50 Years of NWP at DWD

and what to expect in the future?

Ulrich Schättler

Department for Numerical Modeling (FE1)

With contributions from Detlev Majewski, Bodo Ritter, Florian Prill, Harald Anlauf (FE1), Elisabeth Krenzien, Henning Weber (TI15) and many more!

Contents

50 Years of NWP at DWD

Preparations for NWP

All quotations from: H. Reiser: Development of NWP in the Deutscher Wetterdienst, in: "50th Anniversary of Numerical Weather Prediction", Commemorative Symposium, 9-10 March 2000, Book of Lectures.

- ➔ DWD experience in NWP dates back to "the early fifties with manual integration of a quasigeostrophic model".
- ➔ "A group of scientists with K. Hinkelmann as a spiritus rector began around 1950 to consider methods to improve synoptic forecasting practice by numerical integration (…)"
- ➔ "Following the ideas of Rossby, Charney and others (…) this group (…) designed a baroclinic multilevel quasigeostrophic model for manual integration and first experiments with actual weather data were performed in 1952/53".

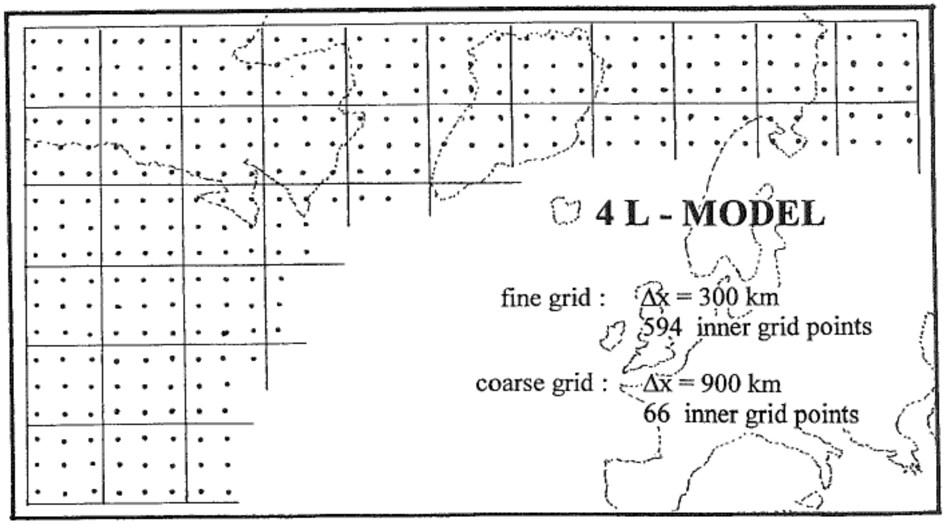

Necessary Hardware

- "(…) these activities were effectively supported by the Air Research and Development Command of the US-Air Force; these projects enabled the employment of additional staff, the acquisition of more electro-mechanical desk calculators and later on renting of machine time".
- "For the purpose of relaxation the inhomogeneous terms (…) were recorded as grid point values in the order (k,i,j) on a large transparent paper (…) of a special relaxation table and serving as a main storage."
- No Flop/s were recorded.

