Deutscher Wetterdienst

Preparing the „Lokal Modell“ for Next Generation Regional Weather Forecasting and Computing

Ulrich Schättler, Elisabeth Krenzien,
Henning Weber

Deutscher Wetterdienst

Contents

- LM and NWP in the last decade
- Next generation regional weather forecasting
- Upcoming events
- Outlook

LM and NWP in the last decade

The „Lokal Modell"

- The nonhydrostatic regional „Lokal Modell" (LM) was developed at DWD in the late 90s
- LM is now used and further developed by the „Consortium for small scale modelling" (COSMO)
- LM is also used by third parties for NWP and other applications (see Outlook)

Meteorological Landscape

CLM Community:

BTU Cottbus
ETH Zürich
FUB Berlin
FZ Karlsruhe
GKSS Geesthacht
Uni Göttingen
Uni Bonn
MPI-M Hamburg (M&D)
MPI-C Mainz
PIK Potsdam
Wegener Center Graz

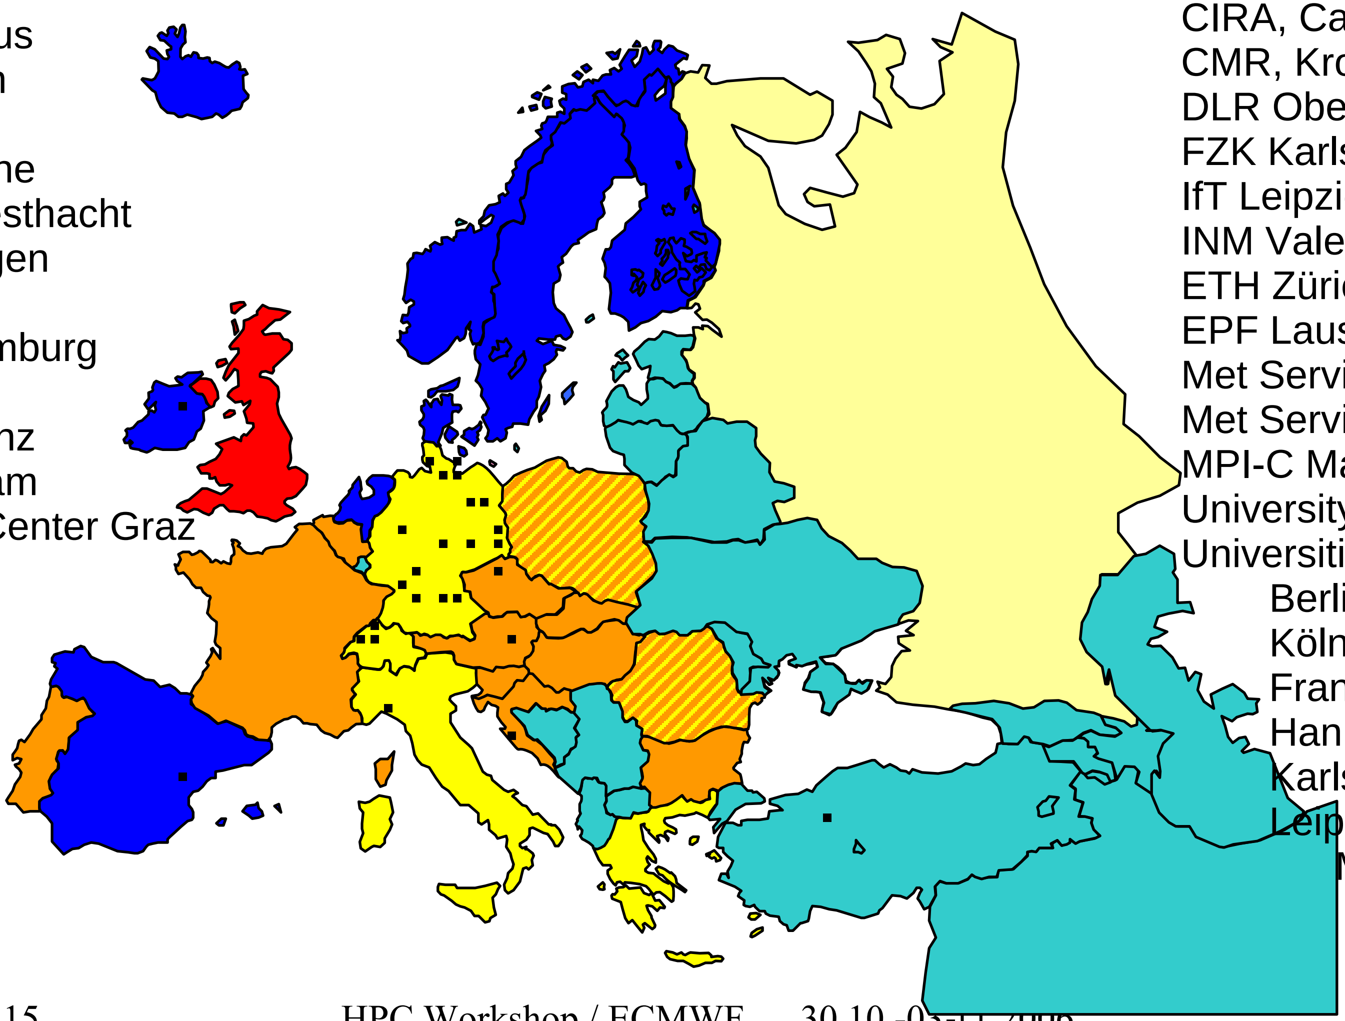

AdW, Prag
AWI Bremerhaven
CIRA, Capua
CMR, Kroatien
DLR Oberpfaffenhofen
FZK Karlsruhe
IfT Leipzig
INM Valencia
ETH Zürich
EPF Lausanne
Met Service Israel
Met Service Turkey
MPI-C Mainz
University College Dublin
Universities:
Berlin, Bern, Bonn, Köln, Dresden, Frankfurt, Genua, Hannover, Hohenheim, Karlsruhe, Kiel, Leipzig, Mainz, München

