Pan-European convection-permitting WRF simulations

Implementation and initial results

Klaus Goergen[1,2], Sebastian Knist[2,3,4], Stefan Kollet[1,2], Harry Vereecken[1,2]

(1) Agrosphere (IBG-3), Jülich Research Centre, Germany; (2) Centre for High-Performance Scientific Computing in Terrestrial Systems, Geoverbund ABC/J, Germany; (3) SimLab TerrSys, Jülich Supercomputing Centre, Jülich Research Centre, Germany; (4) Meteorological Institute, Bonn University, Germany

1. Introduction

Background and motivation

- High-Performance Computing (HPC) developments: Ever-increasing performance and complexity (parallelism, vectorisation, fast interconnects, parallel filesystems, memory hierarchies, etc.) – larger, more complex problems can be simulated
- Climate modelling developments: Higher resolution, multi-physics regional earth system models, growing model domains, long integration times and large ensembles
- Multiple benefits of convection-permitting climate modelling (CPCM), e.g., reduced uncertainties related to convection parameterisation, better representation of convective precipitation, more realistic representation of surface properties or exchange processes
- CPCM is computationally demanding (Prein et al., 2015); however, HPC resources in principle exist, e.g., at institutions like the Jülich Supercomputing Centre (JSC) (Tab.1)

Objective – optimum model performance to allow for extended large CPCM runs

- Use latest standard version of WRF (Skamarock et al., 2008) model for continental convection permitting simulations and optimize usage for our simulations on JSC systems
- Good performance target: 0.25 h wall-clock per simulated day, 30 yrs in about 3.75 months

2. Experiments

Central European model domain, 480 x 456 x 50, ≈3km, WRF v3.6.1 (Fig.1a)

- Simulations: (i) 10 yrs validation + land surface properties and resolution sensitivity runs, (ii) MPI-ESM-LR RCP4.5: 1995-2005, 2040-2050, 2090-2100
- Climate-mode settings: WSM5, RRTMG, YSU PBL, Noah LSM

Pan-European model domain, 1600 x 1552 x 50, ≈3km, WRF v3.5.1/v3.8.1 (Fig.1b)

- Simulations: Selected months, towards seasons and multi-year runs
- Climate-mode settings: WSM6/WSM5, CAM/RRTMG, YSU PBL, Noah LSM (v3.5.1/v3.8.1)

Our interest: Land-atmosphere coupling, regional climate change impacts, hydrological model forcing, added value in precipitation reproduction (talk on Wed., Goergen et al.)

Contribution to EURO- and Med-CORDEX Flagship Pilot Study (talk Thu., Sobolowski)

Fully coupled RESM TerrSysMP (COSMO + CLM + ParFlow + OASIS3-MCT) forecasts at convection permitting resolution (poster Wed., Kollet et al.)

3. Scaling tests

- Currently using highly scalable IBM BG/Q JUQUEEN (testing) and general purpose Linux cluster JURECA (production) in CPU-only setups, accelerators are not yet in use
- 75% efficiency (speedup / ideal speedup) is considered a threshold for efficient code at JSC
- On JUQUEEN (Fig.2a), the 8192 MPI tasks run is about 3.8 times slower than the JURECA simulations (Fig.2b and c), albeit the implementations can hardly be directly compared with each other
- The production runs for the Central Europe domain are done with 10 nodes and 240 MPI tasks, requiring about 0.63 h / day (Fig.2b)
- With 0.2591 h / day of the 3 km pan-EU JURECA run with 3072 MPI tasks (Fig.2c left) and about 83% efficiency (Fig.2c right), the target is met, albeit with very large resource use of 128 compute nodes
- Without using OpenMP/MPI hybrid parallelism, simulation moves from being compute bound to highly communication-bound (Elliot, 2015)
- A runtime improvement is expected from using MICs or GPUs; even if code does not run natively on the accelerator; by offloading e.g. the radiation modules (incl. time-consuming data transfers between the CPU and the accelerator), more frequent calls to these routines may be possible without a performance penalty

Input / Output operations (not shown)

