Towards Global Cloud-Resolving Weather and Climate Prediction at Exascale

Philipp Neumann, Niklas Röber, Joachim Biercamp
Deutsches Klimarechenzentrum (DKRZ)

Luis Kornblueh, Matthias Brück
Max Planck Institute for Meteorology

Daniel Klocke
Deutscher Wetterdienst

Tornados, Disaster & Early Warnings

1. **Weather Extremes: Example Katrina**
   → 1833 deaths, $125 billion damage
   → **accurate prediction via high resolution models**
2. **Climate Change**: Serious threat through by anthropogenic emissions (greenhouse gases)
   → quantitative estimates for change in climate extremes:
   wind storms, flash floods, droughts, …
   → **demand for high resolution modeling** to understand scientific keys such as cloud feedbacks and convective organization in climate[1,2]

1 Bony et al. Nature Geoscience 8:261-268, 2015
2 Schneider et al. Nature Climate Change 7:3-5, 2017

Tornados, Disaster & Early Warnings

1. **We**
2. **Cli**
   ant

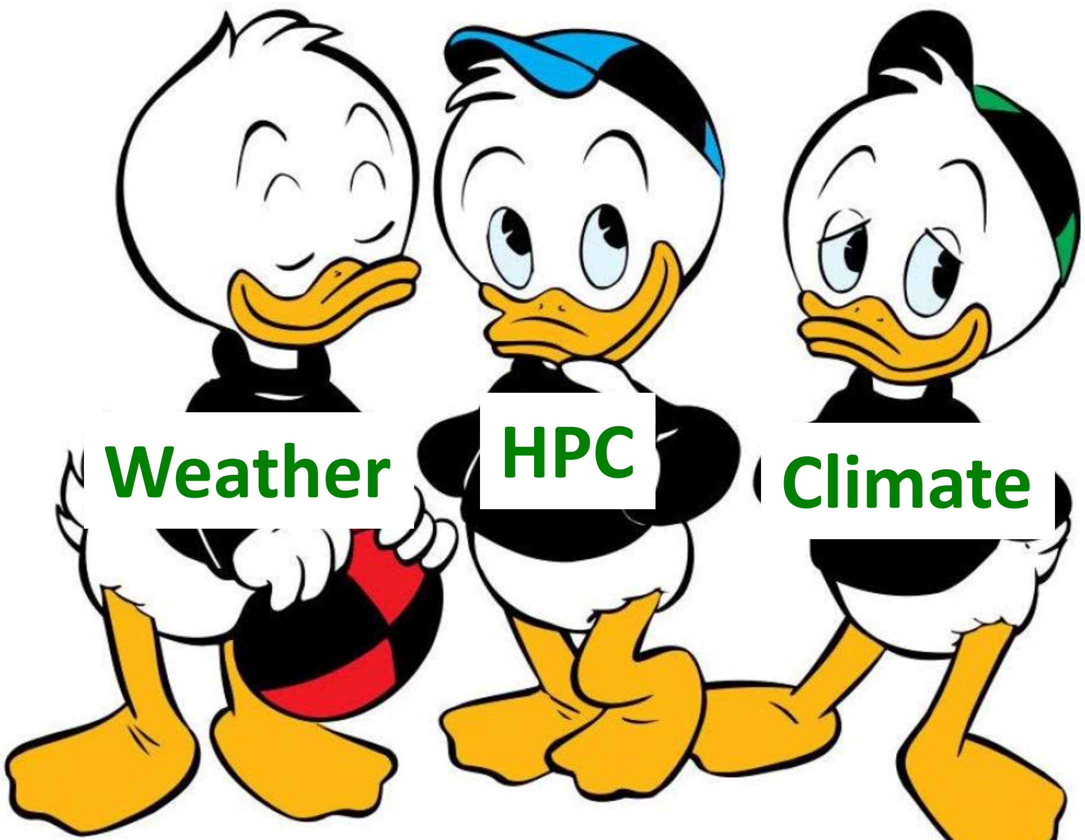

→ **This is where HPC, weather & climate research and communities unite**

1 Bony et al. Nature Geoscience 8:261-268, 2015
2 Schneider et al. Nature Climate Change 7:3-5, 2017

Global High-Resolution Simulations

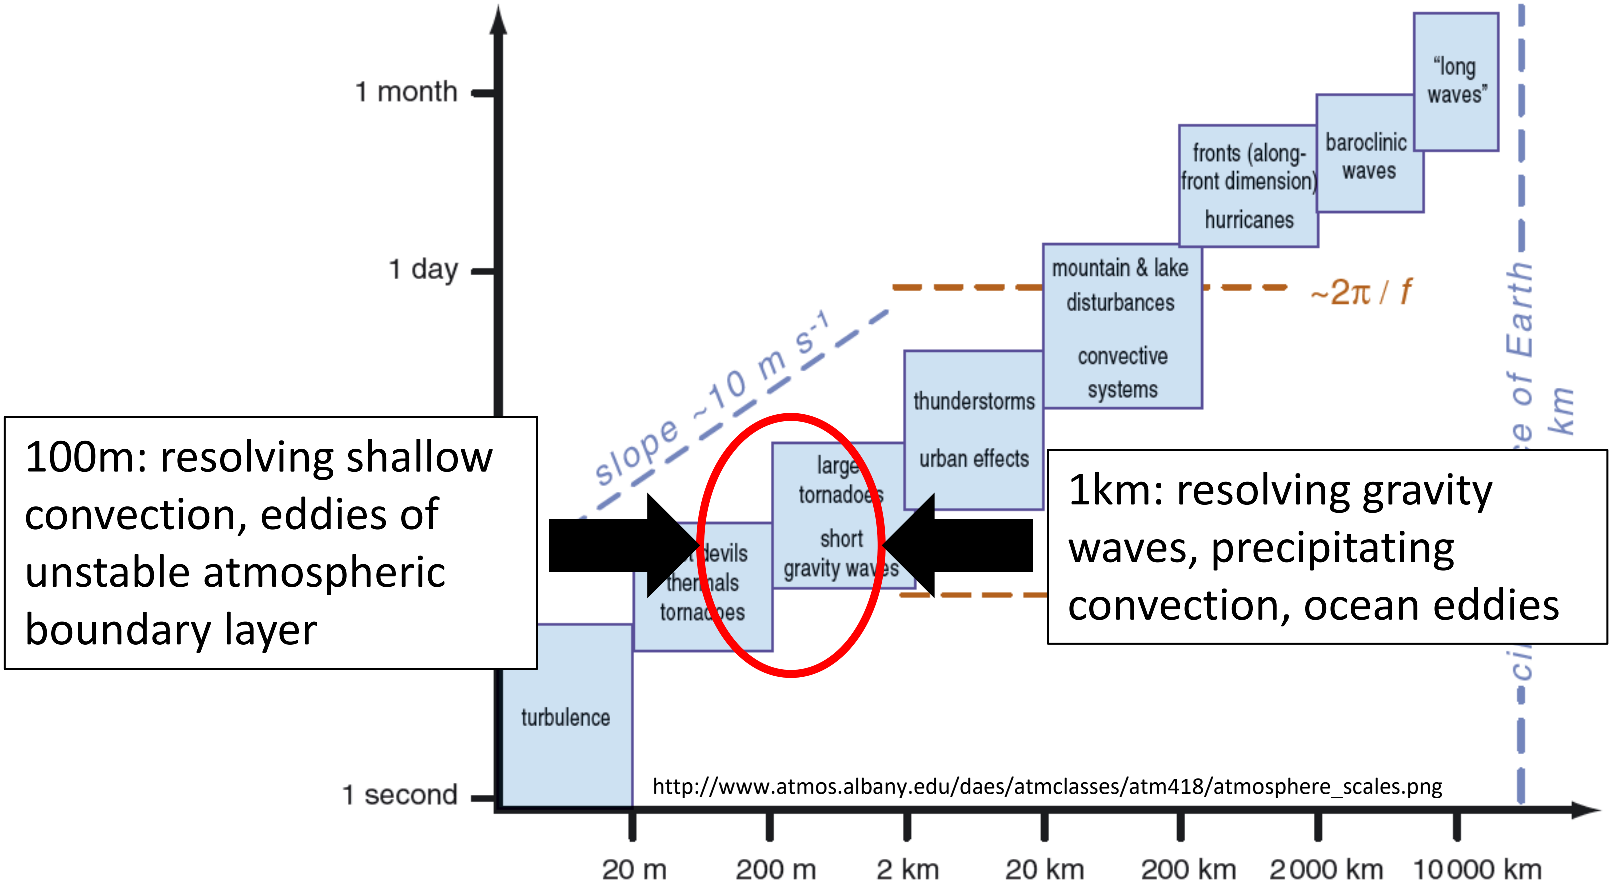

- **Goal: 100-1000m horizontal resolution at global scale**
- Less parametrization → In the limit, we know the equations!
- Challenges: very compute/memory/data intensive
  long-term efforts for model development

Overview

1. Hi-Res Simulations in ESiWACE & DYAMOND
2. Scalability and I/O
3. Performance Prediction via Sparse Grids
4. Summary

Centre of Excellence in Simulation of Weather and Climate in Europe (ESiWACE)

Meet us!

