

45th VI-HPS Tuning Workshop

10 – 13 June 2024 | LRZ

Lecturers



Lecturers

- David Böhme (LLNL)
- Hugo Bolloré (UVSQ)
- Alexander Geiß (TUD)
- Germán Llort (BSC)
- Michele Martone (LRZ)
- Lau Mercadal (BSC)
- Emmanuel Oseret (UVSQ)
- Amir Raofy (LRZ)
- Jan André Reuter (JSC)
- Rudy Shand (Linaro)
- Sameer Shende (University of Oregon)
- Cédric Valensi (UVSQ)
- Josef Weidendorfer (LRZ)
- Brian Wylie (JSC)
- Ilya Zhukhov (JSC)

Organisers

- Cédric Valensi (UVSQ)
- Josef Weidendorfer (LRZ)
- Volker Weinberg (LRZ)
- Helga Tyroller (LRZ)




General Agenda (in CEST)

- 09:00 – 10:30 Lecture
- 10:30 – 11:00 Break
- 11:00 – 12:30 Lecture

- 12:30 – 14:00 Lunch Break

- 14:00 – 15:30 Lecture / Hands-On
- 15:30 – 16:00 Break
- 16:00 – 17:30 Lecture / Hands-On



Guided Tour

- **Date:** Tuesday, 11.06.2024, 17:30 – 18:30 CEST
- **Attention:** You will need a photo ID to enter the compute cubes.



Please register on the participants list on the first day!

Social Event

- **Event:** Self-Paid Dinner
- **Date:** Wednesday, 12.06.2024, 18:30 CEST – Open End

- **Location:**

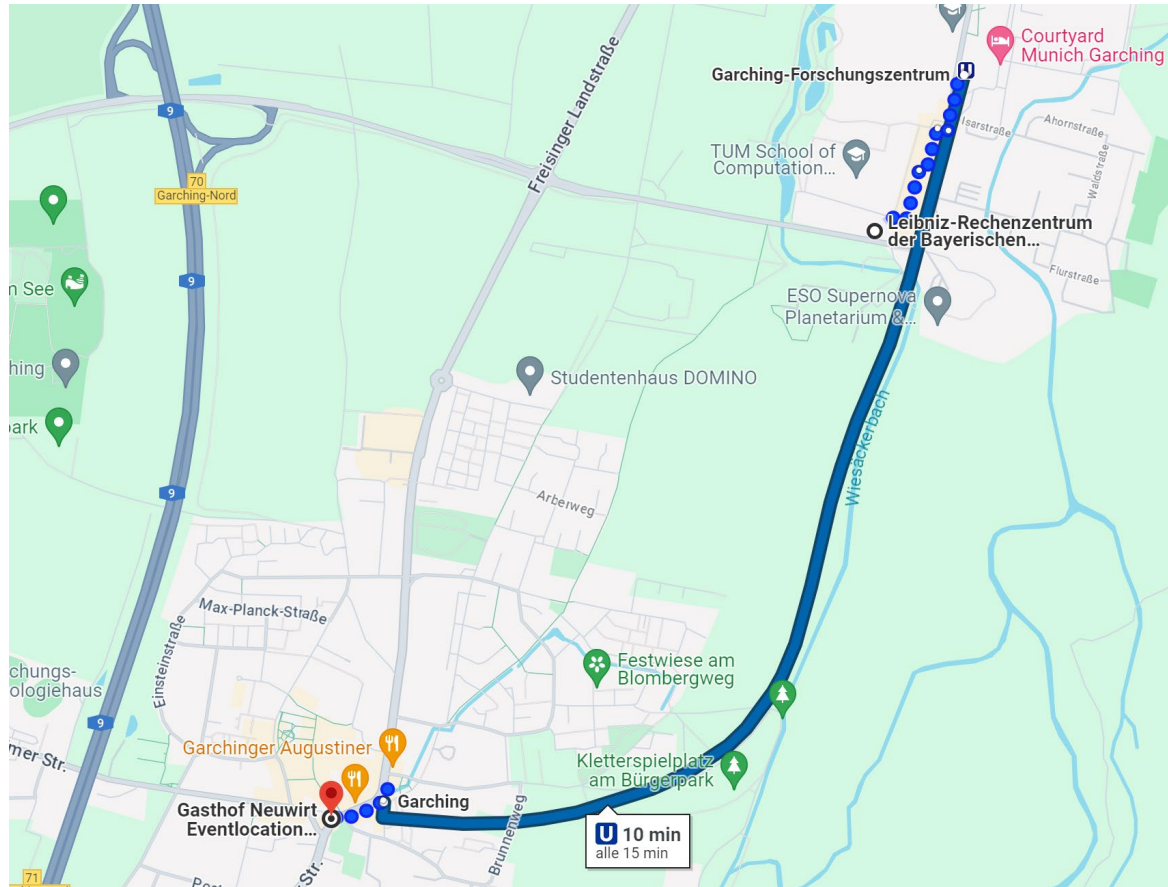
Gasthof Neuwirt
Münchener Str.10
85748 Garching