- CPCM poses a big data challenge: data volume (up to ≈10 TB / month)
- Use of pNetCDF or task local NetCDF over serial I/O (Elliott, 2015)
- Adjust to filesystem characteristics, e.g., via ROMIO hints

4. Next steps

- Improvement of efficiency, use less resources
- Use of hybrid parallelism in WRF (OpenMP/MPI) to reduce communication overhead in combination with optimized tiling strategy in CPU-only and hybrid setups (Elliott, 2015)
- Optimum parallel I/O strategy
- Systematic testing on JSC's heterogeneous HPC architectures (CPUs + GPUs/MICs) based on the built-in functionalities (e.g., RRTMG radiation scheme for GPU and MIC)
- Detailed runtime performance profiling with Score-P/Scalasca, VTune and Darshan
- Testing of in-situ processing and visualisation (VisIt interface)
- "Convection permitting ready" pre- and postprocessing chains
- Use of JURECA's NVIDIA accelerators through CUDA ported physics modules, e.g., for cloud microphysics (WSM5) (Mielikainen et al., 2012); or planetary boundary layer scheme (YSU) (Huang et al., 2015); multiple modules have been ported, e.g., by the Space Science and Engineering Center (SSEC) of University of Wisconsin-Madison

Tab. 1 Production and experimental HPC systems at JSC, one of the three German national supercomputing centres. Possibility to explore capabilities of heterogeneous accelerator systems such as JURECA or DEEP for CPCM. New Xeon Phi KNL deployments are under way (e.g., JUKNIGHT). Central storage on ≈16 PB IBM GPFS. Co-design towards exascale performance (DEEP-ER).

	JURECA	JUQUEEN	JUROPA3	DEEP
C.-nodes CPU cores	1,872 44,928	28,672 458,752	44 704	128 (2,048) 384 (Booster 1) 32 (Booster 2)
CPUs / node	2x Intel Xeon E5-2680 v3 HW, 12 cores, 2.5 GHz	IBM PowerPC A2, 16 cores, 1.6 GHz	2x Intel Xeon E5-2650 SB 8 cores, 2.0 GHz	2x Intel Xeon E5-2680 8 cores, 2.7 GHz
RAM Total RAM	128 / 256 / 512 GB DDR4 277 TB	16 GB DDR3 448 TB	64 GB DDR3 2.8 TB	32 GB 4.0 TB + 6.1 TB + 0.5 TB
Accelerator	75x2 NVIDIA Tesla K80 GPUs	-/-	4x2 NVIDIA Tesla K20X GPUs 4x2Intel Xeon Phi 5110P MICs	Intel Xeon Phi 7120X MICs (B1) Intel Xeon Phi 7120D MICs (B2)
Interconnects	EDR Infiniband	5D-Torus, 40GB/s	FDR Infiniband	QDR Inifiniband, 3D EXTOLL
Peak performance	1.8 (CPU) PFLOPS 0.44 (GPU) PFLOPS	5.9 PFLOPS	0.03 PFLOPS	45 (cluster) + 500 (B1) + 38.4 (B2) TFLOPS

