Atmospheric soundproof model with three-directional MPI parallelization

Zbigniew Piotrowski[1], Andrzej Wyszogrodzki[2]
Piotr Smolarkiewicz[2], Oliver Fuhrer[3]

[1] Institute of Meteorology and Water Management – National Research Institute, Warsaw
[2] National Center for Atmospheric Research, Boulder, CO
[3] MeteoSwiss, Zurich

INNOVATIVE ECONOMY
NATIONAL COHESION STRATEGY

EUROPEAN UNION
EUROPEAN REGIONAL DEVELOPMENT FUND

NCAR

FNP
Foundation for Polish Science

Grants for innovation

IMGW

EULAG model – powerful virtual laboratory

What does the application do?

NCAR

Simulating thermo-fluid flows across a range of scales and physics

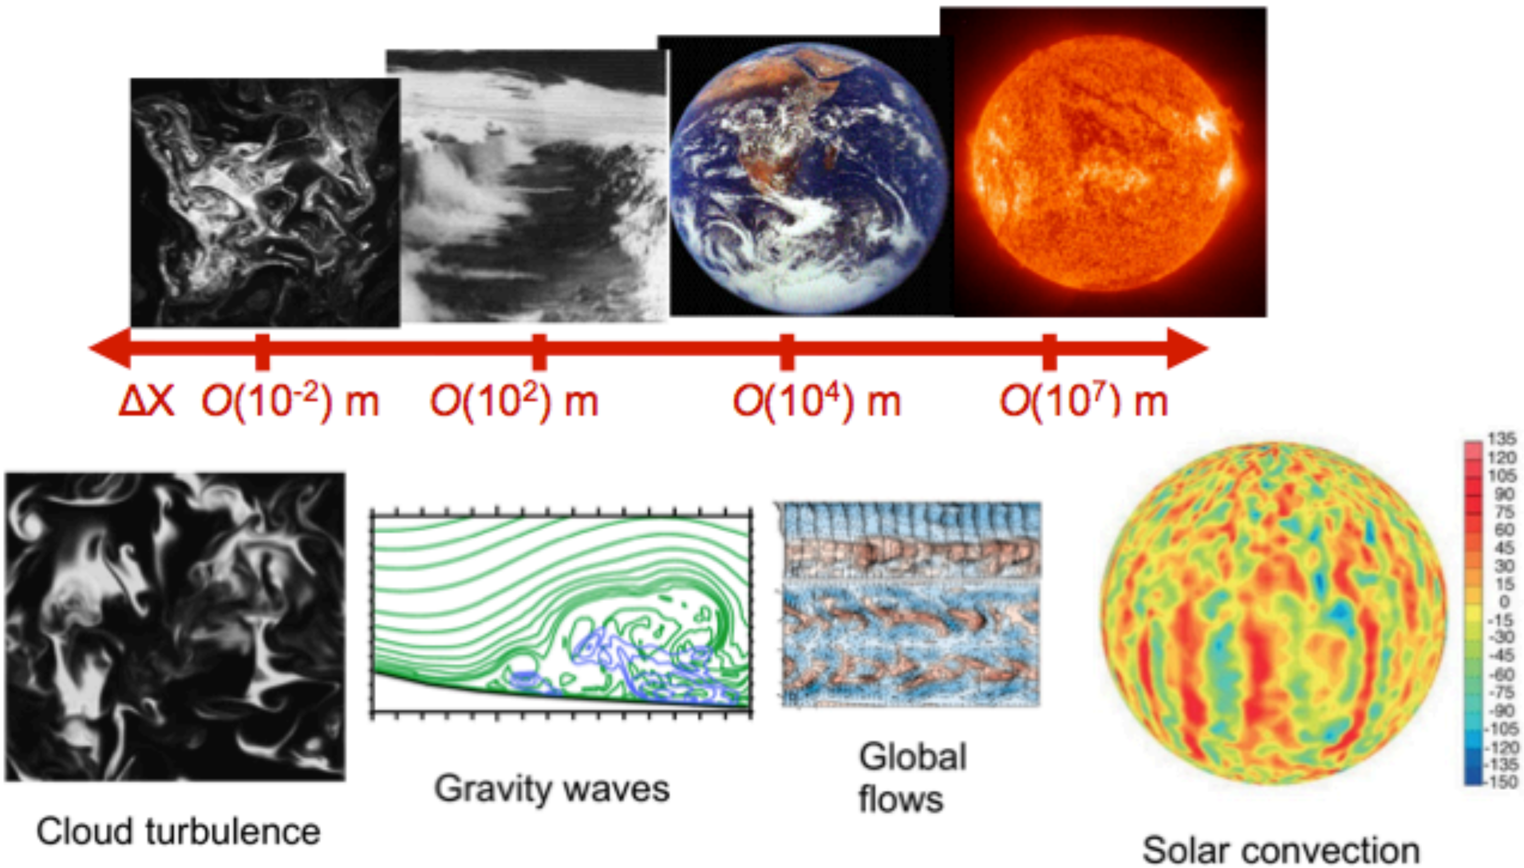

Scalable, multiscale parallel solution necessary !

MERIT

NATIONAL SCIENCE CENTRE

- Prototype COSMO model with anelastic dynamical core based on EULAG has landed,

- Coupled to a limited set of COSMO parameterizations: fluxes, moist physics, radiation, turbulence;

- Capable of forecasting the weather using raw (unsmoothed) SRTM Alpine topography with at least 0.55 km grid spacing;

- Produces results quantitatively comparable to Runge-Kutta dynamical core for large scales, but with more detail in small scales;

- Subject of „CELO" Priority Project of COSMO Consortium, aiming at full operational status in 3 year perspective (data assimilation features still in planning stage);

- Rewrite for GPU and IntelMIC processors within National Science Centre grant "Methods and algorithms for organization of computations in the class of anelastic numerical models for geophysical flows on modern computer architectures with realization in the EULAG model ", 2012-2015, **with 3D-MPI parallelized model as a CPU reference version.**

Numerical design of EULAG

All principal forcings are assumed to be unknown at *n*+1

$$\psi_{\mathbf{i}}^{n+1} = LE_{\mathbf{i}}(\psi^n + 0.5\Delta t R^n) + 0.5\Delta t R_{\mathbf{i}}^{n+1}$$

$\Rightarrow$ system implicit with respect to all dependent variables.

On grids co-located with respect to all prognostic variables, it can be inverted algebraically to produce an elliptic equation for pressure

$$\left\{\frac{\Delta t}{\rho^*}\overline{\nabla}\cdot\rho^*\widetilde{\mathbf{G}}^T\left[\widehat{\widehat{\mathbf{v}}} - (\mathbf{I} - 0.5\Delta t\widehat{\mathbf{R}})^{-1}\widetilde{\mathbf{G}}(\overline{\nabla}\pi'')\right]\right\}_{\mathbf{i}} = 0$$

$$\overline{\mathbf{v}}^s \equiv \overline{\mathbf{v}}^* - \frac{\partial\overline{\mathbf{x}}}{\partial t} \quad \textit{contravariant velocity} \quad \overline{\mathbf{v}}^* \equiv d\overline{\mathbf{x}}/d\overline{t} \equiv \dot{\overline{\mathbf{x}}}$$

$$\widetilde{\mathbf{G}}^T[\widehat{\widehat{\mathbf{v}}} - (\mathbf{I} - 0.5\Delta t\widehat{\mathbf{R}})^{-1}\widetilde{\mathbf{G}}(\overline{\nabla}\pi'')] \equiv \overline{\mathbf{v}}^s$$

Boundary conditions on π'' Imposed on $\overline{\mathbf{v}}^s \bullet \mathbf{n}$ subject to the integrability condition

$$\int_{\partial\Omega}\rho^*\overline{\mathbf{v}}^s \bullet \mathbf{n}\,d\sigma = 0$$

Boundary value problem is solved using nonsymmetric Krylov subspace solver - a preconditioned generalized conjugate residual GCR(*k*) algorithm (Smolarkiewicz and Margolin, 1994; Smolarkiewicz et al., 2004)