- ISC, 24-28 Jun, Frankfurt
- PASC, 2-4 Jul, Basel
- Ciênca, 2-4 Jul, Lisbon
- ICT, 4-6 Dec, Vienna
- …

WP1 Governance & engagement

WP2 Scalability
Global high resolution model demonstrators
→ **ICON, IFS, EC-Earth, NEMO**

WP3 Usability

WP4 Exploitability

WP5 Management & dissemination

Coordinator: DKRZ

WEATHER CLIMATE HPC

DYnamics of the Atmospheric general circulation Modeled On Non-hydrostatic Domains (DYAMOND)

- Goal: Intercomparison of global high-resolution models
- Participation list:
  - ICON/ Luis Kornblueh
  - NICAM/ Ryosuke Shibuya, Chihiro Kodama
  - MPAS/ Falko Judt
  - FV3/ Shiann-Jiann Linn
  - SAM/ Marat Khairoutdinov
  - NASA GEOS5/ William Putman
  - UM/ Pier Luigi Vidale

- Data management and support through DKRZ/ESiWACE
- More information: www.esiwace.eu/services/dyamond

5km Global Simulation with ICON

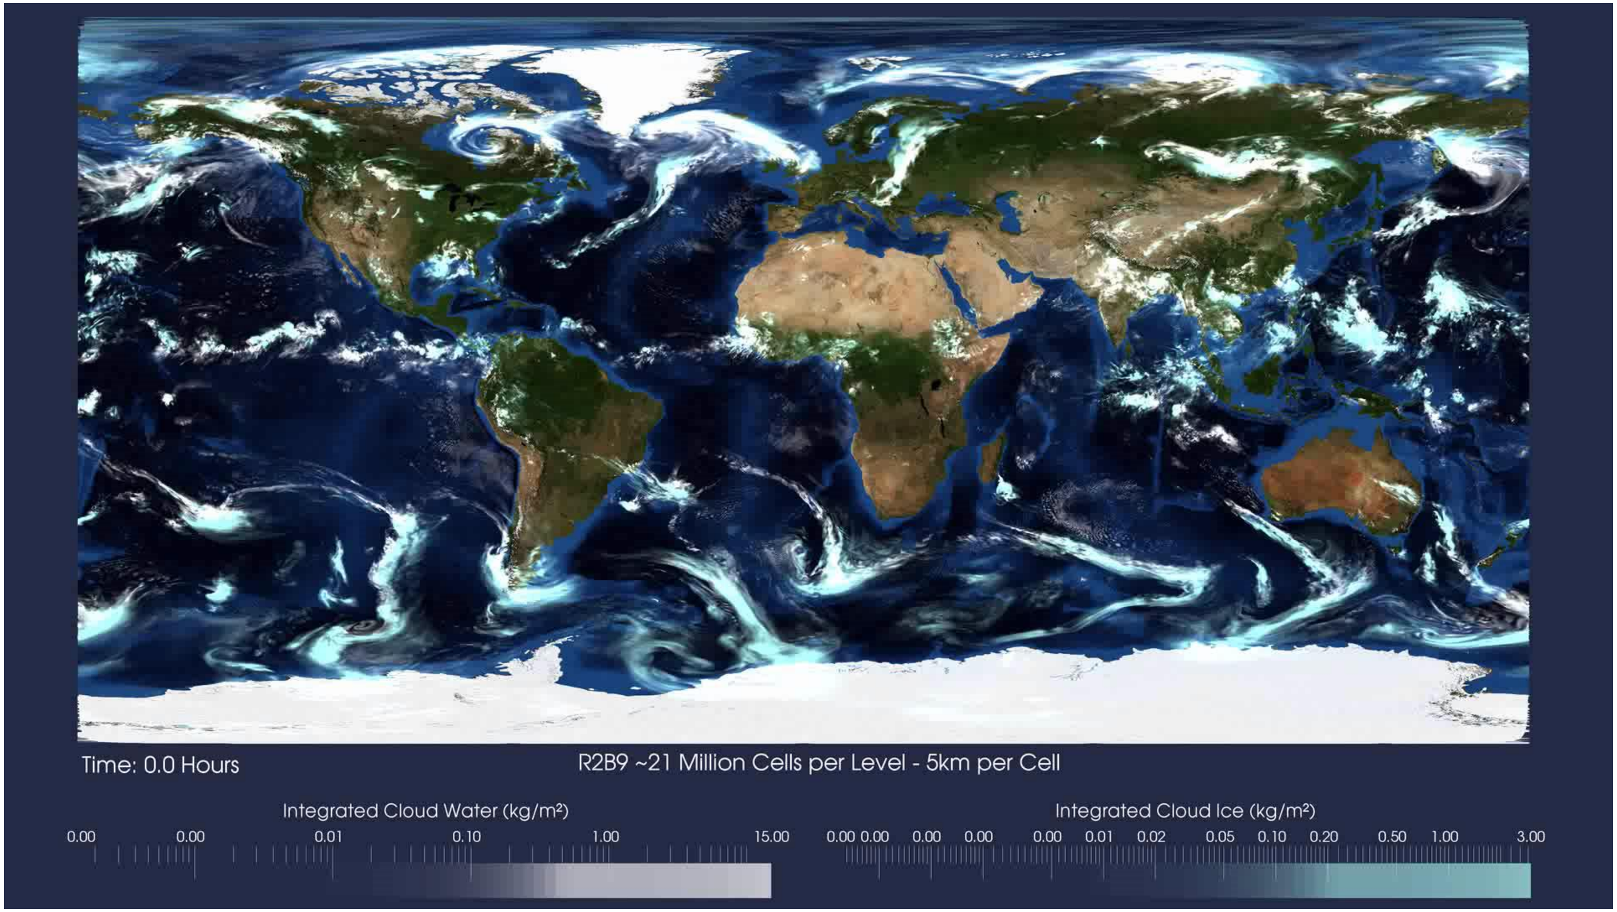

- Icosahedral Non-hydrostatic (ICON) model
- Finite difference-based
- Triangular horizontal grid (21M cells), 137 vertical levels
- 45s time step

