

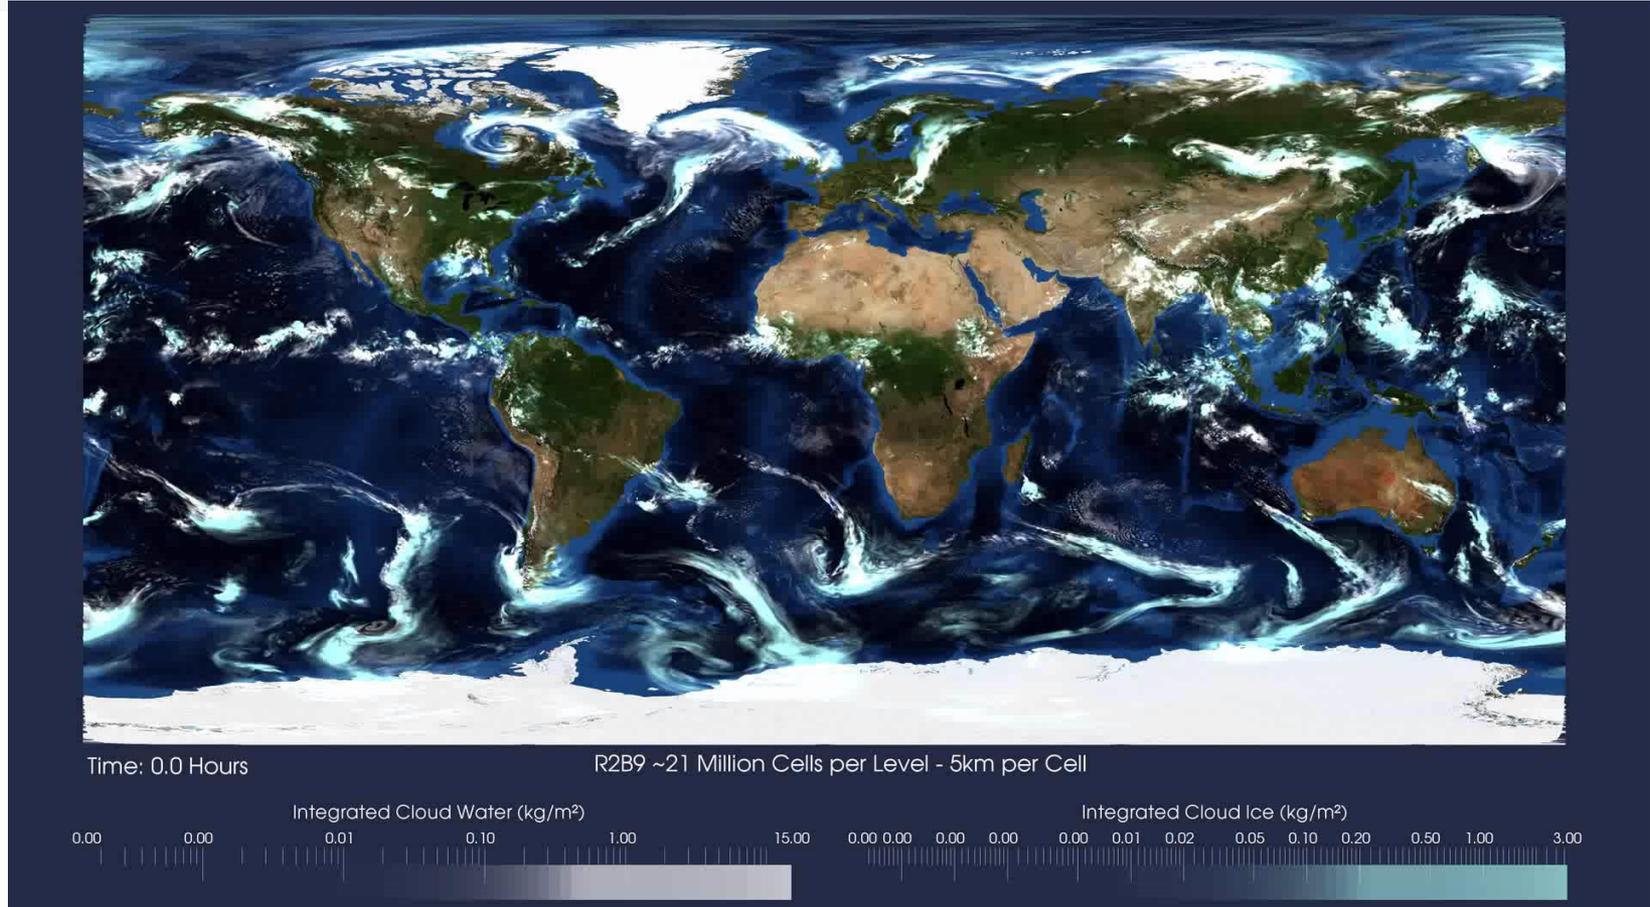
Performance Predictions for Storm-Resolving Simulations of the Climate System