(a) Central-Europe 3 km WRF **(b) Pan-European 3 km WRF**

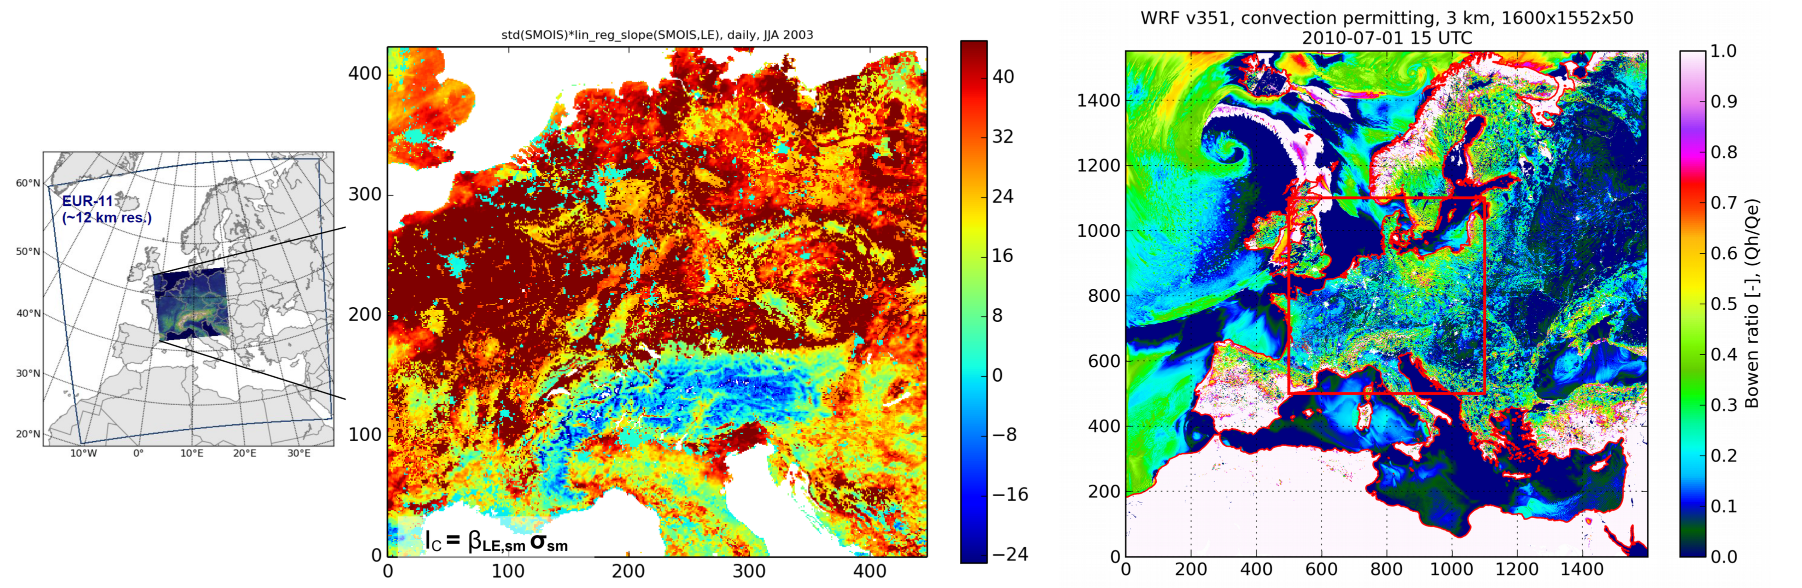

Fig. 1 Model domains at 3 km currently in use. (a) Central Europe domain, here: land-atmosphere coupling index according to Dirmeyer (2011) for heat-wave summer JJA 2003; (b) Pan-European domain, here: Bowen ratio at 2010-07-01. Setup: One-way double nest, ERA-Interim (0.75°) or GCM > EUR-11 (0.11°, ≈12 km) > 0.0275°. The continent-wide model domain allows to investigate scale-dependent processes like atmospheric moisture budgets, convective system evolution, or widespread frontal system dynamics. Processes and their variances may be reproduced across multiple spatial scales and watersheds.

(a) JUQUEEN, strong scaling, EUR-44 (48 km), EUR-11 (12 km), pan-European domain (3 km)

