



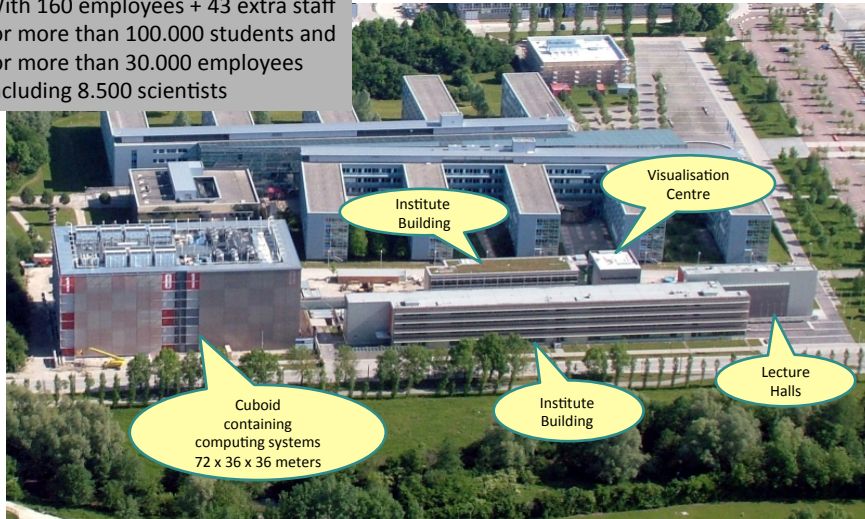
Some Thoughts about Extreme Scaling and Power

Dieter Kranzlmüller

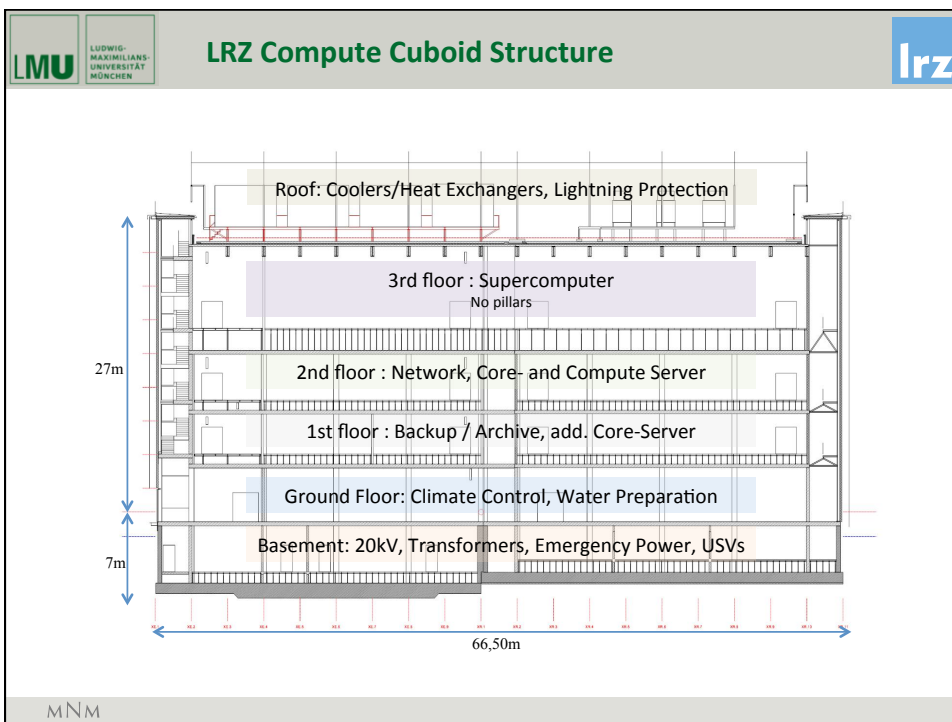

Munich Network Management Team
Ludwig-Maximilians-Universität München (LMU) &
Leibniz Supercomputing Centre (LRZ)
of the Bavarian Academy of Sciences and Humanities