Performance Development at DWD

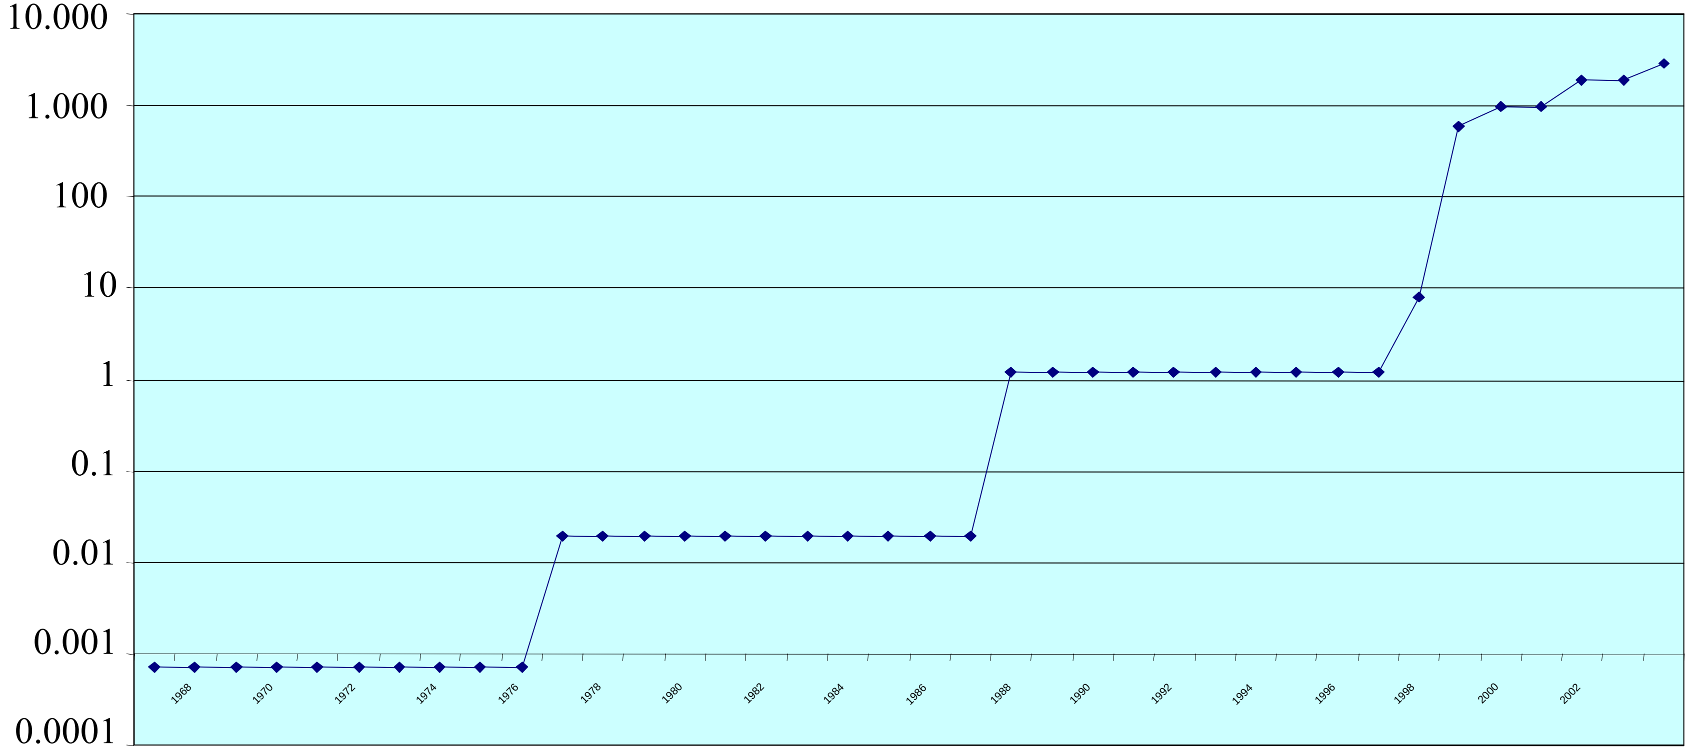

Gflop/s CDC-3800 Cyber 76 Cray YMP/4 C90/8

T3E1200

IBM pwr3/5

Why did we need it

- T3E (1997-2002)
  - LM with 325×325×35 gridpoints, 7 km resolution
  - One of the first operational non-hydrostatic models
  - Further developments, mainly in the physics (prognostic turbulent kinetic energy)
- IBM pwr3 (2002-2005)
  - Introduction of prognostic cloud ice
  - Introduction of prognostic rain and snow
  - Started with the development of a new dynamical core