<https://gasthof-neuwirt.org/>



Please register on the participants list on the first day!

Social Event



Get there by subway / U-Bahn (10 minutes):

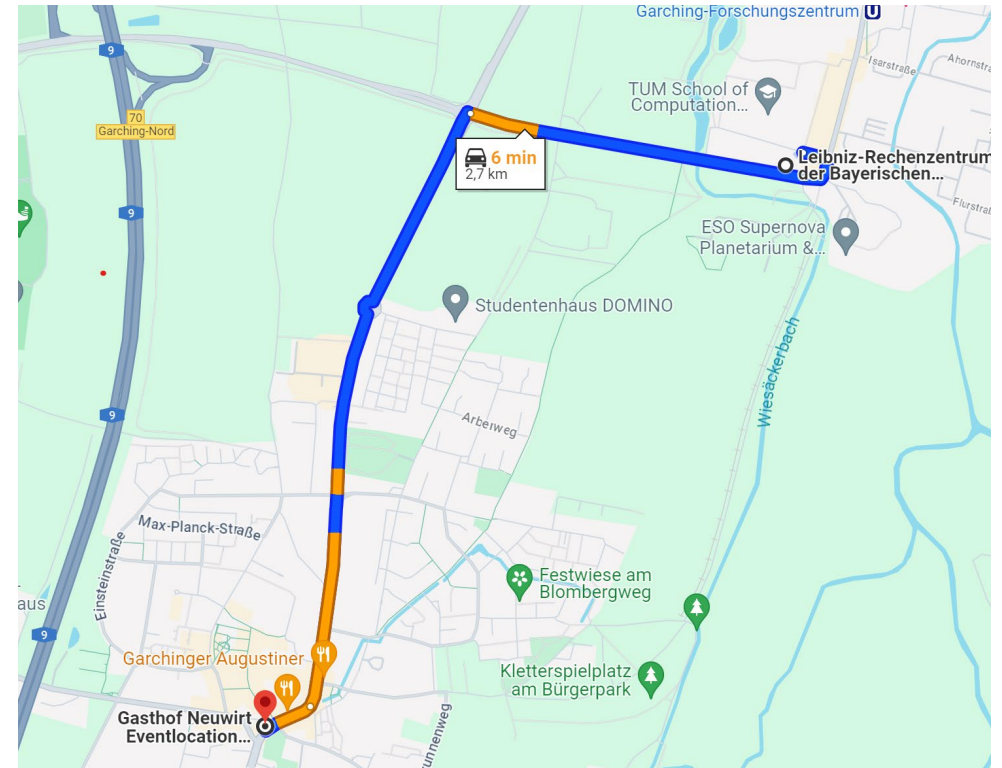
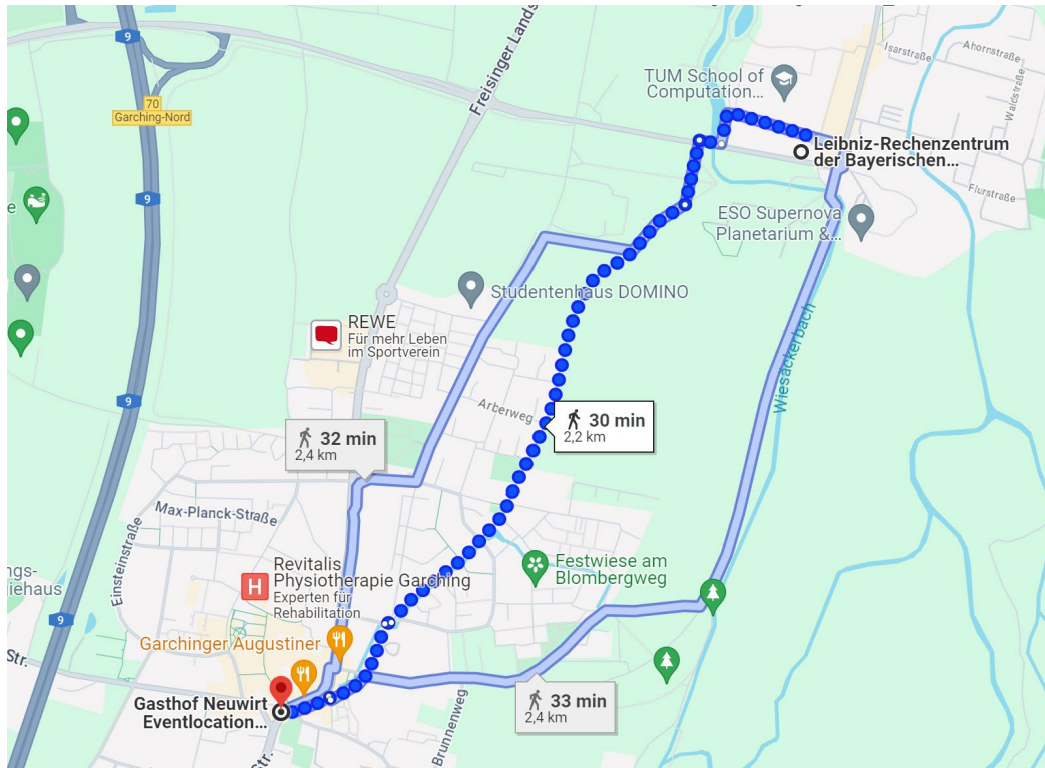
Trains (U6) depart every 15 minutes at Garching Forschungszentrum.