With 160 employees + 43 extra staff for more than 100.000 students and for more than 30.000 employees including 8.500 scientists



Picture: Ernst Graf



LUDWIG-
MAXIMILIANS-
UNIVERSITÄT
MÜNCHEN

SuperMUC Phase 1 + 2

D. Kranzmüller

Energy Efficiency and Extreme Scaling

5

LUDWIG-
MAXIMILIANS-
UNIVERSITÄT
MÜNCHEN

Top 500 Supercomputer List (June 2012)


Rank	Site	Computer/Year Vendor	Cores	R _{max}	R _{peak}	Power
1	DOE/NNSA/LLNL United States	Sequoia - BlueGene/Q, Power BQC 16C 1.60 GHz, Custom / 2011 IBM	1572864	16324.75	20132.66	7890.0
2	RIKEN Advanced Institute for Computational Science (AICS) Japan	K computer , SPARC64 VIIItx 2.0GHz, Tofu interconnect / 2011 Fujitsu	705024	10510.00	11280.38	12659.9
3	DOE/SC/Argonne National Laboratory United States	Mira - BlueGene/Q, Power BQC 16C 1.60GHz, Custom / 2012 IBM	786432	8162.38	10066.33	3945.0
4	Leibniz Rechenzentrum Germany	SuperMUC - iDataPlex DX360M4, Xeon E5-2680 8C 2.70GHz, Infiniband FDR / 2012 IBM	147456	2897.00	3185.05	3422.7
5	National Supercomputing Center in Tianjin China	Tianhe-1A - NUDT YH MPP, Xeon X5670 6C 2.93 GHz, NVIDIA 2050 / 2010 NUDT	186368	2566.00	4701.00	4040.0
6	DOE/SC/Oak Ridge National Laboratory United States	Jaguar - Cray XK6, Opteron 6274 16C 2.200GHz, Cray Gemini interconnect, NVIDIA 2090 / 2009 Cray Inc.	298592	1941.00	2627.61	5142.0
7	CINECA Italy	Fermi - BlueGene/Q, Power BQC 16C 1.60GHz, Custom / 2012 IBM	163840	1725.49	2097.15	821.9
8	Forschungszentrum Juelich (FZJ) Germany	JuQUEEN - BlueGene/Q, Power BQC 16C 1.60GHz, Custom / 2012 IBM	131072	1380.39	1677.72	657.5
9	CEA/TGCC-GENCI France	Curie thin nodes - Bullx B510, Xeon E5- 2680 8C 2.700GHz, Infiniband QDR / 2012 Bull	77184	1359.00	1667.17	2251.0
10	National Supercomputing Centre in Shenzhen (NSCS) China	Nebulae - Dawning TC3600 Blade System, Xeon X5650 6C 2.66GHz, Infiniband QDR, NVIDIA 2050 / 2010 Dawning	120640	1271.00	2984.30	2580.0

www.top500.org

D. Kranzmüller

Energy Efficiency and Extreme Scaling

6



LMU LUDWIG-MAXIMILIANS-UNIVERSITÄT MÜNCHEN **Increasing numbers** **lrz**

Date	System	Flop/s	Cores
2000	HLRB-I	2 Tflop/s	1512
2006	HLRB-II	62 Tflop/s	9728
2012	SuperMUC	3200 Tflop/s	155656
2014	SuperMUC Phase II	3.2 + 3.2 Pflop/s	229960

MNM D. Kranzmüller Energy Efficiency and Extreme Scaling 8



LMU LUDWIG-MAXIMILIANS-UNIVERSITÄT MÜNCHEN **SuperMUC and its predecessors** **lrz**