Scalability of ICON and IFS

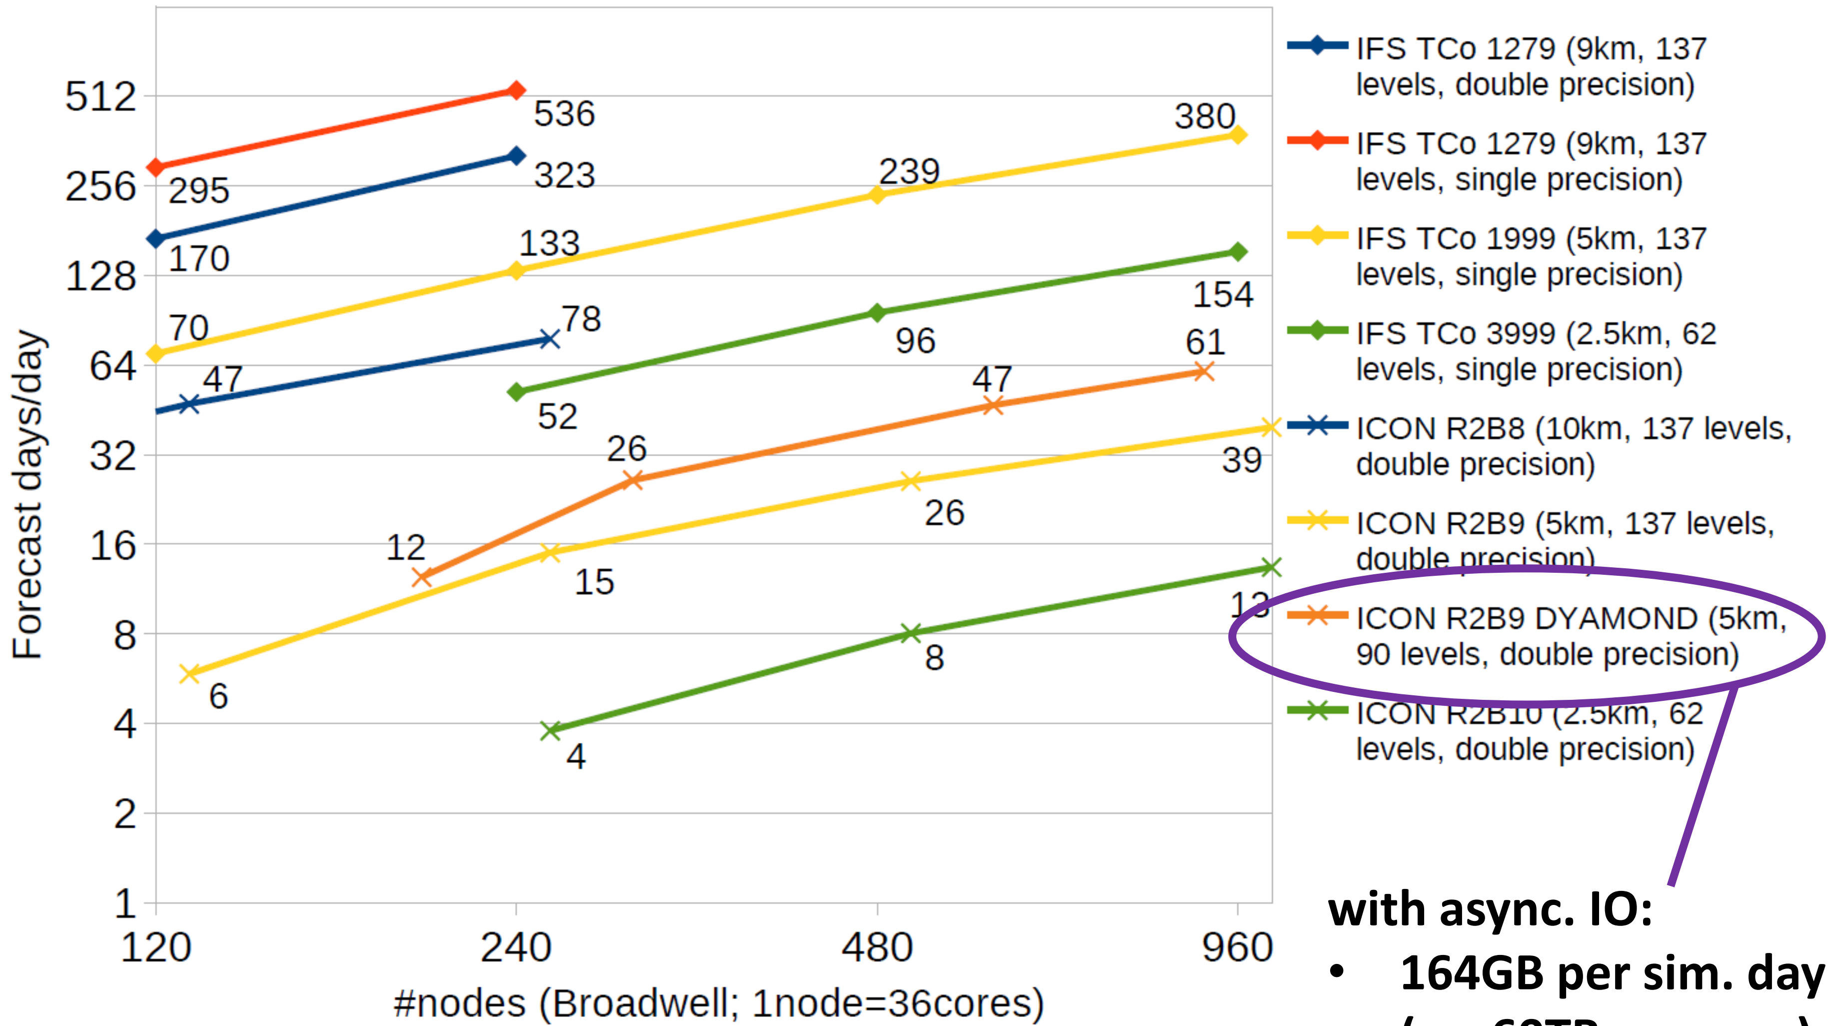

with async. IO:

- **164GB per sim. day (ca. 60TB per year)**
- **682GB checkpoint**

I/O in Numbers: ICON-DYAMOND 5km

Nodes	No I/O procs	wrt_output (s)
150	6	1091
300	6	1332
600	6	1661
600	11	863
900	15	749

→ How to determine optimal splitting?

Large Data Visualisation: Work in Progress @DKRZ

in-situ Visualisation

(ParaView/Catalyst/Cinema)

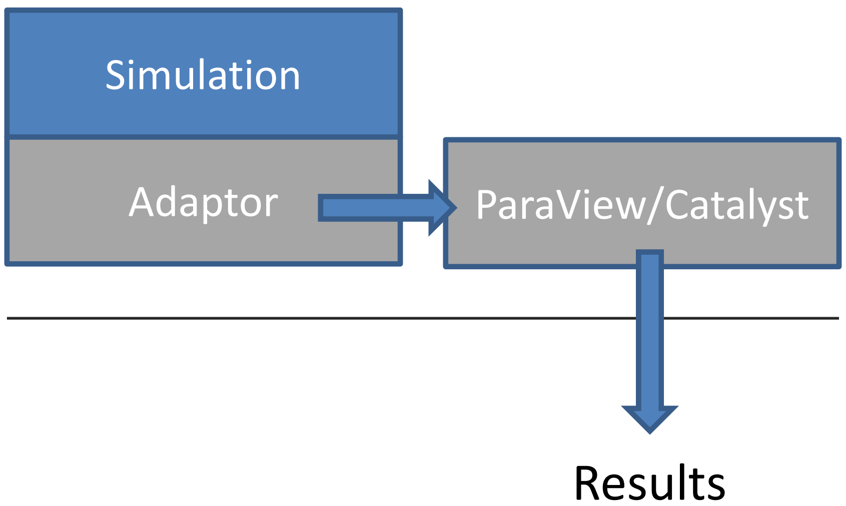

in-situ Compression

(Vapor)