Leave the train after 1 station at Garching.

Social Event

Get there by foot (30 minutes):

Get there by car (6 minutes):



- For coffee and lunch breaks we recommend the bistro, snack/coffee counter and coffee machines of Gastronomie Wilhelm in the TUM School of Computation, Information and Technology (CIT) building just next to the LRZ:



Workshop Webpage

- All slides and materials will be made available under:
- <https://tinyurl.com/vihps-2024>
- Further information on:
 - Agenda
 - Guided Tour
 - Social Event
 - Slides
 - Hands-On Sessions
 - Survey



Using the LRZ Linux Cluster CoolMUC-2

- The LRZ Linux Cluster consists of several segments with different types of interconnect and different sizes of shared memory. All systems have a (virtual) 64 bit address space:
- ➔ • **CoolMUC-2** Cluster with 28-way Haswell-based nodes and FDR14 Infiniband interconnect, used for both serial and parallel processing
- **CoolMUC-3** Cluster with 64-way KNL 7210-F many-core processors and Intel Omnipath OPA1 interconnect, for parallel/vector processing
- **CoolMUC-4** HPDA Cluster with 80-way Ice Lake nodes
- **Large Memory System Teramem** (Intel Broadwell based 6 TByte shared memory server HP DL580)
- **LRZ AI Systems** (DGX A100 , DGX-1 V100, DGX-1 P100 , HPE Intel Skylake + Tesla P100, Tesla V100)