COSMO-EULAG horizontal velocity at 10m – comparison with RK

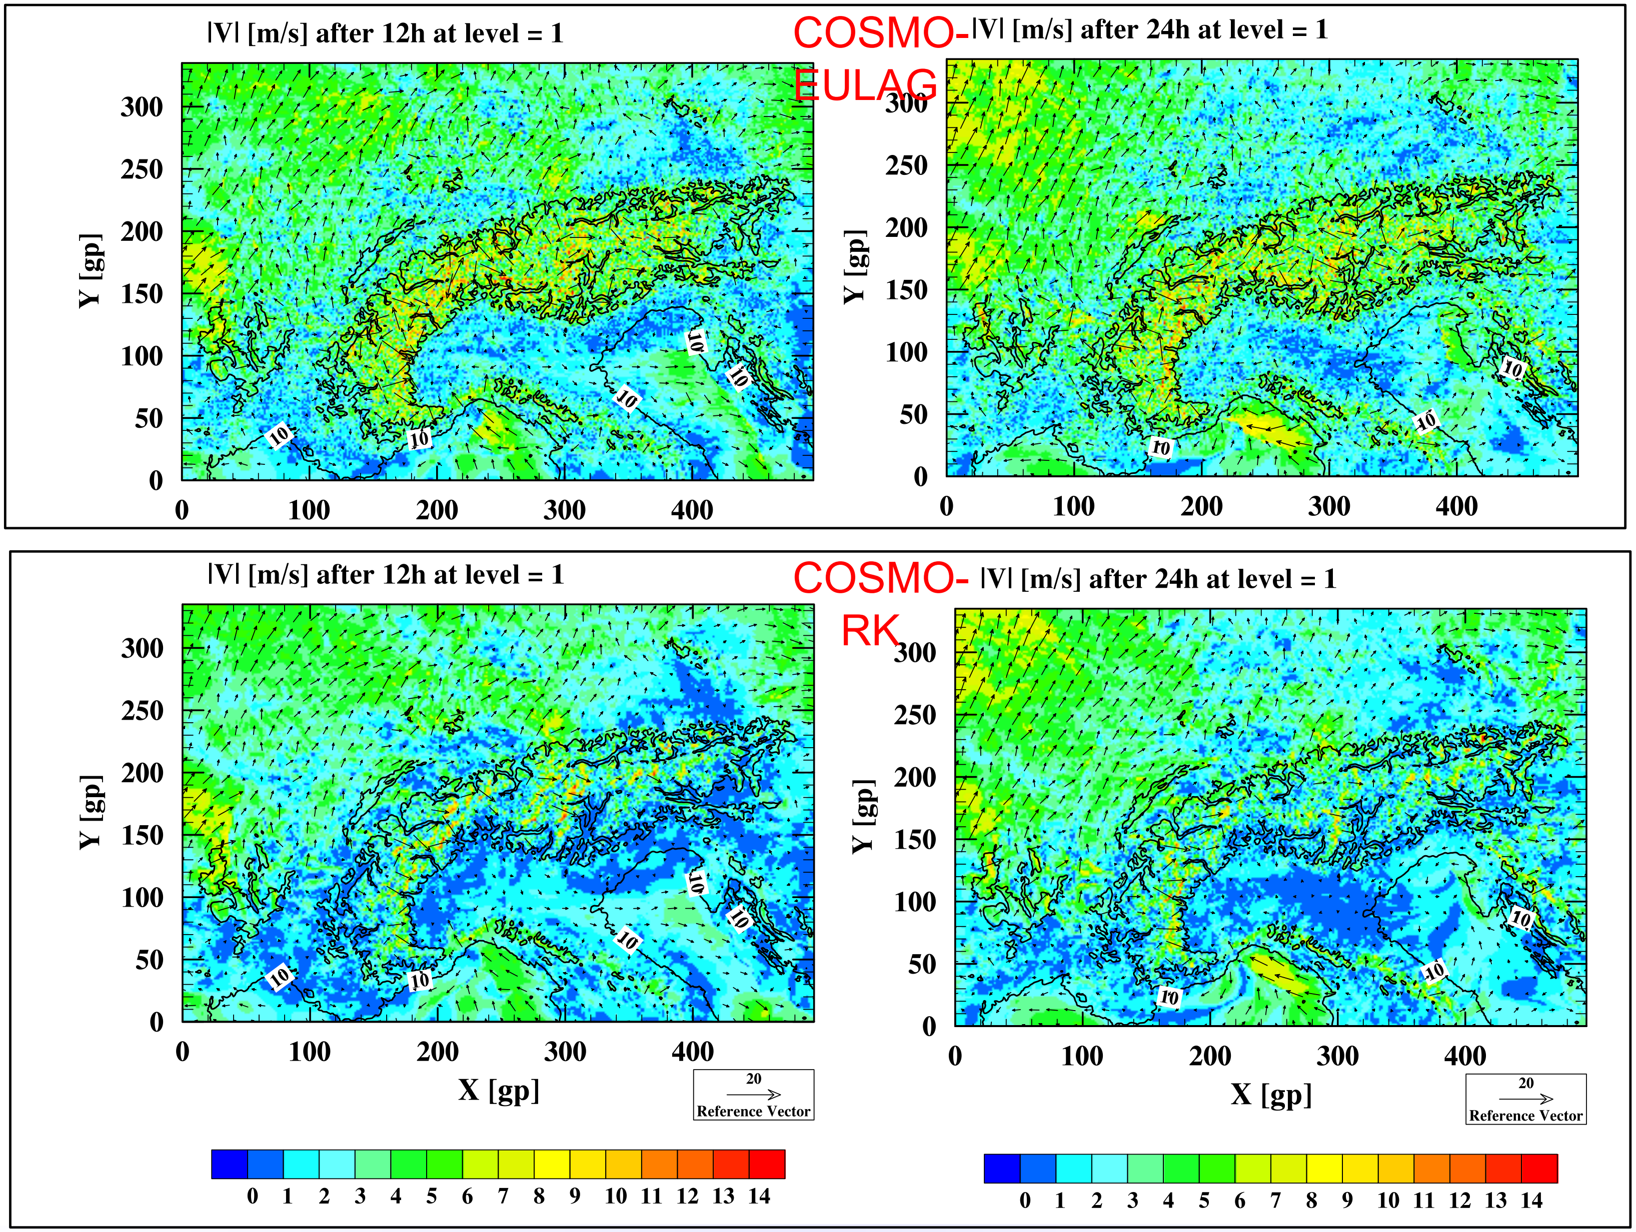

Team of dr B. Rosa, dr M. Kurowski, dr M. Ziemianski, D Wojcik and myself at COSMO General Meeting, Lugano, September 2012

COSMO-EULAG horizontal velocity at 10m – 2.2, 1.1 and .55 km res.

