Max-Planck-Institut für Meteorologie

Deutscher Wetterdienst
Wetter und Klima aus einer Hand

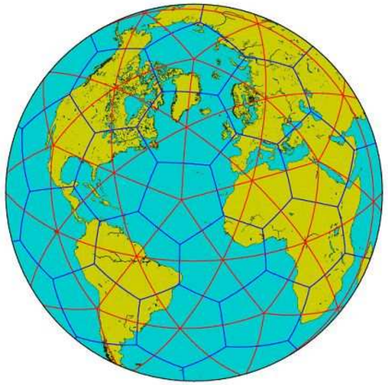

ICON

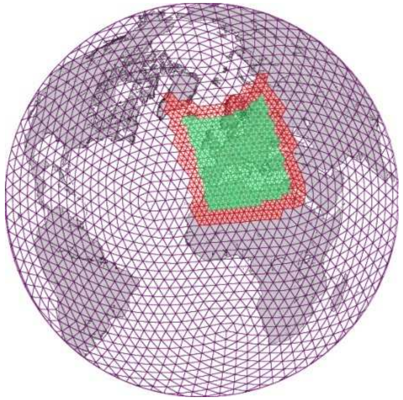

The ICON (Icosahedral Nonhydrostatic) model and its relationship to the ICOMEX project

Günther Zängl
and the ICON and ICOMEX teams

Weather and climate prediction on next-generation supercomputers
23.10.2012

Outline

- ➔ The “levels” of efficiency
- ➔ Overview of the ICON model and its hierarchy of efficiency optimizations
- ➔ The ICOMEX project: key components on the way towards exascale computing
- ➔ Conclusions

The “levels” of efficiency

- Efficiency of the time-stepping scheme and physics-dynamics coupling: time steps vs. computational cost and communication per time step
- “Balancing” of time steps / stability ranges of various model components
- Code-level optimizations to improve usage of memory bandwidth and cache
- Avoid unnecessary computations, copy operations etc.
- Minimization of load imbalance and MPI communication, in particular global communication
- Hierarchy of parallelization levels
- Asynchronous, parallel I/O in order to hide output times
- And last but not least: memory efficiency / memory scaling

Time stepping options in global nonhydrostatic models

- **Fact: at model resolutions permitting breaking gravity waves in the mesosphere, the dynamical core needs to be numerically stable up to wind speeds of at least 250 m/s**

- **Semi-implicit semi-Lagrange can be very efficient due to long time steps, but achieving good scaling properties requires substantial work**
- **Horizontally fully explicit time stepping requires short time steps in general (sound waves), but no further penalty at high wind speeds and no global communication; time splitting is possible between dynamical core and tracer advection / physics parameterizations**
- **Split-explicit or semi-implicit schemes with advective CFL limitation may suffer from the relatively small ratio between sound speed and maximum horizontal wind speed**

ICON = ICOsahedral Nonhydrostatic model

- Joint development project of DWD and Max-Planck-Institute for Meteorology for the next-generation global NWP and climate modeling system
- Nonhydrostatic dynamical core on an icosahedral-triangular C-grid; coupled with (almost) full set of physics parameterizations
- Two-way nesting with capability for multiple nests per nesting level; vertical nesting, one-way nesting mode and limited-area mode are also available

Primary development goals

- ➔ **Better conservation properties (air mass, mass of trace gases and moisture, consistent transport of tracers)**
- ➔ **Grid nesting in order to replace both GME (global forecast model, mesh size 20 km) and COSMO-EU (regional model, mesh size 7 km) in the operational suite of DWD**
- ➔ **Applicability on a wide range of scales in space and time down to mesh sizes that require a nonhydrostatic dynamical core**
- ➔ **Scalability and efficiency on massively parallel computer architectures with O(10^4+) cores**
- ➔ **At MPI-M: Develop an ocean model based on ICON grid structures and operators; Use limited-area mode of ICON to replace regional climate model REMO.**
- ➔ **Later in this decade: participate in the seasonal prediction project EURO-SIP**

Dynamical core

- **Two-time-level predictor-corrector time stepping scheme; for efficiency reasons, not all terms are evaluated in both sub-steps**
- **implicit treatment of vertically propagating sound waves, but explicit time-integration in the horizontal (at sound wave time step; not split-explicit); larger time step (usually 4x or 5x) for tracer advection / fast physics**
- **Unstructured grid with horizontally indirect addressing and halo points collected at the end of the index vector**
- **Computations are grouped according to the fields they are accessing; “nproma”-blocking with long outer loops to optimize cache efficiency**
- **Option for switching the two inner loop indices for operations with indirect addressing**

Tracer advection

- **Finite-volume tracer advection scheme (Miura) with 2nd-order and 3rd-order accuracy for horizontal advection; extension for CFL values slightly larger than 1 available**
- **2nd-order MUSCL and 3rd-order PPM for vertical advection with extension to CFL values much larger than 1 (partial-flux method)**
- **Option to turn off advection of cloud and precipitation variables (and moisture physics) in the stratosphere**
- **Option for (QV) substepping in the stratosphere**

Physics-dynamics coupling

- **Fast-physics processes (saturation adjustment, turbulence, cloud microphysics, surface coupling) are calculated at 4x or 5x the dynamics time step; even longer time step for slow-physics processes (convection, cloud cover diagnosis, radiation, orographic blocking, sub-grid-scale gravity waves)**
- **Reduced grid option for radiation with special domain decomposition such that at most 60% of the grid points are sunlit at the same time**
- **Moist physics can be turned off above the tropopause**

General aspects

- **Full hybrid (MPI/OpenMP) parallelization**
- **Asynchronous output (parallelization still needs to be done)**
- **Domain-merging and processor splitting options for multiple nested domains at the same nesting level and combination of 1-way and 2-way nesting**

- **ICON is 3x – 4x faster than the present operational hydrostatic global model of DWD (GME)**

