  - ↔ Performance profiler: CPU/FPU data, mem. + storage accesses

## Tools at your disposal (III)

### HPC specific tools - Intel

- Intel Advisor
  - ↪ Vectorization + threading advisor: check blockers and opport.
- Intel Inspector
  - ↪ Memory and thread debugger: check leaks/corrupt., data races
- Intel Trace Analyzer and Collector
  - ↪ MPI communications profiler and analyzer: evaluate patterns
- Intel VTune Amplifier
  - ↪ Performance profiler: CPU/FPU data, mem. + storage accesses

### Intel tools are licensed

All come as part of Intel Parallel Studio XE - Cluster edition!

## Tools at your disposal (IV)

### HPC specific tools - Scalasca & friends

- Scalasca
  - ↪ Study behavior of // apps. & identify optimization oport.
- Score-P
  - ↪ Instrumentation tool for profiling, event tracing, online analysis.
- Extra-P
  - ↪ Automatic performance modeling tool for // apps.

## Tools at your disposal (IV)

### HPC specific tools - Scalasca & friends

- Scalasca
  - ↪ Study behavior of // apps. & identify optimization oport.
- Score-P
  - ↪ Instrumentation tool for profiling, event tracing, online analysis.
- Extra-P
  - ↪ Automatic performance modeling tool for // apps.

Free and Open Source!

See other awesome tools at <http://www.vi-hps.org/tools>

## Arm DDT - highlights

### DDT features

- **Parallel debugger:** threads, OpenMP, MPI support
- Controls processes and threads
  - ↳ step code, stop on var. changes, errors, breakpoints
- Deep **memory debugging**
  - ↳ find memory leaks, dangling pointers, beyond-bounds access
- C++ debugging – including STL
- Fortran – including F90/F95/F2008 features
- Python – scripts ran under CPython interpreter, mpi4py
- See vars/arrays **across multiple processes**
- Integrated editing, building and **VCS integration**
- Offline mode for **non-interactive debugging**
  - ↳ record application behavior and state

## Arm DDT - on ULHPC

### Modules

- On Iris: module load tools/ArmForge
- Caution! May behave differently between
  - ↳ Debian+OAR (Gaia, Chaos) and CentOS+SLURM (Iris)

### Debugging with DDT

- 1 Load toolchain, e.g. (for Intel C/C++/Fortran, MPI, MKL):
  - ↳ module load toolchain/intel
- 2 Compile your code, e.g. `mpicc $code.c -o $app`
- 3 Run your code through DDT (GUI version)
  - ↳ iris: `ddt srun ./$app`
  - ↳ gaia/chaos: `ddt mpirun -hostfile $OAR_NODEFILE ./$app`
- 4 Run DDT in batch mode (no GUI, just report):
  - ↳ `ddt --offline -o report.html --mem-debug=thorough ./$app`



## Arm DDT - interface

The screenshot shows the Alinea DDT - Alinea Forge 7.0.3 interface. The main window displays the source code for `dstaticmp.c`. A context menu is open over the code, showing options like "Add breakpoint for All", "Delete breakpoint for All", and "Run to here".

```
82 MPI_Status status; /* Return status for receive */
83
84 t2 = malloc(sizeof(typeThree));
85
86 for(p=0;p<100;p++)
87   bigArray[p]=80000*p;
88
89 MPI_Init(&argc, &argv);
90 MPI_Comm_rank(MPI_COMM_WORLD, &my_rank);
91 MPI_Comm_size(MPI_COMM_WORLD, &pi);
92
93 dynamicArray = malloc(sizeof(int)*100000);
94 sdim = malloc(sizeof(int) * pi);
95
96 for(x=0;x<10000;x++)
97 {
98   dynamicArray[x] = x%10;
99 }
100
101 printf("my rank is %d\n", my_rank);
102
103 for(x=0;x<12;x++)
104 {
105   y = 0;
106   while(y = 12)
107   {
108     tables[1][y] = (x+1)*(y+1);
109     y += my_rank + 1;
110   }
111 }
112
113 if(argc > 1 && my_rank == 0)
114   printf("Rank %d has %d ar
115   printf("They are:");
116   for(x=0; x<10000; x++)
117     printf("%d ", dynamicArray[x]);
118 }
```

