




1

Are we ready for the next level of Big Data?

Prof. Dr. Dieter Kranzlmüller
Leibniz-Rechenzentrum (LRZ) &
Ludwig-Maximilians-Universität München (LMU)

711. WE-Heraeus-Seminar 12.01.2020 | Dieter Kranzlmüller 2

2



3

Science Cooperation

Terra_Byte: Cooperation between DLR and LRZ



- **Storing and Analyzing** Satellite Data (Earth Observation)
- All missions that **DLR** is participating
- **40 to 50 PB** until 2025
- Both, **Data and Compute performance** are required
- **Scale-Up and Scale-Down**
- Usable for all **LRZ partners in Science and Research**



Foto: Alessandro Podo, LRZ

4

Supercomputing - Top in the European Union political agenda



European Commission President
Jean-Claude Juncker

Our ambition is that by 2020, Europe ranks in the top 3 HPC powers worldwide
27 October 2015

04/2016: European Cloud Initiative COM(2016) 178

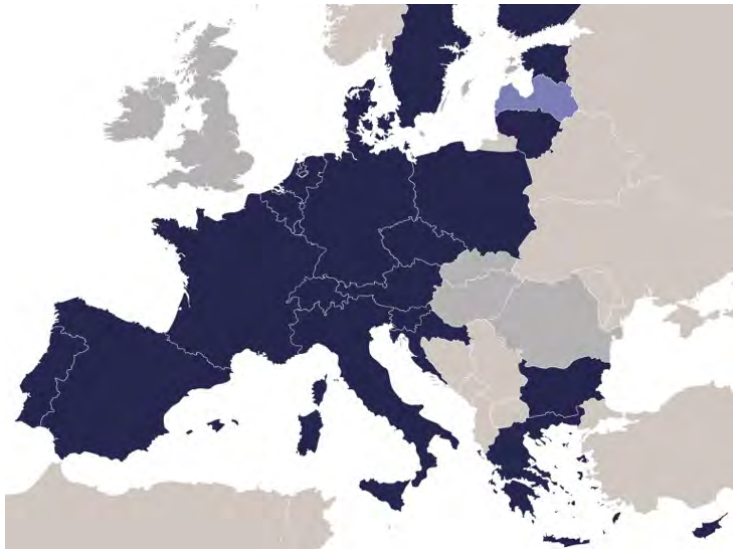
A world-class HPC, data & network infrastructure and a leading HPC and Big Data ecosystem

05/2017: Mid-Term Review of the Digital Single Market Strategy COM(2017) 228

by end-2017, propose a legal instrument providing a procurement framework for an exascale supercomputing & data infrastructure

7

The Future of European Supercomputing The EuroHPC Declaration



Quelle: <https://ec.europa.eu/>
711. WE-Heraeus-Seminar

Primeur Magazine
@primeurmagazine

End of next year Europe will have the number 1, 2, 3 #supercomputers in the world expects #EC @DSMeu Deputy Director-General Khalil Rouhana. Thanks to 1 billion euro investment in #EuroHPC at #ICTPropDay in Helsinki