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

Luis Kornbluh, Matthias Brück
Max-Planck-Institut für Meteorologie

Daniel Klocke
Deutscher Wetterdienst

Global High-Resolution Simulations



- High level of detail → 100-1000m horizontal resolution
- Less parametrisation → In the limit, we know the equations!
- Challenges: very compute/memory/data intensive!

Overview

1. ESiWACE: Overview and Goals
2. ESiWACE, DYAMOND and Perspectives
3. Scalability, IO, Visualisation
4. Performance Prediction
5. Summary

ESiWACE: Overview

- ESiWACE = Centre of **Excellence** in **Simulation** of **Weather** and **Climate** in **Europe**
- Funded by H2020, e-Infrastructures „Centres of Excellence for computing applications“
- ESiWACE leverages two European networks:
 - European Network for Earth System Modelling (ENES)
 - European Centre for Medium-Range Weather Forecasts (ECMWF)

Coordinator:



WEATHER

CLIMATE

HPC



Max-Planck-Institut
für Meteorologie



Bull
atos technologies



Deutscher Wetterdienst
Wetter und Klima aus einer Hand



Institut
Pierre
Simon
Laplace



CERFACS

allinea



SMHI



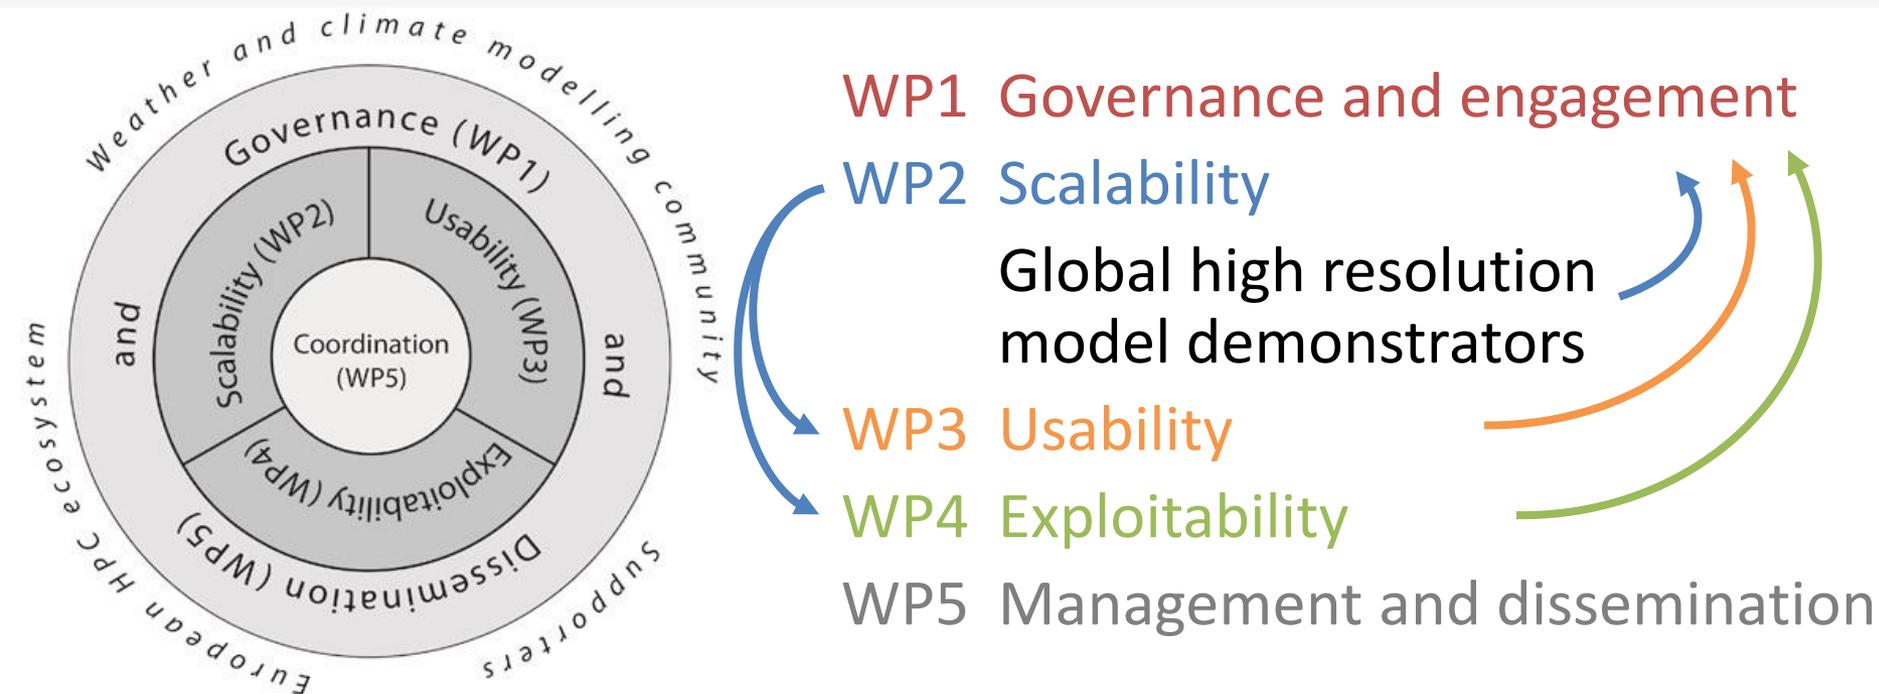
National Centre for
Atmospheric Science
NATURAL ENVIRONMENT RESEARCH COUNCIL