The right sidebar shows the "Locals" window with the following variables:

Variable Name	Value
-argc	1
-argv	0x7ffffff2
-beingWatched	0
-bigArray	
-dest	-134225592
-dynamicArray	0x7ffffe2010
-environ	0x7ffffff28
-i	32767
-message	
-my_rank	5
-p	28
-s	0x0
-sdim	0x235f0
-source	1
-status	
-t2	0x603010
-tables	
-tag	50
-thead	
-troopa	0x0
-x	0
-y	0

The bottom left shows the "Watchpoints" table:

Processes	Scope	Expression	Trigger On	Implemented in
All	#0 main	dynamicArray	writes only	hardware

## Arm MAP - highlights

### MAP features

- Meant to show developers **where&why code is losing perf.**
- **Parallel profiler**, especially made for MPI applications
- Recent Python support, incl. mpi4py, OpenMP, threading
- Effortless profiling
  - ↳ no code modifications needed, may not even need to recompile
- Clear **view of bottlenecks**
  - ↳ in I/O, compute, thread or multi-process activity
- Deep insight in **CPU instructions affecting perf.**
  - ↳ vectorization and memory bandwidth
- **Memory usage over time** – see changes in memory footprint
- Integrated editing and building as for DDT

Full details at Arm HPC Tools: **Forge-MAP**



## Arm MAP - on ULHPC

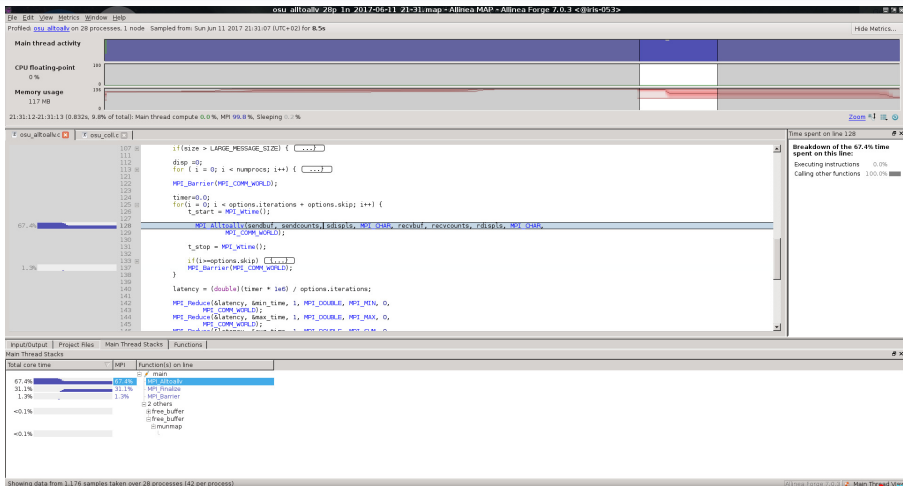
### Modules

- On Iris: `module load tools/ArmForge`
- Caution! May behave differently between:
  - ↳ Debian+OAR (Gaia, Chaos) and CentOS+SLURM (Iris)

### Profiling with MAP

- 1 Load toolchain that built your app., e.g.
  - ↳ `module load toolchain/intel`
- 2 Run your code through MAP (attached, GUI version)
  - ↳ `iris: map srun ./ $app`
  - ↳ `gaia/chaos: map mpirun -hostfile $OAR_NODEFILE ./ $app`
- 3 Run MAP in batch mode (no GUI, create .map file):
  - ↳ `iris: map --profile srun ./ $app`

## Arm MAP - interface



## Arm Perf. Reports - highlights

### Performance Reports features

- Meant to answer **How well do your apps. exploit your hw.?**
- Easy to use, on unmodified applications
  - ↳ outputs HTML, text, CSV, JSON reports
- One-glance view if application is:
  - ↳ **well-optimized** for the underlying hardware
  - ↳ running **optimally at** the given **scale**
  - ↳ **affected by** I/O, networking or threading **bottlenecks**
- Easy to integrate with continuous testing
  - ↳ programmatically improve performance by continuous profiling
- **Energy metric** integrated
  - ↳ using RAPL (CPU) for now on iris
  - ↳ IPMI-based monitoring may be added later