WMO standard verification against IFS analysis: 500 hPa geopotential, NH

blue: GME 40 km with IFS analysis, red: ICON 40 km with IFS analysis

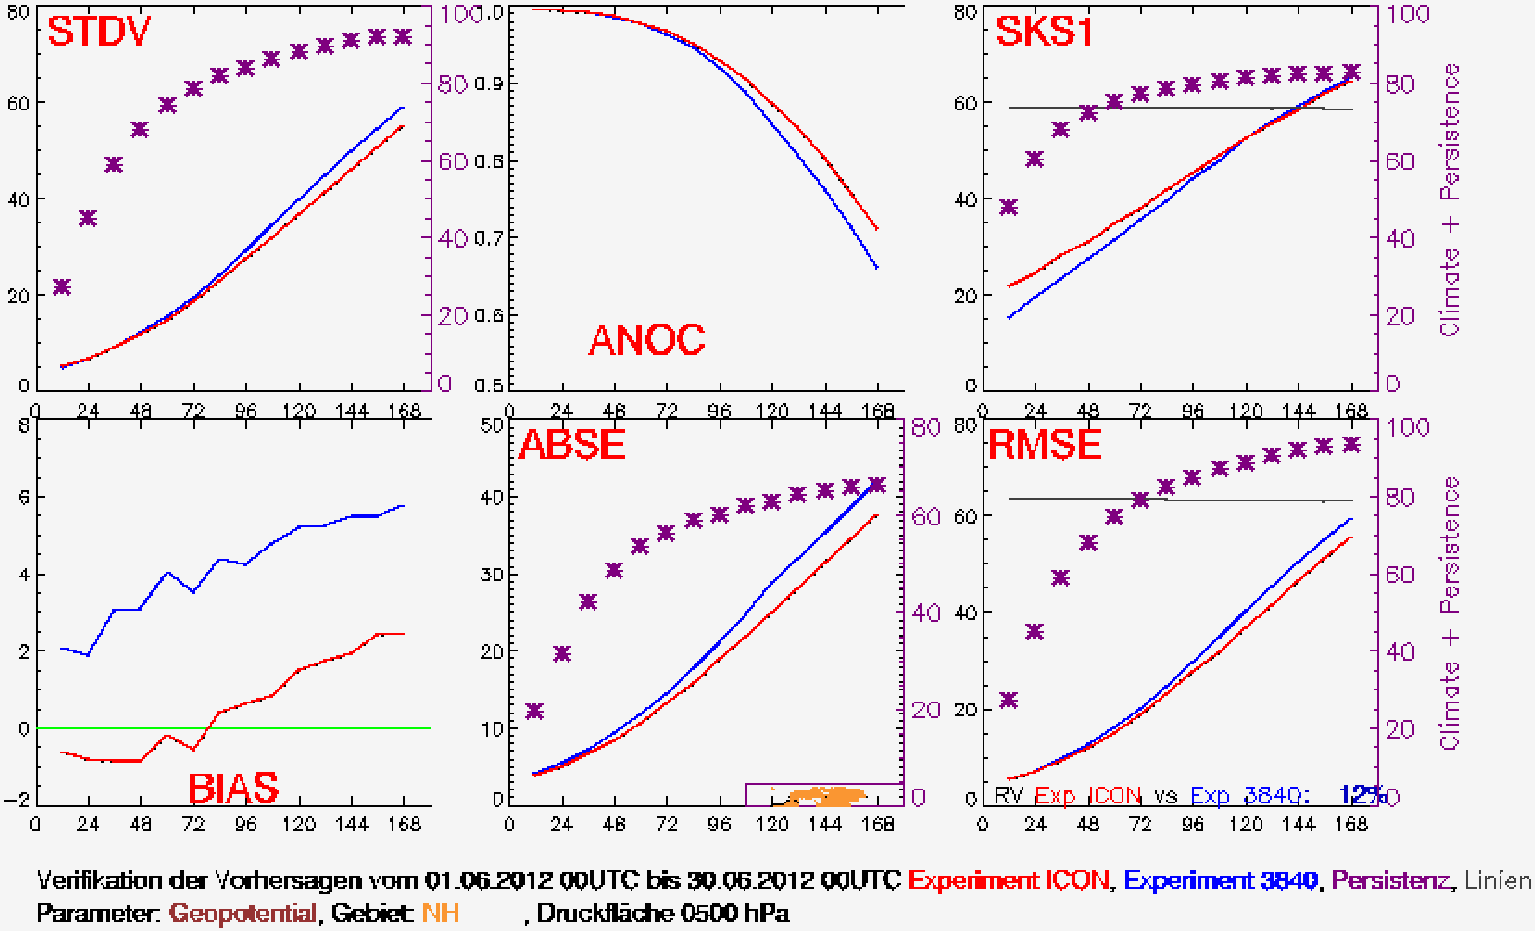

Max-Planck-Institut für Meteorologie

Deutscher Wetterdienst
Wetter und Klima aus einer Hand

WMO standard verification against IFS analysis: 500 hPa geopotential, SH

blue: GME 40 km with IFS analysis, red: ICON 40 km with IFS analysis

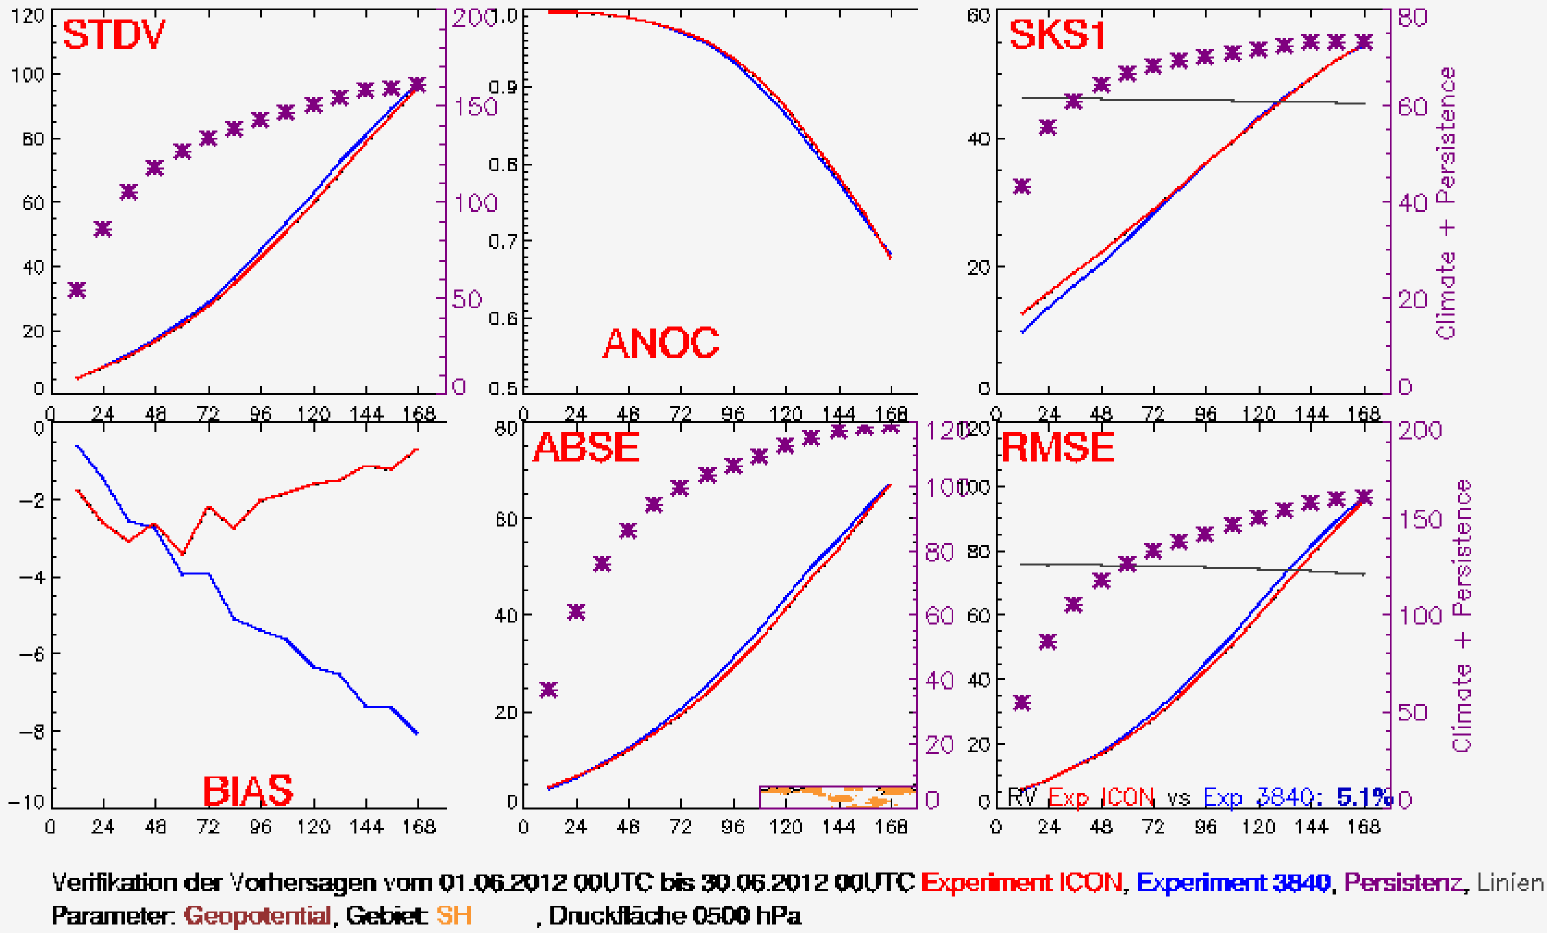