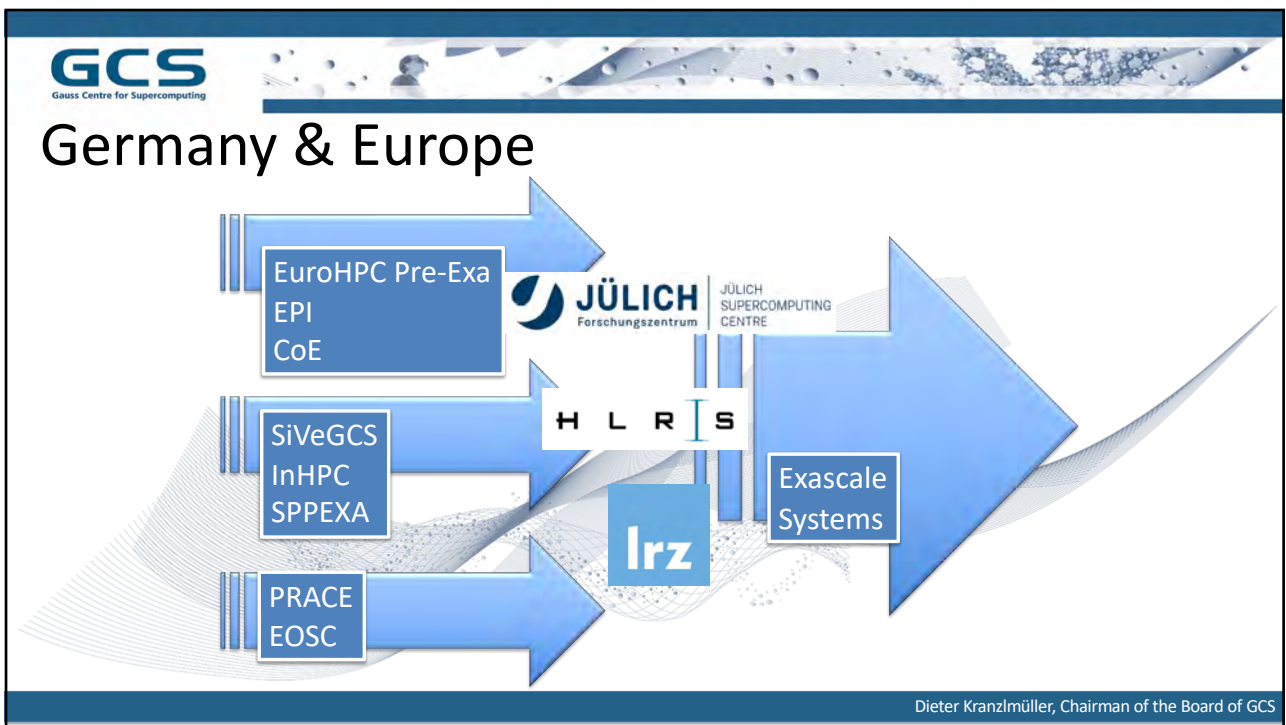
[Tweet übersetzen](#)



8



9



10

German Initiative on Research Data

NFDI – National Research Data Infrastructure

DFG Deutsche
Forschungsgemeinschaft

Aim of the NFDI

... to systematically manage **scientific and research data**, provide long-term data storage, backup and accessibility, and network the data both nationally and internationally.

... by means of a coordinated network of consortia tasked with providing **science-driven data services** to research communities.

- Address the needs of researchers
- Facilitate easy access to data
- Foster sustainability
- Initial funding of approximately **€85 million per year** (incl. 22% overhead)
- Project funding for **10 years**; decision on continued financing in 2026
- Expected to fund approximately 30 consortia, selected in 3 calls

https://www.dfg.de/download/pdf/foerderung/programme/nfdi/welcome_speech_expert_committee_allgoewer.pdf