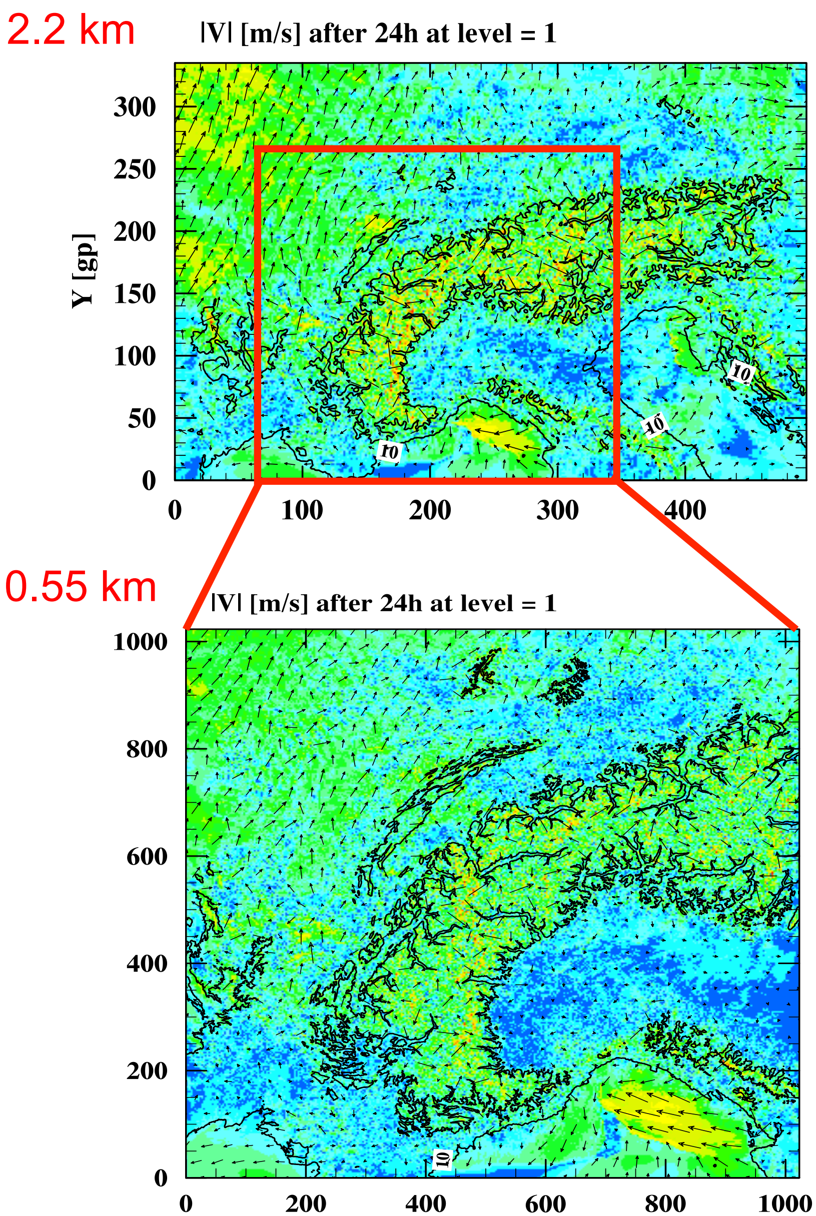

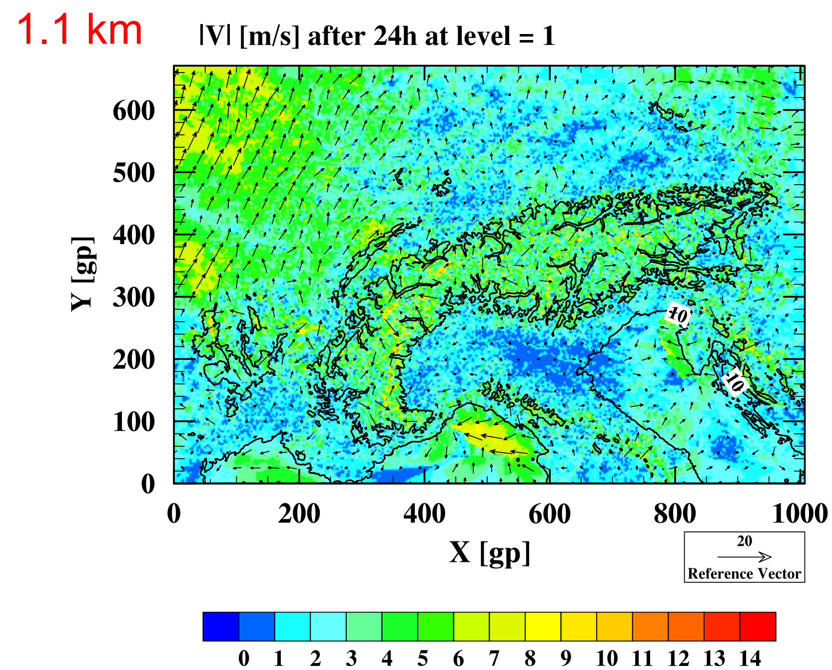

COSMO-EULAG employs traditional horizontal MPI decomposition.

HIGHLIGHT OF TO-DATE PUBLISHED RESULTS ON 3D MPI - EULAG

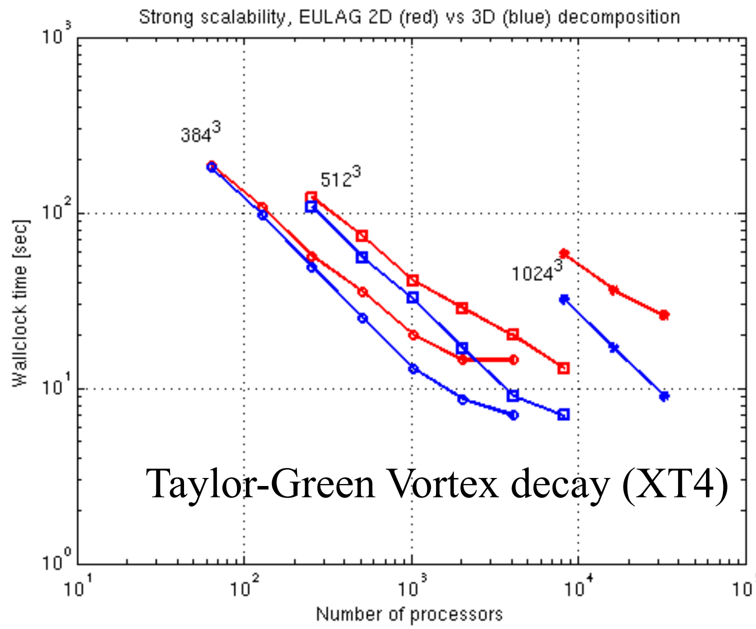

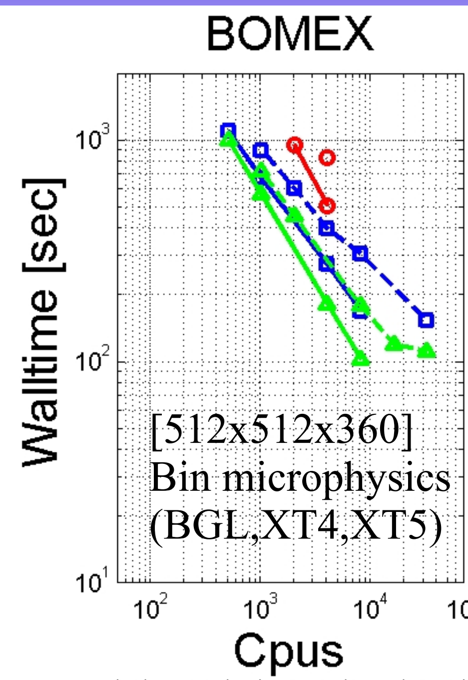

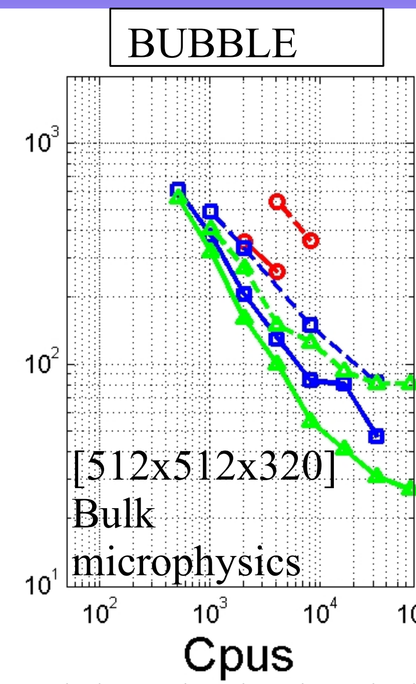

2012 A.A. Wyszogrodzki, Z.P Piotrowski and W.W. Grabowski, "Parallel implementation and scalability of cloud resolving EULAG model", Parallel Processing and Applied Mathematics, Lecture Notes in Computer Science, 7204, 252-261

2011 Z.P. Piotrowski, A.A. Wyszogrodzki, P.K. Smolarkiewicz, "Towards petascale simulation of atmospheric circulations with soundproof equations", Acta Geophys., 59, 2011, 1294 - 1311

NWP – like experiment (using EULAG standalone)

- **Using EULAG as an idealized vehicle carrying a set of physical parameterizations rather than tuning it to provide reasonable forecast product (this delegated to COSMO-EULAG Dynamical Core)**

- **Initial and boundary conditions interpolated from COSMO model 7 km**

- **Grids representing ~2.2 km, ~1.1 km and ~.55 km regional NWP over the Alps, 64 vertical levels, in horizontal: 512x256, 1024x512, 2048x1024 grid points, 1440 timesteps, timestep small enough so it doesn't need to be changed with resolution.**

- **Focus on 16x16 and 8x8 grid points horizontal size of computational subdomain, we check series of 64, 32, 16 and 8 grid points in the vertical and see if we get a better performance and good scaling when adding cores in vertical.**

- **Components tested: MPDATA advection schemes, GCR preconditioned elliptic solver, bulk microphysics with simple ice (Grabowski), simple radiation (Grabowski based on Stevens et al 2005, MWR, p. 1443.), 3D diffusion with simple surface layer model and TKE parameterization.**

We acknowledge that the following results have been achieved using the PRACE Research Infrastructure resource JUQUEEN (BG/Q) and HERMIT (XE6) based in Germany at Juelich and Stuttgart, respectively.

NWP – like experiment (EULAG standalone)

JUQUEEN BG/Q scalability results

NWP – like experiment (EULAG standalone)