Max-Planck-Institut für Meteorologie

Deutscher Wetterdienst
Wetter und Klima aus einer Hand

WMO standard verification against IFS analysis: 500 hPa geopotential, tropics

blue: GME 40 km with IFS analysis, red: ICON 40 km with IFS analysis

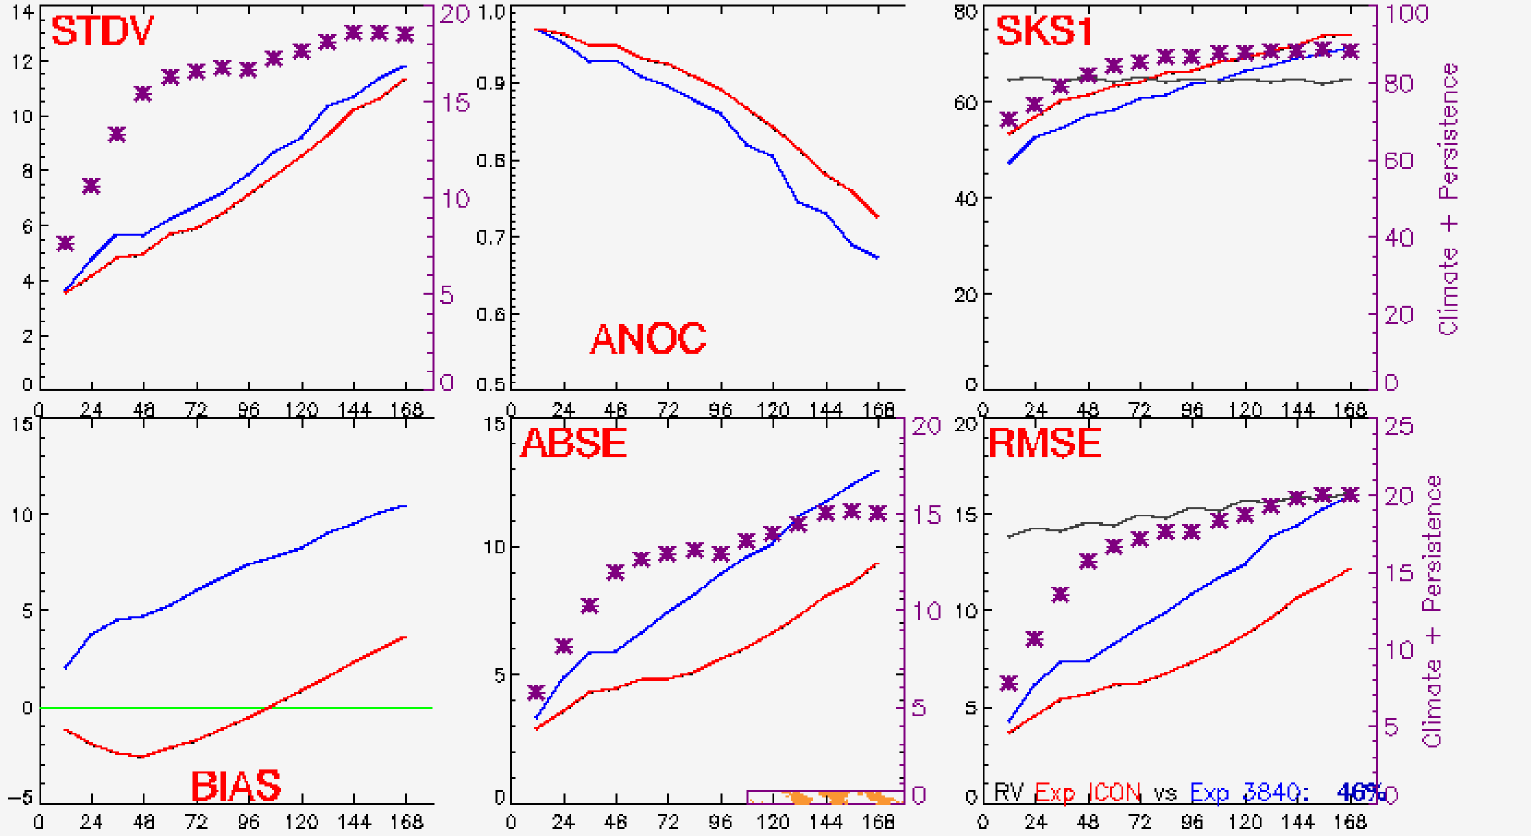

Verifikation der Vorhersagen vom 01.06.2012 00UTC bis 30.06.2012 00UTC Experiment ICON, Experiment 3840, Persistenz, Linien
Parameter: Geopotential, Gebiet: TR , Druckfläche 0500 hPa