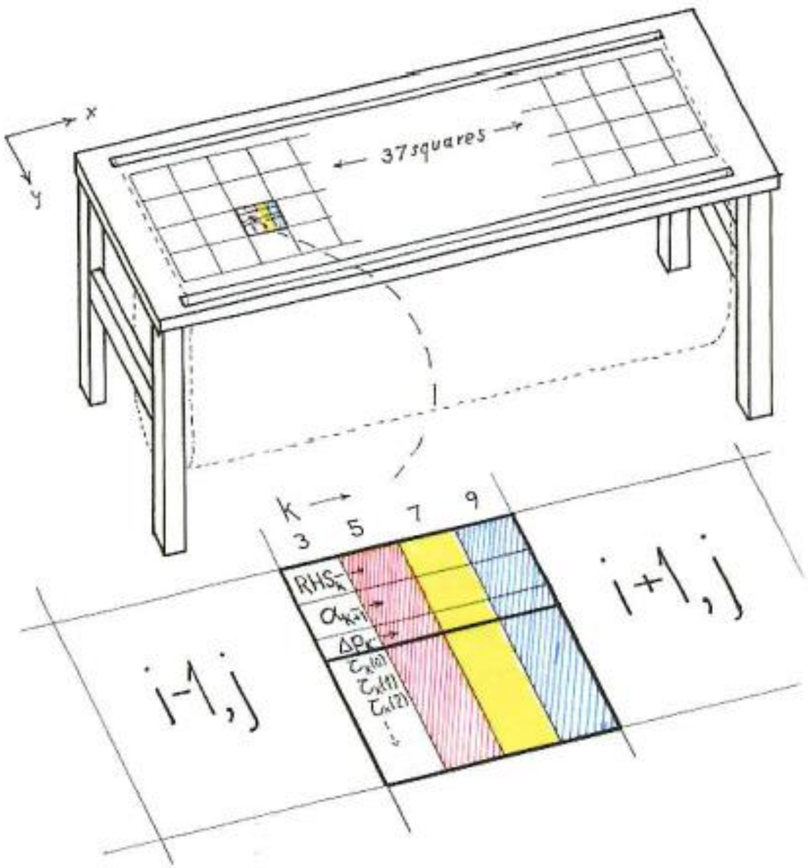

Computers 1966-2016

Deutscher Wetterdienst
Wetter und Klima aus einer Hand

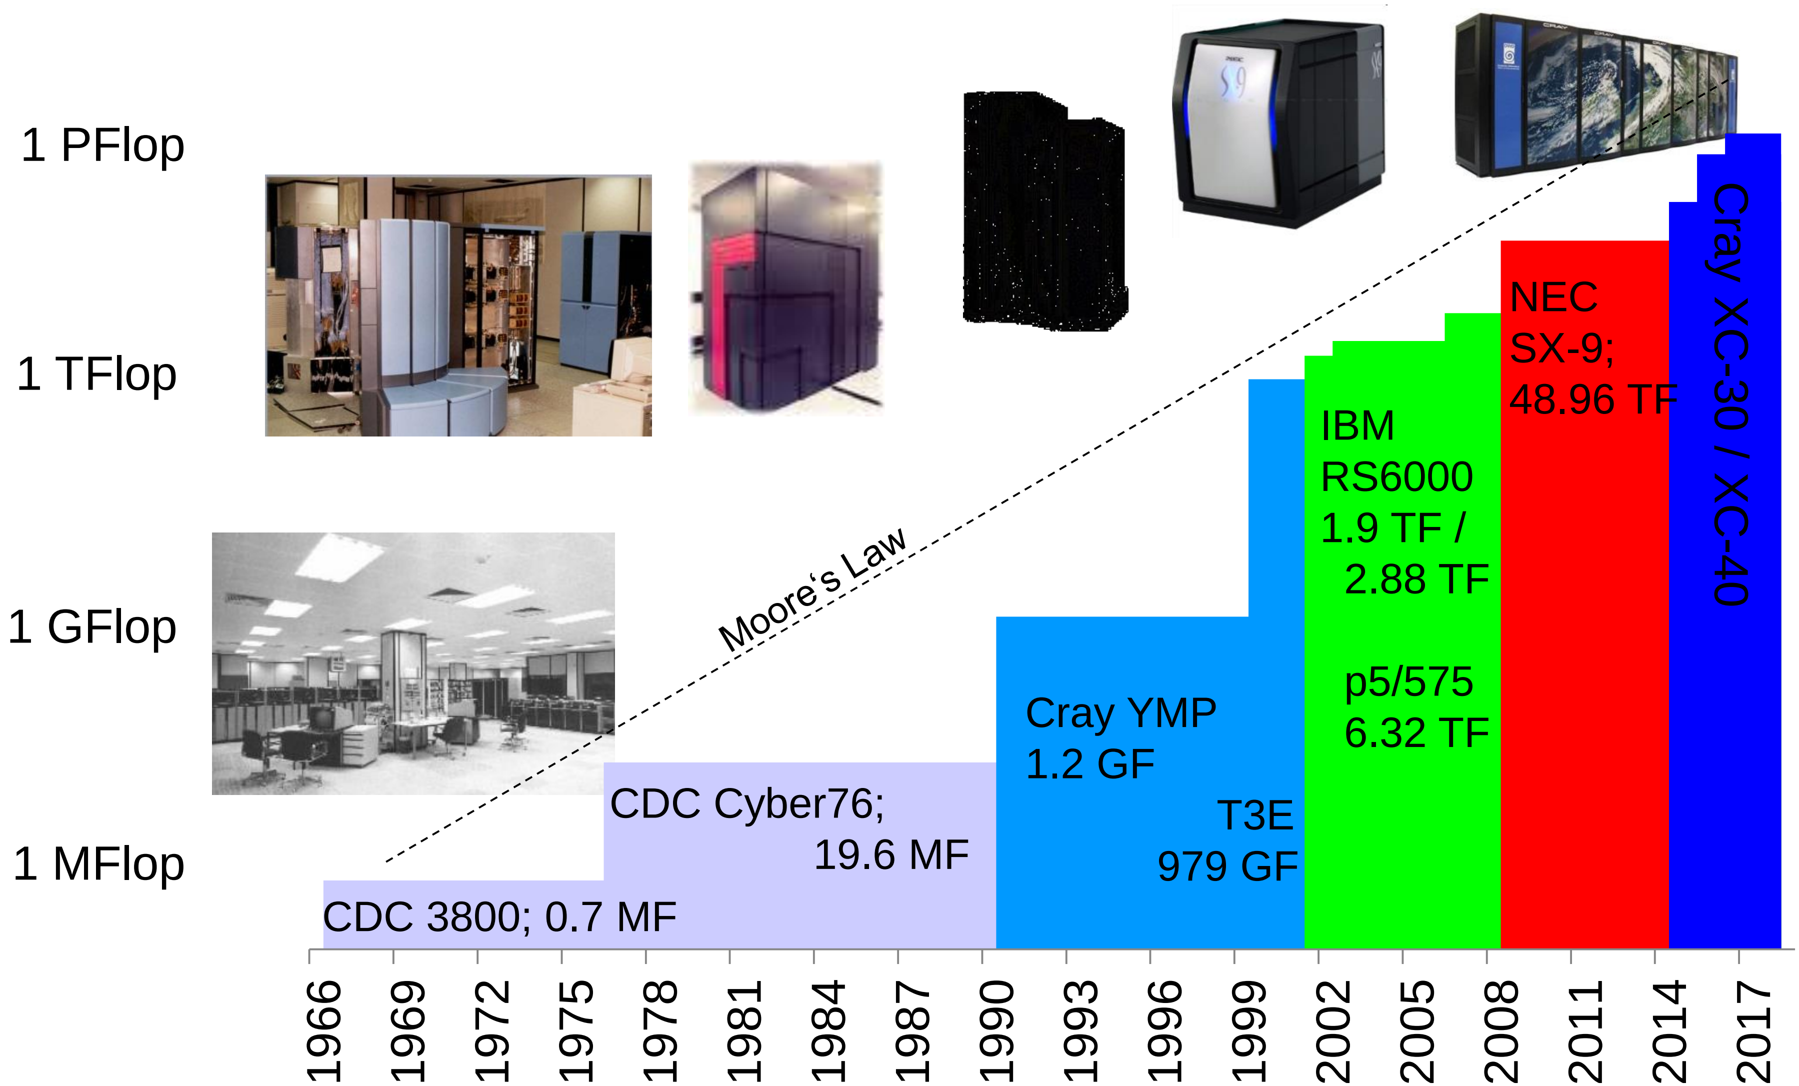

Cray XC-30 / XC-40

- DWD and Cray extended their contract for 2 more years until end of 2018 (with the option of extending it for another year).
- In September 2016 Phase 2 was installed:

Components per cluster	Phase 0	Phase 1	Phase 2
Cabinets, chassis, blades	2 / 6 / 96	5 / 13 / 208	6 / 16 / 256
Max. power consumption (kW)	148	325	407
	IVB	IVB / HSW	HSW / BDW
Compute nodes	364	364 / 432	432 / 544
Cores per node	20	20 / 24	24 / 36
Cores	7280	7280 / 10368	10368 / 19584
Memory per node (GB)	64	64 / 128	128 / 128
Performance peak (TF)	146	146 / 415	415 / 660

NWP generations 1966-2016

Year	Model	Δ	Area	Layers	
1966	BTP	381	145.161	1	barotropic
1967	BKL	381	145.161	5	baroclinic, hemispheric model, dry
1978	BKF	254	64.516	9	barocl., hemispheric, moist
1991	GM	190	36.100	19	global spectral model
1999	GME	60	3.114	31	icosahedral hexagonal grid
2004	GME	40	1.384	40	
2010	GME	30	778	60	
2012	GME	20	346	60	
2015	ICON	13	173	90	non-hydrostatic; triangular grid