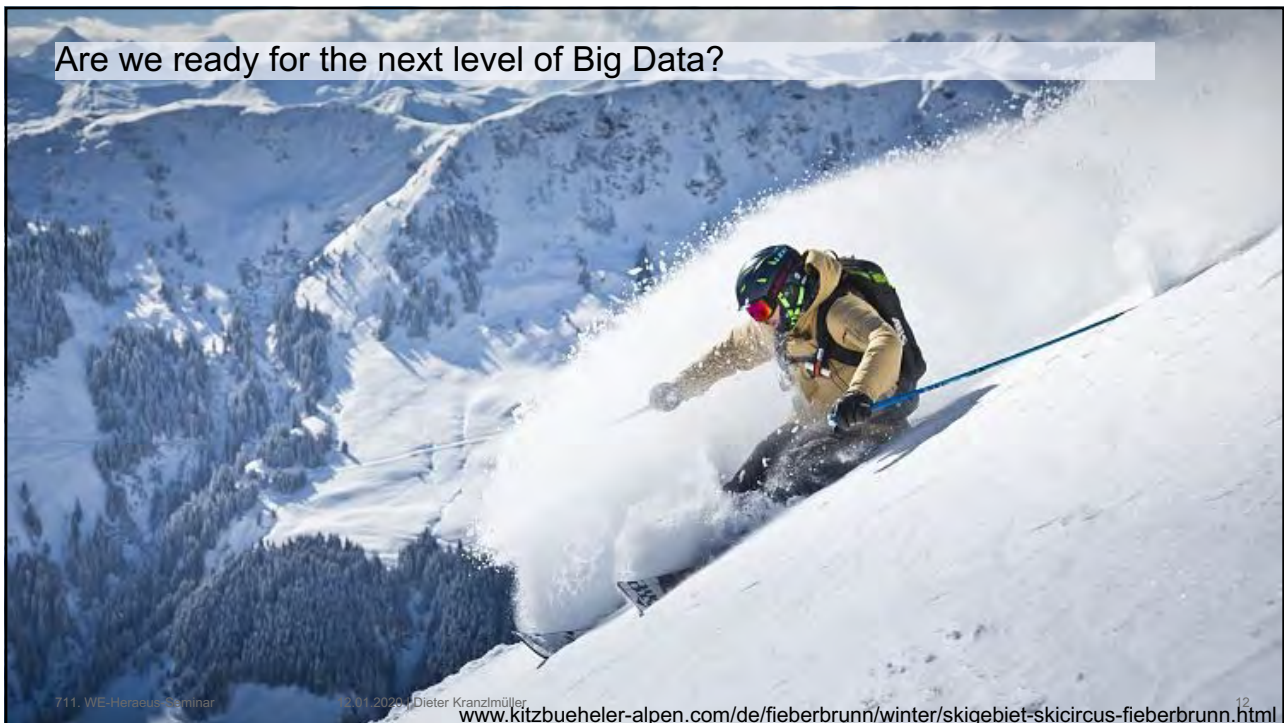
711. WE-Heraeus-Seminar

12.01.2020 | Dieter Kranzlmüller

11

11

Are we ready for the next level of Big Data?



711. WE-Heraeus-Seminar


12.01.2020 | Dieter Kranzlmüller





www.kitzbueheler-alpen.com/de/fieberbrunn/winter/skigebiet-skicircus-fieberbrunn.html

12

12

Top 500 Supercomputers (November 2019)




Rank	System	Cores	Rmax (TFlop/s)	Rpeak (TFlop/s)	Power (kW)
1	 Summit - IBM Power System AC922, IBM POWER9 22C 3.07GHz, NVIDIA Volta GV100, Dual-rail Mellanox EDR Infiniband, IBM DOE/SC/Oak Ridge National Laboratory United States	2,414,592	140,606.0	140,794.9	10,096
2	 Sierra - IBM Power System S922LC, IBM POWER9 22C 3.1GHz, NVIDIA Volta GV100, Dual-rail Mellanox EDR Infiniband, IBM / NVIDIA / Mellanox DOE/NNSA/LLNL United States	1,572,000	94,640.0	125,712.0	7,438
3	 Sunway TaihuLight - Sunway ME6500, Sunway SW26010 280C 1.45GHz, Sunway NRPC National Supercomputing Center in Wuxi China	10,649,600	93,014.6	125,435.9	15,371
4	 Tianhe-2 - TH-IVB-FEP Cluster, Intel Xeon E5-2692v2 12C 2.2GHz, TH Express-2, Matrix-2000, NUDT National Super Computer Center in Guangzhou China	4,981,760	61,444.5	100,678.7	18,482









FAKE NEWS?!

<https://www.top500.org/lists/2019/11/>

711. WE-Heraeus-Seminar 12.01.2020 | Dieter Kranzlmüller 13

13



	Publicly known Exaflop Capacity until 2024	Today's Computing Capacity ²	Increase	Investment: Technology [†] + Ecosystem ^o \$ Billion	GDP: 2017 \$ Trillion	Producer
 China	8 x EF	154 PF	52 x	> 3.2 ^T , ? ^o	12,24	
 USA	3 x EF	238 PF	13 x	1.8 ^T , 1.5 ^o	19,39	
 Japan	2 x EF	33 PF	61 x	> 1.3 ^T ^o	4,87	
 EU	2 x EF	40.7 PF	49 x	1.2 ^T , X ^o	16,46 ¹	

A GLOBAL PRIORITY

High Performance Computing as Key Technology

¹ EuroHPC Countries Total. Source: WorldBank (06.07.18)
² Rmax Top 2 Machines, Top500 Listing, Nov 2018

14

Legend:

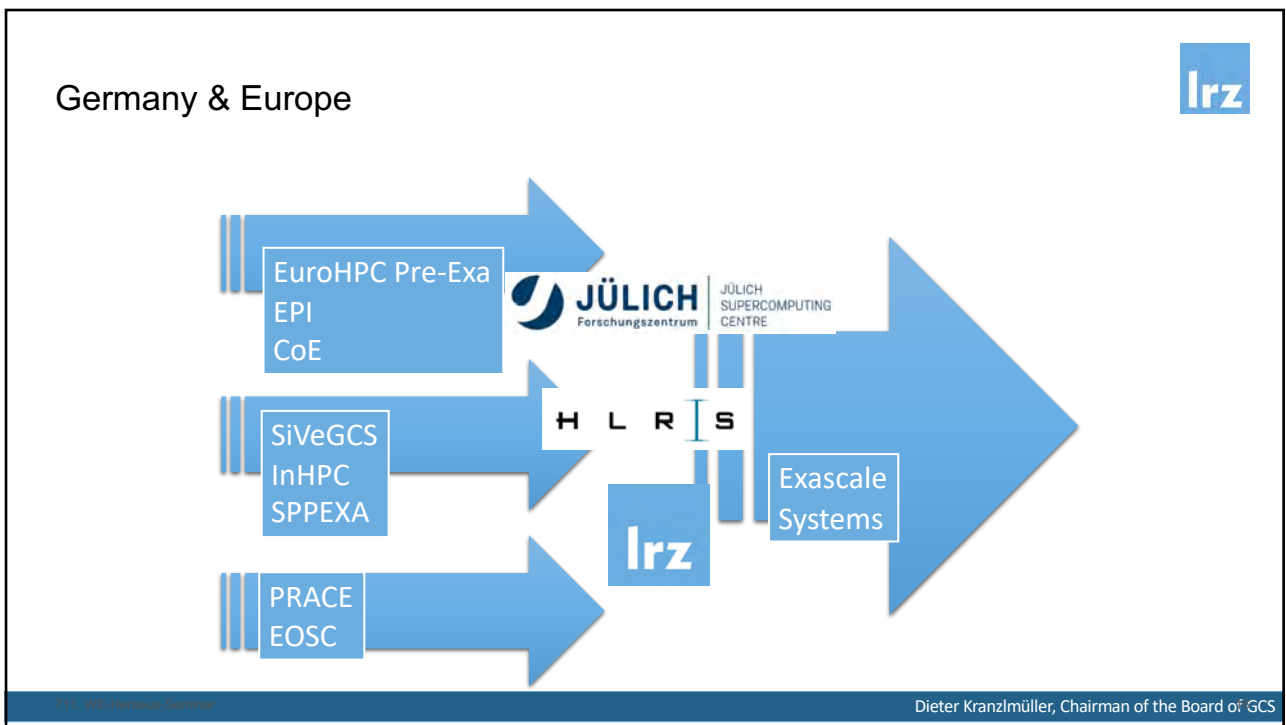
- Pre-exascale – Finland led consortium
- Pre-exascale – Italy led consortium
- Pre-exascale – Spain led consortium
- Exascale – Germany
- Exascale – France
- Other EuroHPC countries

Primeur Magazine @primeurmagine

End of next year Europe will have the number 1, 2, 3 #supercomputers in the world expects #EC @DSMeu Deputy Director-General Khalil Rouhana. Thanks to 1 billion euro investment in #EuroHPC at #ICTPropDay in Helsinki

www.eurohpc.eu

15



16



17

Supercomputing 2018, Dallas, Texas <https://sc18.supercomputing.org>

Rank	Site	Manufacturer	Computer	Country	Cores	Speed (GFLOPS)	Power (MW)
1	Oak Ridge National Laboratory	IBM	Summit IBM Power System, P9 22C 3.87GHz, Mellanox EDR, NVIDIA GV100	USA	2,397,824	143.5	9.8
2	Lawrence Livermore National Laboratory	IBM	Sierra IBM Power System, P9 22C 3.1GHz, Mellanox EDR, NVIDIA GV100	USA	1,572,480	94.6	7.4
3	National Supercomputing Center in Wuxi	NRCPC	Sunway TaihuLight NRCPC Sunway SW26010, 26C 1.5GHz	China	10,649,600	93.0	15.4
4	National University of Defense Technology	NUDT	Tianhe-2A ANUPT TH-1VB-FEP, Xeon E5 12C 2.1GHz, Mellanox EDR	China	4,981,760	61.4	18.5
5	Swiss National Supercomputing Centre (CSCS)	Cray	Phi Dawn Cray XC50, Xeon ES 12C 2.6GHz, Aries, NVIDIA Tesla P100	Switzerland	387,872	21.23	2.38
6	Los Alamos NL / Sandia NL	Cray	Titanium Cray XC40, Intel Xeon Phi 7259 68C 1.4GHz, Aries	USA	979,072	20.16	7.58
7	National Institute of Advanced Industrial Science and Technology	Fujitsu	AI Bridging Cloud Infrastructure (ABCI) PRIMERGY CX2550 M4	Japan	391,680	19.9	1.65
8	Leibniz Rechenzentrum	Lenovo	SuperMUC-NG ThinkSystem S310330, Xeon Platinum 6174 24C 3.1GHz, Intel Cloud Burst	Germany	305,856	19.5	1.5
9	Oak Ridge National Laboratory	Cray	Theta Cray XK7, Xeon E5 12C 2.6GHz, Mellanox EDR, NVIDIA GV100	USA	560,640	17.6	6.3
10	Lawrence Livermore National Laboratory	IBM	Sierra IBM Power System, P9 22C 3.1GHz, Mellanox EDR, NVIDIA GV100	USA	1,572,864	17.2	7.8