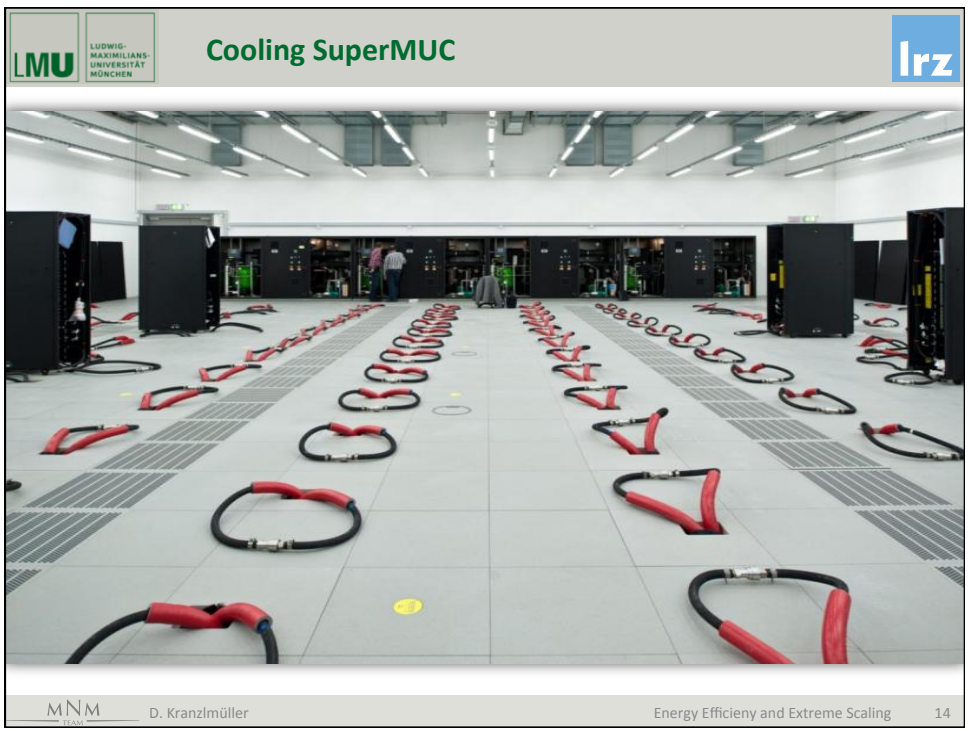
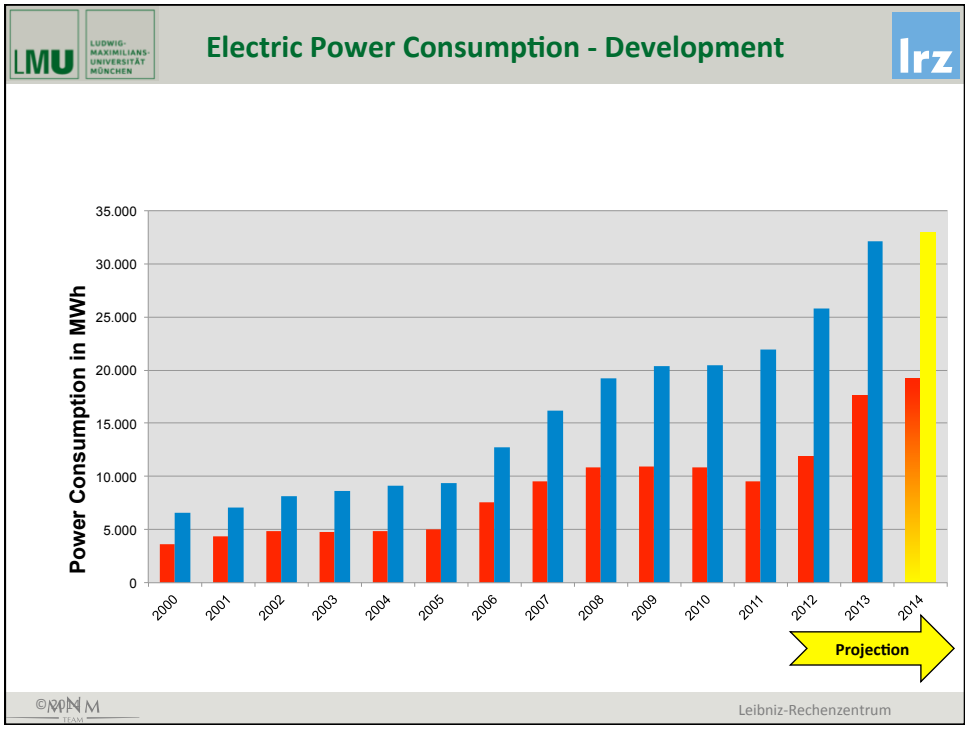
MNM TEAM D. Kranzmüller Energy Efficiency and Extreme Scaling 11