WMO standard verification against IFS analysis: 850 hPa temperature, NH

blue: GME 40 km with IFS analysis, red: ICON 40 km with IFS analysis

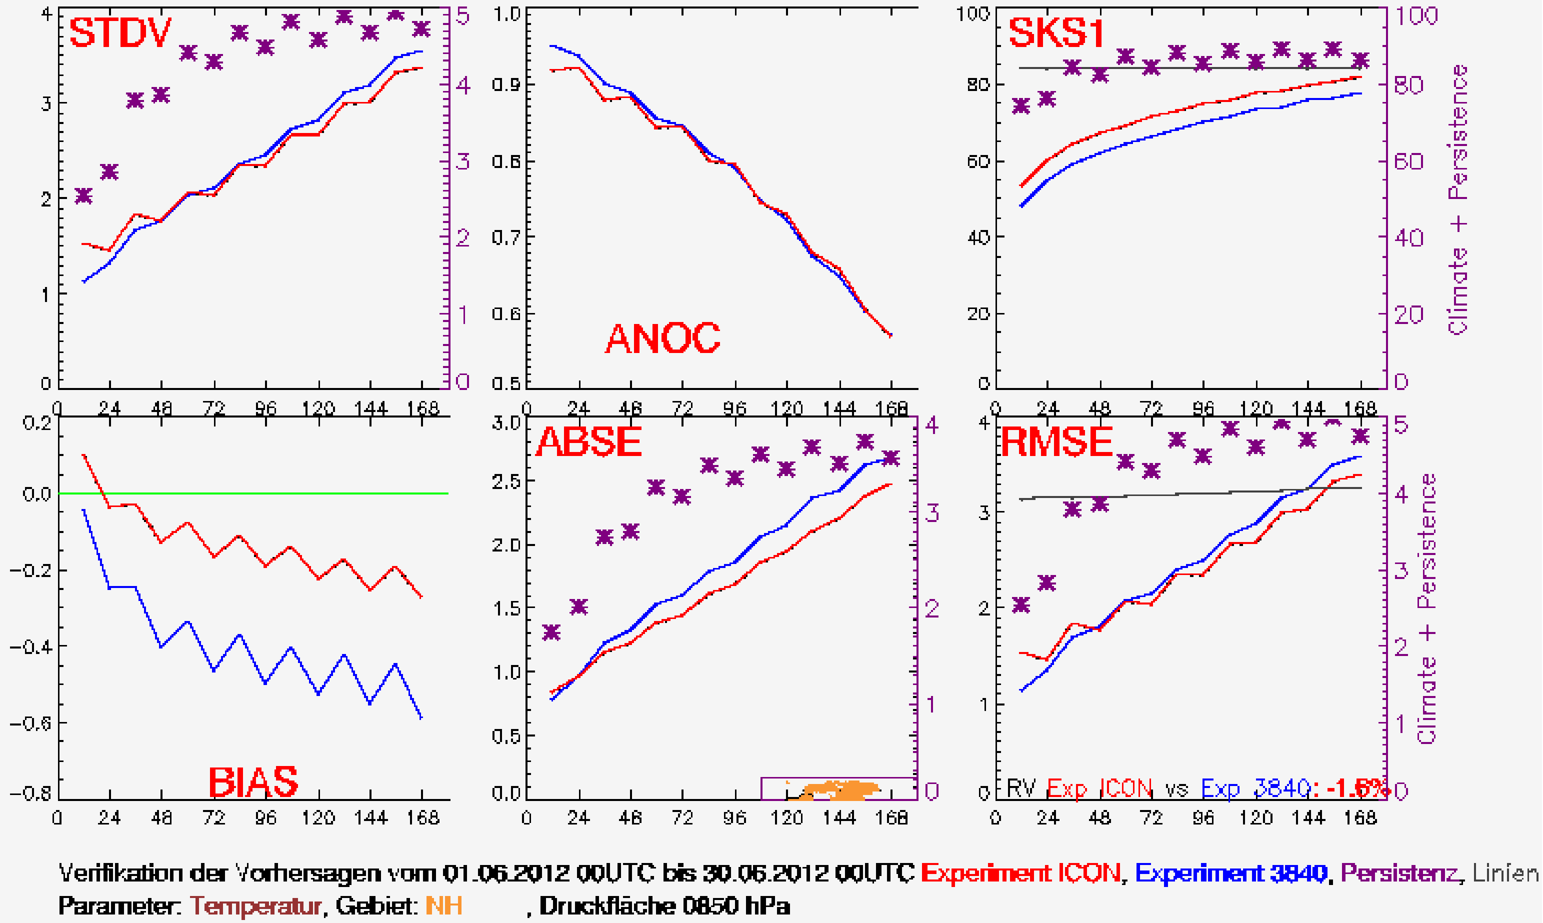

Max-Planck-Institut für Meteorologie

Deutscher Wetterdienst
Wetter und Klima aus einer Hand

WMO standard verification against IFS analysis: 850 hPa temperature, SH

blue: GME 40 km with IFS analysis, red: ICON 40 km with IFS analysis

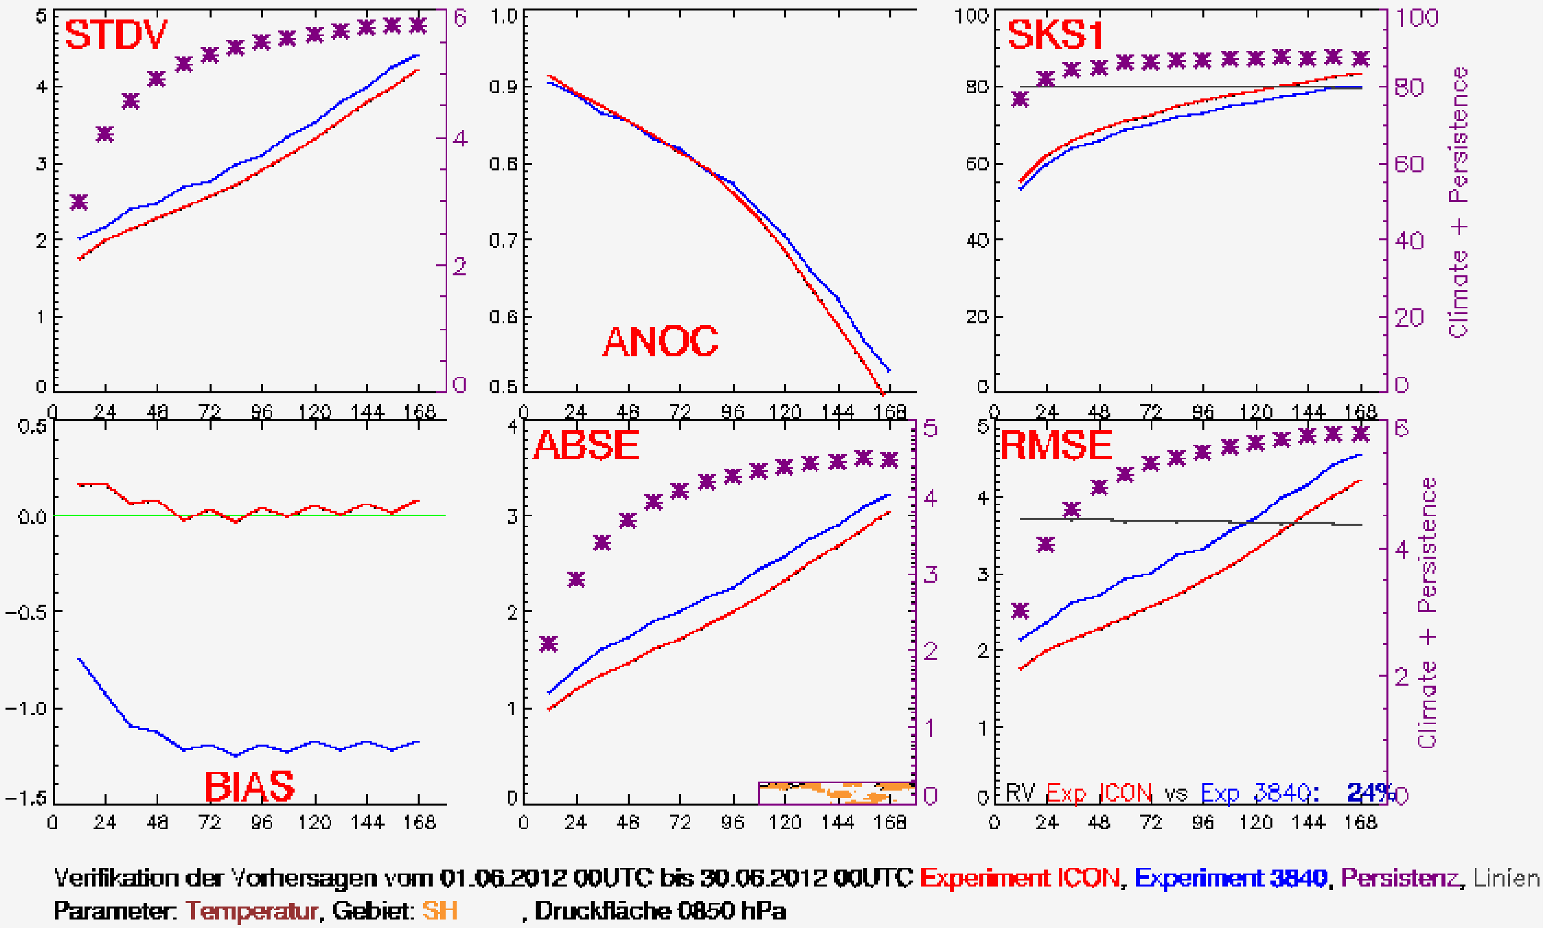

Max-Planck-Institut für Meteorologie

Deutscher Wetterdienst
Wetter und Klima aus einer Hand

Aqua-planet scaling tests, 160/80 km, 96 levels, no output, 32 MPI tasks x 2 OpenMP threads per node; IBM pwr6 @ MPI-M