GRAPH 500

#1 Data-intensive applications **SSSP** benchmark

#5 Data-intensive applications **BFS** benchmark

Top 500 Birds of a Feather Session, 13. November 2018

711. WE-Heraeus-Seminar 12.01.2020 | Dieter Kranzlmüller

18

GCS
Gauss Centre for Supercomputing


JSC "JUWELS" Modular Supercomputer

**JSC "JUWELS Cluster"
Atos BullSequana (2018)**

- 122,448 cores Intel Xeon Skylake
- 12 PetaFlops Peak
- 6.2 PetaFlops Sustained
- 273 TeraByte Main Memory
- 75 PetaByte Disk

**JSC "JUWELS Booster"
TBD (2020)**

- TBD cores
- >60 PetaFlops Peak
- TBD PetaFlops Sustained
- TBD TeraByte Main Memory
- 75 PetaByte Disk



711. WE-Heraeus-Seminar 19

19

GCS
Gauss Centre for Supercomputing

HLRS HPE "Hawk"

HPE BADGER (2020)

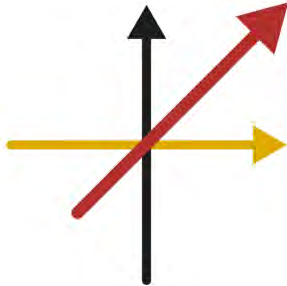
- 720,896 cores AMD Rome
- 27,11 PetaFlops Peak
- 746 TeraByte Main Memory
- ~26 PetaByte Disk