LMU LUDWIG-MAXIMILIANS-UNIVERSITÄT MÜNCHEN **LRZ Building Extension** **lrz**

Picture: Horst-Dieter Steinhöfer

Figure: Herzog+Partner für StBAM2 (staatl. Hochbauamt München 2) Picture: Ernst A. Graf

MNM TEAM D. Kranzmüller Energy Efficiency and Extreme Scaling 12







iDataPlex DWC Rack
w/ water cooled nodes
(front view)



iDataPlex DWC Rack
w/ water cooled nodes
(rear view of water manifolds)

Torsten Bloth, IBM Lab Services - © IBM Corporation





Photos: StBAM2 (staatl. Hochbauamt München 2)

LMU LUDWIG-MAXIMILIANS-UNIVERSITÄT MÜNCHEN **Energy Consumption in Data Centres** lrz

■ Data Centers are “Heaters with integrated logic”

Electrical Power IN **IT Equipment 47%** INDOOR DATA CENTER HEAT Heat OUT

- Chiller 23%
- Humidifier 3%
- CRAC/CRAH 15%
- PDU 3%
- UPS 6%
- Lighting / aux devices 2%
- Switchgear / generator 1%

Torsten Bloth, IBM Lab Services - © IBM Corporation

MNM D. Kranzmüller Energieeffizientes Supercomputing 19



LMU LUDWIG-MAXIMILIANS-UNIVERSITÄT MÜNCHEN **ICT and Global Warming** lrz

“In 2012, global CO2 emissions were 31.7 GtCO2. This represents a 1.2% year-on-year increase in emissions”
International Energy Agency
<https://www.iea.org/publications/freepublications/publication/CO2EmissionsFromFuelCombustionHighlights2014.pdf>

“ICT industry contributes to about 2 percent of global carbon dioxide emissions”
“Emission of carbon by the ICT sector is likely to double by the year 2020”
<http://www.natureworldnews.com/articles/467/20130107/ict-sector-account-2-percent-global-carbon.htm>

Picture © Flickr User johnb

MNM D. Kranzmüller Energieeffizientes Supercomputing 20



LMU LUDWIG-MAXIMILIANS-UNIVERSITÄT MÜNCHEN **SuperMUC Phase 1 + 2** lrz

Phase 1
3.2 PFLOP/s

SuperMUC

Phase 2
3.2 PFLOP/s


MNM D. Kranzmüller Extreme Scale Computing 23

LMU LUDWIG-MAXIMILIANS-UNIVERSITÄT MÜNCHEN **LRZ Application Mix** lrz

- Computational Fluid Dynamics: Optimisation of turbines and wings, noise reduction, air conditioning in trains
- 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
- Material Science: Semiconductors
- Chemistry: Catalytic reactions
- Medicine and Medical Engineering: Blood flow, aneurysms, air conditioning of operating theatres
- Biophysics: Properties of viruses, genome analysis
- Climate research: Currents in oceans

> 200 Applications on SuperMUC
> 14 Applications use full machine (Phase 1 or 2)




MNM D. Kranzmüller Energy Efficiency and Extreme Scaling 24




Results (Sustained TFlop/s on 128000 cores)




Name	MPI	# cores	Description	TFlop/s/island	TFlop/s max
Linpack	IBM	★ 128000	TOP500	161	2560
Vertex	IBM	★ 128000	Plasma Physics	15	245
GROMACS	IBM, Intel	★ 64000	Molecular Modelling	40	110
Seissol	IBM	★ 64000	Geophysics	31	95
waLBerla	IBM	★ 128000	Lattice Boltzmann	5.6	90
LAMMPS	IBM	★ 128000	Molecular Modelling	5.6	90
APES	IBM	★ 64000	CFD	6	47
BQCD	Intel	★ 128000	Quantum Physics	10	27