# Arm Perf. Reports - on ULHPC

## Modules

- On Iris: `module load tools/ArmReports`
- Caution! May behave differently between:
  - ↪ Debian+OAR (Gaia, Chaos) and CentOS+SLURM (Iris)
  - ↪ Gaia: can collect GPU metrics
  - ↪ Iris: can collect GPU & energy metrics

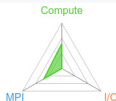
## Using Performance Reports

- 1 Load toolchain that you run your app. with, e.g.
  - ↪ `module load toolchain/intel`
- 2 Run your application through Perf. Reports
  - ↪ `iris: perf-report srun ./$app`
  - ↪ `gaia/chaos: perf-report mpirun -hostfile $OAR_NODEFILE ./$app`
- 3 Analysis by default in `.html` and `.txt` indicating also run config.

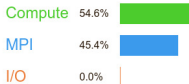
## Arm Perf. Reports - output (I)



Command: `srun gmx_mpi mdrun -s bench_mase_cubic.tpr -nsteps 10000`  
 Resources: 1 node (28 physical, 28 logical cores per node)  
 Memory: 126 GiB per node  
 Tasks: 28 processes, OMP\_NUM\_THREADS was 0  
 Machine: iris-053  
 Start time: Sun Jun 11 2017 20:13:59 (UTC+02)  
 Total time: 19 seconds  
 Full path: `/mnt/irisgpfps/apps/resif/data/production/v0.1-20170602/default/software/bio/GROMACS/2016.3-intel-2017a-hybrid/bin`



Summary: `gmx_mpi` is **Compute-bound** in this configuration



Time spent running application code. High values are usually good. This is **average**; check the CPU performance section for advice

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 **negligible**; there's no need to investigate I/O performance

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

### CPU

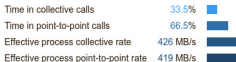
A breakdown of the **54.6%** CPU time:



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

### MPI

A breakdown of the **45.4%** MPI time:



Most of the time is spent in **point-to-point calls** with an average transfer rate. Using larger messages and overlapping communication and computation may increase the effective transfer rate.

## Arm Perf. Reports - output (II)

### CPU

A breakdown of the **54.6%** CPU time:

Single-core code	5.5%	
OpenMP regions	94.5%	████████████████████
Scalar numeric ops	5.2%	
Vector numeric ops	44.2%	████████
Memory accesses	50.6%	████████

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

### I/O

A breakdown of the **0.0%** I/O time:

Time in reads	0.0%	
Time in writes	0.0%	
Effective process read rate	0.00 bytes/s	
Effective process write rate	0.00 bytes/s	

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

### Memory

Per-process memory usage may also affect scaling:

Mean process memory usage	75.6 MiB	████████
Peak process memory usage	86.6 MiB	████████████████
Peak node memory usage	11.0%	

The peak node memory usage is very low. Running with fewer MPI processes and more data on each process may be more efficient.

### MPI

A breakdown of the **45.4%** MPI time:

Time in collective calls	33.5%	█
Time in point-to-point calls	66.5%	█
Effective process collective rate	426 MB/s	█
Effective process point-to-point rate	419 MB/s	█

Most of the time is spent in **point-to-point calls** with an average transfer rate. Using larger messages and overlapping communication and computation may increase the effective transfer rate.

### OpenMP

A breakdown of the **94.5%** time in OpenMP regions:

Computation	99.5%	████████████████████
Synchronization	0.5%	
Physical core utilization	100.0%	████████
System load	101.9%	████████

OpenMP thread performance looks good. Check the CPU breakdown for advice on improving code efficiency.

### Energy

A breakdown of how the **0.899 Wh** was used:

CPU	100.0%	████████
System	not supported %	
Mean node power	not supported W	
Peak node power	not supported W	

The **whole system energy** has been calculated using the CPU energy usage.

System power metrics: No Allinea IPMI Energy Agent config file found in (null). Did you start the Allinea IPMI Energy Agent?