1^{st} and 9^{th} Generation: 1966 vs. 2016

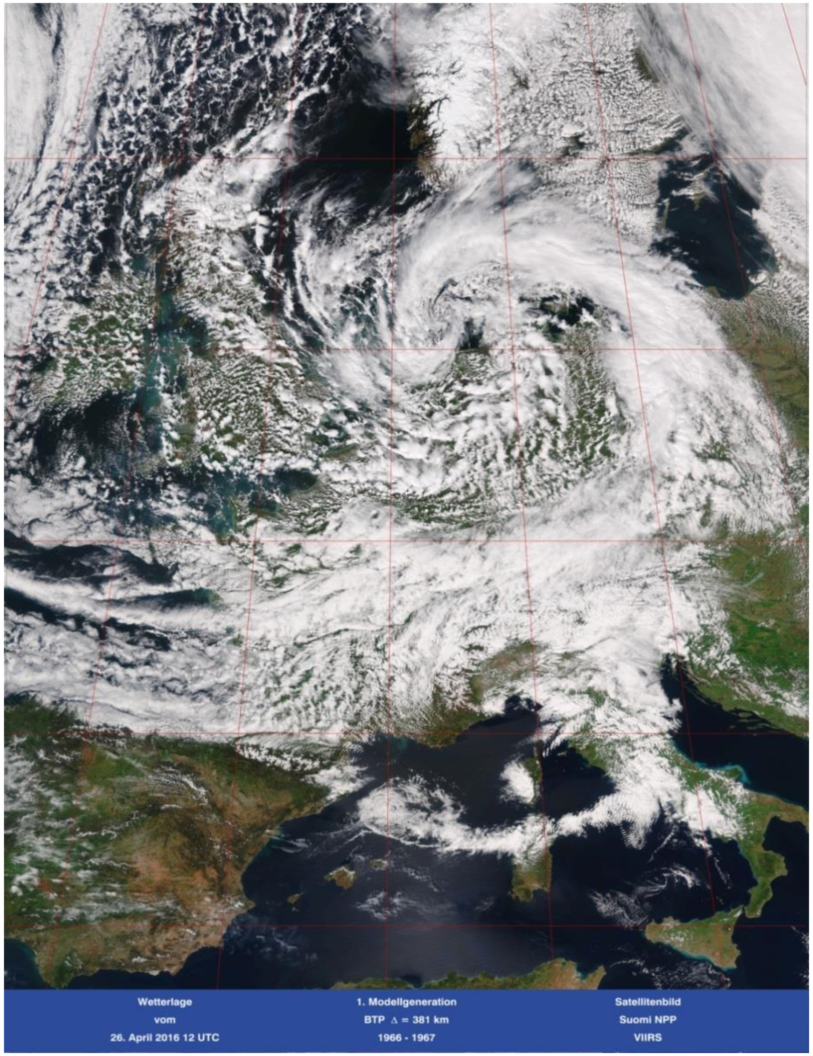

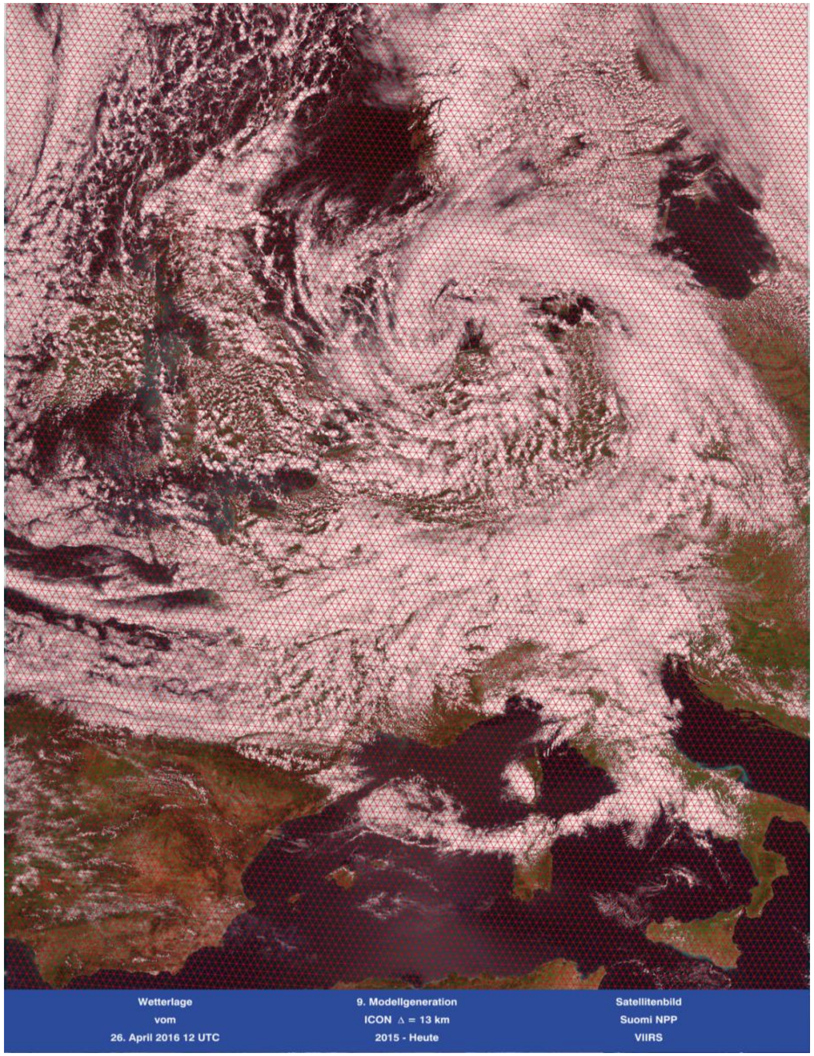

~ 3 - 4 grid points for Germany, 1 layer

~ 2000 grid points for Germany, 90 layers

Regional Models 1991-2016

Deutscher Wetterdienst
Wetter und Klima aus einer Hand

Year	Model	Δ	Area	
1991	Europa	55	Europe	hydrostatic
1993	Deutschland	14	Germany	in addition to EM
1999	COSMO	7	Central Europe	non-hydrostatic
2002	COSMO-EU	7	Europe	
2007	COSMO-DE	2.8	Germany	in addition to COSMO-EU
2011	COSMO-DE-EPS	2.8	Germany	convection permitting ensemble system; 20 members

The NWP generations are also characterized by increasing complexity of the physical parameterizations, numerical methods and software design.

Current Models and Future Plans

Current Models

Deutscher Wetterdienst
Wetter und Klima aus einer Hand

Since	Model / Analysis	Δ	Start Times	Comments
01/2015	deterministic ICON	13	00/12 +180h	90 layers
			06/18 +120h	
			03/09/15/21 + 30h	
01/2016	Ensemble variational (EnVar) for high res. analysis, LETKF for ensemble	13 40	3-hourly update	40 members
07/2016	EU-Nest in ICON	6.5		60 layers

COSMO-EU has been replaced by ICON-EU Nest and will be switched off.

Plans to extend the area and resolution for COSMO-DE and COSMO-DE-EPS have been postponed due to the considerable changes related to ICON.

Updates were also necessary for the database and archive server.