Science & Technology
Facilities Council



ESiWACE: Goals and Structure

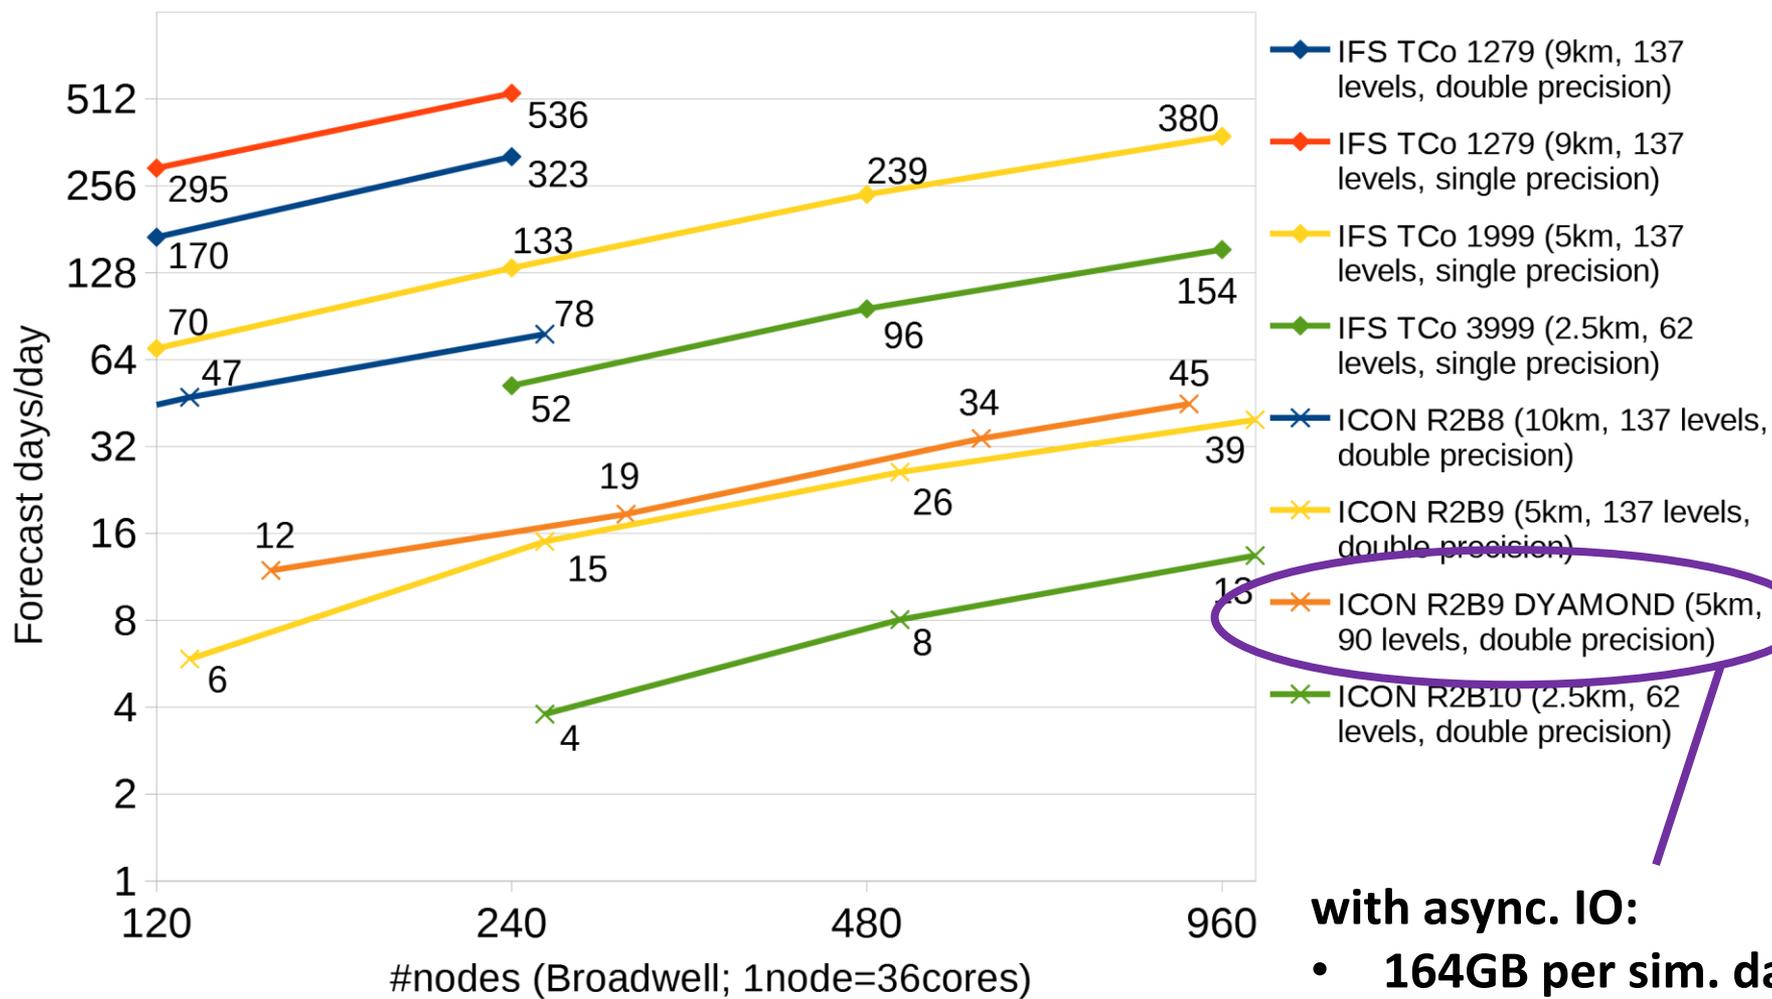


ESiWACE substantially improves efficiency and productivity of numerical weather and climate simulation on high-performance computing platforms by supporting the end-to-end workflow of global Earth system modelling.

ESiWACE, DYAMOND and Perspectives

- ESiWACE1: Establishment of global high-resolution simulations to demonstrate the computability of weather/climate predictions with (at least)
 - 1km resolution (atmosphere only)
 - 10km resolution (atmosphere-ocean)
 - Codes: IFS, NEMO, ICON, EC-EARTH
 - Explicit representation of deep convection, convective clouds, small-scale ocean eddies
 - Decreasing errors and uncertainties in weather and climate prediction
- DYAMOND: Intercomparison of global high-resolution models
- ESiWACE2
- Long-term: Extreme-scale robust high-resolution ensemble simulations at 1 simulated year per day (SYPD)

Scalability of ICON: Global High-resolution Simulations



with async. IO:

- 164GB per sim. day
- 682GB checkpoint

I/O in Numbers: Outtakes from DYAMOND runs

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?