## Intel Advisor - highlights

### Advisor features

- Vectorization Optimization and Thread Prototyping
- Analyze vectorization opportunities
  - ↳ for code compiled either with Intel and GNU compilers
  - ↳ SIMD, AVX\* (incl. AVX-512) instructions
- Multiple data collection possibilities
  - ↳ loop iteration statistics
  - ↳ data dependencies
  - ↳ memory access patterns
- Suitability report - predict max. speed-up
  - ↳ based on app. modeling

Full details at [software.intel.com/en-us/intel-advisor-xe](https://software.intel.com/en-us/intel-advisor-xe)

## Intel Advisor - on ULHPC

### Modules

- On iris/gaia/chaos: `module load perf/Advisor`

### Using Intel Advisor

- 1 Load toolchain: `module load toolchain/intel`
- 2 Compile your code, e.g. `mpicc $code.c -o $app`
- 3 Collect data e.g. on gaia:

```
mpirun -n 1 -gtool "advixe-cl -collect survey \  
-project-dir ./advisortest:0" ./$app
```

- 4 Visualise results with `advixe-gui $HOME/advisortest`





## Intel Advisor - interface

The screenshot shows the Intel Advisor interface with the following components:

- Left Panel:** Navigation menu with sections: Vectorization Workflow, Threading Workflow, Run Rooftline, 1. Survey Target, 1.1 Find Trip Counts and FLOPS, Mark Loops for Deeper Analysis, 2.1 Check Dependencies, and 2.2 Check Memory Access Patterns.
- Top Bar:** Application title and window controls.
- Main Content Area:**
  - Summary:** Elapsed time: 247.46s, Vectorized: Not Vectorized, FILTER: All Modules, All Sources.
  - Vectorization Advisor:** Description of the toolset.
  - Program metrics:** Elapsed Time: 247.46s, Vector Instruction Set: SSE, SSE2, Number of CPU Threads: 1.
  - Loop metrics:** Bar chart showing Total CPU time (246.84s, 100.0%), Time in 8 vectorized loops (176.79s, 71.6%), and Time in scalar code (70.04s, 28.4%).
  - Vectorization Gain/Efficiency:** Vectorized Loops Gain/Efficiency: 1.81x (90% bar), Program Approximate Gain: 1.58x.
  - Top time-consuming loops:** Table with columns: Loop, Self Time, Total Time.
  - Collection details:** Section for further analysis.
  - Platform information:** MPI rank: 0, CPU Name: Intel(R) Xeon(R) CPU X5670 @ 2.93GHz, Frequency: 2.93 GHz, Logical CPU Count: 12, Operating System: Linux, Computer Name: gaia-100.gaia-cluster.uni.lu.

Loop	Self Time	Total Time
loop in ComputeSPMV_ref at ComputeSPMV_ref.cpp:67	70.624s	70.624s
loop in ComputeSYMGs_ref at ComputeSYMGs_ref.cpp:93	56.756s	56.756s
loop in ComputeSYMGs_ref at ComputeSYMGs_ref.cpp:74	44.903s	44.903s
loop in ComputeSYMGs_ref at ComputeSYMGs_ref.cpp:67	28.317s	73.220s
loop in ComputeSYMGs_ref at computeSYMGs_ref.cpp:86	20.331s	77.087s

# Scalasca & friends - highlights

## Scalasca features

- Scalable performance analysis toolset
  - ↳ for large scale // applications on 100.000s of cores
- Support for C/C++/Fortran code with MPI, OpenMP, hybrid
- 3 stage workflow: instrument, measure, analyze
  - ↳ at compile time, run time and resp. postmortem
- Score-P for instrumentation + measurement, Cube for vis.
  - ↳ Score-P can also be used with Periscope, Vampir and Tau
- Facilities for measurement optimization to min. overhead
  - ↳ by selective recording, runtime filtering

Full details at <http://www.scalasca.org/about/about.html>

## Scalasca - on ULHPC

### Modules

- On iris/gaia/chaos:

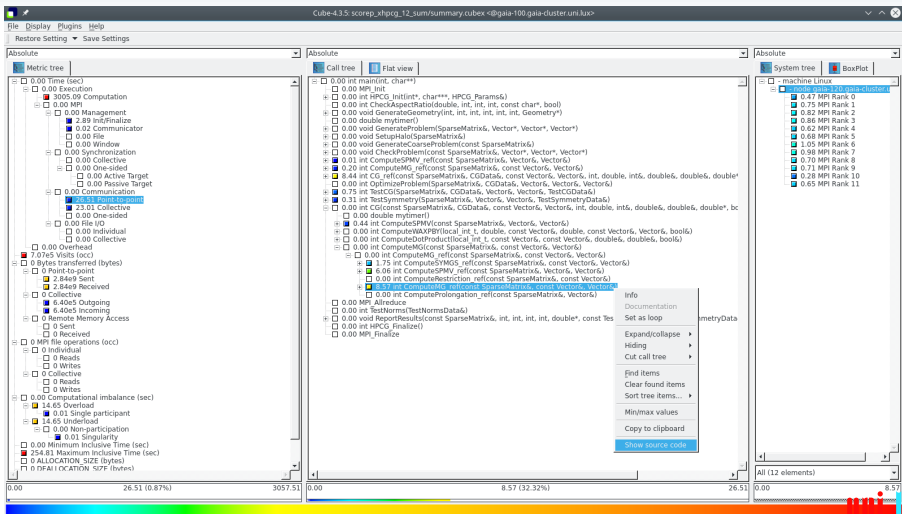
```
module load perf/Scalasca perf/Score-P
```

### Using Scalasca

- 1 Load toolchain: `module load toolchain/foss`
- 2 Compile your code, e.g. `scorep mpicc $code.c -o $app`
- 3 Collect data e.g. on gaia: `scan -s mpirun -n 12 ./$app`
- 4 Visualise results with `square scorep_$app_12_sum`
  - ↪ or generate text report: `square -s scorep_$app_12_sum`
  - ↪ ... and print it: `cat scorep_$app_12_sum/scorep.score`



## Scalasca visualisation with Cube-P





# Summary

- 1 Introduction
- 2 Debugging and profiling tools
- 3 Conclusion**



# Now it's up to you

Easy right?



## Now it's up to you

Easy right?

Well not exactly.



## Now it's up to you

Easy right?

**Well not exactly.  
Debugging always takes effort and real applications are  
never trivial.**





## Now it's up to you

**Easy right?**

**Well not exactly.  
Debugging always takes effort and real applications are  
never trivial.**

**But we do guarantee it'll be /easier/ with these tools.**



# Conclusion and Practical Session start

## We've discussed

- A couple of small utilities that can be of big help
- HPC oriented tools available for you on UL HPC

## And now..

**Short DEMO time!**



# Conclusion and Practical Session start

## We've discussed

- A couple of small utilities that can be of big help
- HPC oriented tools available for you on UL HPC

## And now..

**Short DEMO time!**

Your Turn!



## Hands-on start

- We will first start with running HPCG (unmodified) as per:

[ulhpc-tutorials.rtfld.io/en/latest/advanced/HPCG/](http://ulhpc-tutorials.rtfld.io/en/latest/advanced/HPCG/)

- ... your tasks:

- 1 perform a timed first run using unmodified HPCG v3.0 (MPI only)
  - ✓ use `sacct -j $JOBID -l` to get details
  - ✓ single node, use `≥ 80 80 80` for input params (`hpcg.dat`)
- 2 run HPCG (timed) through Allinea Perf. Report
  - ✓ use `perf-report` for diff. // configurations: 1 vs 2 nodes all cores
- 3 instrument and measure HPCG execution with Scalasca
  - ✓ check MPI comm. measurements, especially collective operations

- Reservations for the workshop on Iris:

→ regular nodes (batch partition): `--reservation=hpcschool`

→ accel. nodes (gpu partition): `--reservation=hpcschool-gpu`

# Questions?

<http://hpc.uni.lu>

## High Performance Computing @ uni.lu

Prof. Pascal Bouvry  
Dr. Sebastien Varrette  
Valentin Plugaru  
Sarah Peter  
Hyacinthe Cartiaux  
Clement Parisot  
Dr. Frédéric Pinel  
Dr. Emmanuel Kieffer

University of Luxembourg, Belval Campus  
Maison du Nombre, 4th floor  
2, avenue de l'Université  
L-4365 Esch-sur-Alzette  
*mail: hpc@uni.lu*



1 Introduction

2 Debugging and profiling tools  
3 Conclusion