



Center of Excellence for Exascale in Solid Earth

The HLRS team in ChEESE-2

José Gracia

High Performance Computing Center (HLRS)
University of Stuttgart



This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 823844

Outline

- HLRS
- SPMT group and involvement in ChEESE-2






Höchstleistungsrechenzentrum Stuttgart (HLRS)

High Performance Computing Center Stuttgart

- Central unit of the U Stuttgart
 - origin going back to 1968
- 1st German National Supercomputing Center
 - founded 1996
 - compute resources for German researchers
- Gauss Center for Supercomputing
 - founded 2007
 - HLRS, JSC/Jülich and LRZ/Garching
- PRACE RI hosting member
 - founded ~2012
 - open for EU researchers since 2004
- Staff
 - ~100 scientific staff
 - ~40 technical & admin staff
 - ~10 student assistants
- Broad experience in collaborative research projects on national, European, and international level
- Significant share of compute resources is used by industrial customers








Research Projects



EXCELLERAT



CIRCE



Computational Immediate Response Center for Emergencies

exaFOAM



SDC4LIT



CASTIEL



BWDATA
ARCHIV

MONT-BLANC

Coe GSS



Centre of excellence



DreamCloud



myThOS

Many Threads Operating Syst

HPC
Europa

MIKELANGEL



olca



PAA AGE



dash

CREST



ORCHESTRA



BEAM-ME



PRACE



EXCESS



HONK



FORTISSIMO

Center of Excellence for Exascale in Solid Earth

PHANTOM

Industrial Collaborations

RECOM
SERVICES

Lauer & Weiss



ELECTRIC ANT LAB
Scientific Modeling & Simulations



STELBA
Hydro

MANN+HUMMEL

OPTIMA
EXCELLENCE IN PHARMA



DÜRR

HSG-IMIT



Putzmeister





Departments and Research Directions

- **Numerical Methods & Libraries**
- **Training**
- **Philosophy of Computer Simulation**
- **Visualization**
Virtual & augmented reality
- **Scalable Programming Models and Tools**
Hybrid HPC models, debugging tools, performance analysis & optimization
- **Service Management & Business Processes**
Cloud computing, virtualization, data analytics, ML & AI



Current HLRS Flagship System

“Hawk” HPE Apollo 9000:

- 720.896 cores AMD EPYC Rome
- 5632 nodes, 2 sockets, 128GB memory
- Performance:
 - 25 PFs *Peak* performance
 - 19.3 PFs *Linpack* performance rank 30 on current TOP500 list
- ~1.44 PB Main Memory
- ~42 PB Disk Storage





Scalable Programming models and Tools (SPMT)

- established in 2012, 6 - 10 members

Research directions

- User productivity:
 - Model-centric debugging for task-based HPC models
 - High-level programming models/idioms for expressing concurrency, data distribution and communication hiding
 - Understanding common performance issues in real codes
- Smart runtime systems:
 - Overlapping computation with communication and I/O
 - Tuning runtime parameters for energy-efficiency





Who is José Gracia

- Studied physics at U Heidelberg, Germany
- 2002 PhD in theoretical astrophysics
- 7 years PostDoc in Heidelberg, Athens and Dublin

Large scale numerical simulations on clusters

HPC from a *computational science* perspective

- Joined HLRS in 2009; evaluation of novel HPC models
- Group leader since 2012

HPC from a *computer science* perspective





Who is Christoph Niethammer

- Studied physics and mathematics at U Stuttgart, Germany
- 2022 PhD in HPC
- Joined HLRS as student assistant in 2009
- Research interests:
 - Scalable Molecular Dynamics
 - MPI, in particular topologies & collectives
 - Scientific Workflows
 - Performance analysis

HPC as a tool for scientist





Role in ChEESE-2

- WP2:
 - Code audits (T2.1)
- WP4: WP lead
 - Workflow implementation and optimization
 - Lightweight performance analysis of workflow components
- WP7:
 - Liaison activities (T7.2)





Thanks for your attention!

Questions?