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

900 nodes (Mistral, compute2), 15 IO procs,

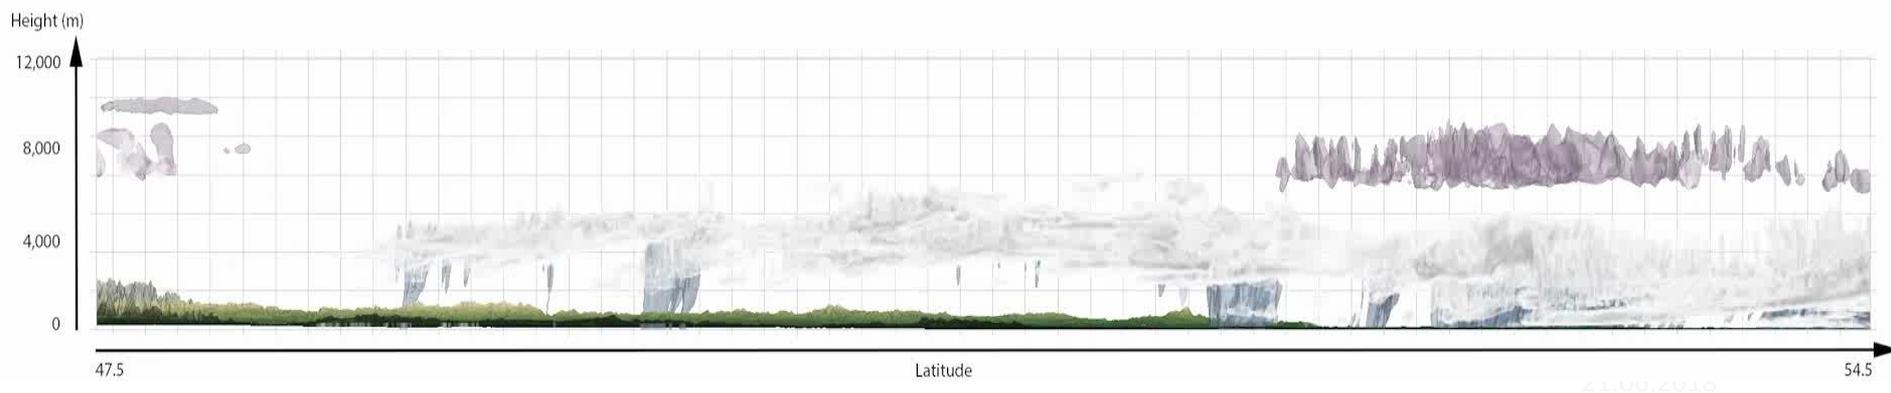
1 simulated day, 2D/3D/rh,omega output every 15min/3h/15min

filename	variables	grb (GB)	nc (GB)	ratio nc/grb
atm1_2d_ml_20160801T000000Z	tqv_dia, tqc_dia, tqi_dia, tqg, tq_s	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

The Purpose of Simulation is Insight...

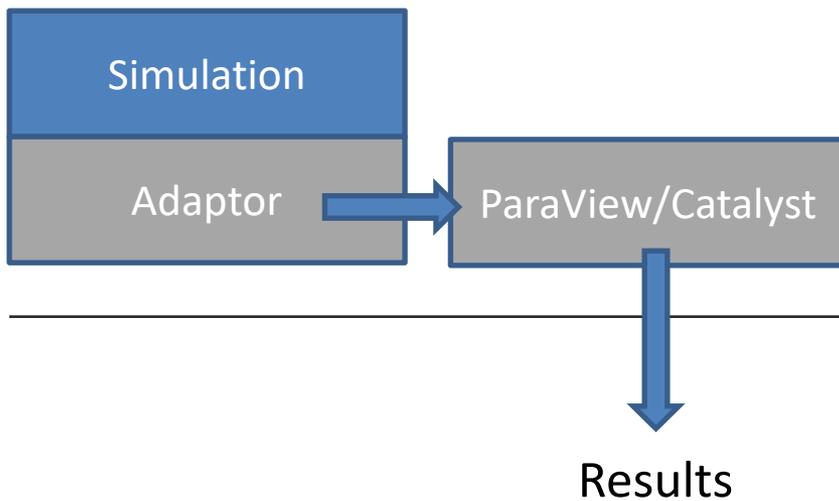
See, understand, learn, communicate ...