711. WE-Heraeus-Seminar 20

20

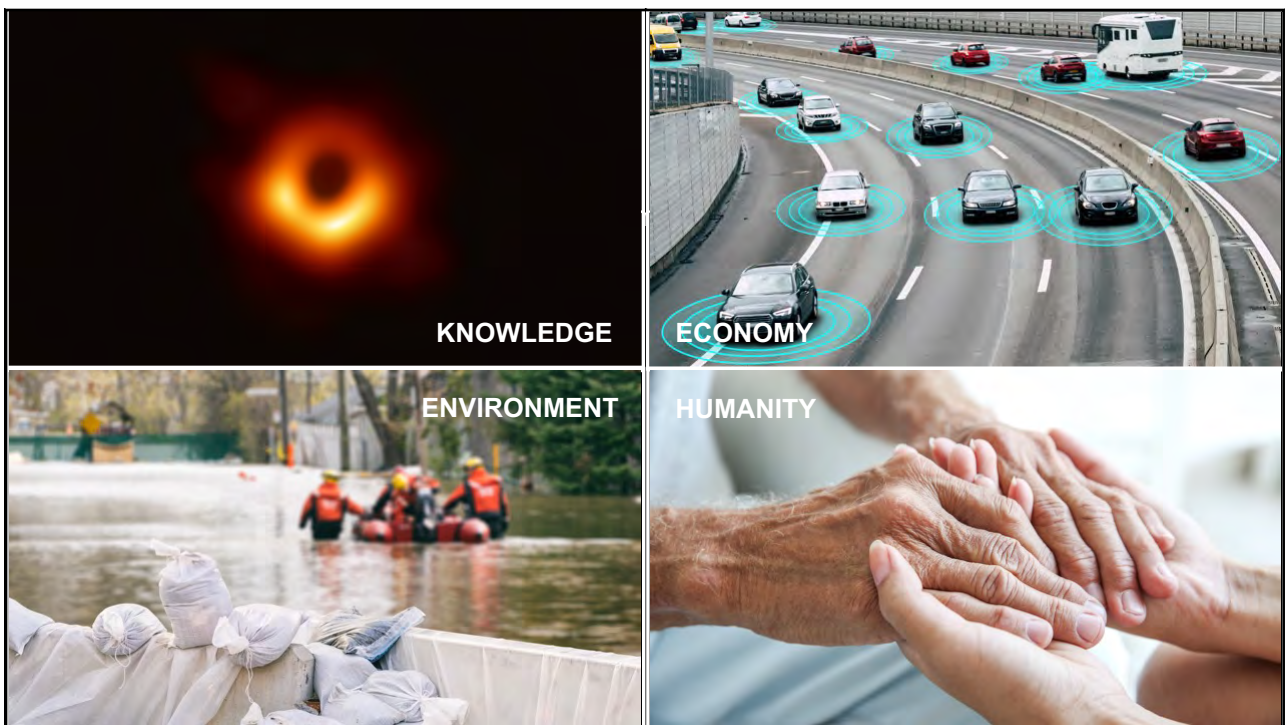
GCS Smart Scaling Strategie



Prof. Dr. Dieter Kranzlmüller (LRZ)
Prof. Dr. Thomas Lippert (FZJ)
Prof. Dr. Michael Resch (HLRS)

711. WE-Heraeus-Seminar

21



KNOWLEDGE

ECONOMY

ENVIRONMENT

HUMANITY

22



23



24

LRZ User Application Profile

Application Mix – General Purpose, Broad Use

Computational Fluid Dynamics
Optimisation of turbines/wings, noise reduction

Fusion
Plasma in a future fusion reactor (ITER)


Astrophysics
Origin and evolution of stars and galaxies

Solid State Physics
Superconductivity, surface properties


Geophysics
Earth quake scenarios

...


- > 9.5 billion compute hours consumed
- > 6.3 million jobs processed
- > 820 research projects carried out
- > 2,230 researchers as clients



SuperMUC-NG Science Symposium
<https://www.lrz.de/presse/ereignisse/SuperMUC-NG-Science-Symposium/>



Get your digital copy
of the results book on SuperMUC projects



25



26

Results of the computation obtained on SuperMUC



Alexandros Stamatakis | Scientific Computing Group, Heidelberg Institute for Theoretical Studies (HITS)/Exelixis Lab
711. WE-Heraeus-Seminar 12.01.2020 | Dieter Kranzlmüller


27

27

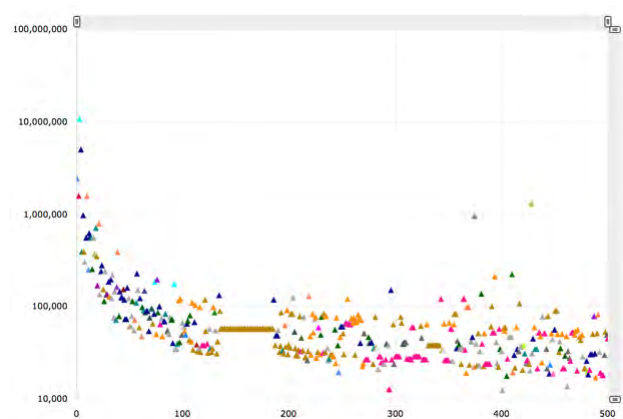
<https://www.top500.org/statistics/efficiency-power-cores/> Scaling – Number of Cores



1993:



2018:




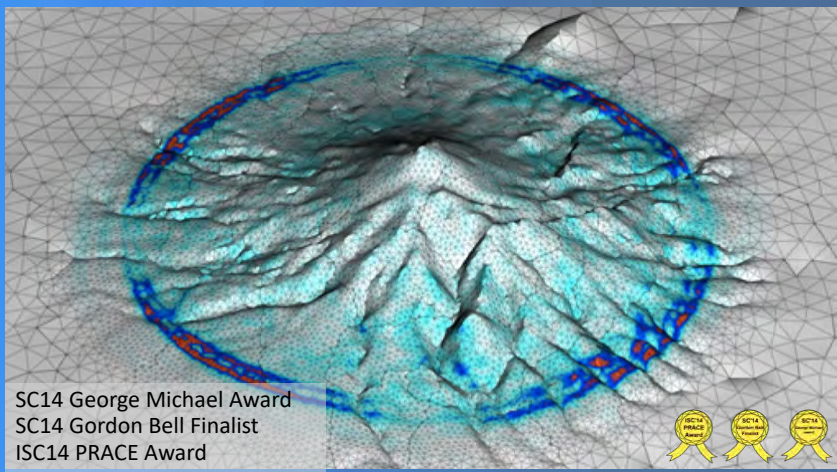
711. WE-Heraeus-Seminar 12.01.2020 | Dieter Kranzlmüller