Prognostic Rain and Snow

without | Observations | with

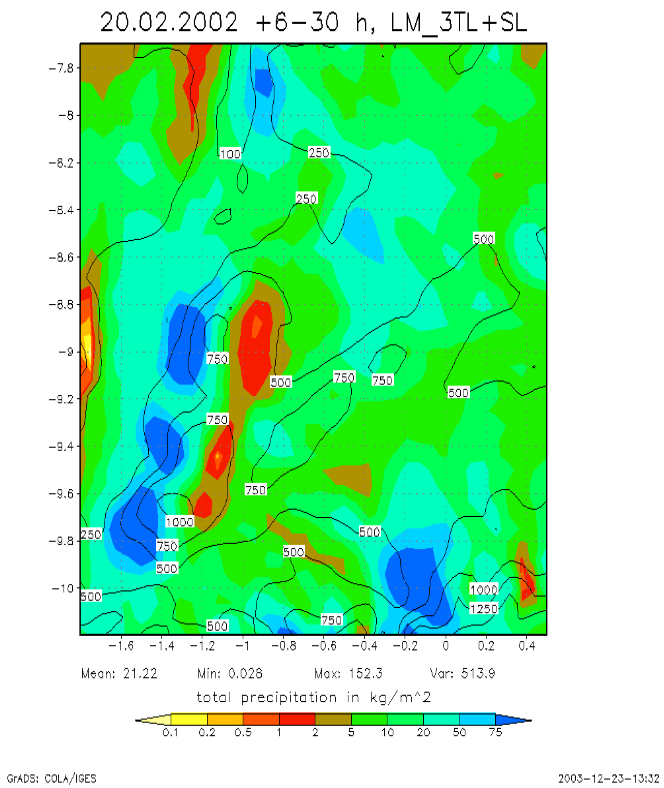

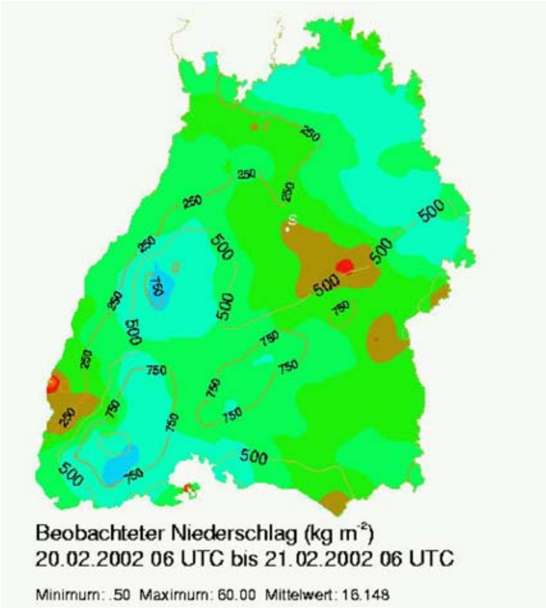

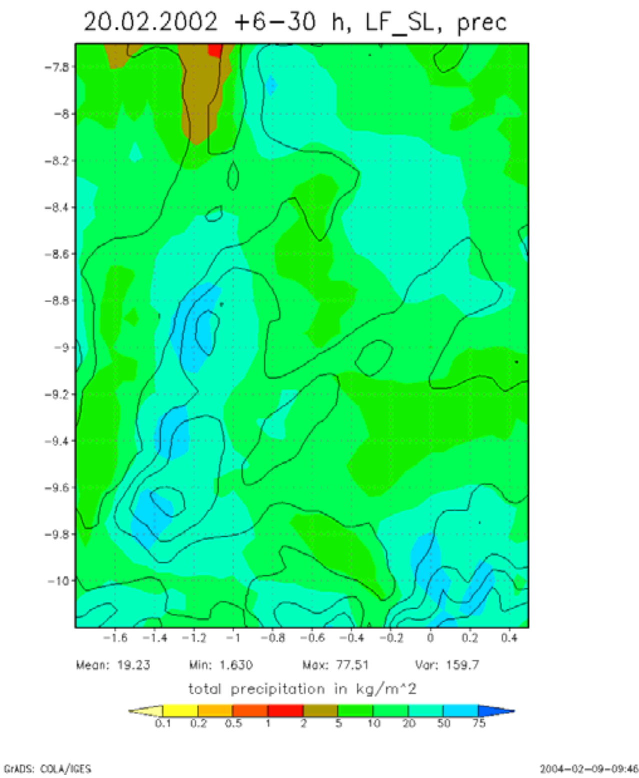

Why did we need it (II)

- IBM pwr5 (2005-2006)
  - LM-E with 665 × 657 × 40 grid points, 7 km resolution
  - Development of a „nowcasting“ version of the LM: LM-K
  - New dynamical core based on Runge-Kutta methods with higher order in space and time
    - For 7 km: dt_{RK}=72s vs. dt_{LF}=40s
    - For 2.8 km: dt_{RK}=30s vs. $dt_{LF}\approx$10s
  - Latent Heat Nudging of Radar Data
  - Graupel scheme, Lake model, shallow convection