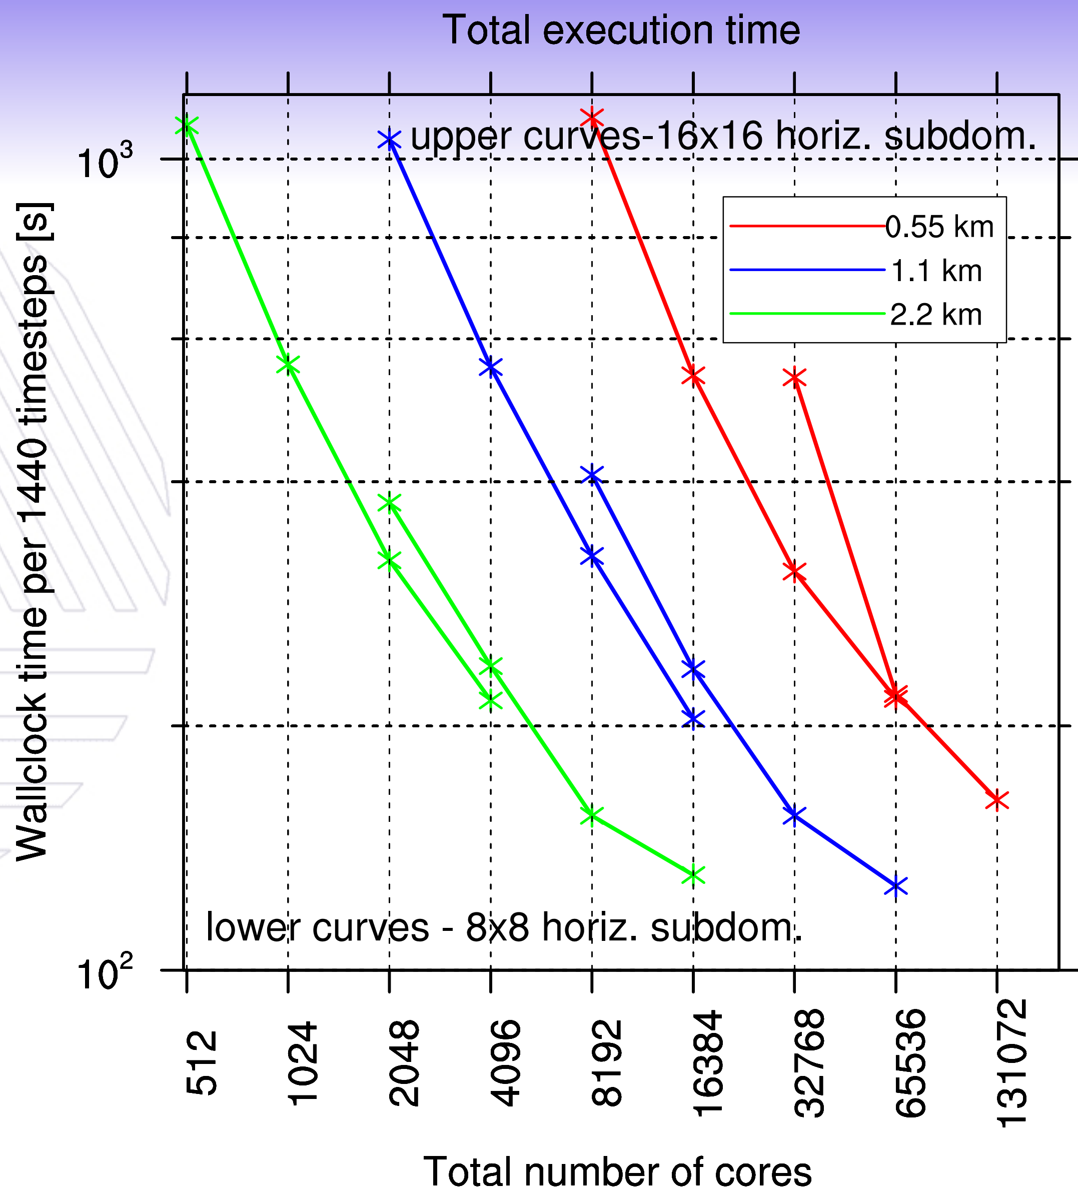

Disclaimer:

Total time does NOT represent time-to-solution from the client (weather service) point of view.

Setup is too simplified, only mimicking numerical weather prediction.

Work in progress.

NWP – like experiment (EULAG standalone)

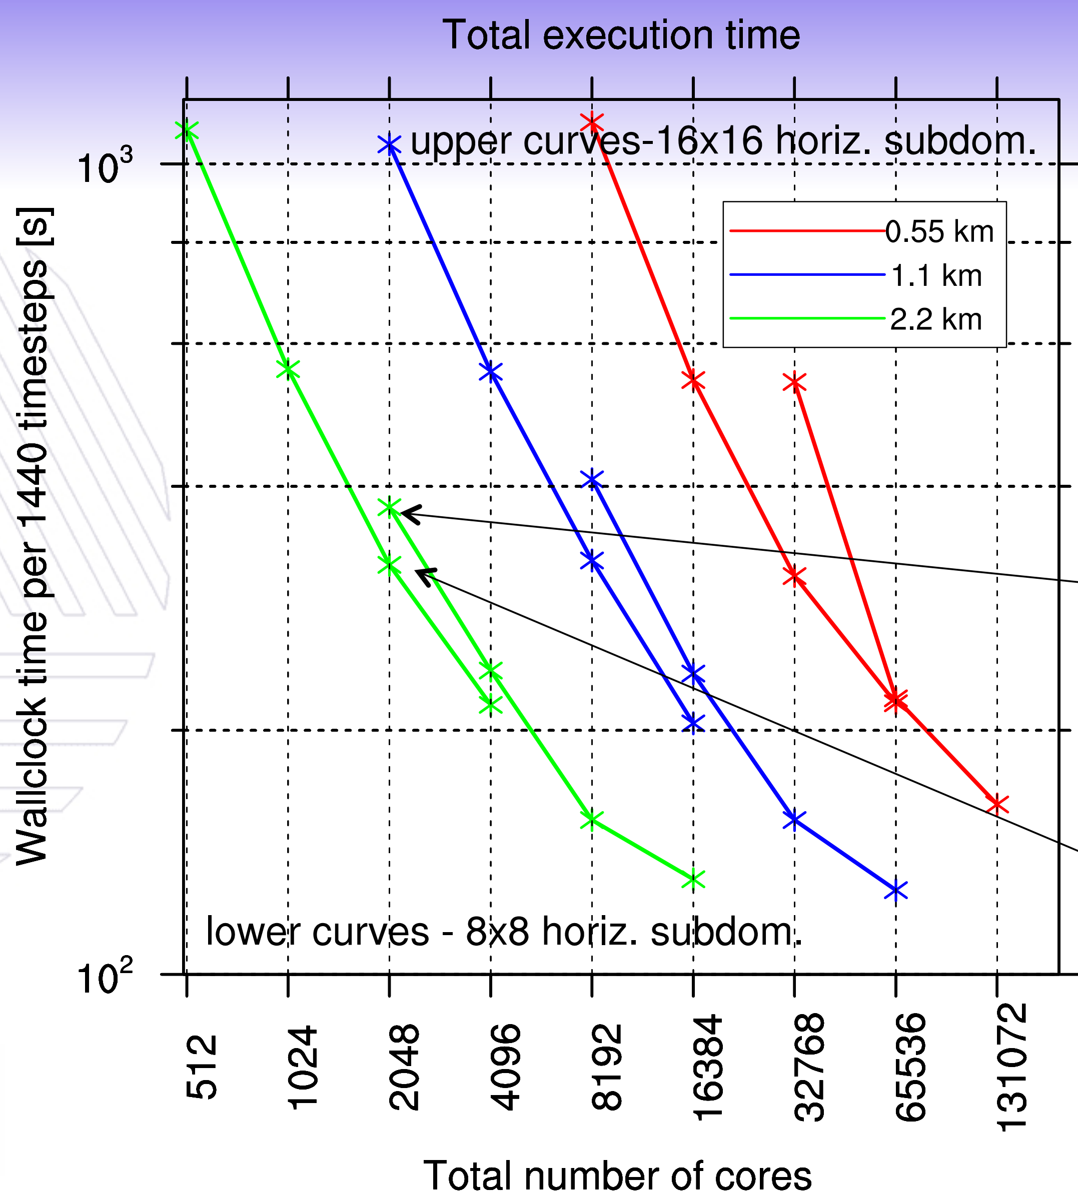

64x32x1 processor grid,
8x8x64 grid points subdomain

Note lower execution time when decomposing in vertical for the same number of cores used !

32x16x4 processor grid,
16x16x16 points subdomain

NWP – like experiment (EULAG standalone)