WMO TEMP-Verification, Europe, +24/72

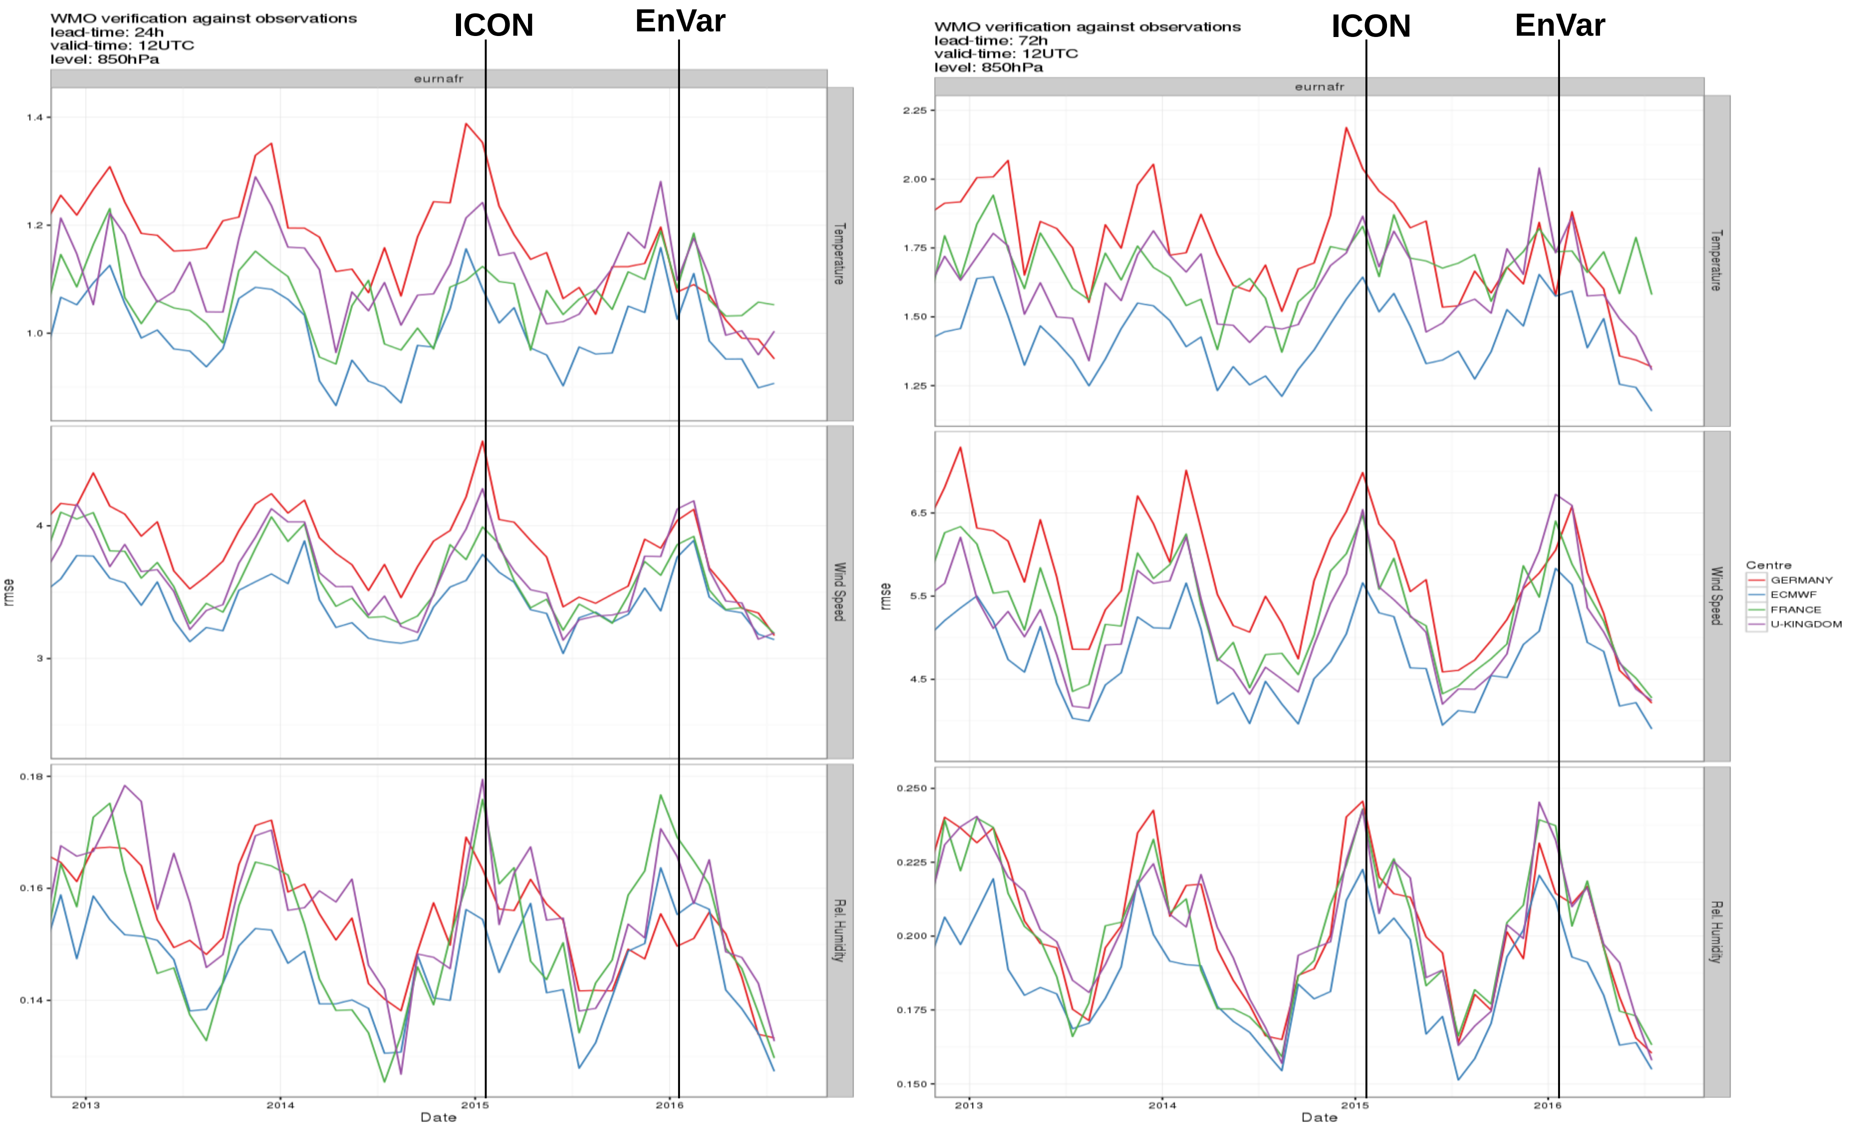

Future Plans for Cray Phase 2

- ➔ Ensemble data assimilation (LETKF) for COSMO-DE
- ➔ COSMO-D2: extended domain: 651×761 grid points, 2.2 km; 65 layers
- ➔ COSMO-D2-EPS with 40 members