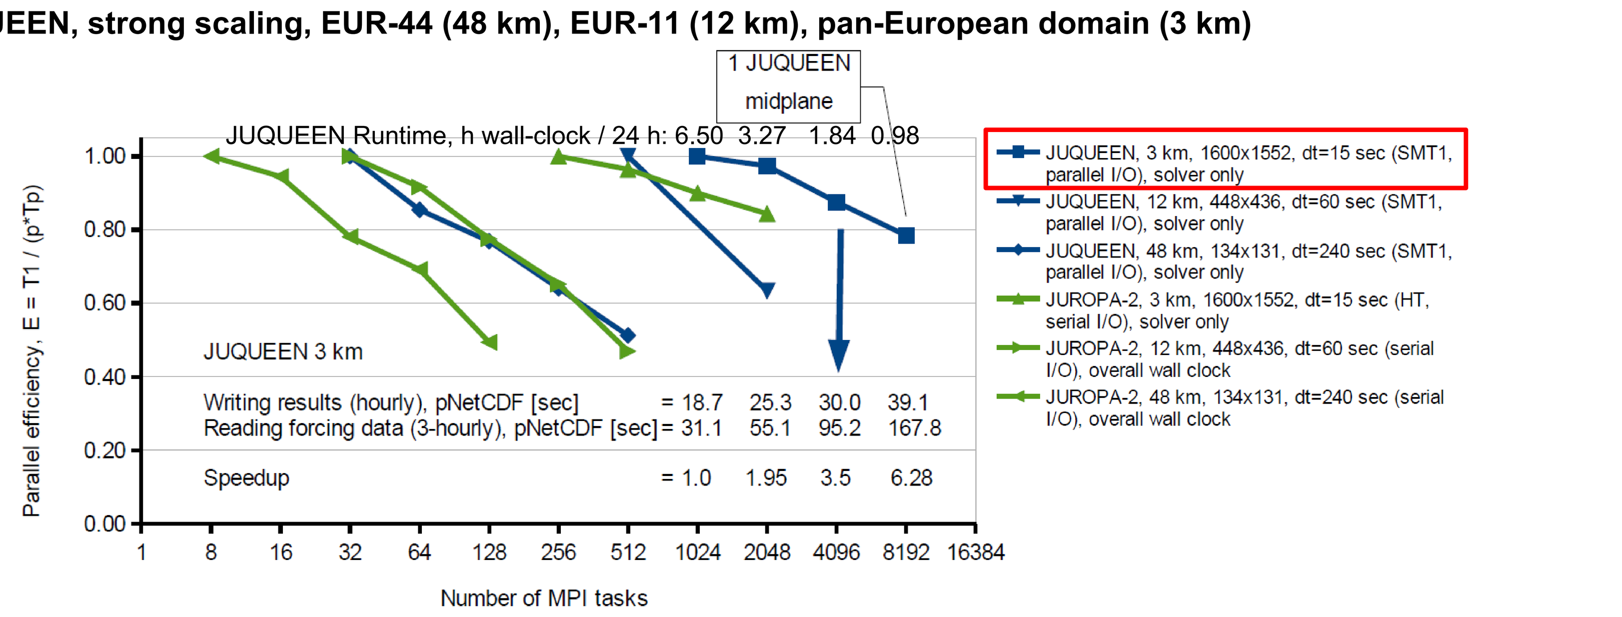

(b) JURECA, strong scaling, EUR-11 (12 km), Central Europe domain (3 km)