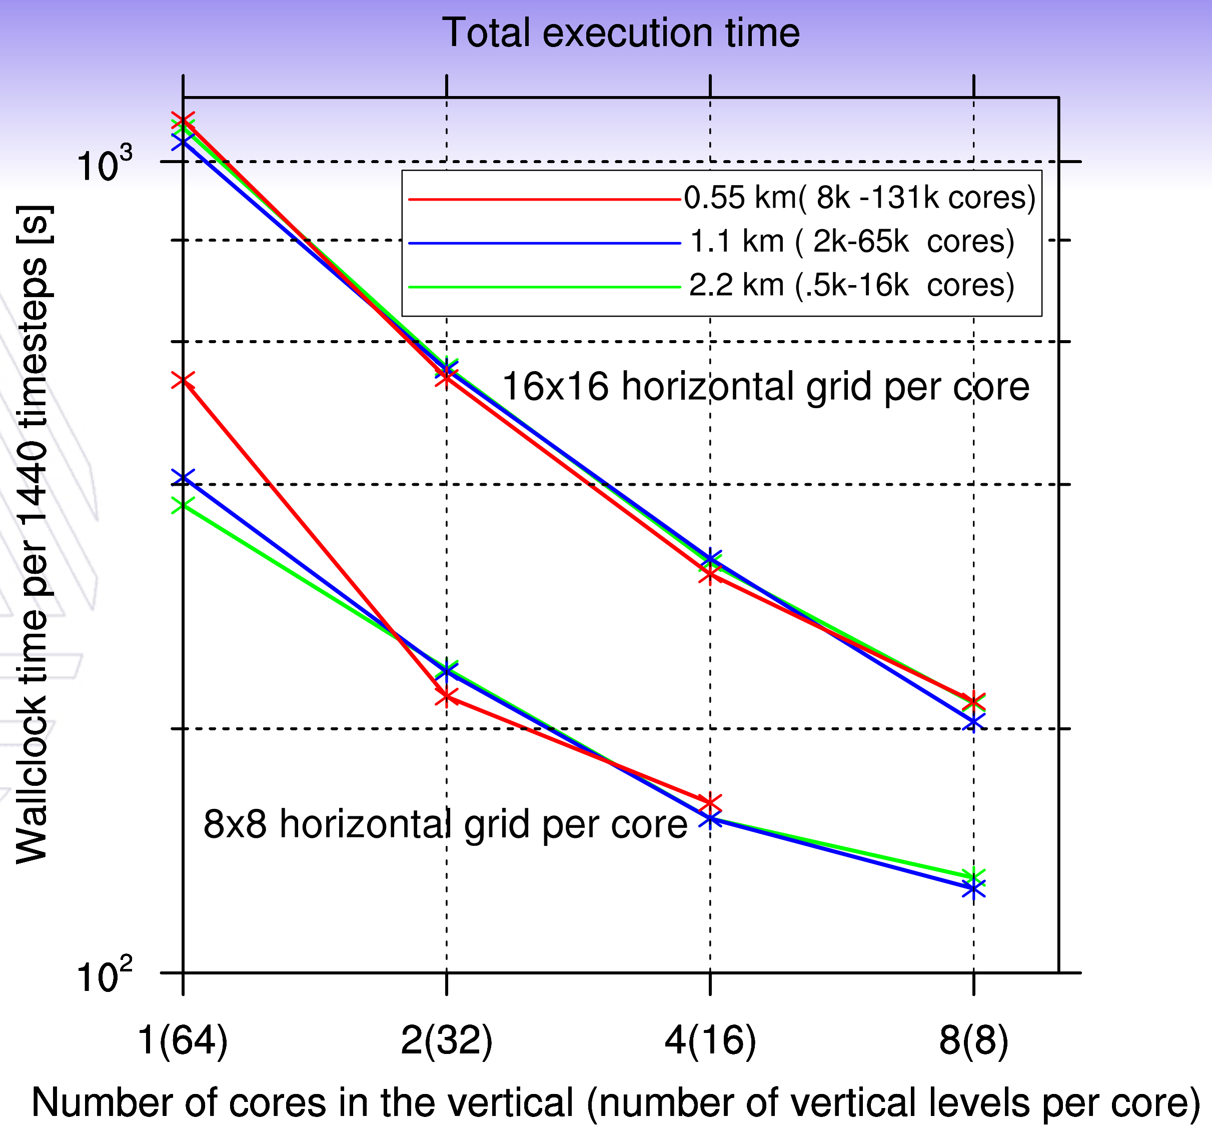

NWP – like experiment (EULAG standalone)

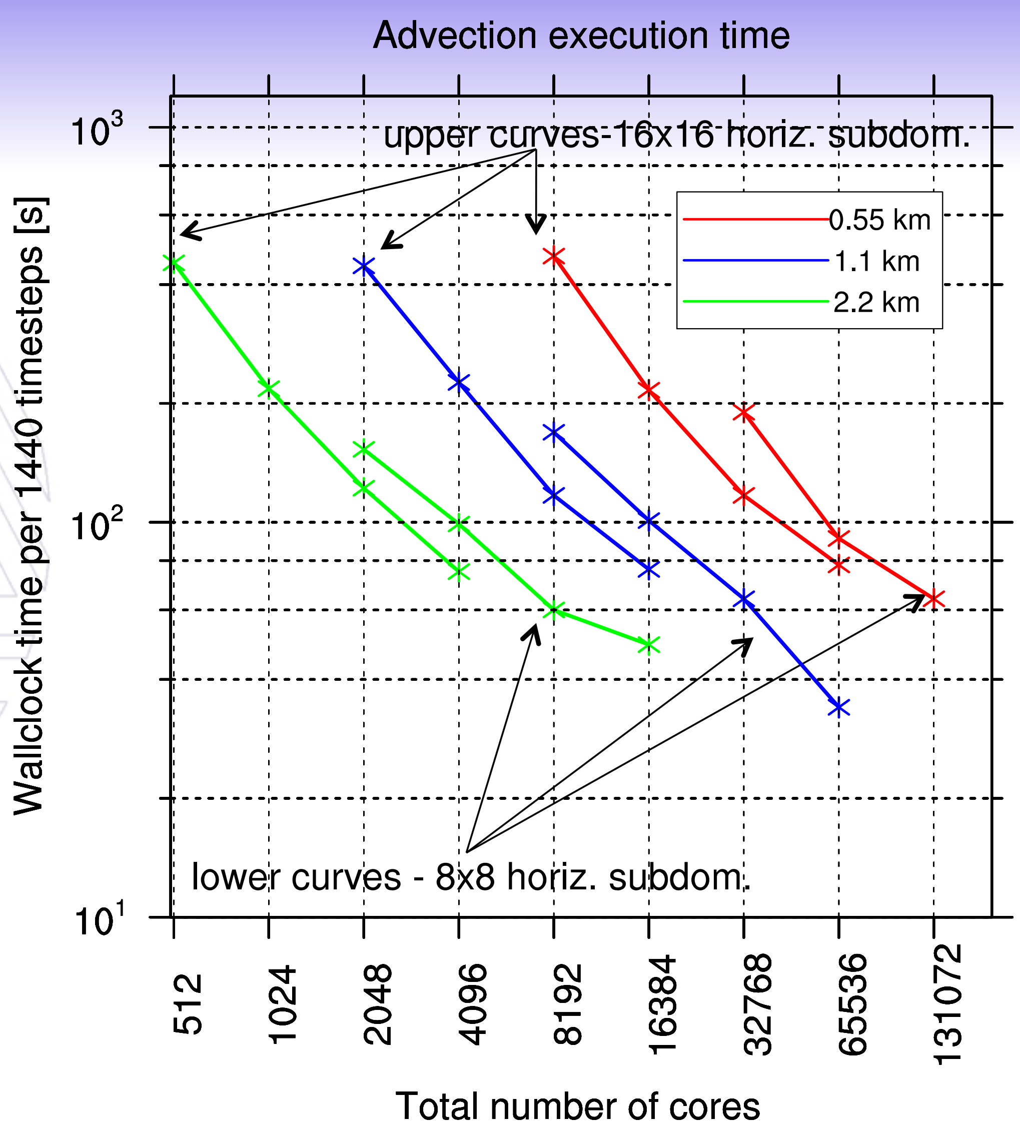

NWP – like experiment (EULAG standalone)

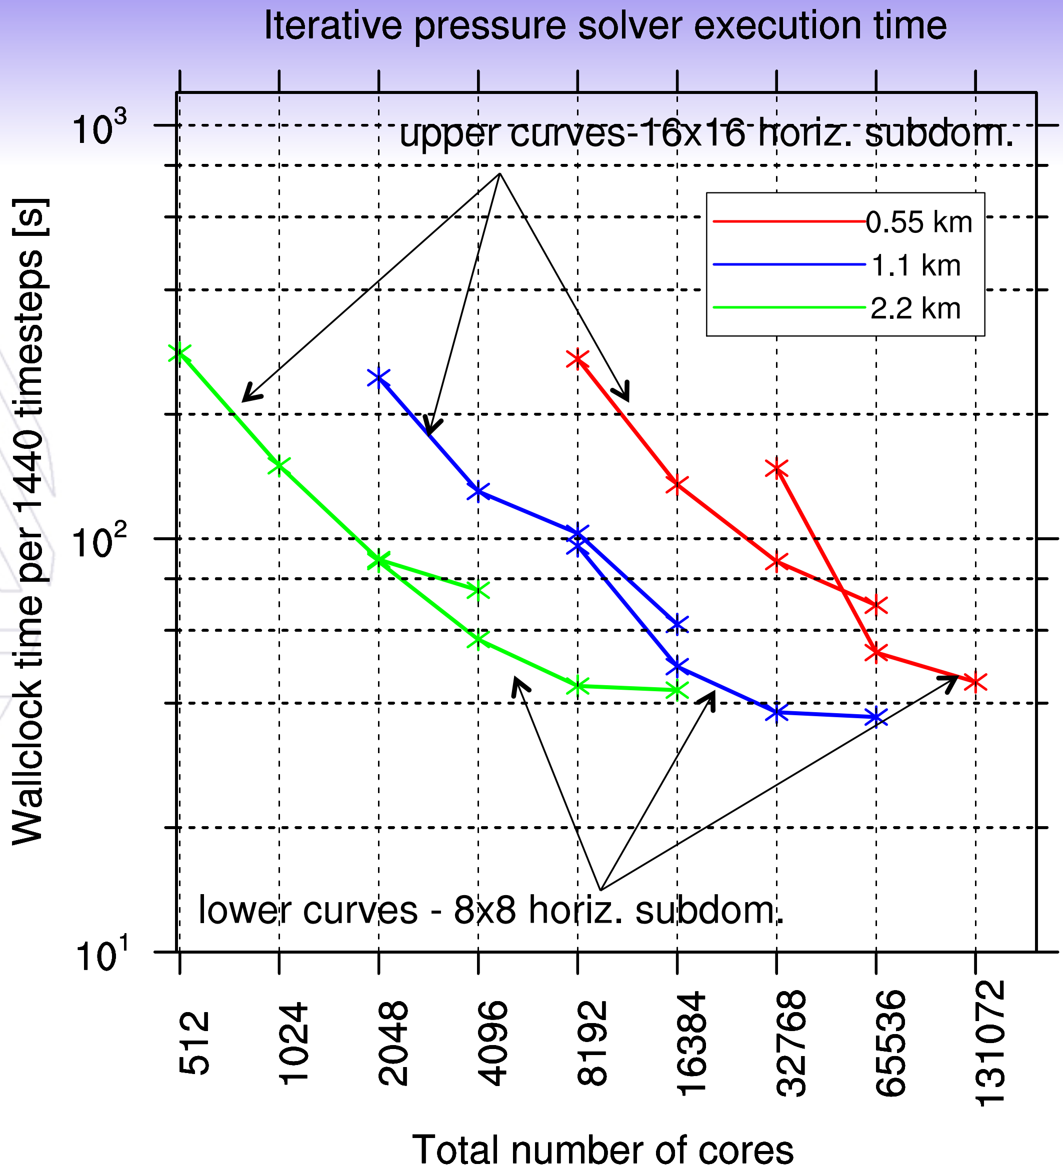

Problems with the reproducibility of the results (possible I/O problem), so these results only preliminary.