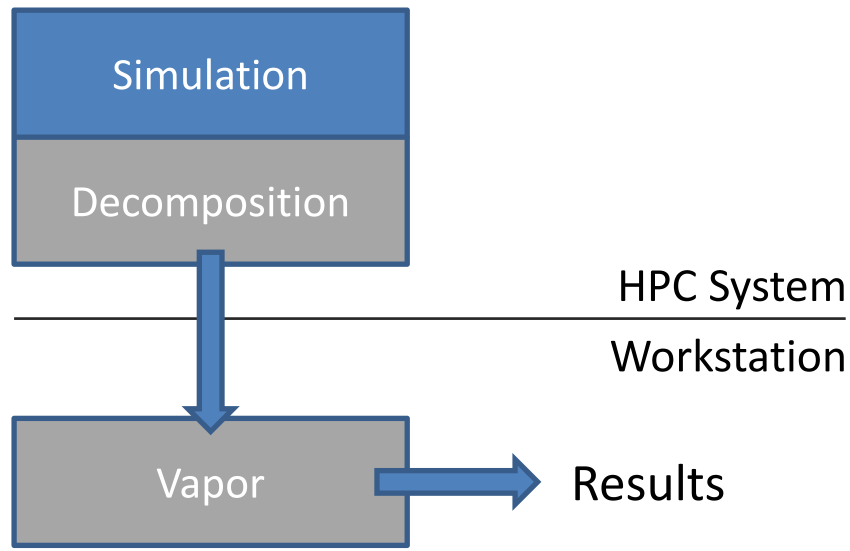

Perspectives

- **0.1-0.3 SYPD at 5km, atmosphere-only, O(30k Broadwell cores)**
- Other groups:
  - FV3/GFDL: 0.4 SYPD at 6.5km, 55k cores (Cray XC40)
  - NICAM/Japan: 0.6 SYPD at 14km, 20k cores (K computer)
- Long-term goal:
  1km atmosphere-ocean simulations at 1 SYPD
  3km ensemble atmosphere-ocean simulations
  → **this will require at least exascale computing** and corresponding **data capacity/handling capability**
  → challenge: model=long-term development
- **ESiWACE: The evolutionary path**
  → optimise current (production) models
- **ESCAPE/ESCAPE-2: The revolutionary path**
  - DSLs to enhance programmability/portability
  - Mixed precision arithmetics
  - Increasing the levels of concurrency (dwarf concept)

Performance Prediction: Objective

- **Multi-parameter influence on computational performance**

  → **computational:** OpenMP/MPI decomposition, loop-blocking, vector lengths, …

  → **algorithmic:** time step, number of iterations, error control/tolerance,…

  → **all aforementioned categories for every model subcomponent**

  → high-dimensional parameter space

- **Objective: performance estimate for Earth System Models…**

  …to gain insight into (wanted or unwanted) hotspots

  …to improve scheduling (relevant to workflows?)

- **Approach:** **Regression on high-dimensional parameter space via adaptive sparse grids**