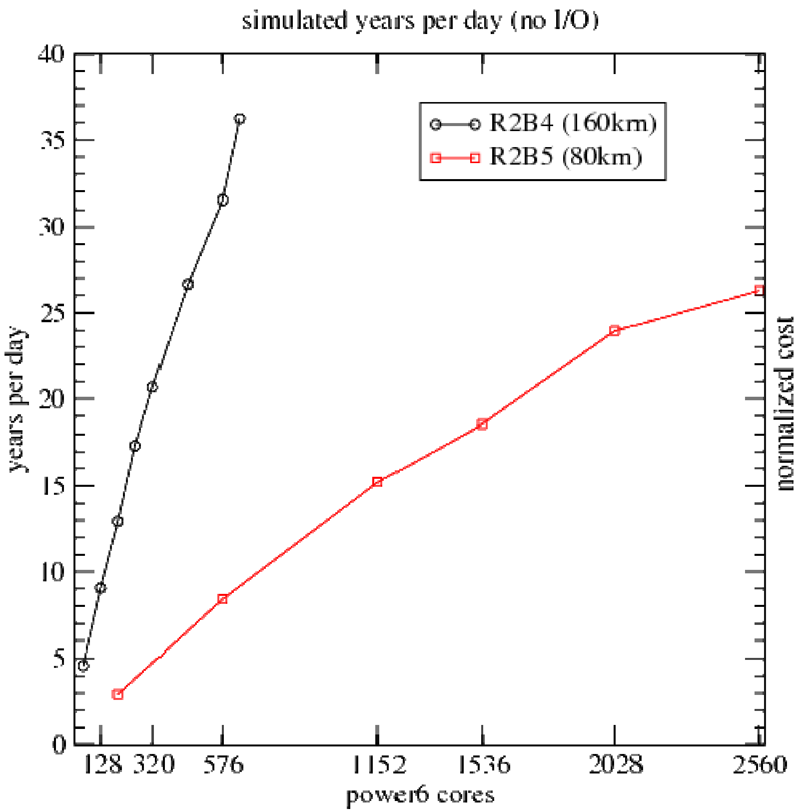

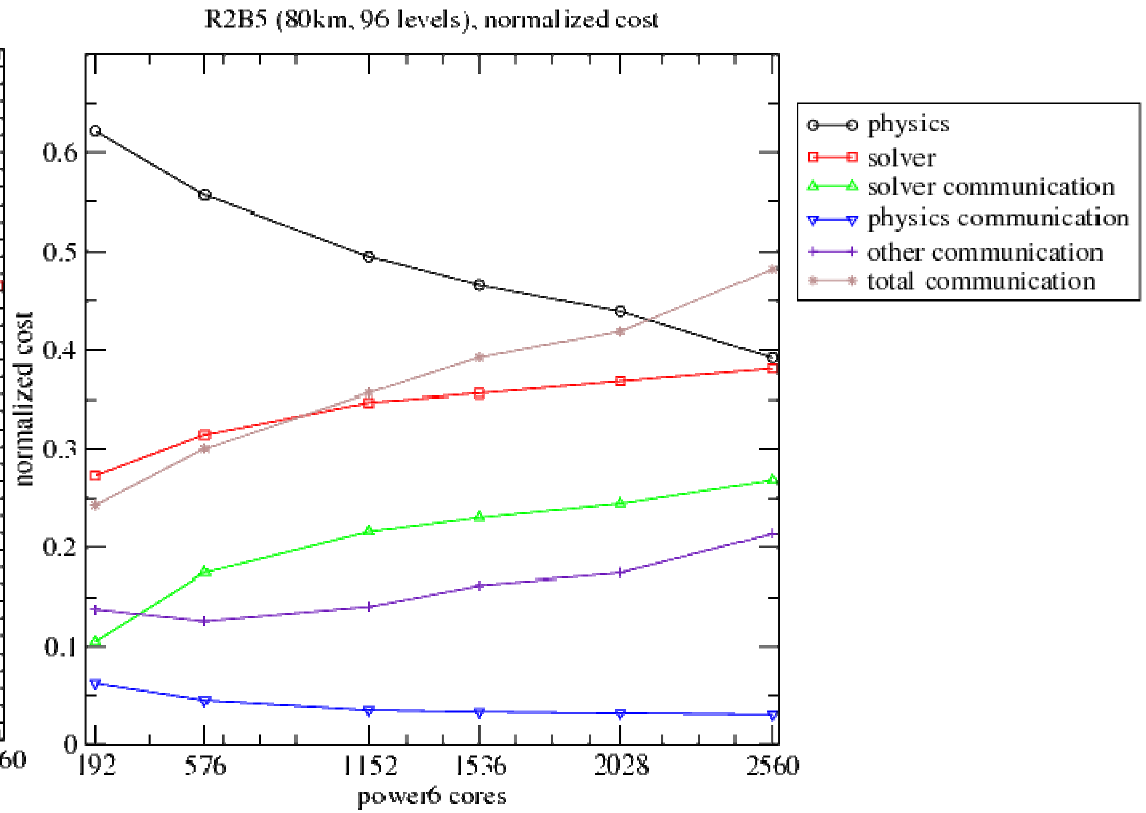

L. Linardakis, MPI-M

Max-Planck-Institut für Meteorologie

Deutscher Wetterdienst
Wetter und Klima aus einer Hand

Real-data scaling test with 40 km / 90 levels, 6-hourly output, 32 MPI tasks x 2 OpenMP threads per node; IBM pwr6 @ ECMWF