NWP – like experiment (EULAG standalone)

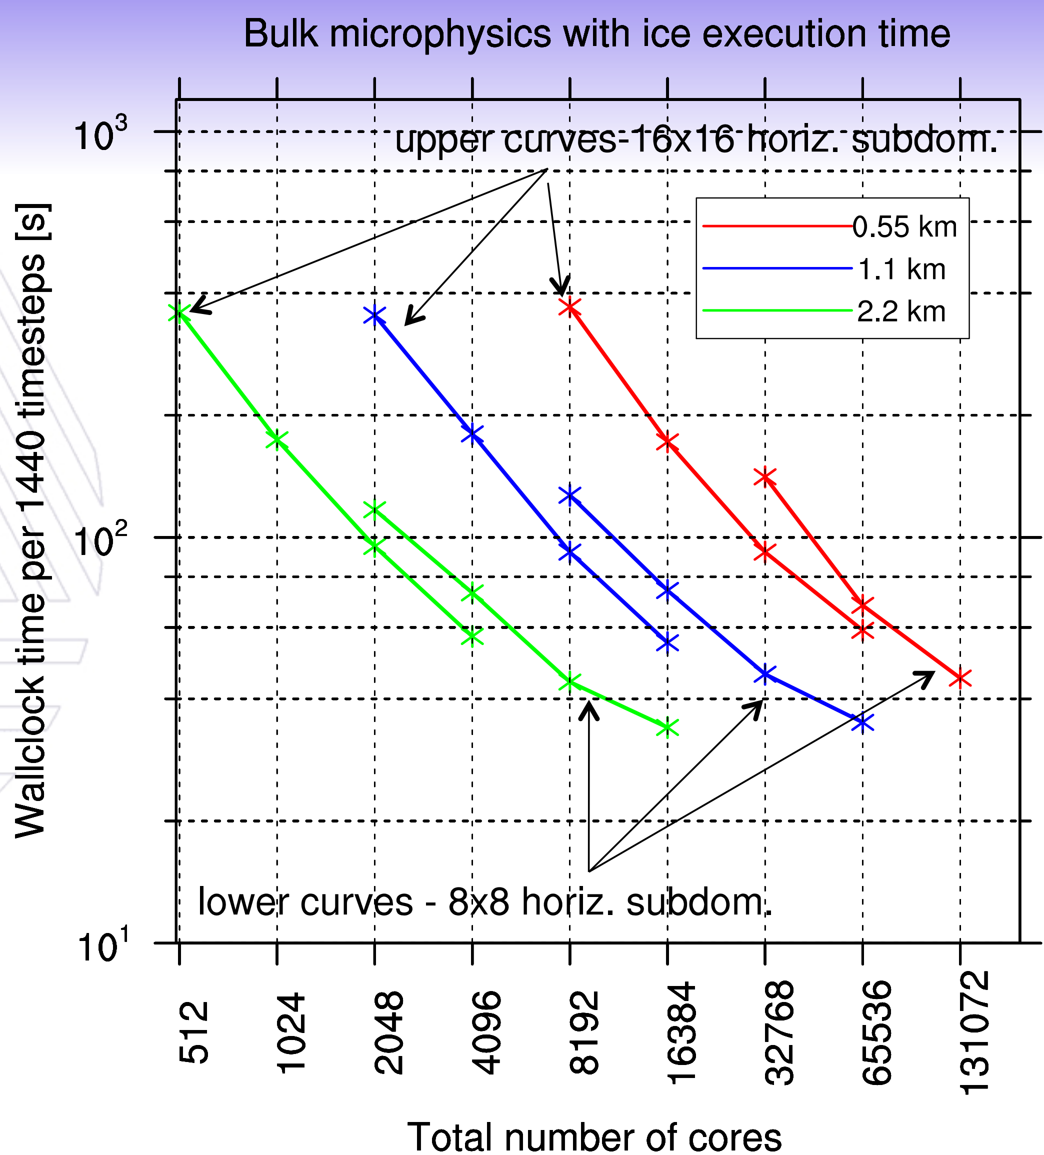

NWP – like experiment (EULAG standalone)

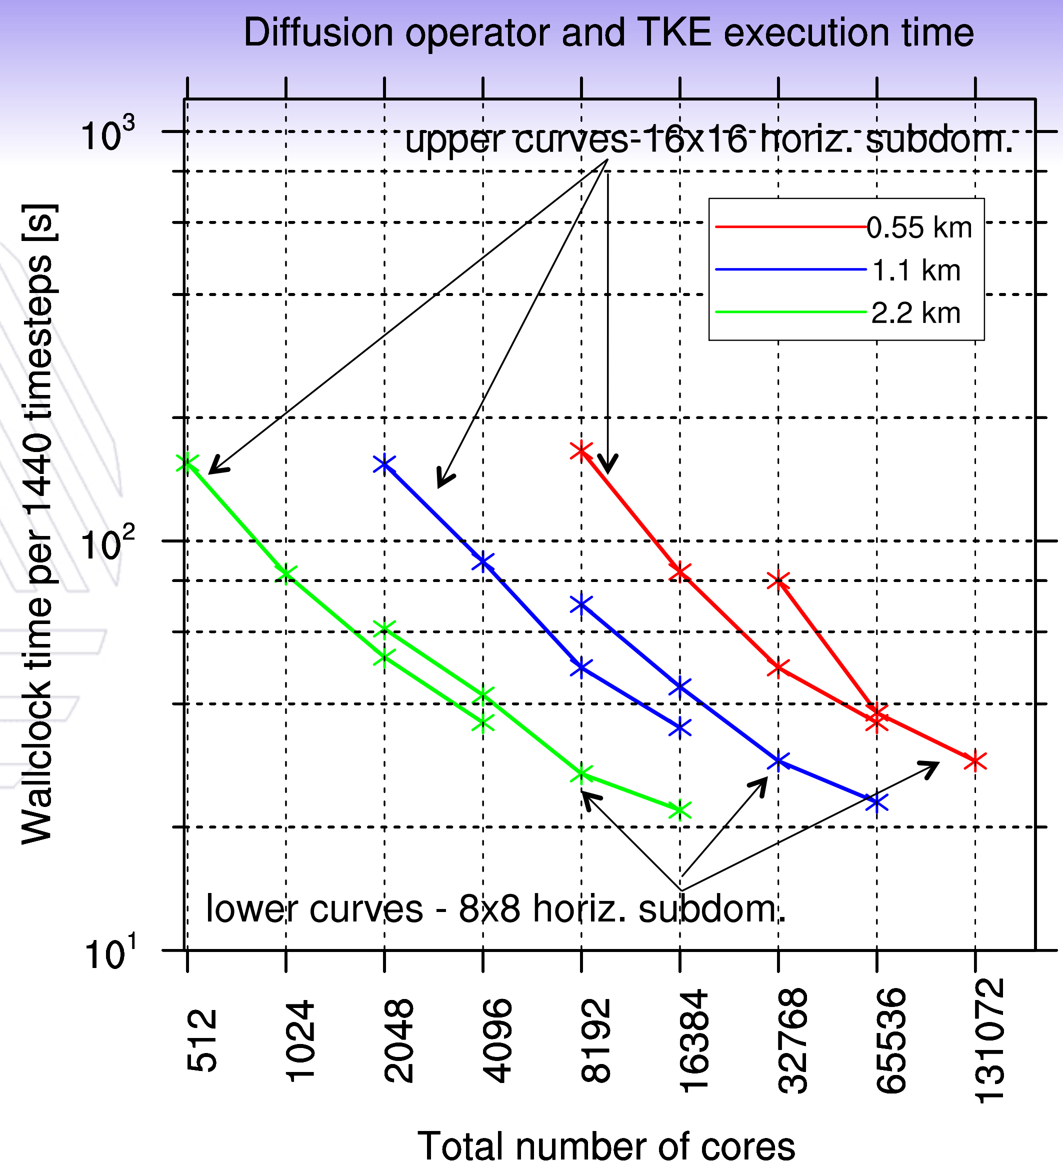

Non-parallel algorithms:

- Tridiagonal solver in preconditioner of GCR solver (many calls per timestep).