- ➔ ICON-EPS: 40 km, 80 members (EU-Nest with 28 km)

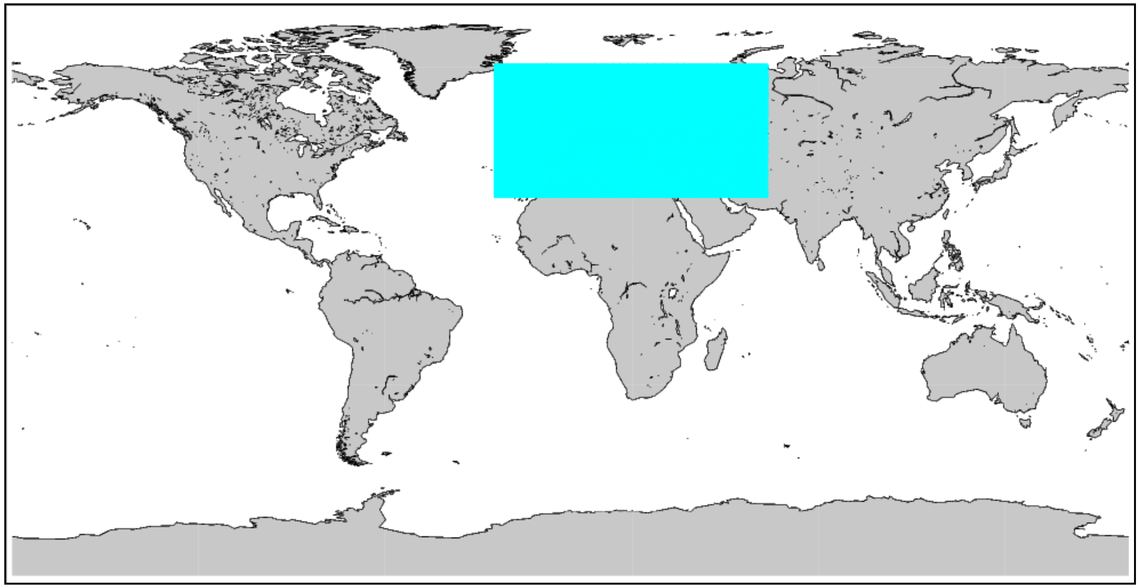

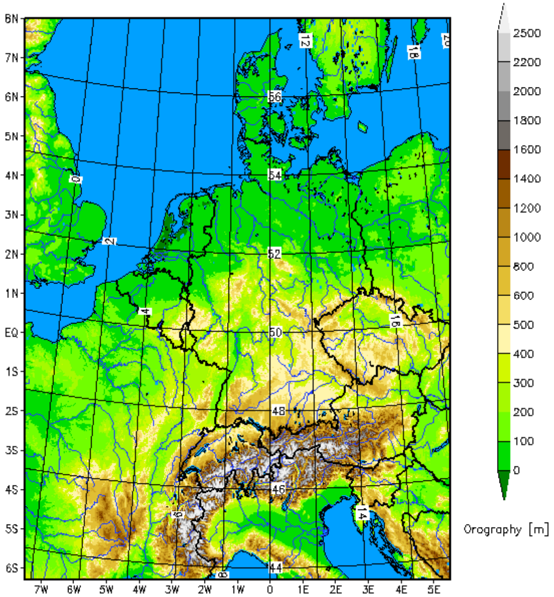

Future Development of Convective Scale NWP

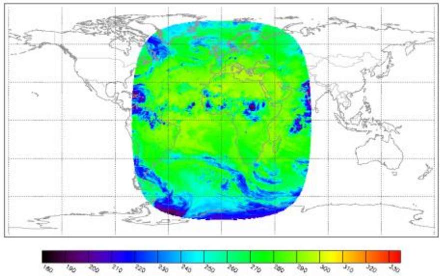

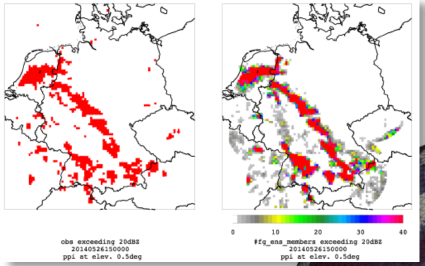

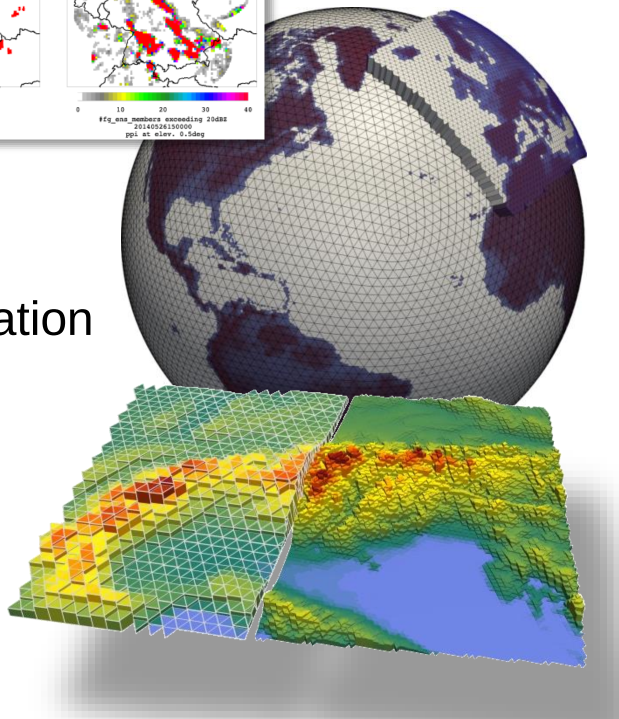

- ➔ Assimilation of new observations
  - ➔ Radar volume data, consistence with LHN
  - ➔ SEVIRI Infrared and SEVIRI Near-VIS
  - ➔ Lightning (LPI – lightning potential index)
- ➔ Model physics / dynamics
  - ➔ E.g. two-moment cloud microphysics
  - ➔ Interaction of physics and dynamics
  - ➔ Model grid spacing ~1km (or less)
- ➔ Development of new EPS-based products and verification
  - ➔ Detection of HIW, precipitation, gusts, turbulence
  - ➔ Consistency between global and regional EPS
- ➔ Seamless prediction with a focus on shortest range

Developments in Software Design

- Unification of ICON and COSMO physical parameterizations (is nearly ready!)
  - Both models use the same schemes for cloud microphysics, turbulence, SSO and the soil and surface, but in different implementations
  - Because ICON only uses a one-dimensional vector to store horizontal fields, we had to change the data structure in COSMO for the parameterizations.
  - A "copy-in/copy-out" mechanism has been implemented to transform all necessary fields between the parameterizations and the rest of the model (which still is ijk-structure)

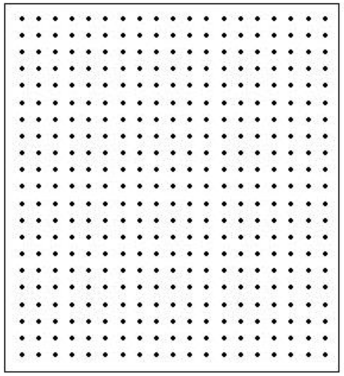

(i,j,k) data format

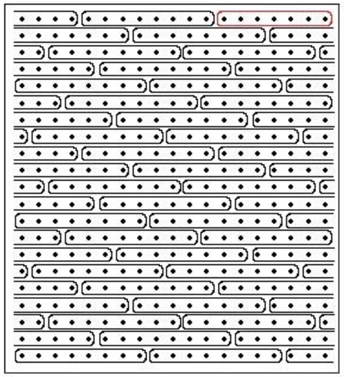

(nproma,k) data format

Replacement of COSMO-Model