Precipitation

LM | RADAR | LM-K Testsuite

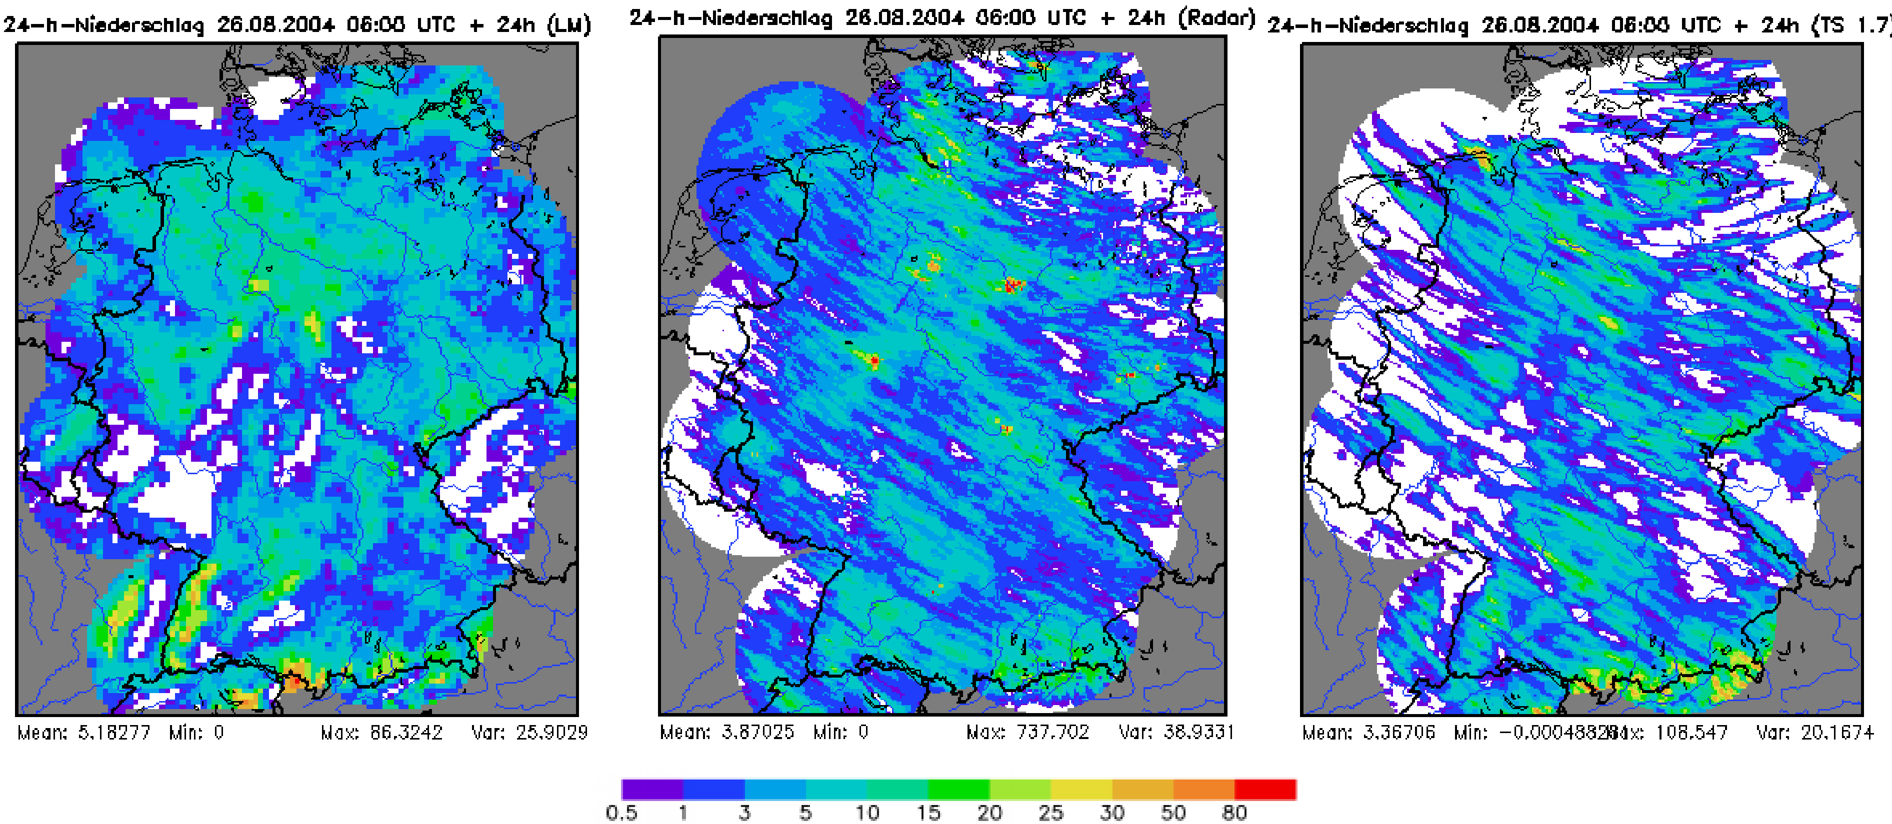

Why did we need it (III)

- 2 × IBM pwr5 (from 2006 on)
  - Divided operations and development
  - Pre-operational LM-K runs:
    - 421 × 461 × 50 grid points, 2.8 km resolution
    - Running 8 times a day for 21 hours
  - And also development work has increased

Historical Summary: Additional compute performance was used for bigger domains, higher resolution and more expensive algorithms

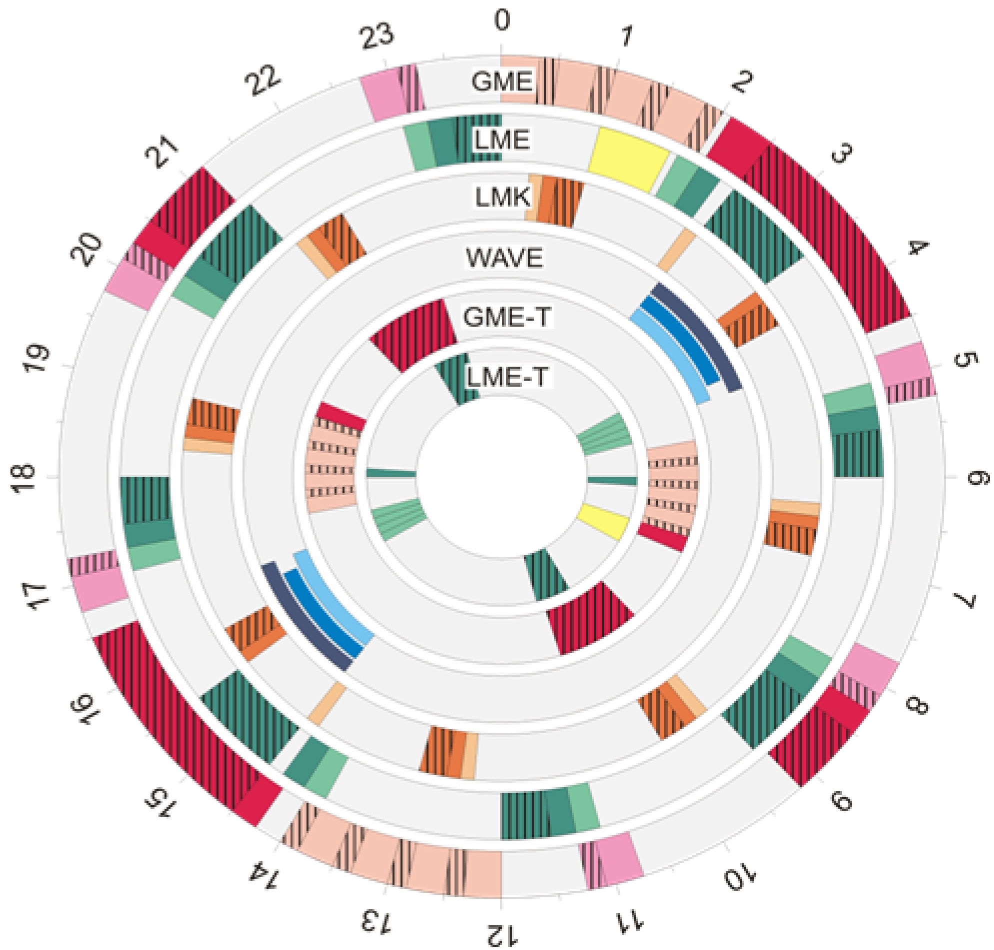

Scheduling of operational Jobs

Next generation regional weather forecasting

Model Development in the next years

- Higher resolution:
  - Possible; but uncertain
  - Perhaps need totally different physics
- Bigger domains
  - possible; what does the customer need?
- More sophisticated algorithms
  - Definitely
- And: Ensemble Prediction Systems

Why Ensemble Predictions?