- Radiation (very simple but called every timestep, one sequential recurrence + one broadcast operation).

- Implicit 1D advection in moist model (one call per moist timestep).

- In the examples provided, parallelized in the naive (fully sequential, column waits) way.

NWP – like experiment (EULAG standalone)

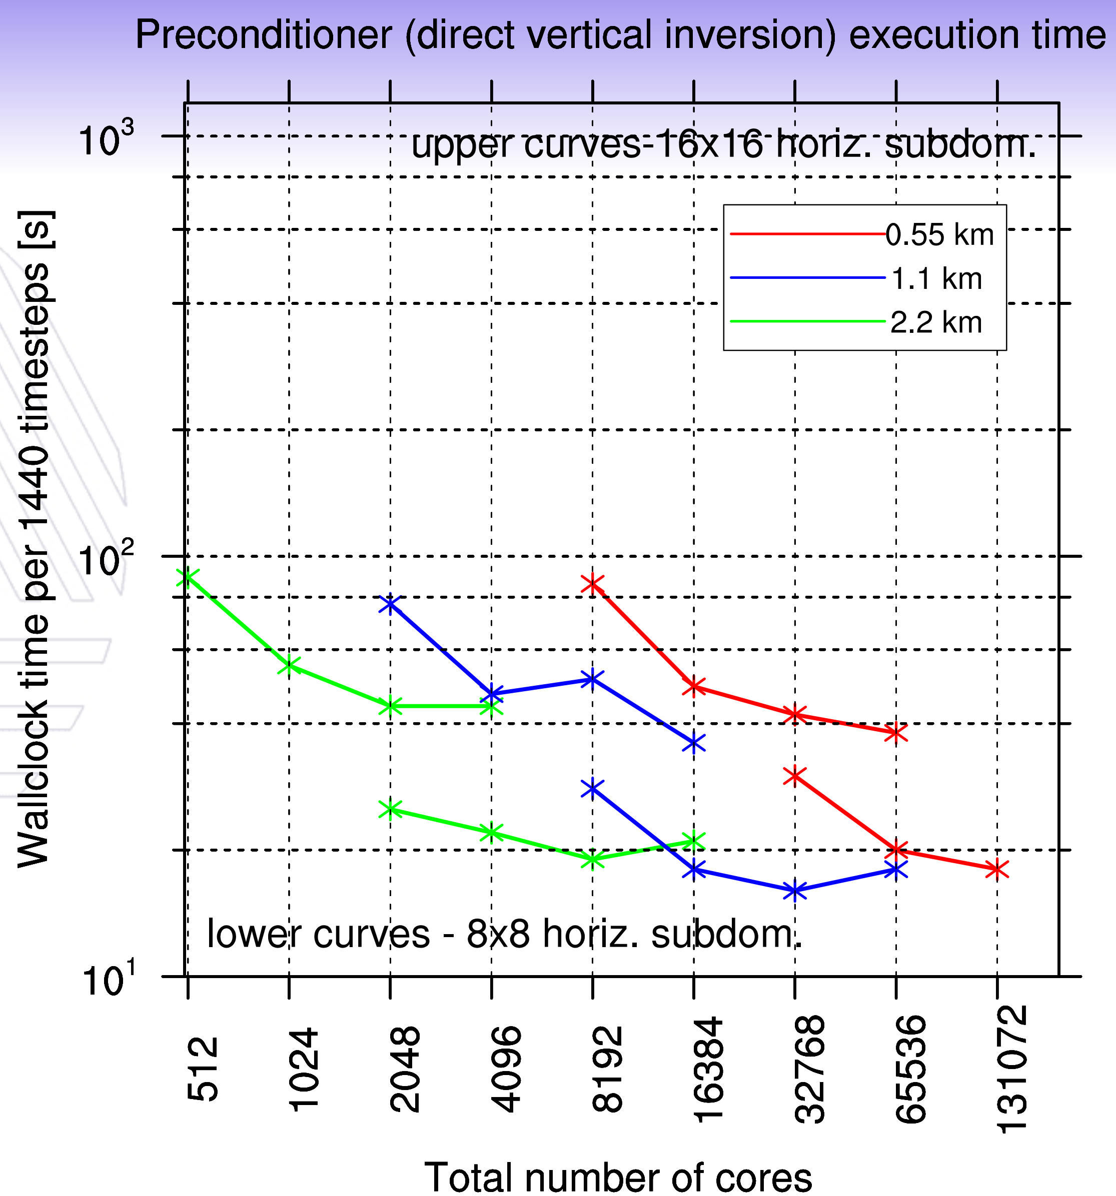

NWP – like experiment (EULAG standalone)

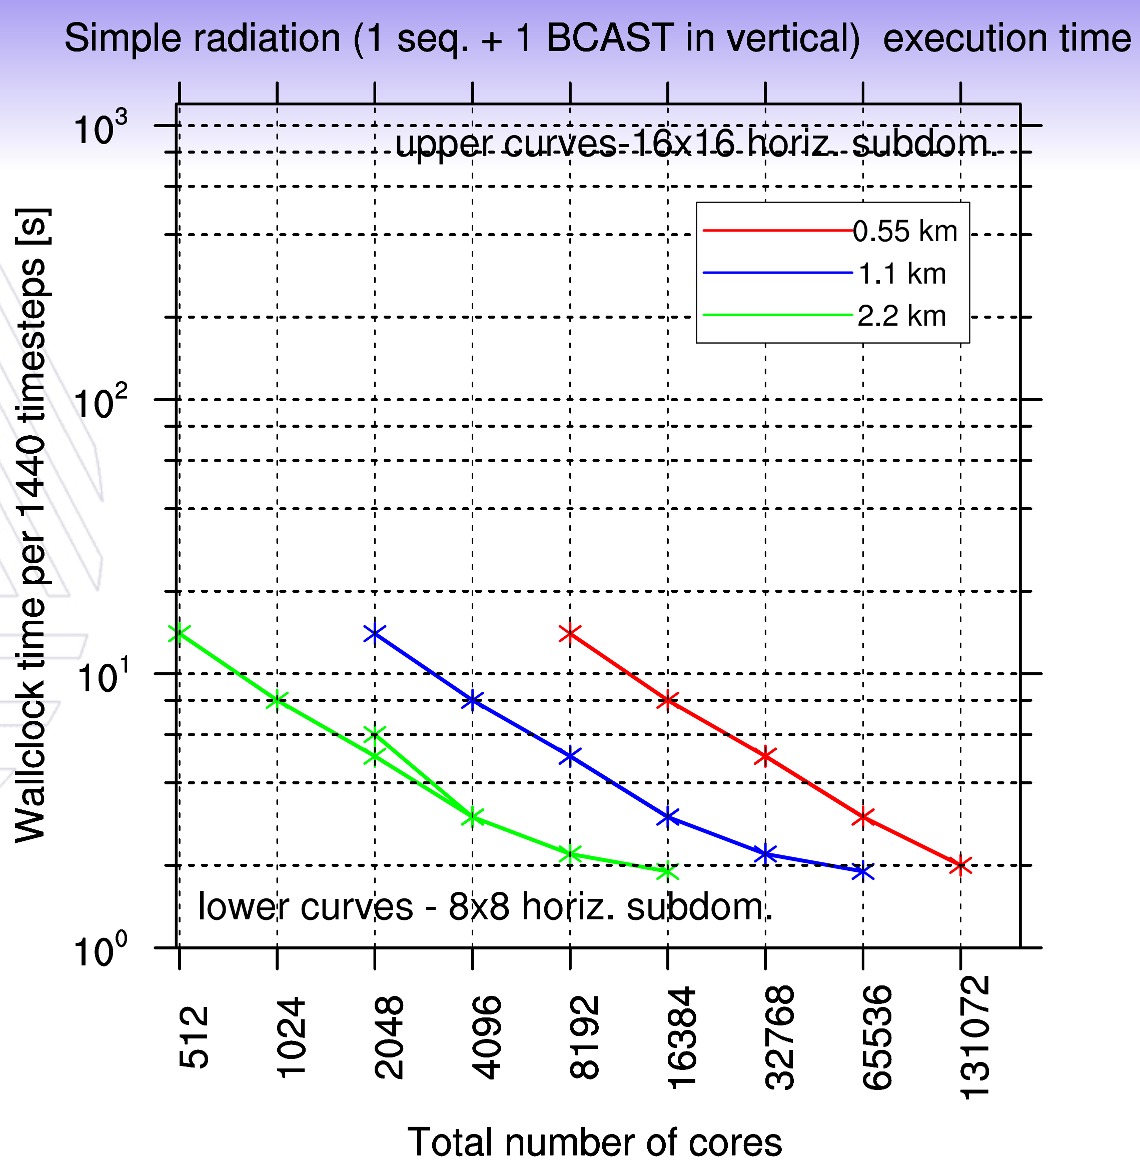

NWP – like experiment (EULAG standalone)

HERMIT XE 6 scalability results

NWP – like experiment (EULAG standalone)