Sparse Grids

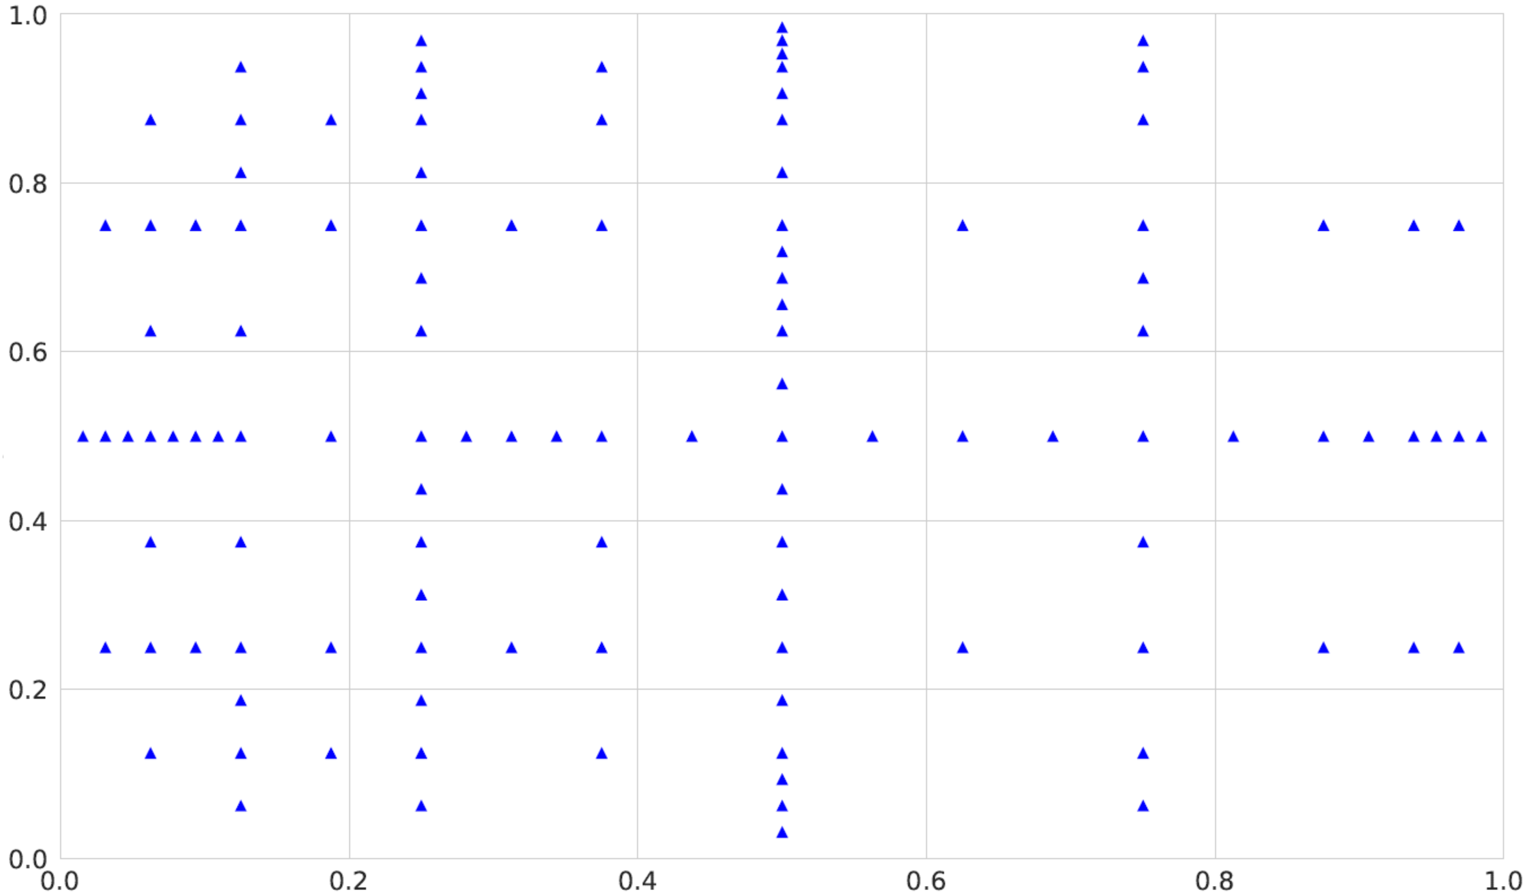

Theorem 1 *For the interpolation error of a function $f \in H^2_{0,mix}$ in the sparse grid space $V^s_{0,n}$ holds*

$$||f - f^s_n||_2 = \mathcal{O}(h^2_n \log(h^{-1}_n)^{d-1}).$$

SG: O(N(log N)$^{d-1}$) points

Full grid: O(N^d) points

J. Garcke.
Sparse grids in a nutshell