- Weather predictions are NOT determinstic!
- Small fluctuations in the initial conditions lead to large differences in the prediction
- An ensemble of many predictions with different initial conditions gives us the probability for a certain weather situation

Example: Temperature „plume“ for Mainz

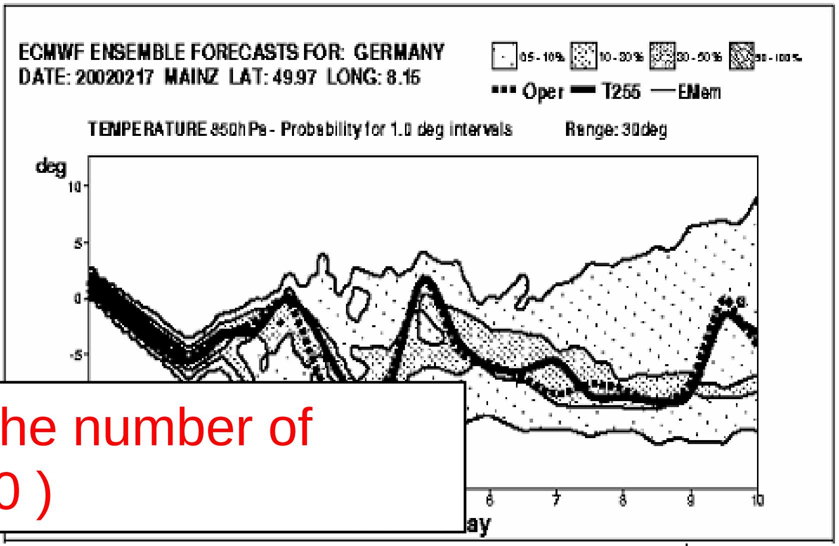

but: computing requirements scale with the number of ensemble members! (factor of 20-40)

Development of a regional EPS

- Work has started in COSMO (there is some experience by COSMO-LEPS)
- Scientific question: how to make the members different
  - Disturbing initial and boundary conditions
  - Changing the physics
  - Using different external parameters
- How many members do we need?

Development of a regional EPS (II)

- We definitely need more computing power
- And the boring part of the story:
  - This is embarrassingly parallel
  - The program paradigm used (MPI) is still sufficient
  - We do not go towards a Peta-Flop computer (at least not in the near future)

Upcoming Events

Replacement of Computing Center

- To run 20-40 members of a LMK ensemble, we aim at a performance enhancement by a factor of 15, giving about 100 Tflops peak
- An invitation to tender will be issued later this year
- The new machines will be put in a new building
- There will be 2 phases for the procurement (with performance(phase2) $\geq$ performance(phase1))

Compute Server Requirements

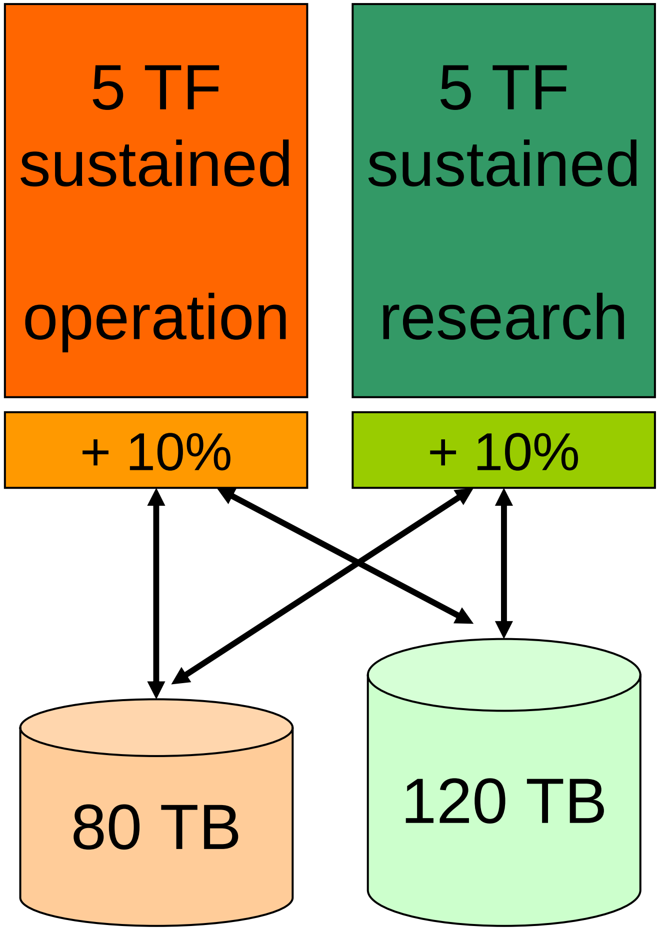

- Two independent compute servers (one for operations, one for research)
- Research machine is backup for the operational server
- Each 5 Teraflop/s sustained performance (measured with LM-RAPS benchmark)
- Additional 10% of peak for serial and interactive applications
- 200 TB storage (I/O: 2 GB/s in, 2 GB/s out)

Availability of the operational server:
at least 99.8% per month (1,5h downtime)