- Confirmatory visualisation
- Interactive visualisation
- Animations & stills for communication

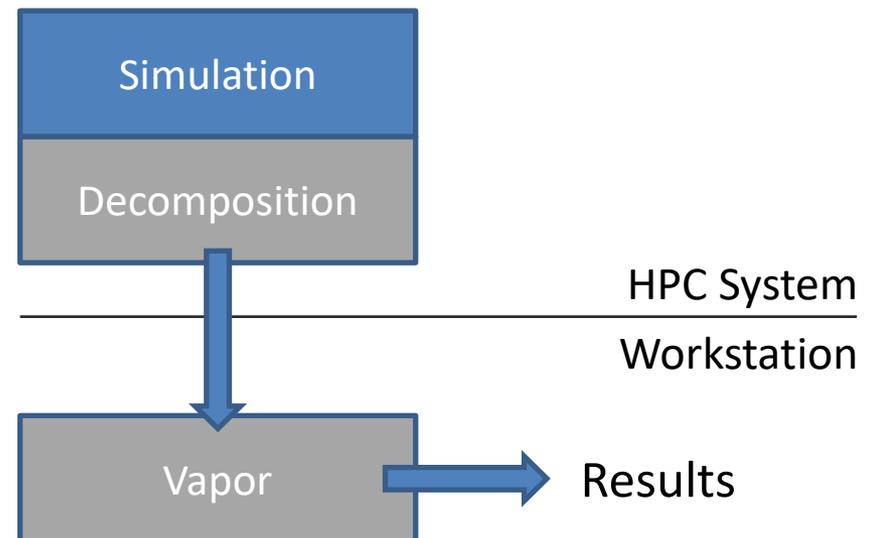


Large Data Visualisation

in-situ Visualisation (ParaView/Catalyst/Cinema)

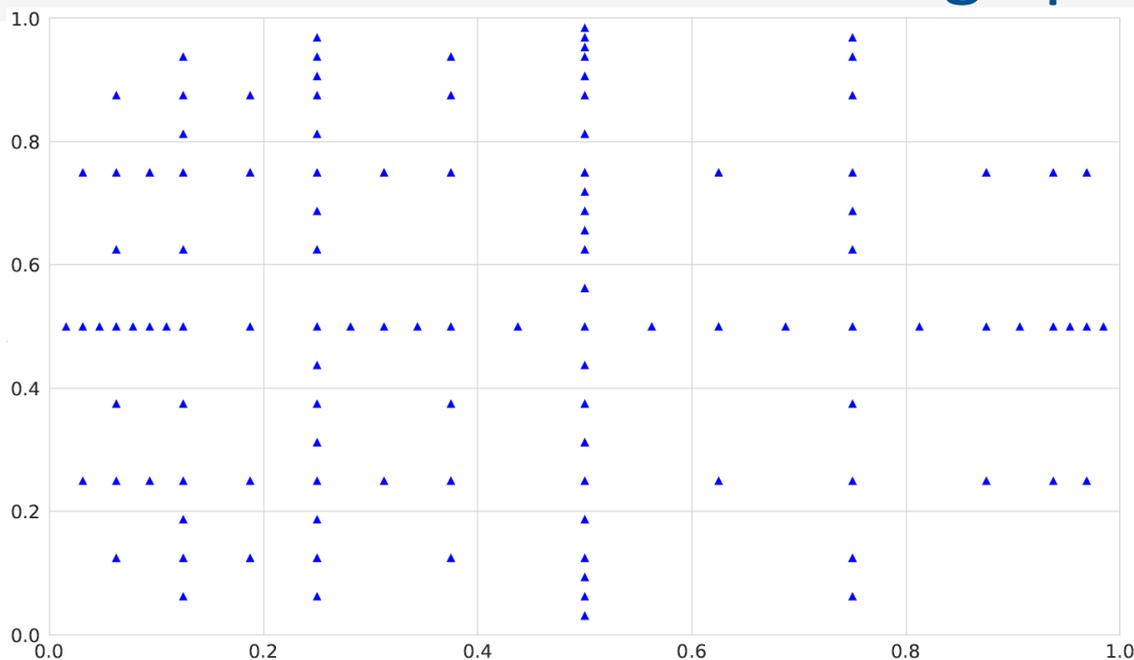


in-situ Compression (Vapor)



