Implementation of elliptic solvers within ParCS parallel framework

Gordey Goyman, Vladimir Shashkin
Marchuk Institute of Numerical Mathematics RAS (INM RAS)
Hydrometeorological Center of Russia (HMCR)
Moscow Institute of Physics and Technology (MIPT)
gordeygoyman@gmail.com

Russian Supercomputing Days 2021

Outline

- Motivation
- Model problem description
- Linear solvers implementation details
- Scalability test results
- Conclusion

Motivation

The next generation global atmospheric dynamical core is under development at INM RAS and Hydrometcentre of Russia

Program complex that are

- **Scalable** and **numerically efficient**

  model horizontal resolution up to 1 km (10^{10} degrees of freedom)

  model runtime < 20 min for 1 day forecast

- **Flexible**

  Wide range of applications (operational forecast, climate modeling, long range probabilistic prediction)

  Possibility to switch between grids, numerical methods and equations set

Motivation

Atmospheric hydro-thermodynamics

- **stiff system** with phenomena time scales from seconds to months
- severe restriction on **explicit** methods time step

Semi-implicit time stepping:

- allows to use larger integration time steps
- requires solution of **elliptic Helmholtz problem** at every time step
- **Scalable** and **numerically efficient** solver needed

Semi-implicit methods

Main idea: Implicit approximation only for terms describing fast phenomena

Example

$$\frac{\partial u}{\partial t} + c_{adv}\frac{\partial u}{\partial x} = -g\frac{\partial h}{\partial x},$$
$$\frac{\partial h}{\partial t} + c_{adv}\frac{\partial h}{\partial x} = -H\frac{\partial u}{\partial x}.$$

Advection with the speed c_{adv} and gravity waves propagation with $c_g = \sqrt{gH}$

Assuming $c_{adv} \ll c_g$

$$\frac{u^{n+1} - u^{n-1}}{2\Delta t} + c_{adv}\frac{\partial u^n}{\partial x} = -\frac{g}{2}\left(\frac{\partial h}{\partial x}^{n+1} + \frac{\partial h}{\partial x}^{n-1}\right)$$
$$\frac{h^{n+1} - h^{n-1}}{2\Delta t} + c_{adv}\frac{\partial h^n}{\partial x} = -\frac{H}{2}\left(\frac{\partial u}{\partial x}^{n+1} + \frac{\partial u}{\partial x}^{n-1}\right)$$

$\Rightarrow$

Helmholtz type equation

$$\left(I - gH\Delta t\frac{\partial^2}{\partial x^2}\right)h^{n+1} = \tilde{R}_h$$

Analogous equation is obtained in case of 3D atmospheric equations!

Model problem

$$\left(\frac{1}{\gamma_k^2} - \nabla \cdot \nabla\right)\psi_k = f_k, k = 1..N_z$$

After discretizing in space we obtain set of systems of linear equations

$$\left(\frac{1}{\gamma_k^2} I - L\right) x = b$$

Condition number determined by $\frac{\gamma_k^2}{\Delta x^2} \sim \left(\frac{c_s \Delta t}{4 \max \Delta x}\right)^2 = \frac{4}{CFL^2}$

Atmospheric models horizontal waves $CFL \sim 4 - 10$

Model problem. Spatial discretization

Equiangular gnomic cubed sphere grid

- Grid points on cube's faces determined by

$$x = \tan\alpha,\ \ y = \tan\beta,\ \ \alpha,\beta \in [-\pi/4, \pi/4]$$

- Differential operators

$$\nabla f = \left(\frac{\partial f}{\partial \alpha}, \frac{\partial f}{\partial \beta}\right) \qquad \nabla \cdot \mathbf{v} = \frac{1}{G}\left(\frac{\partial G\tilde{u}}{\partial \alpha} + \frac{\partial G\tilde{v}}{\partial \beta}\right)$$