- Based on the various node types the LRZ Linux cluster offers a wide span of capabilities:
 - mixed **shared and distributed memory**
 - **large software portfolio**
 - flexible usage due to **various available memory sizes**
 - parallelization by **message passing (MPI)**
 - shared memory parallelization with **OpenMP or pthreads**
 - **mixed (hybrid) programming** with MPI and OpenMP
 - **secure shell** based logins and data transfer to generally accessible front end nodes
 - **development environment** with compilers, tools and libraries available on front end nodes, run time environments and applications available on batch nodes.
 - resource assignment via **SLURM scheduler**
 - **data management:**
 - SCRATCH space for short lifetime data (removal is forced)
 - DSS/HOME area with small quota for program and configuration data
 - DSS/PROJECT area (max. 10 TByte) upon request for long lifetime data

First Self-Assembled Linux Cluster (1999-2002)



Cluster Components (2012)



CoolMUC-2 (2015): The six racks to the left



CoolMUC-3 (2017)



Hardware	
Number of nodes	812
Cores per node	28
Hyperthreads per core	2
Core nominal frequency	2.6 GHz
Memory (DDR4) per node	64 GB (Bandwidth 120 GB/s - STREAM)
Bandwidth to interconnect per node	13,64 GB/s (1 Link)
Bisection bandwidth of interconnect (per island)	3.5 TB/s
Latency of interconnect	2.3 μ s
Peak performance of system	1400 TFlop/s
Infrastructure	
Electric power of fully loaded system	290 kVA
Percentage of waste heat to warm water	97%
Inlet temperature range for water cooling	30 ... 50 °C
Temperature difference between outlet and inlet	4 ... 6 °C
Software (OS and development environment)	
Operating system	SLES15 SP1 Linux
MPI	Intel MPI 2019, alternatively OpenMPI
Compilers	Intel icc, icpc, ifort 2019
Performance libraries	MKL, TBB, IPP
Tools for performance and correctness analysis	Intel Cluster Tools

```
ssh -Y lxlogin1.lrz.de -l xxyyzzz  
ssh -Y lxlogin2.lrz.de -l xxyyzzz  
ssh -Y lxlogin3.lrz.de -l xxyyzzz  
ssh -Y lxlogin4.lrz.de -l xxyyzzz
```

Haswell (CoolMUC-2) login node
Haswell (CoolMUC-2) login node
Haswell (CoolMUC-2) login node
Haswell (CoolMUC-2) login node

- The **Intel OneAPI software stack** is automatically loaded at login.
- **Intel Fortran and C/C++ Compilers** comprise both traditional and new, LLVM-based compiler drivers:
 - traditional: `ifort`, `icc`, `icpc`
 - LLVM-based ("next-gen"): `ifx`, `icx`, `icpx`, `dpcpp`
- See <https://doku.lrz.de/intel-compilers-11481687.html>
<https://doku.lrz.de/intel-oneapi-11481686.html>
- By default, **OpenMP directives** in your code are ignored. Use the `-qopenmp` option to activate OpenMP.
- Use `mpirexec -n #tasks` to run **MPI programs**. The compiler wrappers' names follow the usual `mpicc`, `mpif90`, `mpiCC` pattern.

How to use the CoolMUC-2 System

- **Submit a job:**
`sbatch --reservation=hhps1s24 job.sh`
- **List own jobs:**
`squeue -M cm2_tiny -u hpckurs??`
- **Cancel jobs:**
`scancel -M cm2_tiny jobid`
- **Interactive Access:**
 - `module load salloc_conf/cm2_tiny`
 - `salloc --nodes=1 --time=02:00:00 --reservation=hhps1s24 --partition=cm2_tiny`
 - **or:** `srun --reservation= hhps1s24 --pty bash`

CoolMUC-2 SLURM OpenMP Batch File



→ /lrz/sys/courses/vihps/job-omp.sh

```
#!/bin/bash
#SBATCH -o /dss/dsshome1/08/hpckurs00/coolmuc.%j.%N.out
#SBATCH -D /dss/dsshome1/08/hpckurs00/
#SBATCH -J coolmuctest
#SBATCH --clusters=cm2_tiny
#SBATCH --nodes=1
#SBATCH --cpus-per-task=28
#SBATCH --get-user-env
#SBATCH --reservation=hhps1s24
#SBATCH --time=02:00:00

module load slurm_setup
export OMP_NUM_THREADS=$SLURM_CPUS_PER_TASK
./myprog.exe
```

Please do not use
#SBATCH --mail-type=...
as this could be considered as a denial
of service attack on the LRZ mail hub!