- In the long run, COSMO will be replaced by a limited-area-mode of ICON (ICON-LAM).
- Will have to maintain only one code in the future.
- Tentative schedule for COSMO to ICON-LAM migration
  - 2017/2018: Testing in the framework of a COSMO Priority Project.
  - 2019/2020: Test of EDA for ICON-LAM
  - 2021 and beyond: Gradual replacement of the COSMO-Model by ICON-LAM by COSMO partners and licensees

- ICON will be the main tool for NWP at DWD in the future.

Further Hardware

- DWD plans to launch the next ITT in 2018 to replace the current hardware in 2019. Preparations have already started. (*).
- Future plans for the operational suite from 2020 on (which resolutions, ensemble sizes) are still under discussion and depend on the available money for new hardware.
- Besides money, also DWD now experiences the limitations of power and performance.
- How can new architectures help?

(*) Information subject to change without further notice!

Testing Novel Architectures

Testbeds at DWD

- For GP GPUs:
  - DWD purchased cluster licence for PGI compiler to use on Linux workstations.
  - Available GPU is NVS 315 with 1 GByte of memory and 32 cores („Flexible and energy efficient low profile solution“).
  - Which is technically working!
- For Intel KNL:
  - Also included in Phase 2 of the Cray contract is a Cray CS400 cluster (1 rack, 3 chassis) with 12 Intel KNL- and 1 login-node.
  - Available since beginning of October.
  - Login node: Intel® Xeon® CPU E5-2690 v3: Haswell.
  - KNL nodes: Intel® Xeon Phi™ CPU 7230 with 64 cores.

- Would also be possible to include GPUs in this cluster.