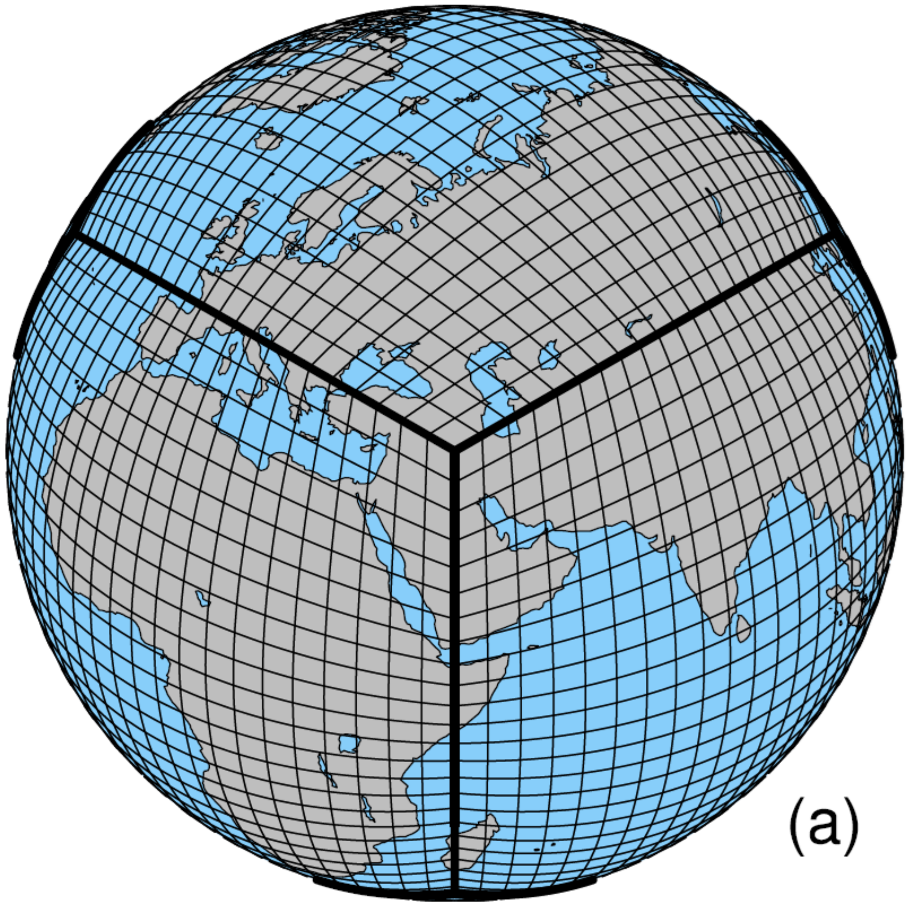

(a)

Finite-difference approximation

- Arakawa C-type variables staggering
- Standard 2nd order formulae
- Bilinear interpolation for transformation covariant->contravariant vector components
- Halo interpolation procedure near cube's panels edges

Solvers description

BiCGstab

- Krylov subspace method. Searches approximate solution

$$x^m \in span\{b, Ab, A^2b, \dots, A^{m-1}b\}$$

- Allows to solve equations with non-symmetric matrices
- Completely matrix-free

Each iteration requires

- 2 matrix-vector products $y = Ax$
- 4 axpy type operations $y = \alpha x + p$
- 4 dot products (x, y) or (x, x) – scalability bottleneck

Solvers description

Geometric Multigrid

Eliminates error on sequence of coarser grids $\Delta x, 2\Delta x, 4\Delta x, \ldots$

- Smoother – weighted Jacobi method
- Restriction operator – 4 point cell average
- Prolongation – bilinear interpolation
- Coarse grid matrices – discretization of initial operator

Condition number determined by $\frac{\gamma_k^2}{\Delta x^2} \sim \left(\frac{c_s \Delta t}{4 \max \Delta x}\right)^2 = \frac{4}{CFL^2}$

At coarse levels $CFL \to \frac{CFL}{2^{k-1}}$, k – number of MG level

For $CFL \sim 4 - 10$, it is sufficient to use 4 MG levels

ParCS parallel framework

Along with the development of the model, we also work on our own software infrastructure

ParCS (Parallel Cubed Sphere)

Fortran 2008 object oriented library providing:

- Domain decomposition
- Distributed data storage
- Parallel exchanges
- Input/Output