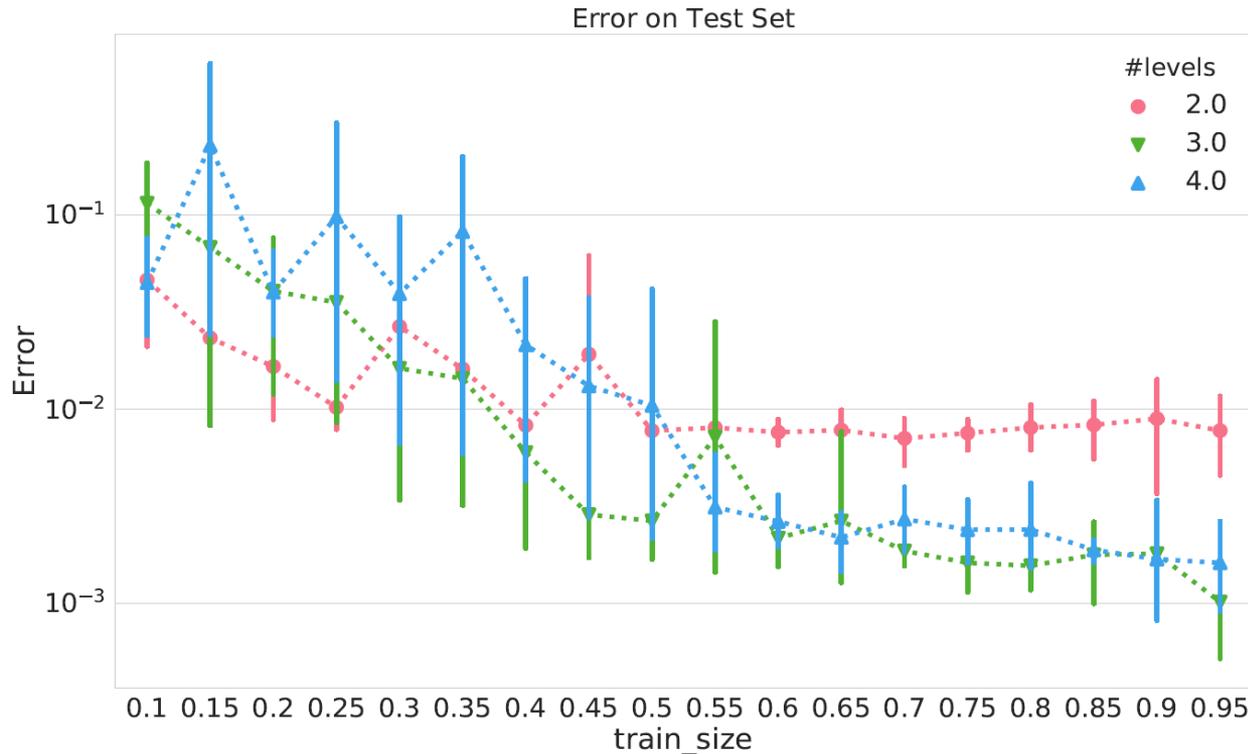
21.06.2018

Towards Performance Prediction using Sparse Grids



- **Objective: performance estimate for complex ESMs...**
 ...to gain insight into (wanted or unwanted) hotspots
 ...to improve scheduling (relevant to workflows?)
- **Multi-parameter influence on computational performance**
 → OpenMP/MPI decomposition, loop-blocking, vector lengths, ...
- **Approach: Regression on (potentially) high-dimensional parameter space via adaptive sparse grids**

Towards Performance Prediction using Sparse Grids



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

Summary

- ESiWACE – Joining forces to explore computability of extreme-scale weather and climate simulations
 - ENES HPC Workshop, May 17-18 2018, Lecce/Italy
 - European Seminar on Computing, June 3-8 2018, Pilsen/Czech Rep.
 - Teratec Forum, June 19-20 2018, Palaiseau/France
 - ISC, June 24-28 2018, Frankfurt/Germany
 - PASC, July 2-4 2018, Basel/Switzerland
- 5km global resolution simulations incl. IO at O(45) forecast days/day
→ still some way to 1 SYPD...
- Performance prediction with sparse grids delivers accurate run time estimates in various applications

Contacts: neumann@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**.