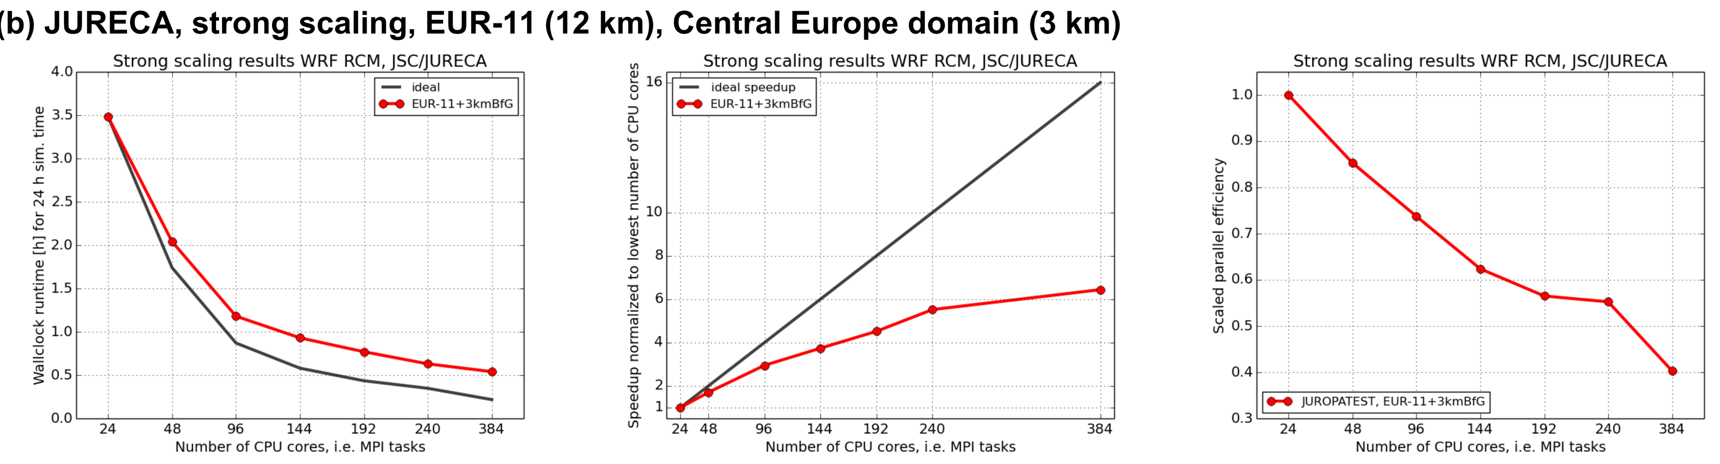

(c) JURECA, strong scaling, pan-European domain (3 km)