Infrastructure: DWD's new building at Offenbach 2004-2008

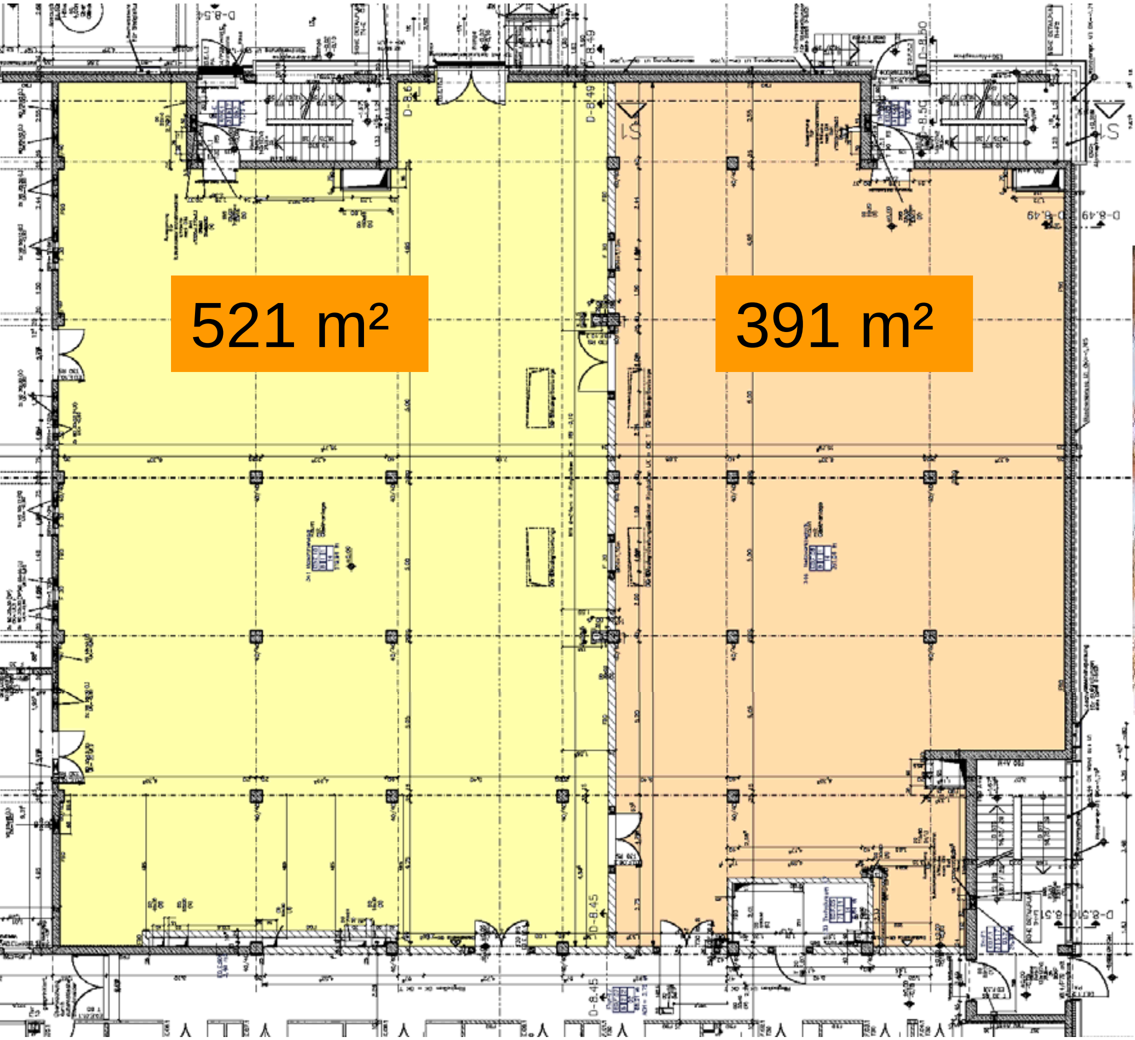

with new central

computer facilities

Peak TFlops at DWD

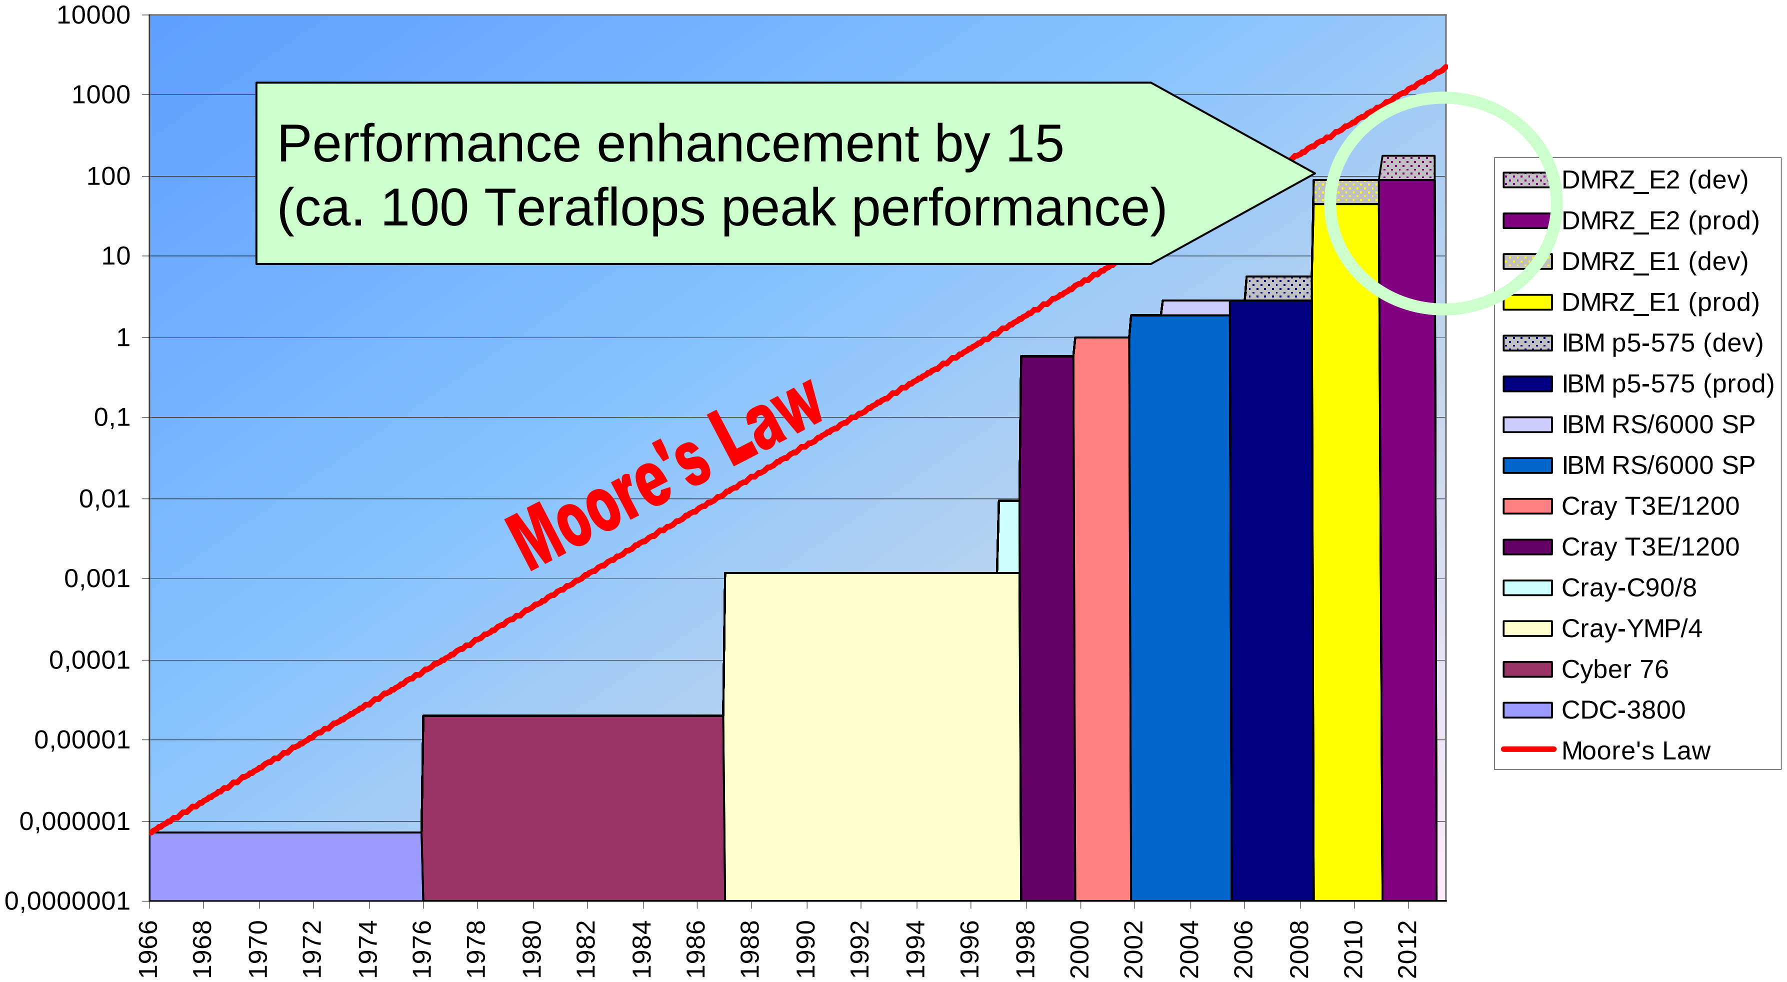

LM_RAPS_4.0

- A new LM_RAPS distribution will be used to measure the performance
- New features of LM_RAPS_4.0
  - Consolidated Runge-Kutta dynamical core
  - Modifications to physical parameterizations
  - Bigger amount of data for output (some variables written every 15 minutes)
- LM-K will be the main test job

LM_RAPS_4.0 – First Results

LM-K 421 × 461 × 50 grid points; 6 h forecast on pwr5 (DWD), pwr5+ (ECMWF), NEC SX6

# Nodes	4	8 4		16 8		32 16	
# (log.) CPUs	1×4	8×8	8×16	8×16	16×16	16×16	16×32
Total Time	3805.06	1556.76 1624.83	1099.94 1244.53	783.24 889.04	589.04 699.87	494.49	405.10
Computations	3432.55	1096.06 1181.63	878.78 1011.24	523.80 566.11	418.97 489.07	292.75	234.15
Communi-cation	304.78	411.46 406.31	163.06 181.81	218.06 292.61	109.22 164.17	170.10	116.07
I/O	37.03 only Inp.	41.38 33.20	50.06 45.37	34.98 26.31	50.92 40.49	27.84	48.43

LM_RAPS_4.0 – Some Flop/s

LM-K 421 × 461 × 50 grid points; 6 h forecast on pwr5 (DWD), pwr5+ (ECMWF), NEC SX6

# Nodes	4	8 4		16 8		32 16	
# (log.) CPUs	1×4	8×8	8×16	8×16	16×16	16×16	16×32
Flop (10^{12})	45.9	46.9 45.6	47.7 46.2	47.8 46.5	48.7 47.4	47.7	49.5
Flop per grid point and step	6581.45	6723.19 6527.39	6827.65 6624.65	6842.67 6664.70	6972.29 6792.09	6835.44	7087.05
GFlop/s	12.3	30.1 28.0	43.3 37.2	60.9 52.3	82.5 67.7	96.4	121.8