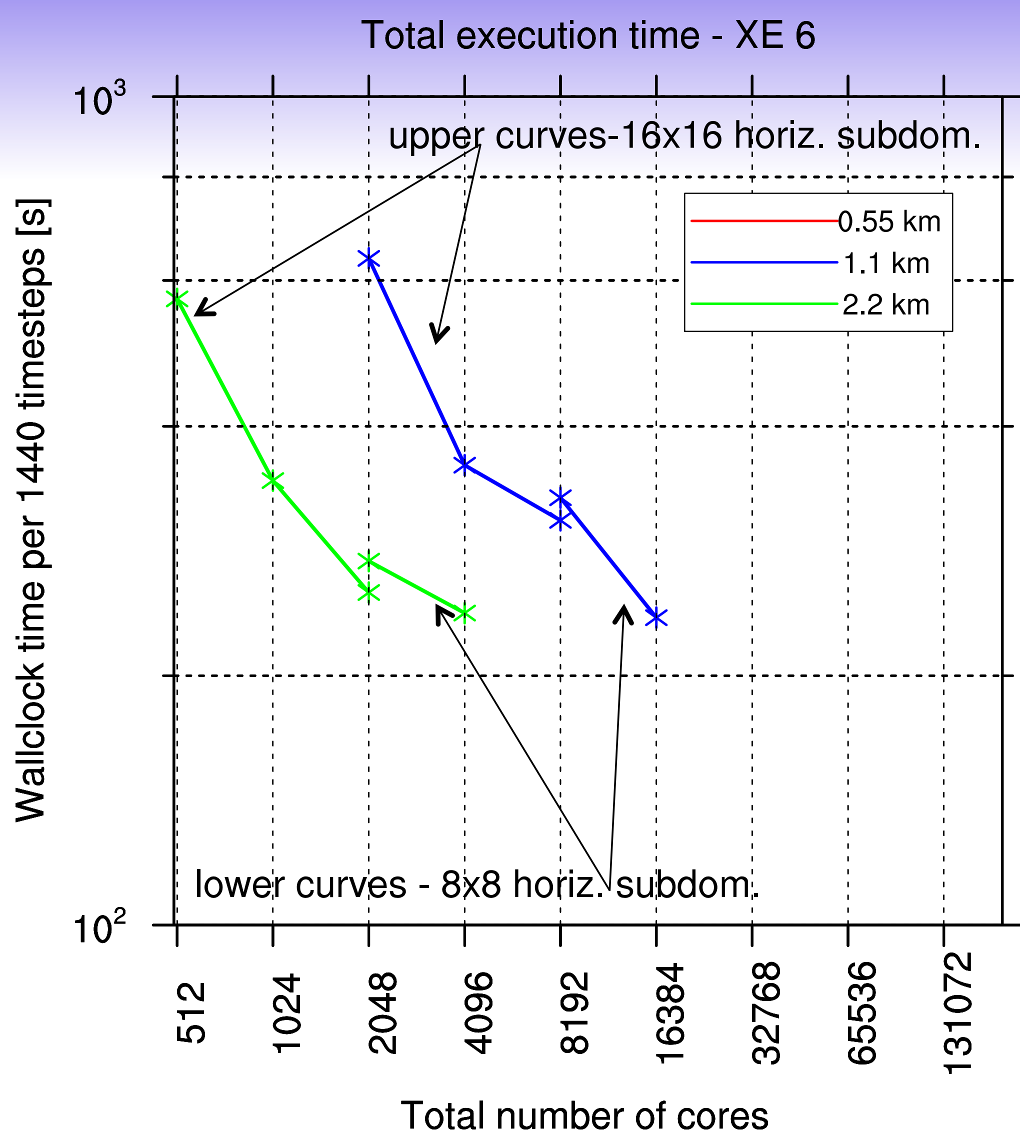

On Cray XE 6, 8x8x32 gridpoints subdomain size seems to be a scalability limit.

NWP – like experiment (EULAG standalone)

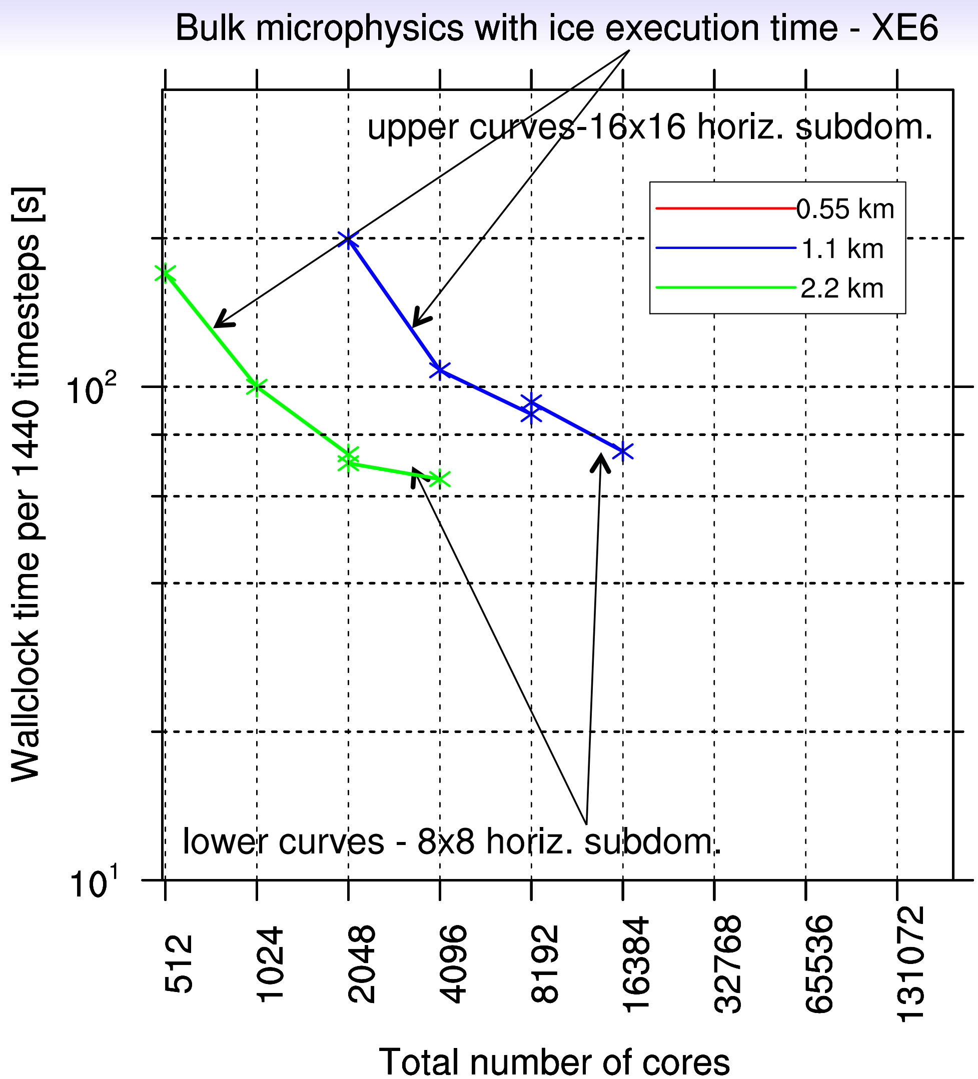

Moist microphysics scales in vertical similar to total program execution time (note different vertical axis range than in the previous slide).

Remarks

- Sequential algorithm (parameterization) could be often split into parallel and non-parallel part. In this case the non-parallel part may be reduced to single MPI Gather + single MPI Scatter operation in the column + some computations. Pipelining of the sequential parts may also help.

- Performance improvement is possible, speedup dependent on the particular setup.

- Scaling for the NWP case may worsen/improve given larger stiffnes of the elliptic problem (larger number of preconditioner iterations)/ implementation of the alternative preconditioners, respectively.

- Elliptic solvers need (increasingly costly) global communication; communication-avoiding methods probably need to be investigated towards the sustained petascale.

- On the fast core machines, it might be difficult to obtain a significant reduction of time-to-solution when using parallelization in the vertical of the NWP setup because of communication wall and not neccessarily the sequential nature of physical processes in the vertical.

- Nevertheless, vertical parallelization adds flexibility and may benefit in easier vectorization (e.g. long innermost loops)

Conclusions

- Three dimensional MPI parallelization in EULAG allows for significant reduction of time-to-solution with using many more cores in simulations alluding to regional NWP on BG/Q.

- Such parallelization also reduces time-to-solution for fixed number of cores.

- It might be worth to estimate if penalty resulting from sequential nature of recurrences resulting from physical processes or tri-diagonal solvers is significant enough to prevent parallelization in the vertical.

- Three dimensional MPI parallelization contributes to the symmetry of coding and may result in better readibility.

- Might be an alternative to tedious effective hybrid parallelization of large atmospheric codes that usually necessitates significant rewrite of the code.

- Further shared memory extensions are possible, hybrid 3D-MPI decomposition and OpenMP might be the optimal solution.

FNP
Foundation for Polish Science

The " Towards peta-scale numerical weather prediction for Europe " project is realized within the Homing Plus programme of Foundation for Polish Science,
co-financed from European Union, Regional Development Fund

IMGW