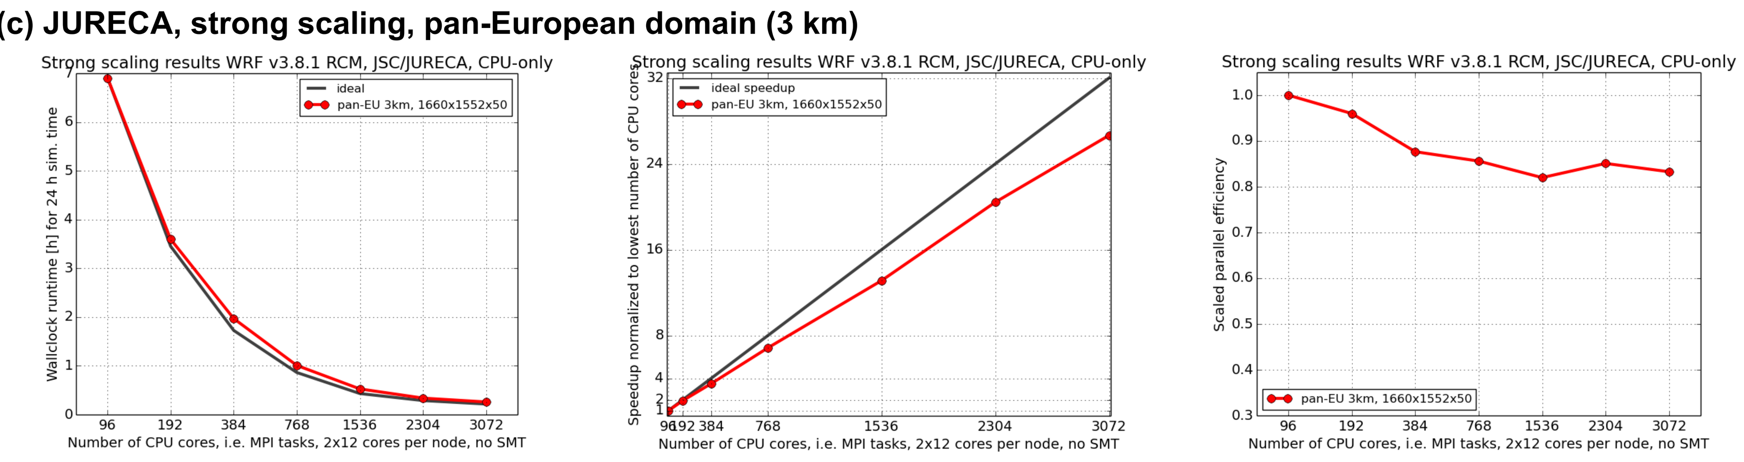

Fig. 2 Strong scaling results for CPCM setups on different JSC HPC systems. Real data test case (6 hr ERA-Interim driven): 2017-07-01, test run for 24 hrs.
(a) Different model domains: efficiency; WRF v3.5.1, JUQUEEN, pNetCDF parallel I/O, please note the model runs with the squares and the triangles, these can be compared to the results from experiment (c), the timings are without I/O and initialisation.
(b) Double-nest EUR-11 (12 km, 448 x 436 x 50) and Central Europe (3 km, 480 x 456 x 50) domains: wall-clock runtime (left), speedup (middle), efficiency (right); WRF v3.6.1, JURECA, serial NetCDF I/O, -O2 -xAVX optimisation.
(c) Pan-European (3 km, 1600 x 1552 x 50) domain: wall-clock runtime (left), speedup (middle), efficiency (right); WRF v3.8.1, JURECA, without I/O and initialisation, -O3 -xCORE-AVX2 optimisation.

Acknowledgements

Simulations for this study took place on the HPC systems of the Jülich Supercomputing Centre as part of different compute time grants: The authors gratefully acknowledge the computing time granted on the supercomputer JURECA at the Jülich Supercomputing Centre through VSR (project ID JJSC1500, Goergen). The compute time allocation of the Simulation Laboratory Terrestrial Systems at the Jülich Supercomputing Centre for the JUQUEEN and JURECA HPC systems has been substantially increased through a share of the Research Centre Jülich Board of Directors' quota for the SLTS simulation projects. Part of the work has been funded through a research and development cooperation on hydro-meteorology between the Federal Institute of Hydrology, Koblenz, Germany, and the Meteorological Institute, Bonn University.

Poster 1.2.1, GEWEX Convection-Permitting Climate Modeling Workshop, Boulder, Co, USA, 6-8 September 2016

References

Dirmeyer, P. A. (2011), The terrestrial segment of soil moisture–climate coupling, Gephys. Res. Lett., 38(16), L16702.

Elliott, S. (2015), Performance Analysis and Optimization of the Weather Research and Forecasting Model (WRF) on Intel Multicore and Manycore Architectures, International Conference for High Performance Computing, Networking, Storage and Analysis (SC15), 15-20 November 2015, Austin, Texas, USA.

Huang, M., J. Mielikainen, B. Huang, H. Chen, H.-L. A. Huang, and M. D. Goldberg (2015), Development of efficient GPU parallelization of WRF Yonsei University planetary boundary layer scheme, Geosci. Model Dev., 8(9), 2977-2990.

Mielikainen, J., B. Huang, H.-L. A. Huang, and M. D. Goldberg (2012), Improved GPU/CUDA Based Parallel Weather and Research Forecast (WRF) Single Moment 5-Class (WSM5) Cloud Microphysics, IEEE J. Sel. Top. Appl. Earth Obs. Remote Sens., 5(4), 1256-1265.

Prein, A. F. et al. (2015), A review on regional convection-permitting climate modeling: demonstrations, prospects, and challenges, Rev. Geophys., 53(2), 323-361.

Skamarock, W., J. Klemp, J. Dudhia, D. Gill, D. Barker, M. Duda, W. Wang, and J. Powers (2008), A description of the Advanced Research WRF Version 3, Tech. rep., NCAR.

Contact

Klaus Goergen
k.goergen@fz-juelich.de

Agrosphere (IBG-3)
Institute of Bio- and Geosciences
Jülich Research Centre (FZJ)
Jülich, Germany

Centre for High-Performance Scientific Computing in Terrestrial Systems (HPSC TerrSys)
Geoverbund ABC/J (Germany)

http://www.fz-juelich.de/ibg/ibg-3
http://www.fz-juelich.de/ias/jsc/slts
http://www.hpsc-terrsys.de

Follow us on Twitter @HPSCTerrSys

HPSC TerrSys YouTube Channel

Member of the Helmholtz Association