At multiblock logically-rectangular computational grids

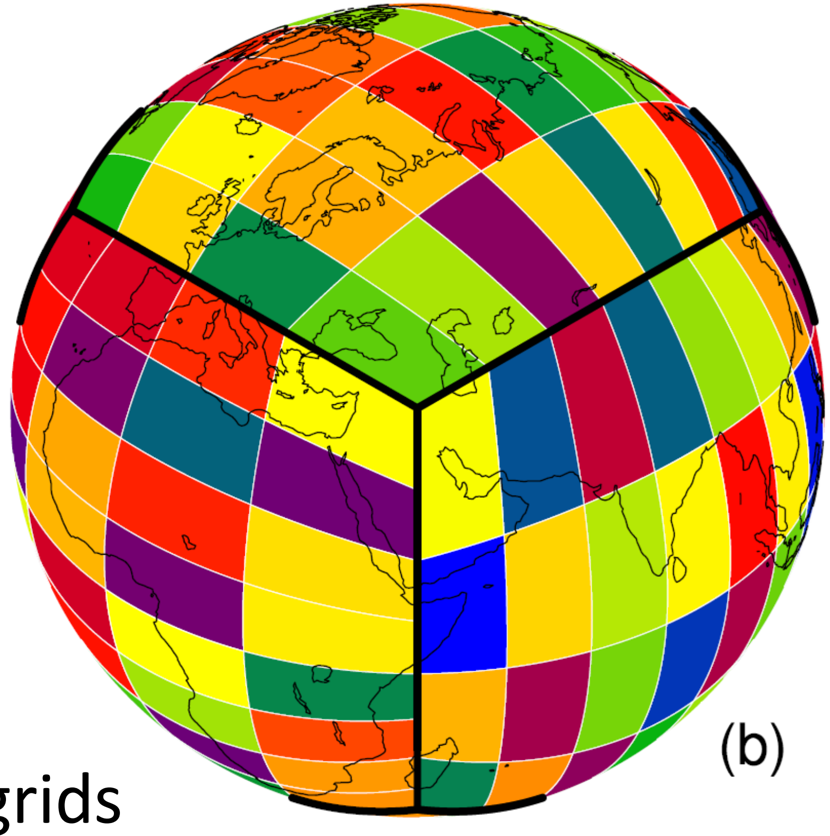

(b)

Implementation details

Pure MPI parallel implementation

- Domain decomposed into set of tiles (rectangular section of the grid)
- Number of MPI-processes = 6*M*N = number of tiles. Ideally M=N
- Halo zones (overlap regions between tiles) exchanges strategy

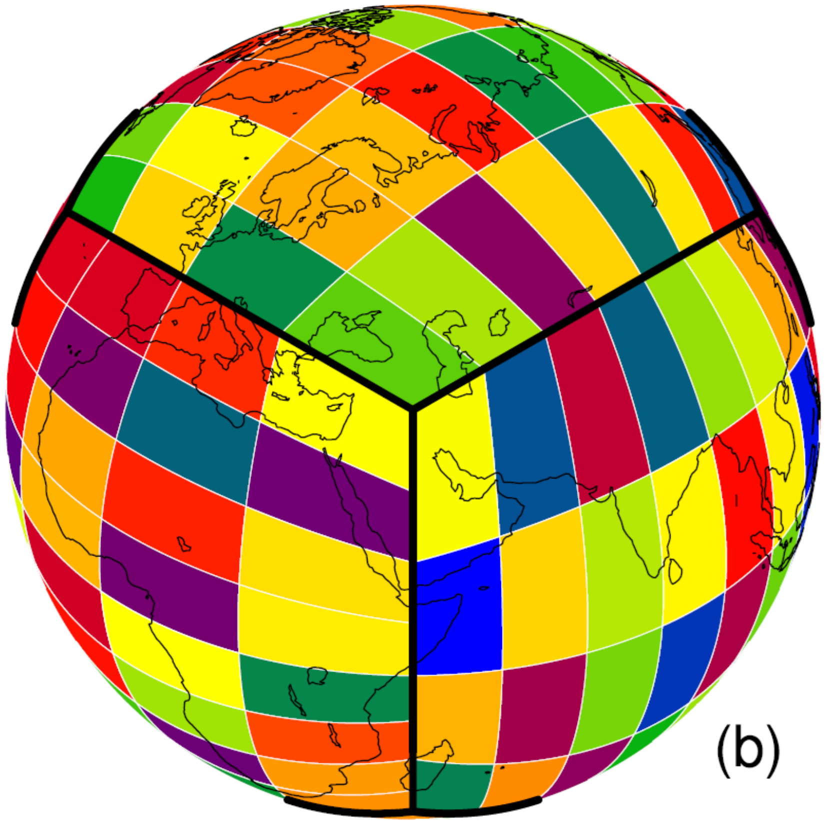

(b)