COSMO-Model on GP GPUs

Deutscher Wetterdienst
Wetter und Klima aus einer Hand

Porting Strategy

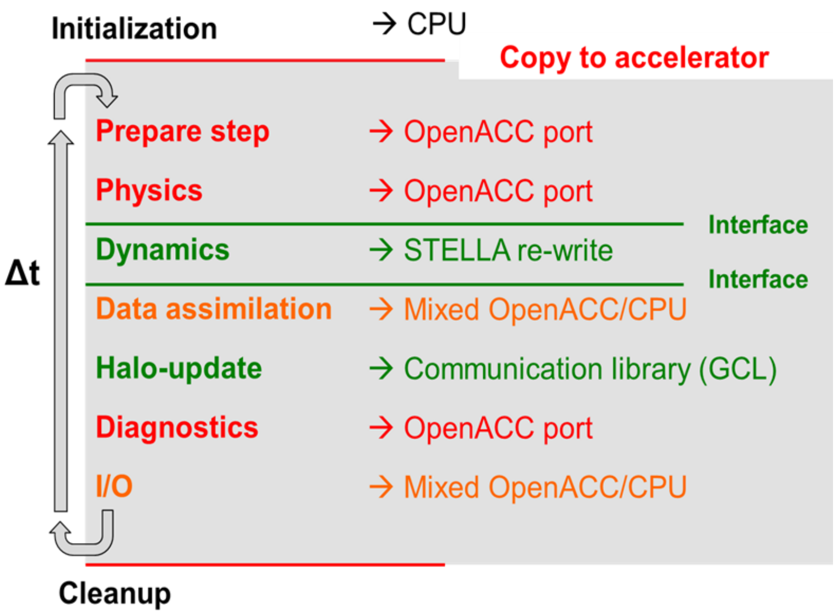

- MeteoSwiss already ported full COSMO-Model to GPUs
- End of March 2016 they started operational runs with this version (which is based on COSMO-Model 4.19, now we have 5.03 with several significant changes)
- Process has started to implement GPU changes to the official COSMO-Model version
- The future of the STELLA (new code name: gridTools) re-write is not clear yet.
- Do not miss the presentation by Carlos Osuna on Thursday!

My First Steps on a GPU

Deutscher Wetterdienst
Wetter und Klima aus einer Hand

- Task: Implement the radiation interface between ijk- and blocked data structure and compute necessary input for radiation scheme
- The routines from the radiation scheme have been ported by Xavier Lapillonne from MCH
- Besides porting the loops (see right), you have to get all

```
!$acc data create
!$acc copyin
!$acc update device / host
!$acc delete
```

correct.

```
! Temperatures at layer boundaries
!$acc parallel
!$acc loop gang vector collapse(3)
DO  k = 2, ke
 DO jp = 1, nradcoarse
  DO ip = 1, ipdim
   !get ij indices for blocked structure
   i = mind_ilon_rad(ip,jp,ib)
   j = mind_jlat_rad(ip,jp,ib)
       …
   zti(ip,k,jp) =                          &
   (t(i,j,k-1,ntl)*zphfo*(zphf - zpnf )&
  + t(i,j,k  ,ntl)*zphf *(zpnf - zphfo))&
    *  (1.0_wp/(zpnf *(zphf - zphfo)))
  ENDDO
 ENDDO
ENDDO
!$acc end parallel
```

And after a few trials and errors: It worked ☺

My First Steps on a GPU (II)

Deutscher Wetterdienst
Wetter und Klima aus einer Hand

Scheme	CPU	GPU	GPU	GPU	GPU
nproma	16	16	32	128	1024
Total Time	15.46	132.26	36.12	21.43	18.24
Radiation	1.82	107.24	18.68	5.72	3.10
Update Device / Host	-	1.59	1.59	1.63	1.62

Tests with a bigger domain showed the same behaviour, but bigger `nproma` then lead (on my "low profile" GPU) to

```
Out of memory allocating 7045760 bytes of device memory
Failing in Thread:1
total/free CUDA memory: 1068171264/6311936
```

Very First Tests on KNL

Deutscher Wetterdienst
Wetter und Klima aus einer Hand

- COSMO still is a flat MPI implementation, which is considered to be not beneficial on KNL.
- The new implementation of the COSMO-ICON physics with a blocked data structure allows for a fast OpenMP parallelization.
- Tried to investigate the difference between flat MPI and hybrid MPI-OpenMP
- But be careful: these are really the first tests!

```
DO ib = 1, nblocks
  ! do the copy-in
  (...)

  CALL radiation_interface

  CALL turb_interface

  CALL conv_interface

  CALL soil interface

  ! do the copy-out
  (...)
ENDDO
```

Timings for the Physics

Deutscher Wetterdienst
Wetter und Klima aus einer Hand

- There were two problems to be solved to make reasonable tests:
  1. OpenMP parallelization of the block-loop over physical parameterizations
  2. Find a correct way for cpu- / memory-binding and thread-affinity.
- The following timings were obtained with
  - I_MPI_PIN_PROCESSOR_LIST allcores
  - KMP_AFFINITY=verbose,compact

- Compiler used: Intel 2017.0.014 with –O3 –xMIC-AVX512
- COSMO domain size: 201×201×40: should fit into 16 GB MVDRAM
- No questions are taken (I myself have more than I can answer right now!)

MPI tasks / Threads	**8x8 / 1**	**8x8 / 2**	**8x8 / 4**
Radiation	2.46	2.33	2.21
Turbulence	8.08	6.68	3.54
Convection	2.06	2.03	1.22
Soil-Model	1.44	0.86	1.90

Conclusions

- 50 years and still have to learn new tricks.
- Changes take time and need preparations!
- Testbeds are now available at DWD to test new architectures and programming models.
- It is not yet clear to us, how NWP codes will look like in the future.
- Forecasts are always difficult, but most probably our next computer will not be a pure GPU or KNL machine.

- Regarding software development, we want to maintain a single source code for every model / application.


Thank you
very much
for your
attention