CoolMUC-2 SLURM MPI Batch File



→ /lrz/sys/courses/vihps/job-mpi.sh

```
#!/bin/bash
#SBATCH -o /dss/dsshhome1/08/hpckurs00/coolmuc.%j.%N.out
#SBATCH -D /dss/dsshhome1/08/hpckurs00
#SBATCH -J coolmuctest
#SBATCH --clusters=cm2_tiny
#SBATCH --nodes=2
#SBATCH --ntasks-per-node=28
#SBATCH --get-user-env
#SBATCH --reservation=hhps1s24
#SBATCH --time=02:00:00

module load slurm_setup
mpiexec -n $SLURM_NTASKS ./myprog.exe
```

Please do not use
#SBATCH --mail-type=...
as this could be considered as a denial
of service attack on the LRZ mail hub!

And now enjoy the workshop!

Survey

- Please fill out the online survey under



<https://tinyurl.com/hhps1s24-survey>




- This helps us to
 - increase the quality of the courses and
 - design the future training programme at LRZ and GCS according to your needs and wishes.



Upcoming Courses and Workshops



AI Training Series - Introductory AI Workshop by Intel	17.06.2024 – 17.06.2024	ONLINE	0.00 EUR	94	10.06.2024	Register now
EuroCC2 AI for Science Bootcamp (apply via openhackathons.org)	24.06.2024 – 26.06.2024	ONLINE	0.00 EUR	80	14.05.2024	
AI Training Series - High Performance Data Analytics Using R at LRZ	10.07.2024 – 10.07.2024	ONLINE	0.00 EUR	89	05.07.2024	Register now
Deep Learning and GPU Programming Workshop	09.09.2024 – 12.09.2024	ONLINE	0.00 EUR	37	02.09.2024	Register now
Iterative Solvers for Linear Systems	17.09.2024 – 18.09.2024	Leibniz Rechenzentrum	25.00 EUR – 420.00 EUR	27	03.09.2024	Register now
Introduction to C++	22.10.2024 – 24.10.2024	ONLINE	30.00 EUR – 600.00 EUR	63	15.10.2024	Register now
Node-Level Performance Engineering	03.12.2024 – 05.12.2024	ONLINE	30.00 EUR – 600.00 EUR	58	26.11.2024	Register now
Modern C++ Software Design	09.12.2024 – 11.12.2024	ONLINE	30.00 EUR – 600.00 EUR	53	27.11.2024	Register now



- **Information on further HPC courses:**

- by LRZ: <http://www.lrz.de/services/compute/courses/>
- by the Gauss Centre of Supercomputing (GCS): <https://www.gauss-centre.eu/trainingsworkshops>
- by German Centres (collected by the Gauß-Allianz): <https://hpc-calendar.gauss-allianz.de/>
- by the EuroCC National Competence Centres: <https://www.eurocc-access.eu/services/training/>
- by European Centres: <https://hpc-portal.eu/upcoming-events-courses>



- Subscribe via <https://lists.lrz.de/mailman/listinfo/training-announce> to get regular information on Education and Training events (co-) organised by LRZ.
- Subscribe via <https://www.lrz.de/wir/newsletter/> for the general LRZ newsletter.
- Subscribe via <https://lists.rwth-aachen.de/postorius/lists/vi-hps-news.lists.rwth-aachen.de/> for the VI-HPS newsletter.

Thank you!
We hope to see you again in a future training event!