MNM D. Kranzmüller Energy Efficiency and Extreme Scaling 26

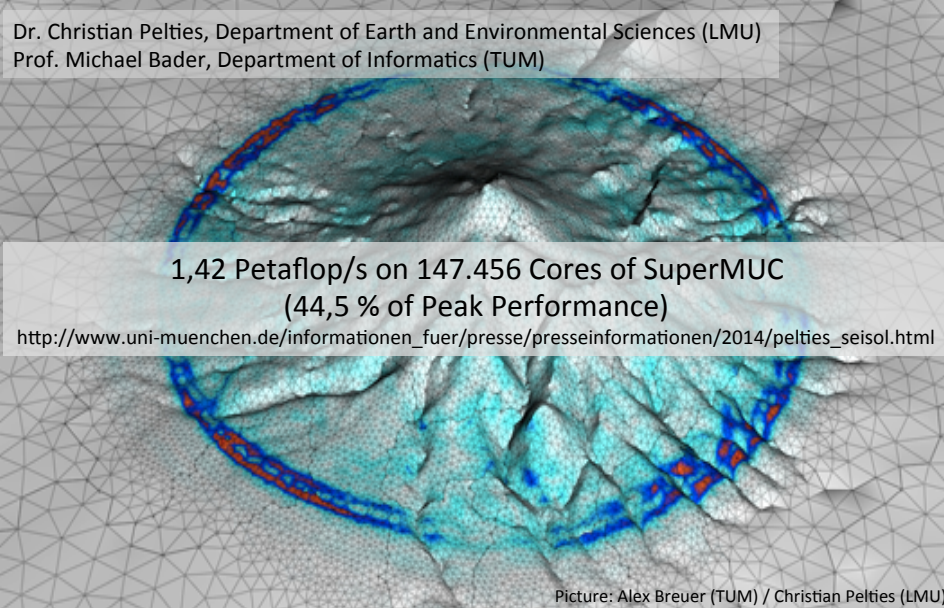


Partnership Initiative
Computational Sciences π CS


- **Individualized services** for selected scientific groups – flagship role
 - Dedicated point-of-contact
 - Individual support and guidance and targeted training & education
 - Planning dependability for use case specific optimized IT infrastructures
 - Early access to latest IT infrastructure (hard- and software) developments and specification of future requirements
 - Access to IT competence network and expertise at Computer Science and Mathematics departments
- **Partner contribution**
 - Embedding IT experts in user groups
 - Joint research projects (including funding)
 - Scientific partnership – joint publications
- **LRZ benefits**
 - Understanding the (current and future) needs and requirements of the respective scientific domain
 - Developing future services for all user groups


 D. Kranzmüller
 Energy Efficiency and Extreme Scaling 27




SeisSol - Numerical Simulation of Seismic Wave Phenomena





Dr. Christian Pelties, Department of Earth and Environmental Sciences (LMU)
 Prof. Michael Bader, Department of Informatics (TUM)



1,42 Petaflop/s on 147.456 Cores of SuperMUC
(44,5 % of Peak Performance)
http://www.uni-muenchen.de/informationen_fuer/presse/presseinformationen/2014/pelties_seisol.html


Picture: Alex Breuer (TUM) / Christian Pelties (LMU)




 D. Kranzmüller
 Energy Efficiency and Extreme Scaling 28



**Partnerschaftsinitiative
Computational Sciences πCS**


Goals for LRZ:

- Thematic focusing – **Environmental Computing**
- Strengthening science through innovative, high performance IT technologies and modern IT infrastructures and IT services
- Interdisciplinary integration (technical and personnel) of scientists and (international) research groups
- Novel requirements and research results at the interface of scientific computing and computer-based sciences
- Increased prospects for attracting research funding through established IT expertise as contribution to application projects
- Outreach and exploitation