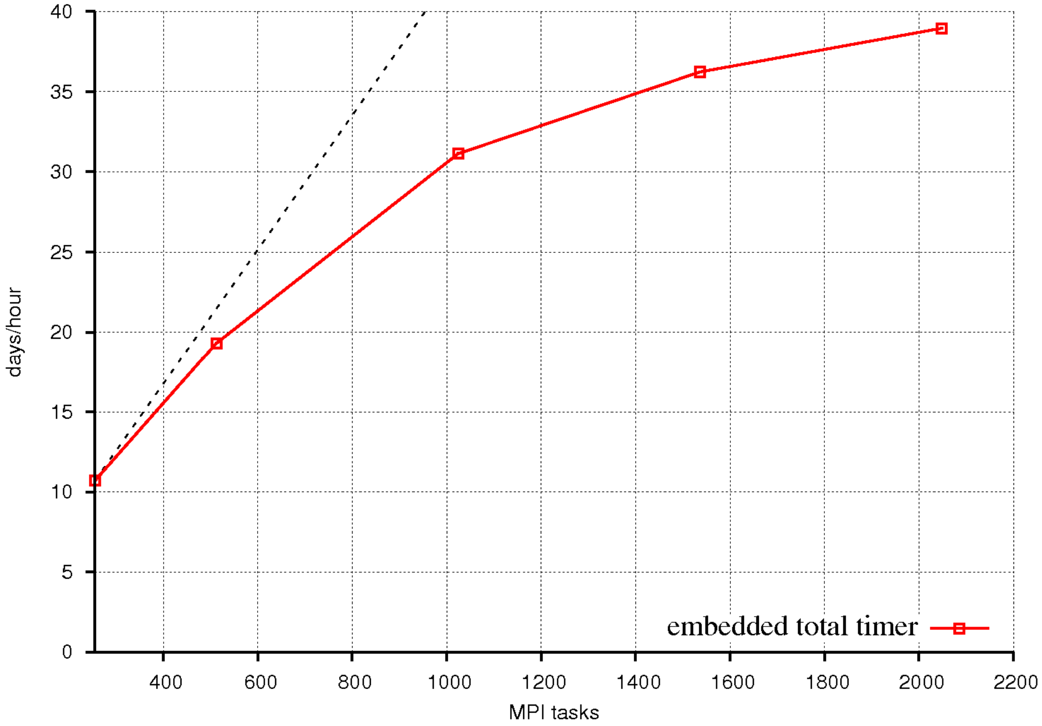

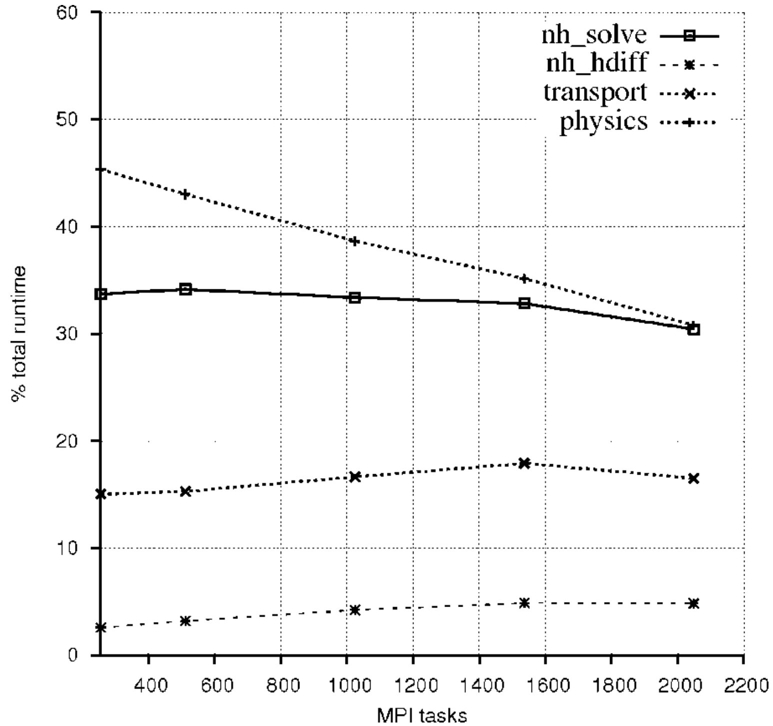

F. Prill, DWD

Real-data test with 20/10/5 km, 90/60/54 levels, 6-hourly output, 32 MPI tasks x 2 OpenMP threads per node; IBM pwr6 @ ECMWF

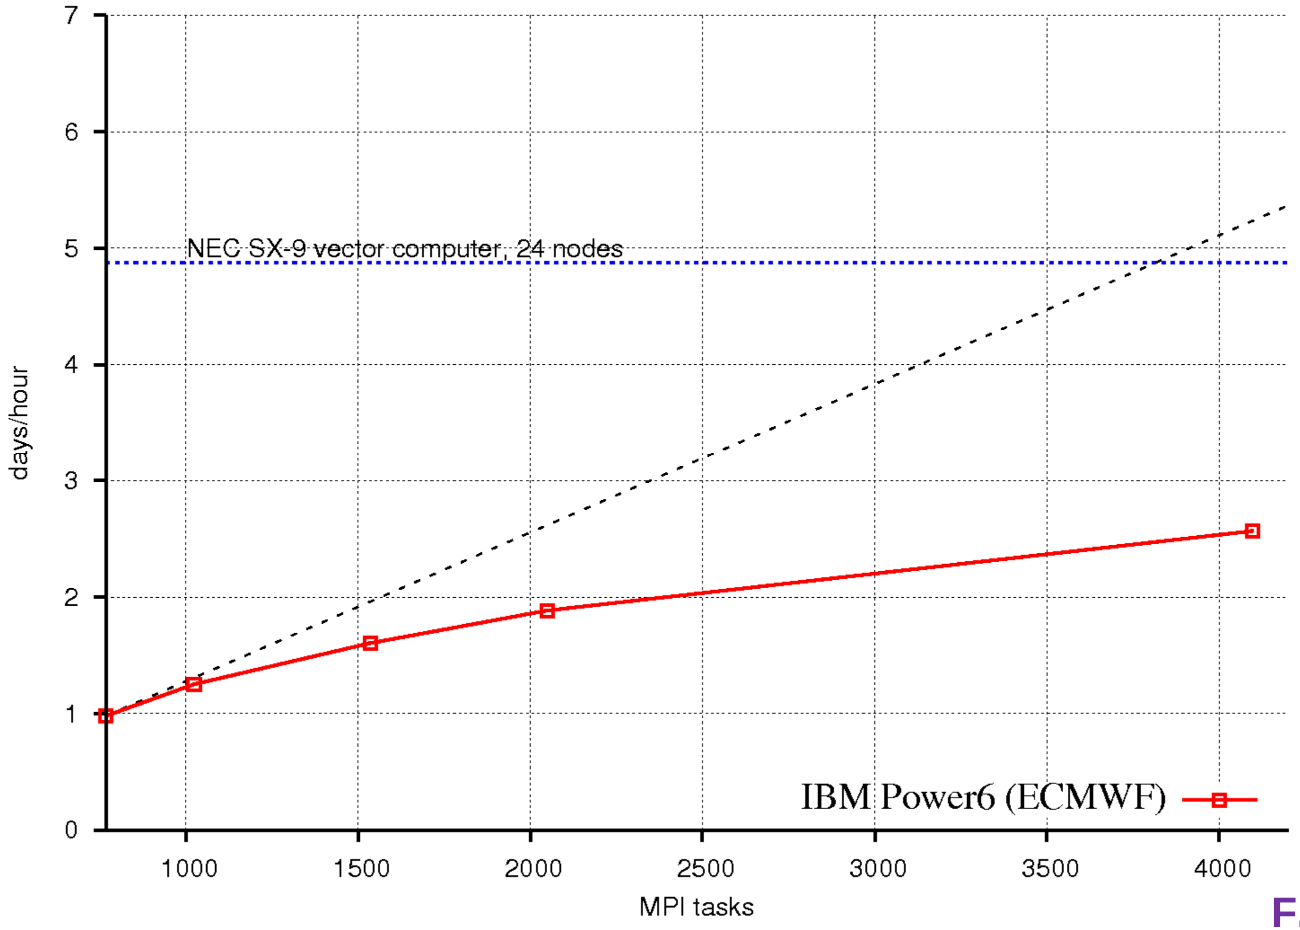

F. Prill, DWD

ICOMEX

Deutscher Wetterdienst
Wetter und Klima aus einer Hand

ICOsahedral-grid Models for EXascale Earth-system simulations

- One of six successful projects of the G8 call issued in 2009
- Six work packages representing four global modeling groups
- Main strategic goal: address a selection of key issues for future cutting-edge computing and develop generic solutions

Project partners:

Günther Zängl (Lead PI, Deutscher Wetterdienst)
Hirofumi Tomita (RIKEN/AICS)
Masaki Satoh (U. Tokyo)
Thomas Ludwig (U. Hamburg/DKRZ)
Leonidas Linardakis(MPI-M)
John Thuburn(U. Exeter/MetOffice)
Thomas Dubos (IPSL/École Polytechnique)

ICOMEX: Participating models

Deutscher Wetterdienst
Wetter und Klima aus einer Hand

- NICAM : Hirofumi Tomita (RIKEN/AICS), Masaki Satoh (U. Tokyo)
- Development started >10 years ago
- world's first global convection-permitting simulations
- Structured icosahedral-hexagonal A-grid

- ICON : Günther Zängl (DWD), Marco Giorgetta (MPI-M)
- Development started ~8 years ago
- Unstructured icosahedral-triangular C-grid
- Two-way nesting and limited-area option

- MPAS : John Thuburn (U. Exeter)
- Development started a few years ago
- Unstructured icosahedral-hexagonal C-grid
- Special care of energy and vorticity budgets

- DYNAMICO : T. Dubos (IPSL/École Polytechnique)
- Work really started only two years ago
- Structured icosahedral-hexagonal C-grid
- So far hydrostatic version only

ICOMEX: Work packages

Deutscher Wetterdienst
Wetter und Klima aus einer Hand

- WP1: Model intercomparison and evaluation (Satoh/Tomita)
- WP2: Abstract model description scheme / domain-specific language (Linardakis)
- WP3: Feasibility study for using GPUs (Dubos)
- WP4: Implicit solvers for massively parallel computing platforms (Thuburn)
- WP5: Parallel internal postprocessing (Thuburn/Dubos)
- WP6: Parallel I/O (Ludwig)

plus

- WP7: Collaboration with hardware vendors (coordinated by Ludwig)

Specific tasks

Deutscher Wetterdienst
Wetter und Klima aus einer Hand

WP1:

- ➔ **Intercomparison of the participating models with respect to scientific and computational aspects**
- ➔ **Test cases: Jablonowski-Williamson baroclinic wave test, Aqua-planet experiments, 30-year AMIP runs**
- ➔ **Make high-performance computing platforms of the project partners mutually accessible**

WP2:

- ➔ **Use domain-specific language (DSL) with automatic Fortran code generation in order to optimize memory layout and loop orders for a variety of platforms**
- ➔ **Testbed: nonhydrostatic dynamical core of ICON**

Specific tasks

WP3:

- ➔ GPU implementation of the hydrostatic dynamical core of DYNAMICO (besides conventional CPU implementation, which is done outside ICOMEX)
- ➔ Comparative testing and benchmarking

WP4:

- ➔ Development of a multigrid elliptic solver for semi-implicit time integration schemes
- ➔ Testbed code: MPAS

Specific tasks

WP5:

- **Development of a parallel internal postprocessing library**
- **First specific task (currently in progress): conservative remapping**
- **Later: combine with asynchronous parallel I/O**

WP6:

- **Development of output benchmarks and systematic testing of exisisting platforms, file systems and libraries**
- **Identification of bottlenecks and development of solutions/workarounds**
- **Development and testing of more efficient file formats**

First results

Scaling intercomparison between NICAM and ICON (sorry, 1 node only)

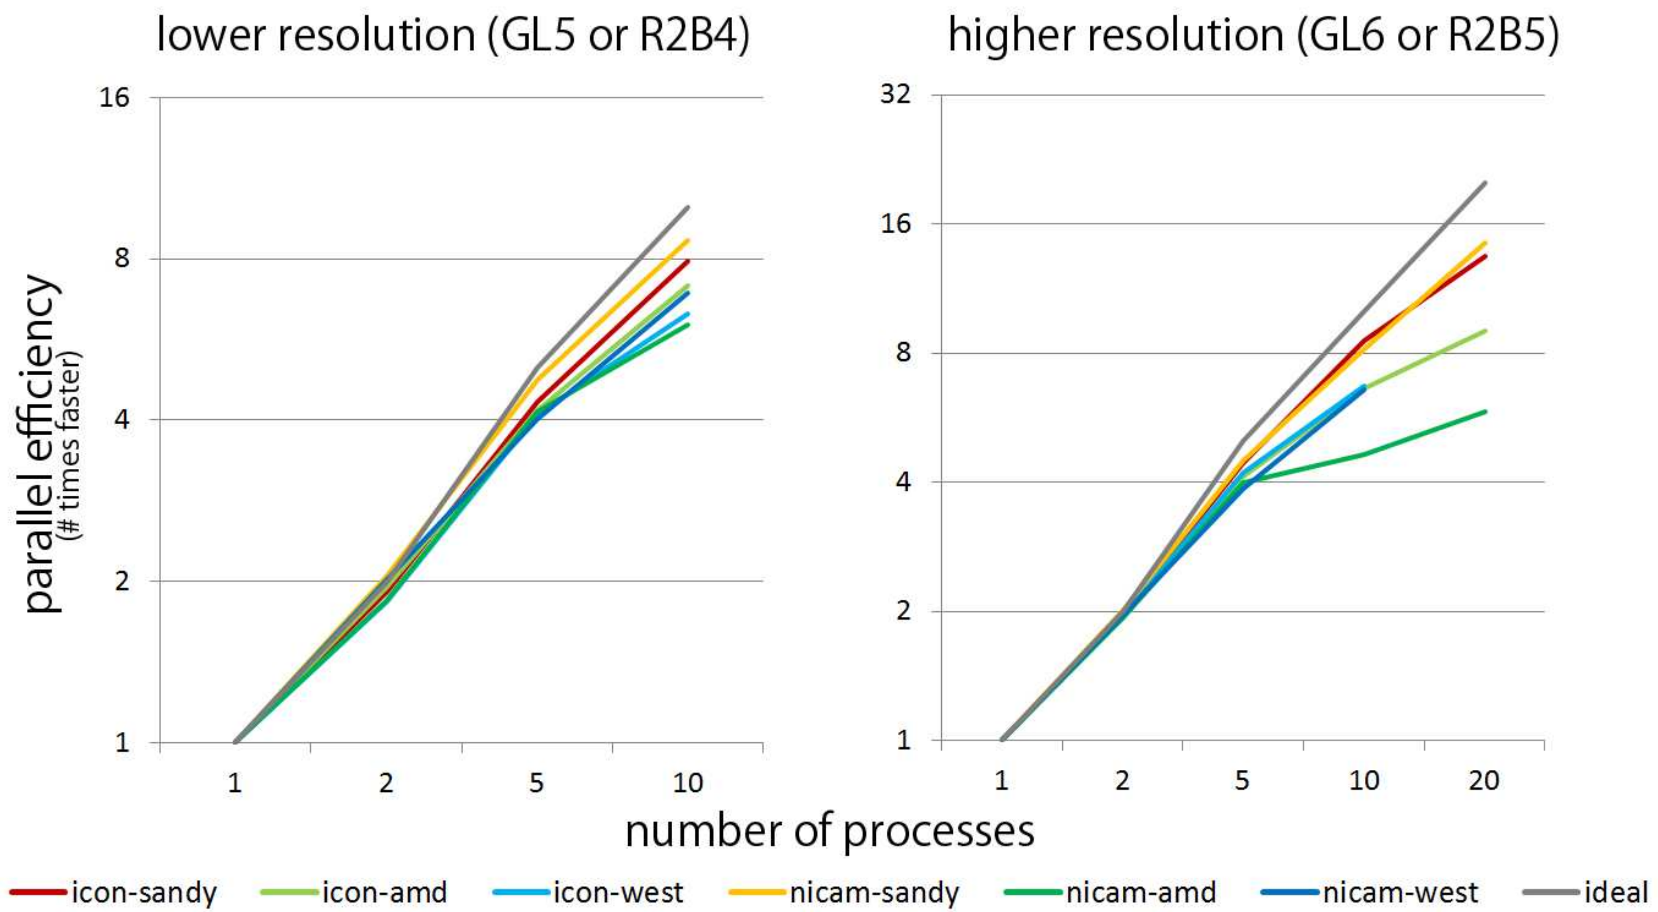

First results

Domain-specific language for ICON: synthetic test for the dynamical core, 20480 cells x 35 levels