28


28

Extreme Scaling Case Study
SeiSol – Numerical Simulation of Seismic Wave Phenomena





SC14 George Michael Award
 SC14 Gordon Bell Finalist
 ISC14 PRACE Award




**1,42 Petaflop/s
 on 147.456 Cores
 of SuperMUC**
 44,5 % of Peak
 Performance

Dr. Christian Pelties
 Department of Earth
 and Environmental
 Sciences (LMU)

Prof. Michael Bader
 Department of
 Informatics (TUM)

711. WE-Heraeus-Seminar
12.01.2020 | Dieter Kranzlmüller
29

29



PARTNERSHIP FOR ADVANCED COMPUTING IN EUROPE

- ▶ About PRACE
- ▶ HPC Access
- ▶ Industry
- ▶ PRACE Projects and Outcomes
- ▶ Training and Documentation
- ▶ Events
- ▶ Media
- ▶ Outreach
- ▶ Contact

Home page » Events » PRACEdays » [Dr. Alice-Agnes Gabriel wins 2020 PRACE Ada Lovelace Award for HPC](#)


Dr. Alice-Agnes Gabriel wins 2020 PRACE Ada Lovelace Award for HPC

Tuesday 7 January 2020


PRACE is delighted to announce that Dr. Alice-Agnes Gabriel, [Department of Earth and Environmental Sciences Geophysics](#), Ludwig-Maximilians-University of Munich (LMU) is the winner of the 2020 PRACE Ada Lovelace Award for HPC for her outstanding contributions to and impact on HPC in Europe.

Dr. Alice-Agnes Gabriel, is a lecturer in Professor Heiner Igel's Chair of Seismology in the Institute of Geophysics, LMU. Her research is focusing on understanding the physics of earthquakes using theoretical analysis, physics-based forward models, innovative observation techniques and High-Performance Computing to bridge spatio-temporal scales.


Dr. Gabriel's career is characterized by first-rate earthquake scenarios realised on some of the largest supercomputers worldwide. Her research work was nominated and awarded by several prizes such as one of the Finalist by Association for Computing Machinery (ACM) [Gordon Bell Prize in 2014](#), at International Conference for High-Performance Computing, Networking, Storage, and Analysis (SC14), New Orleans, USA; and [Best Paper Award](#), IEEE/ACM International Conference for High-Performance




Press releases



30





31

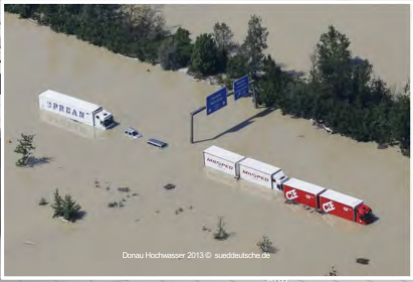


32

StMUV Project ClimEx



Assessment of the effects of climate change on hydrological extreme events such as floodings and droughts



Coupling HPC models for Climate and Hydrology, workflow and data management

Titel:	Klimawandel und Hydrologische Extremereignisse – Risiken und Perspektiven für die Wasserwirtschaft in Bayern und Québec
Partner:	Prof. Ralf Ludwig, LMU, LfU (DE), OURANOS, CEHQ, ETS (CA)
Duration:	2015 - 2019