Experiments setup

- Cray XC40 (Roshydromet), Cray ARIES interconnect, Intel Xeon E2697v4 18-core CPUs, i.e. 36 CPU cores per node.
- Tests for problems at 3 grids:
  C216L30 – 46 km horizontal resolution, 8.39×10^6 grid points
  C432L30 – 23 km, 3.36×10^7 grid points
  C1080L30 – 9 km, 2.1×10^8 grid points
- $\max CFL = 7.5$ for all grids
- Stopping criterion $\frac{\|r^k\|}{\|r^0\|} = 10^{-5}$

Strong scaling results

46 km problem scales up to 1944 cores 9 km problem at least to 4860 CPU cores

Multigrid method $\approx 4-6$ times faster than BiCGstab

24 hour forecast

9km model, 100s time step, 4680 cores
BiCGstab – 3.6 min,
MG method – 0.6 min,
of the runtime in solver block

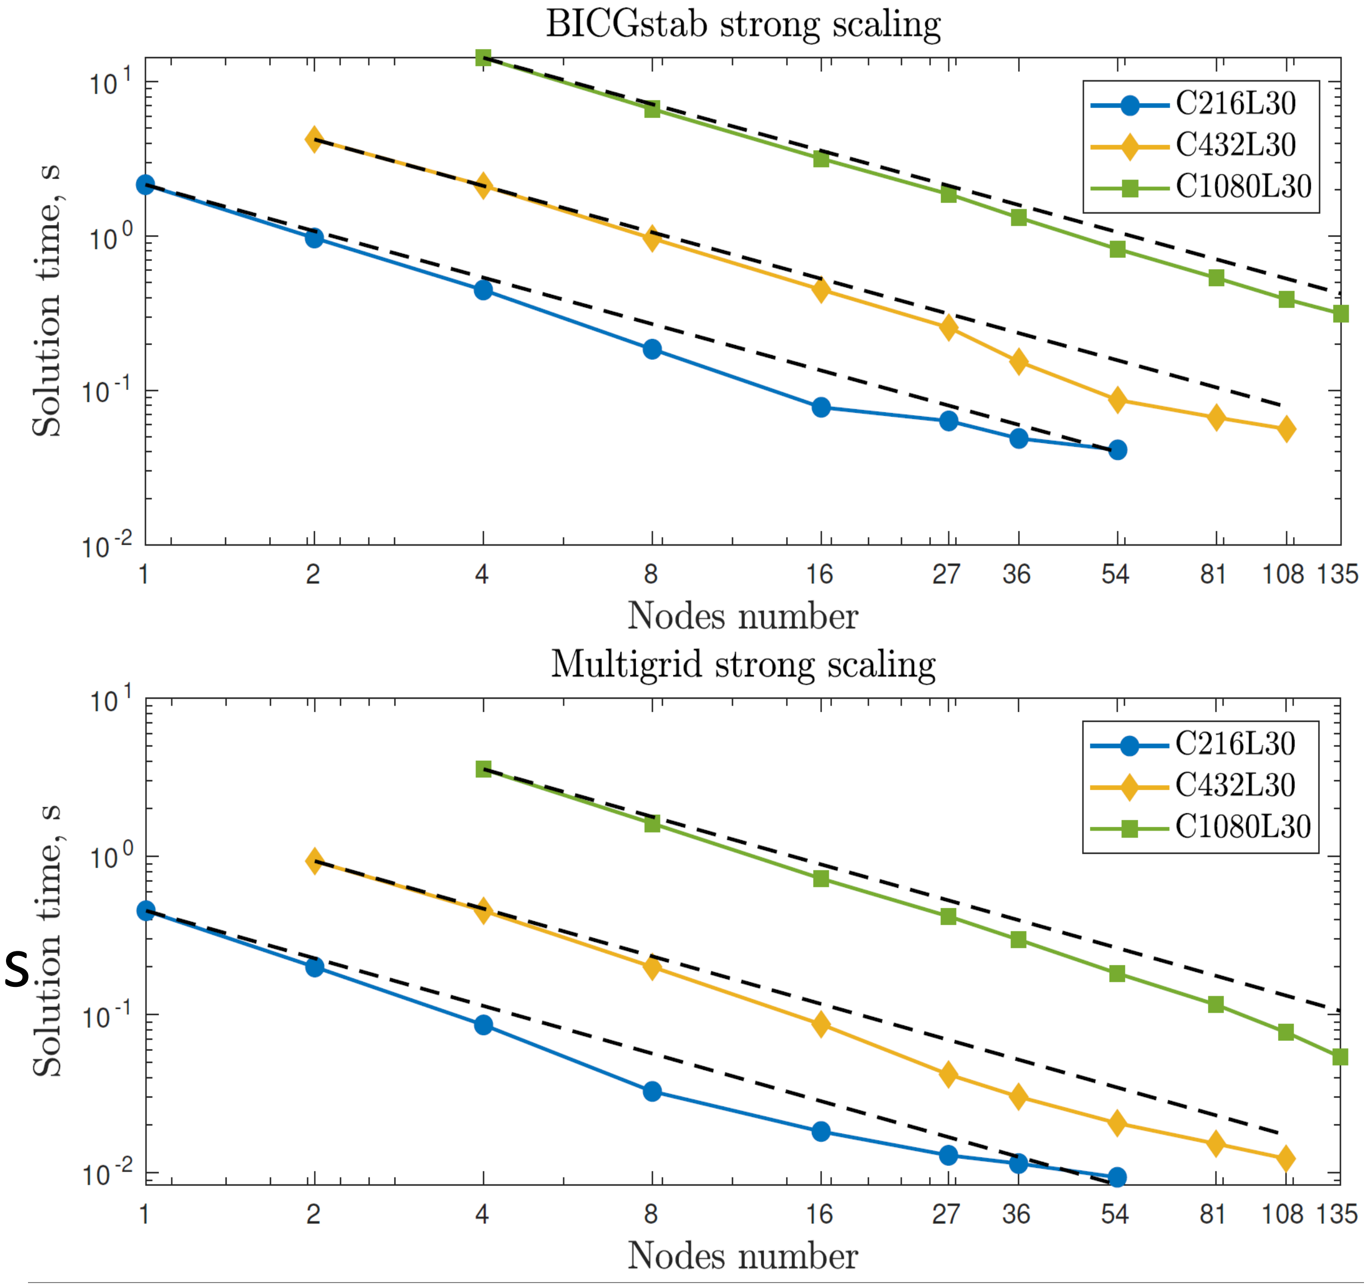

Strong scaling. BiCGstab components

Strong scaling of the BiCGstab components for the problem at C216L30 grid

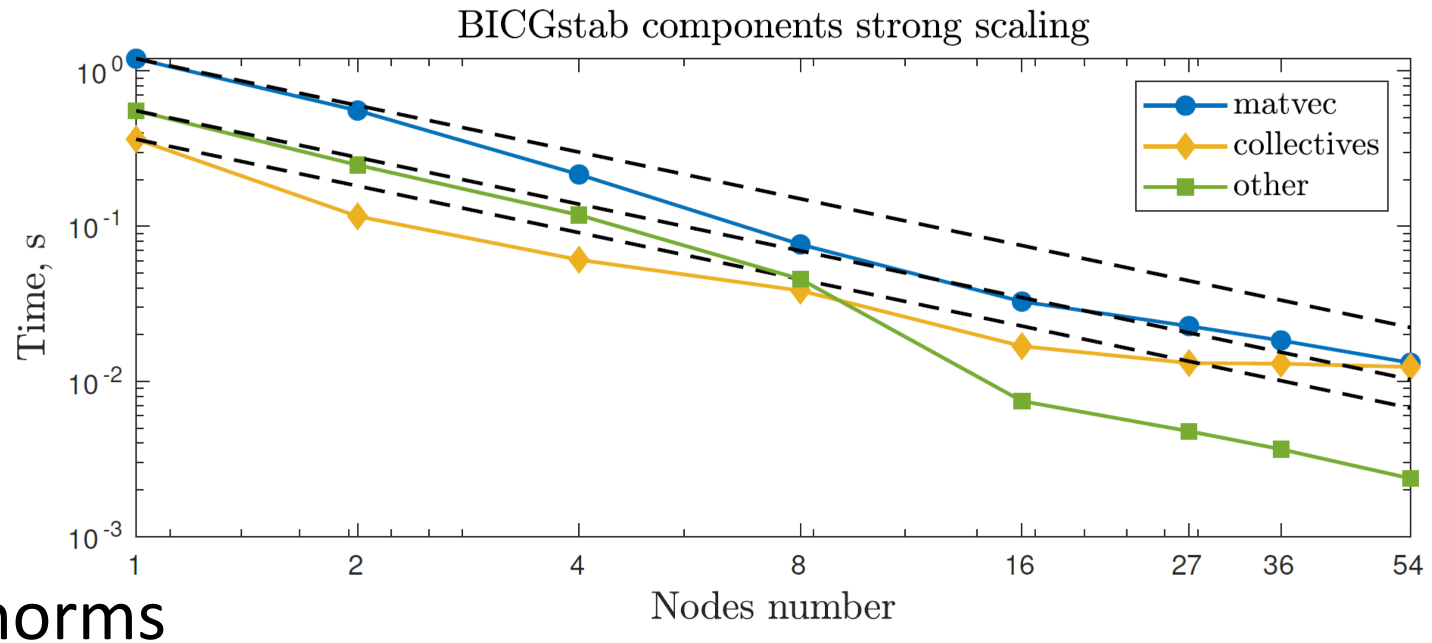

“matvec” – matrix vector products

“collectives” – vector dot products/norms

“other” – axpy type operations

“collectives” scales only up to 576 CPU cores

Strong scaling. Multigrid levels

Strong scaling of the MG computations at different levels. C216L30 grid