Specific issue considered here: optimization of memory layout for IBM pwr6

Cores	32	64	128	192
No DSL time/(cell*iter)	1.574e-06	7.012e-07	3.574e-07	2.777e-07
DSL time /(cell*iter)	1.390e-06	6.008e-07	6.008e-07	2.504e-07
No DSL iterations/sec	635479	1426037	2798150	3601217
DSL iterations/sec	719527	1664402	3096318	3993947
Speed-up	13%	17%	11%	11%

First results

NetCDF write performance of different library builds on the IBM pwr6 @ DKRZ
in parallel runs with 4 processes (MiB/s)

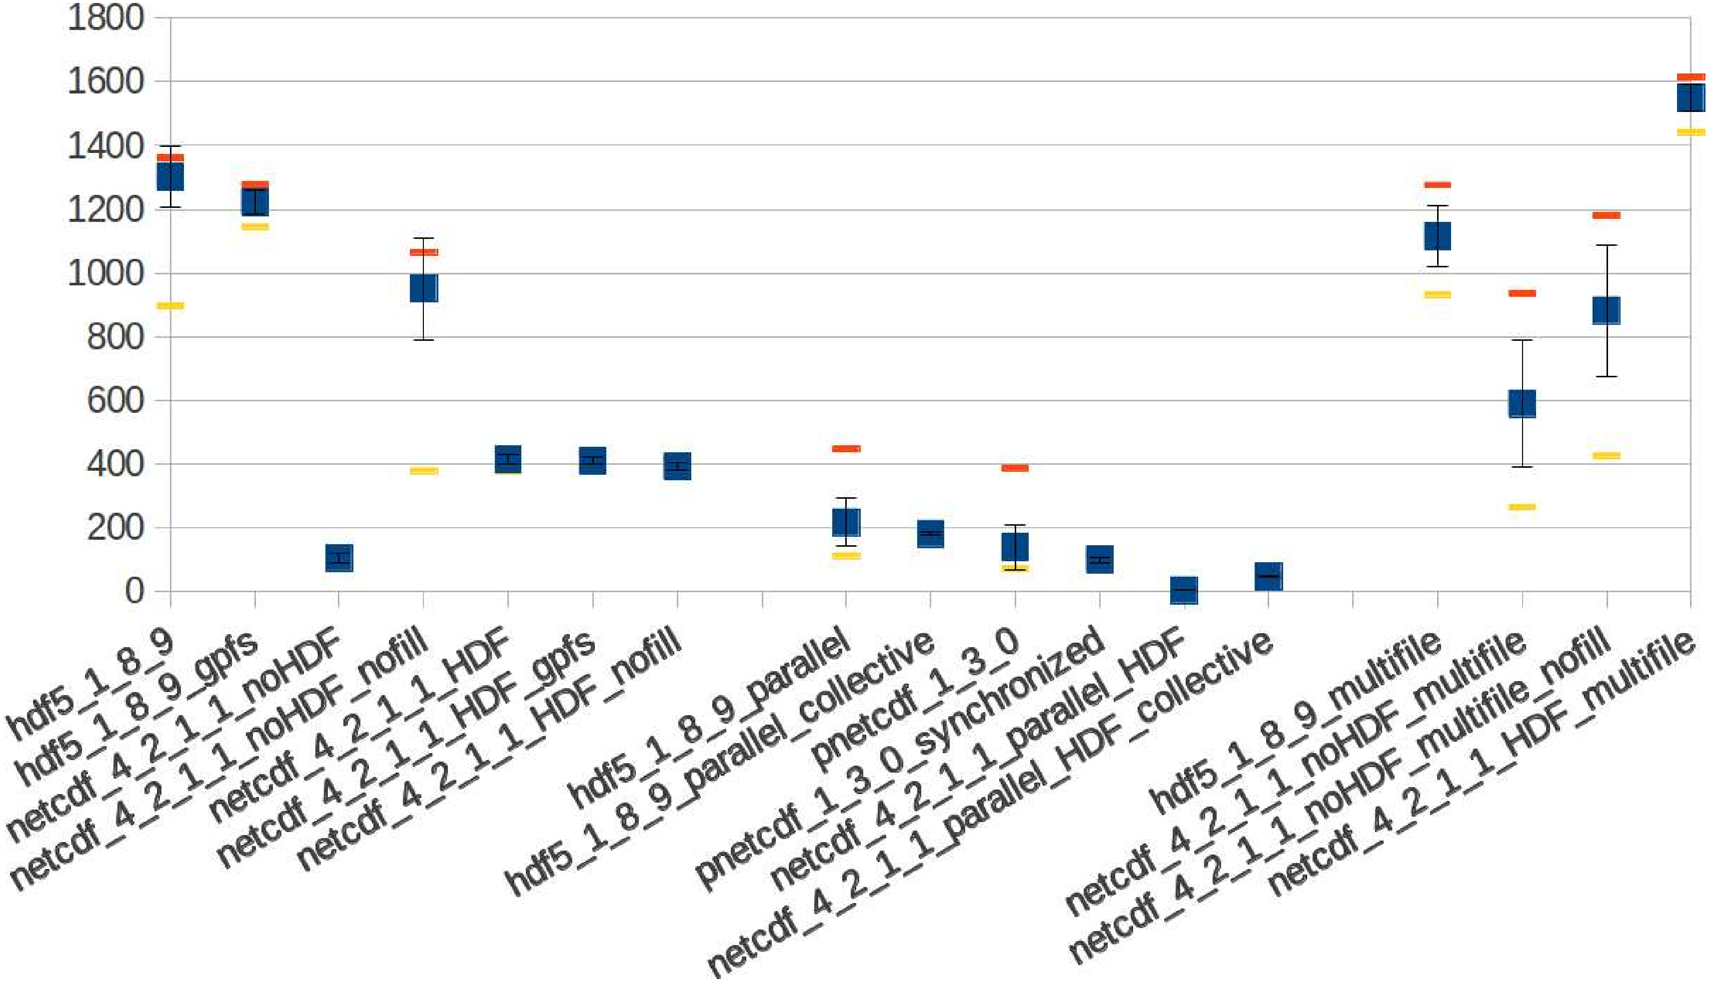

Summary and conclusions

ICON

- **The efficiency of ICON constitutes a substantial progress over the operational hydrostatic GME. Nevertheless, further improvements are needed to its scaling behavior (both memory and compute scaling)**
- **Forecast quality with full physics coupling is comparable with the operational GME even though systematic testing and tuning is only in its initial phase**
- **Next major step: coupling with data assimilation**

ICOMEX

- **Delays in recruiting project scientists in some WP's, therefore varying level of progress**
- **First interesting results are becoming available, indicating significant potential for optimization in several components of our models**