Procurement Schedule

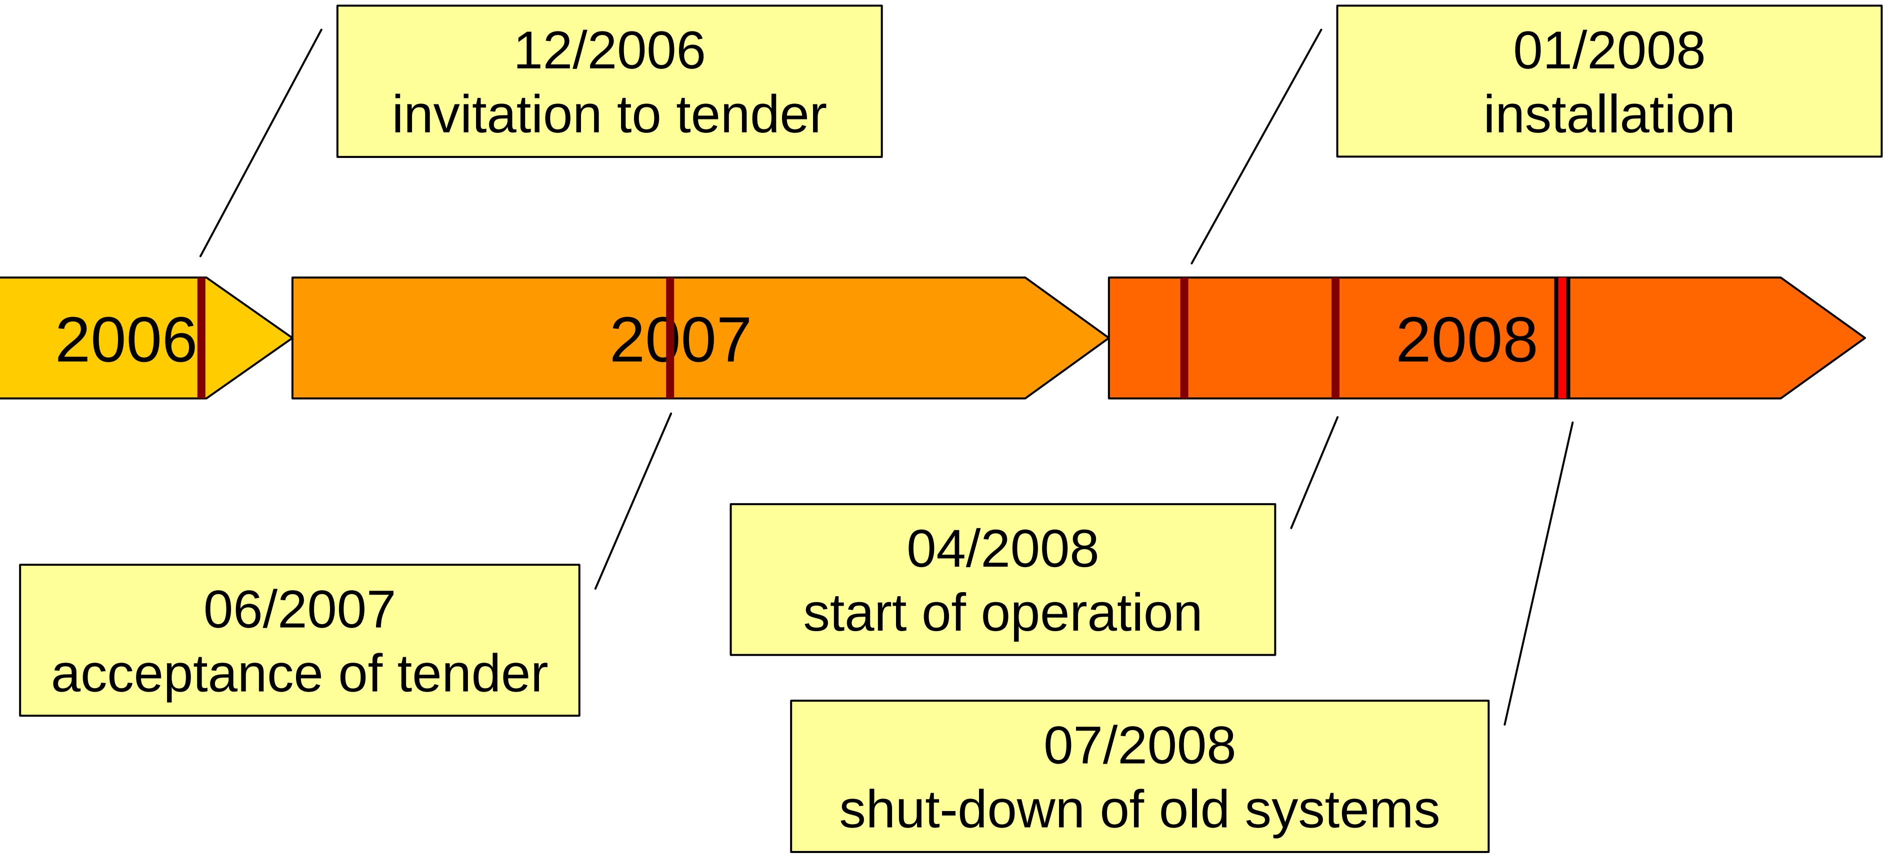

Outlook

Other LM Activities

- DWD will not need a petaflop computer in the near future, but:
  - CLM: LM is also used as a regional climate model (see: http://www.clm-community.eu)
  - For example: 150 years for several (!) climate scenarios. One run takes about 300 days on a NEC SX6 node
  - LM and chemistry: The Karlsruhe Research Institute will provide a chemistry model (LM_ART) for online coupling to the LM. This adds about 100 additional prognostic scalar variables (up to now: 5 dynamical and 6 scalar variables)

More Activities

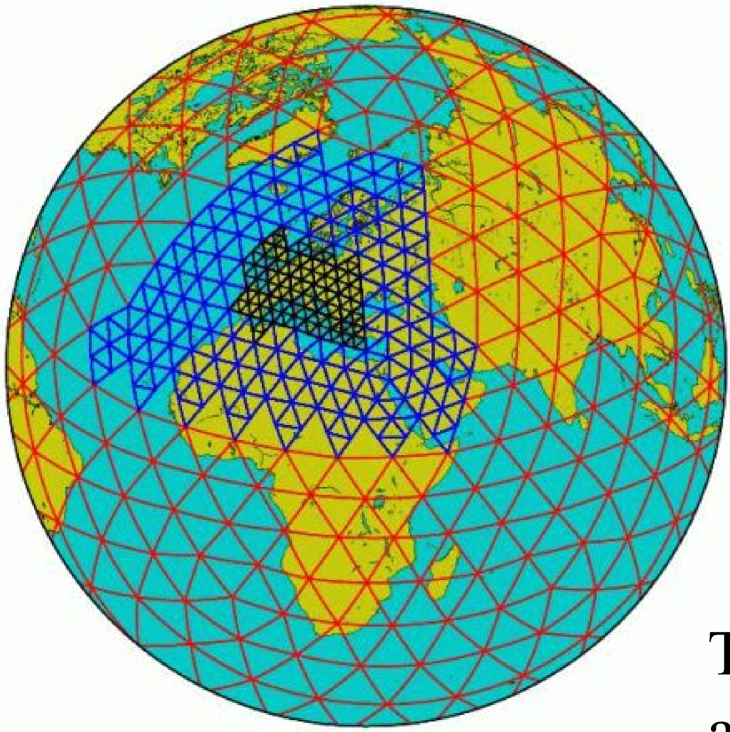

ICON (ICOsahedral Nonhydrostatic model development of DWD and MPI-Met in Hamburg) with local zooming option

There will be “petaflop applications” also in NWP in some years time.

Are we ready then?

Thank You