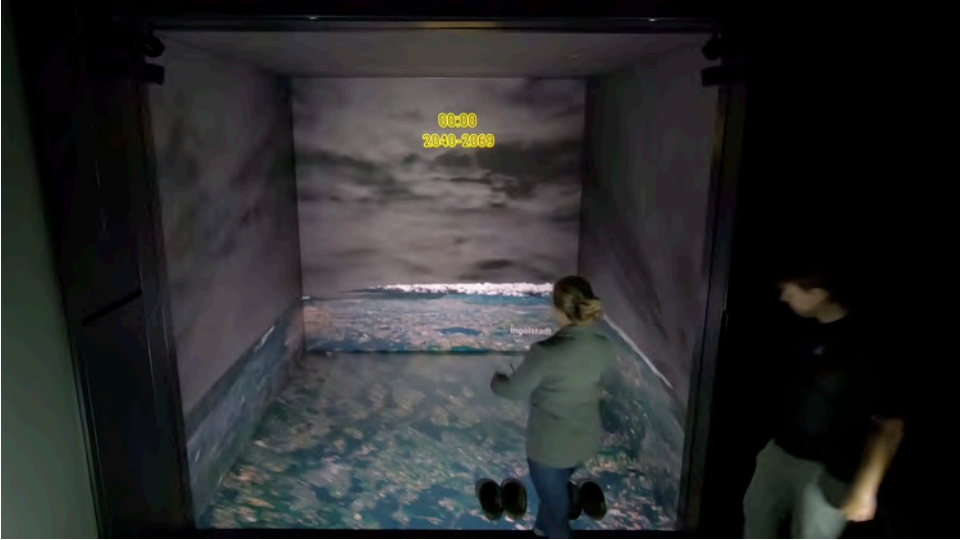

www.climex-project.org

711. WE-Heraeus-Seminar 12.01.2020 | Dieter Kranzlmüller 33

33

Effects of extreme meteorological events on hydrology in Bavaria and Québec

StMUV Project ClimEx - Influence of Climate Change



711. WE-Heraeus-Seminar 12.01.2020 | Dieter Kranzlmüller 34

34

From Long-Tail to Power-Users

High-Performance Computing for a Broad Range of Users

Requirements and big demands from Universities
(from world-renowned scientist to high potentials and aspiring students)

Applications with different requirements on hardware and software

Users with different levels of expertise in HPC
from different application domains

Broad Architecture Portfolio

Focus on Usability & Workflow

Integration of new domains

Expertise in Consulting & Partnership

Extensive Education & Training

Trusted bidirect. Communication with Users

Support for Academic Startups

Digitalisation Environment Biodiversity Chemistry Geophysics CFD Particle Physics Medicine AI Quantum Physics

lrz

35



36

Best Practices for Collaborations LRZ Partnership Model



Scientific partnership must be on equal footing

1. Invite interesting groups for **partnership workshops**
 - Introduce yourself and invite them to introduce themselves
 - Discuss open questions, pain points and potential collaboration ideas
2. Choose scientific groups for **flagship roles**
3. Establish a **dedicated point-of-contact**
 - Provide individual support and guidance, including targeted training & education
 - Inform about use case specific (optimized) IT infrastructures and possibilities
 - Offer early access to latest IT infrastructure (hard- and software) developments
 - Collect specifications of future requirements
 - Bridge to other partners and experts at CS and Math departments
4. Embed IT experts in user groups
5. Joint research projects (including funding), joint publications

37

Are we ready for the next level of Big Data?



- We need the **best computing infrastructure** for science and research
- We need to **emphasize collaboration** between domain scientists and computer scientists
- We need to **integrate HPC, AI, QC, ...** and **Research Data** into tools for science

Contact:

Prof. Dr. Dieter Kranzlmüller
kranzlmuller@lrz.de / @Kranzlmuller



38

THE INNOVATION OF SUPERCOMPUTING

HPC
Modeling & Simulation (M&S)

Quantum Computing

AI & Machine Learning

Big Data

711. WE-Heraeus-Seminar 39

The slide features a dark blue background with a central circular flow diagram of three arrows. On the left is a server rack icon. On the right is a brain icon with circuitry. At the bottom center is a bar chart with three bars of increasing height. The text 'THE INNOVATION OF SUPERCOMPUTING' is in white and yellow. 'HPC Modeling & Simulation (M&S)' is in white. 'Quantum Computing' is in white. 'AI & Machine Learning' is in white. 'Big Data' is in white. The GCS logo is in the top left. The text '711. WE-Heraeus-Seminar' and '39' are in the bottom left and right respectively.

39

THE INNOVATION OF GERMAN SUPERCOMPUTING

TIER 0/1 Centres

TIER 2 Centres

Data Infrastructure

711. WE-Heraeus-Seminar 40

The slide features a dark blue background with a central circular flow diagram of three arrows. On the left is a server rack icon. On the right is a server rack icon with three horizontal bars. At the bottom center is a bar chart with three bars of increasing height. The text 'THE INNOVATION OF GERMAN SUPERCOMPUTING' is in white and yellow. 'TIER 0/1 Centres' is in white. 'TIER 2 Centres' is in white. 'Data Infrastructure' is in white. The GCS logo is in the top left. A German flag icon is in the top right. The text '711. WE-Heraeus-Seminar' and '40' are in the bottom left and right respectively.

40



THE INNOVATION OF GERMAN SUPERCOMPUTING

National Tier-2 High Performance Computing (NHR)



Gauss Alliance



National Research Data Infrastructure (NFDI)

711. WE-Heraeus-Seminar 41

41



kranzmueller@lrz.de / @Kranzmueller

42