 D. Kranzlmüller
 Energy Efficiency and Extreme Scaling 29



Motivations


- Severe storms, and floods/flash-floods are highly impacting on human society and economical activities

Flood		
	2013	2003-2012
Occurrence	145	172
Killed	9 545	5 689
Affected	31 350 240	106 139 202
Damage (US\$ million)	53.90	25.35

Storm		
	2013	2003-2012
Occurrence	95	99
Killed	9 197	17 527
Affected	48 629 303	30 638 289
Damage (US\$ million)	47.84	58.90

Earthquake (incl. Tsunami)		
	2013	2003-2012
Occurrence	27	28
Killed	1 120	67 882
Affected	7 029 162	8 111 667
Damage (US\$ million)	12.01	46.01

Mass movement wet		
	2013	2003-2012
Occurrence	11	18
Killed	235	896
Affected	1 031	344 166
Damage (US\$ million)	—	0.15


Drought		
	2013	2003-2012
Occurrence	9	15
Killed	—	39
Affected	7 955 904	36 407 747
Damage (US\$ million)	1.08	4.93

Wild fires		
	2013	2003-2012
Occurrence	9	10
Killed	34	72
Affected	8 381	211 967
Damage (US\$ million)	—	2.55

Extreme temperature		
	2013	2003-2012
Occurrence	15	25
Killed	2 102	14 435
Affected	270 016	9 011 290
Damage (US\$ million)	—	3.83

Mass movement dry		
	2013	2003-2012
Occurrence	1	1
Killed	46	23
Affected	2	408
Damage (US\$ million)	0.048	—


Volcano		
	2013	2003-2012
Occurrence	3	6
Killed	—	36
Affected	105 106	116 207
Damage (US\$ million)	—	0.015






DRIHM





DRIHM2US











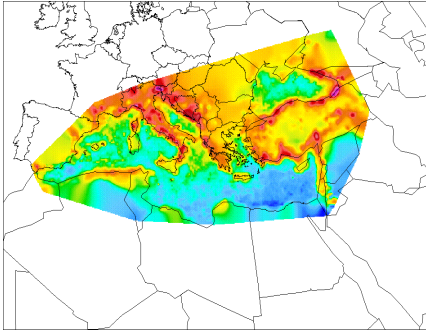
EM-DAT
The International Disaster Database
Centre for Research on the Epidemiology of Disasters - CRED


30

LUDWIG-
MAXIMILIANS-
UNIVERSITÄT
MÜNCHEN

The Mediterranean region

- The FLASH project estimated over 29 billion euros the material damages produced by floods in the Mediterranean region during the 1990-2006 period
- The total number of casualties has been estimated over 4,500, concentrating in the Mediterranean African countries.



SSMI and raingauge observations
1978-1994

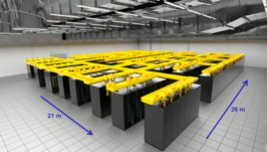
MNM

31


LUDWIG-
MAXIMILIANS-
UNIVERSITÄT
MÜNCHEN

DRIHM Model Chains


HPC



HTC



Cloud Computing



Rainfall

RainFARM

WRF-ARW

WRF-NMM

MesoNH

Arome

Raingauge observations

P-Interface – NetCDF 1.6

DRIfT

RIBS

HBV

Streamflow observation

Q-Interface – WaterML 2.0

Water level Flow & Impact

Mascaret

Delft-3D

RFSM

Property Damage

↓ Gridded data
NetCDF 1.6

↓ Point series data
WaterML 2.0

↓ OpenMI 2

↓ Single

↓ Single/Ensemble

MNM

32

e-Science 2015 in Munich

Ludwig-Maximilians-Universität (LMU), Munich

31 August – 04 September 2015

<http://escience2015.mnm-team.org/>



LUDWIG-
MAXIMILIANS-
UNIVERSITÄT
MÜNCHEN



der Bundeswehr
Universität  München