Performance Prediction: Sparse Grid Regression

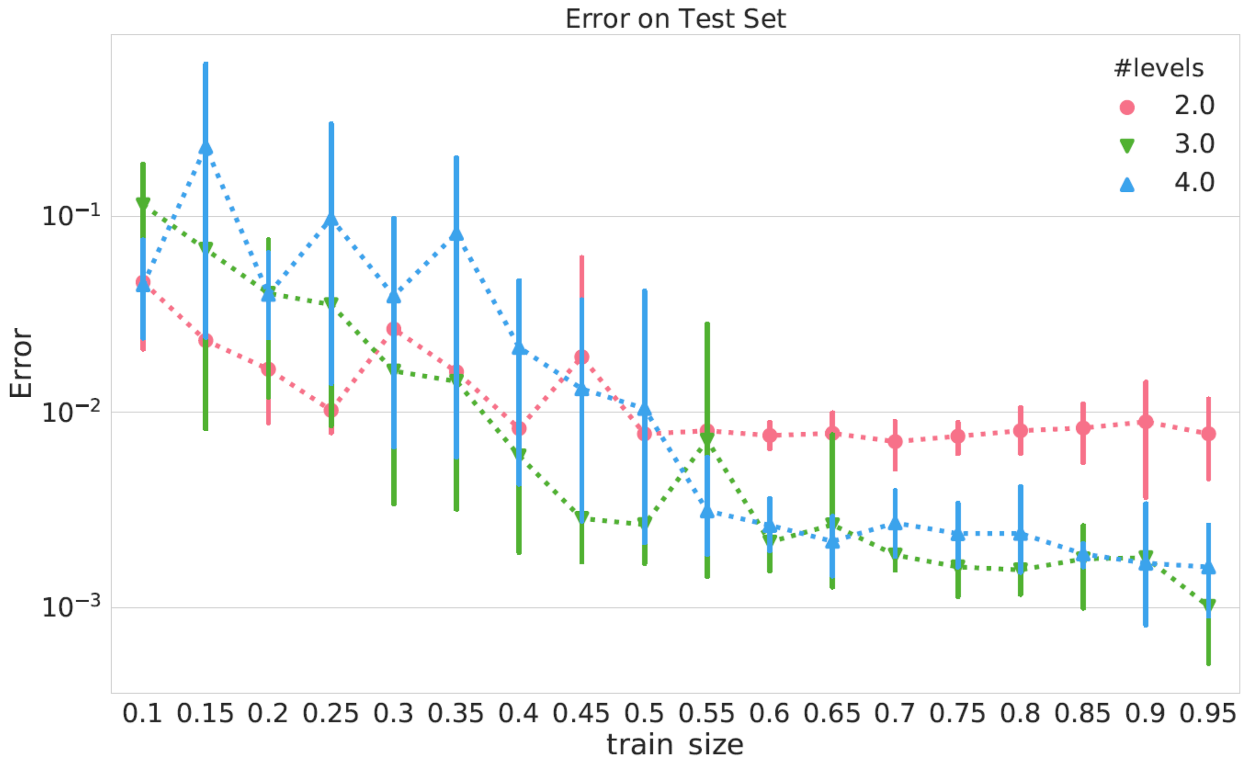

- Configuration: ICON-DYAMOND R2B4 (160km global res.), no I/O
- Run times on single-node (dual-socket Broadwell)
- Parameters: number OpenMP threads/MPI tasks
  loop-blocking (nproma)
- Courtesy by Paula Harder, DKRZ