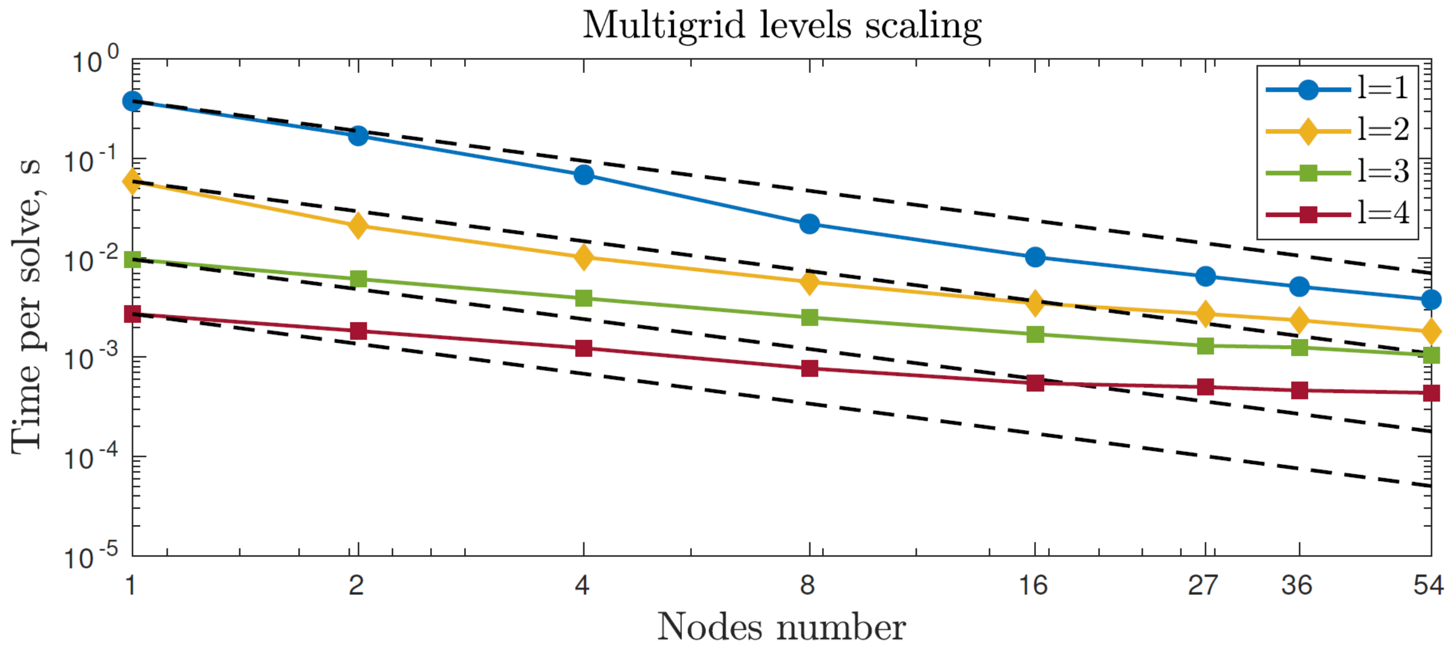

Parallel efficiency decrease at the last two levels.

Hybrid MPI-OpenMP parallelization

Use of single precision for exchanges/computations at coarse levels.

Conclusions

- BiCGstab and geometric multigrid solvers for the solution of Helmholtz type problem at the cubed sphere grid within ParCS parallel framework

- Both algorithms scales at least up to 4680 CPU cores

- Good starting point for further testing and optimization of the implemented solvers and ParCS library within a non-hydrostatic atmospheric dynamics model

Thank you for attention!