Summary

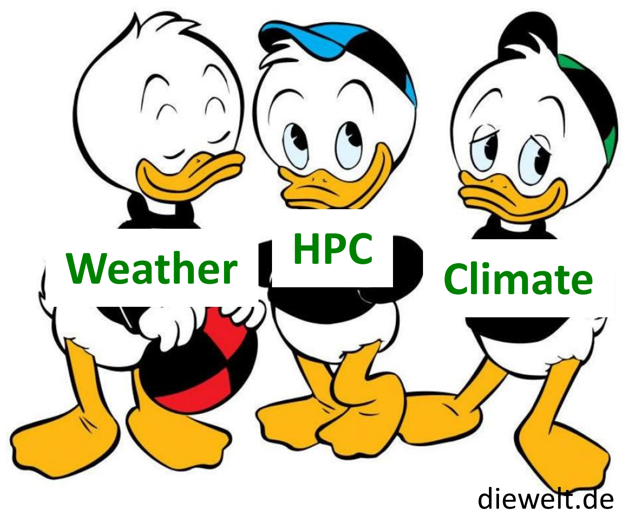

- ESiWACE – Joining forces to explore computability of extreme-scale weather and climate simulations
- DYAMOND: A project to intercompare global high-resolution models

- 5km global resolution simulations incl. I/O at O(0.1-0.3) SYPD

→ long way to 1 SYPD, requiring co-design efforts far beyond the duck trio ☺

→ compute/data hardware, HPC, software engineering, weather, climate, mathematics, ...

- Performance predictions with sparse grids deliver accurate run time estimates in various applications (weather simulation, particle systems)

Contacts: neumann@dkrz.de, esiwace@dkrz.de, www.esiwace.eu

Acknowledgement:
ESiWACE has received funding from the **European Union's Horizon 2020 research and innovation programme** under grant agreement No 675191. The authors gratefully acknowledge the computing time granted by the **German Climate Computing Centre**.

BACKUP Top 10 Reasons…Why 1km Would Be A Great Leap

Bjorn Stevens. ExtremeEarth. Presentation @EGU 2018, Vienna:

1. Convection is resolved (rainbands).
2. Surface orographic effects and gravity waves are resolved (storm tracks).
3. Shallow circulations (clouds and convection, feedbacks and forcing).
4. Ocean eddies are resolved (southern ocean stratification).
5. Tropopause dynamics are resolved (storm tracks and stratospheric water vapor).
6. Bathymetric effects on water mass formation are resolved (variability).
7. Allows a native representation of land surface (land use changes).
8. Remaining problems, such as microphysics & turbulence become tractable (parameterization).
9. Simulates observables (brings different science communities to the same table).
10. Direct link to impacts (connects directly to application communities).

BACKUP I/O in Numbers: Grib vs Netcdf (ICON-DYAMOND 5km)

900 nodes (Mistral, compute2), 15 IO procs,
1 simulated day, 2D/3D/rh,omega output every 15min/3h/15min

filename	variables	grb (GB)	nc (GB)	ratio nc/grb
atm1_2d_ml_20160801T000000Z	tqv_dia, tqc_dia, tqi_dia, tqg, tqs	5.7	38	6.7
atm2_2d_ml_20160801T000000Z	clct, lhfl_s, shfl_s, pres_sfc, tot_prec, cape_ml	9.9	46	4.6
atm_2d_avg_ml_20160801T000000Z	asob_s, athb_s, asob_t, athb_t, asou_t, asodifu_s, athd_s, athu_s	16	61	3.8
atm3_2d_ml_20160801T000000Z	u_10m, v_10m, t_2m, qv_2m, tqr	9.1	38	4.2
atm_3d_pres_ml_20160801T000000Z	pres	9.1	49	5.4
atm_3d_qv_ml_20160801T000000Z	qv	13	49	3.8
atm_3d_t_ml_20160801T000000Z	temp	13	49	3.8
atm_3d_tot_qc_dia_ml_20160801T000000Z	tot_qc_dia	1.4	49	35.0
atm_3d_tot_qi_dia_ml_20160801T000000Z	tot_qi_dia	0.96	49	51.0
atm_3d_u_ml_20160801T000000Z	u	14	49	3.5
atm_3d_v_ml_20160801T000000Z	v	14	49	3.5
atm_3d_w_ml_20160801T000000Z	w	12	49	4.1
atm4_2d_ml_20160801T000000Z	cin_ml, t_g, qv_s, umfl_s, vmfl_s	8	38	4.8
atm_omega_3d_pl_20160801T000000Z	omega	9	38	4.2
atm_rh_3d_pl_20160801T000000Z	rh	14	38	2.7