Aladin Workshop/Hirlam All Staff Meeting
Lisbon, 4–7 April 2016

# The way out of the scalability gridlock: ESCAPing dwarfs

Daan Degrauwe, Steven Caluwaerts, Piet Termonia

RMI Belgium

# Content

- **Why it's bad** (recap from last year)
  - scalability of spectral methods
  - non-spectral methods have their own problems
- **Why it's even worse**
  - some 3D results
- **The way out**
  - The ESCAPE project
- **Conclusions**

# The spectral method



- From the *accuracy* point of view, spectral methods are unsurpassable: their order of accuracy is infinite!
- Moreover, the calculation of derivatives and solving the Helmholz equation are trivial. This allows for (semi-)implicit timestepping and large timesteps.
  So our spectral dynamics are also quite *efficient*.

# The spectral method



- From the *accuracy* point of view, spectral methods are unsurpassable: their order of accuracy is infinite!
- Moreover, the calculation of derivatives and solving the Helmholz equation are trivial. This allows for (semi-)implicit timestepping and large timesteps. So our spectral dynamics are also quite *efficient*.
- . . . but they require spectral transforms (FFT or Legendre transform for the global). These are nonlocal, i.e. they require domain-wide communication. This makes their use problematic on massively parallel machines.
- Another disadvantage of a spectral model is the requirement of a homogeneous reference state for the semi-implicit timestepping. e.g. orography cannot be treated implicitly.
- *But at what point do the costs no longer justify the accuracy?*
  - to answer this question, we must closely investigate the alternatives.

# Local methods



- When considering alternatives for the spectral horizontal discretization, we try to keep as much as possible of the model intact:
  - only way to make a clean comparison
  - limited development cost (no need to modify physics, . . . )

  So for the time being, we stick to a semi-implicit time discretization and a semi-Lagrangian advection scheme.

- When considering alternatives for the spectral horizontal discretization, we try to keep as much as possible of the model intact:
  - only way to make a clean comparison
  - limited development cost (no need to modify physics, . . . )

  So for the time being, we stick to a semi-implicit time discretization and a semi-Lagrangian advection scheme.

- Finite-difference discretizations are considered on the following grids:

A-grid

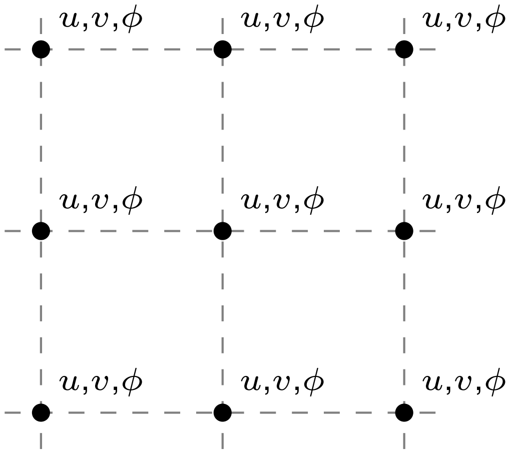


wind speed
not staggered

C-grid

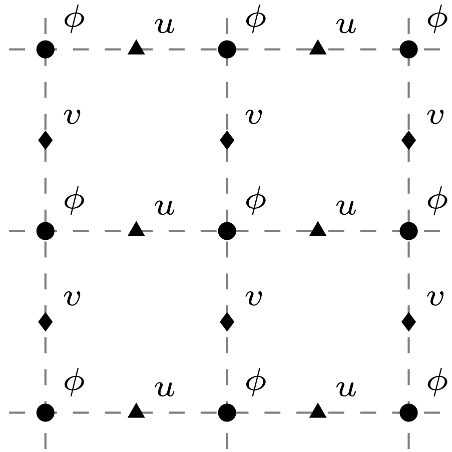


wind speed
staggered

Z-grid

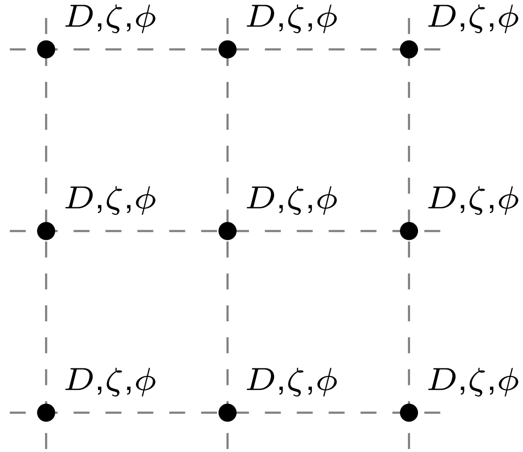


vorticity/divergence
not staggered

# Local methods



- Shallow-water toy model tests show that each grid has its own problems:
  - C-grid is staggered and requires heavy modifications
  - A-grid has bad propagation of short waves (negative group velocity)
  - Z-grid has bad projection of short waves, resulting in very noisy wind fields

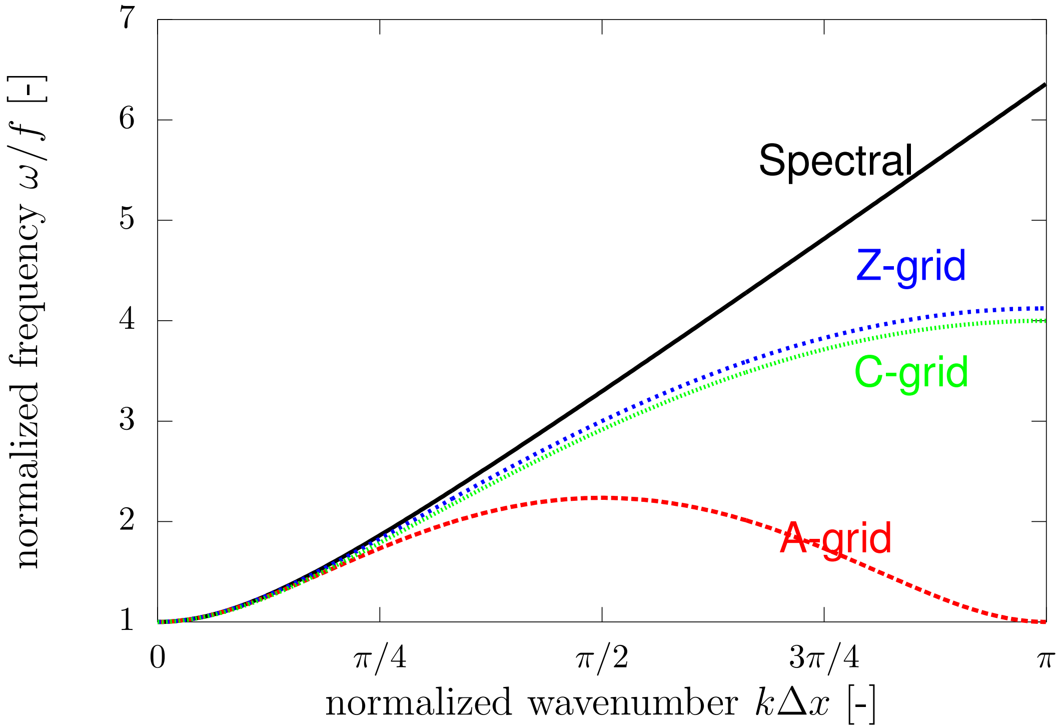

# Local methods



- Shallow-water toy model tests show that each grid has its own problems:
  - C-grid is staggered and requires heavy modifications
  - A-grid has bad propagation of short waves (negative group velocity)
  - Z-grid has bad projection of short waves, resulting in very noisy wind fields

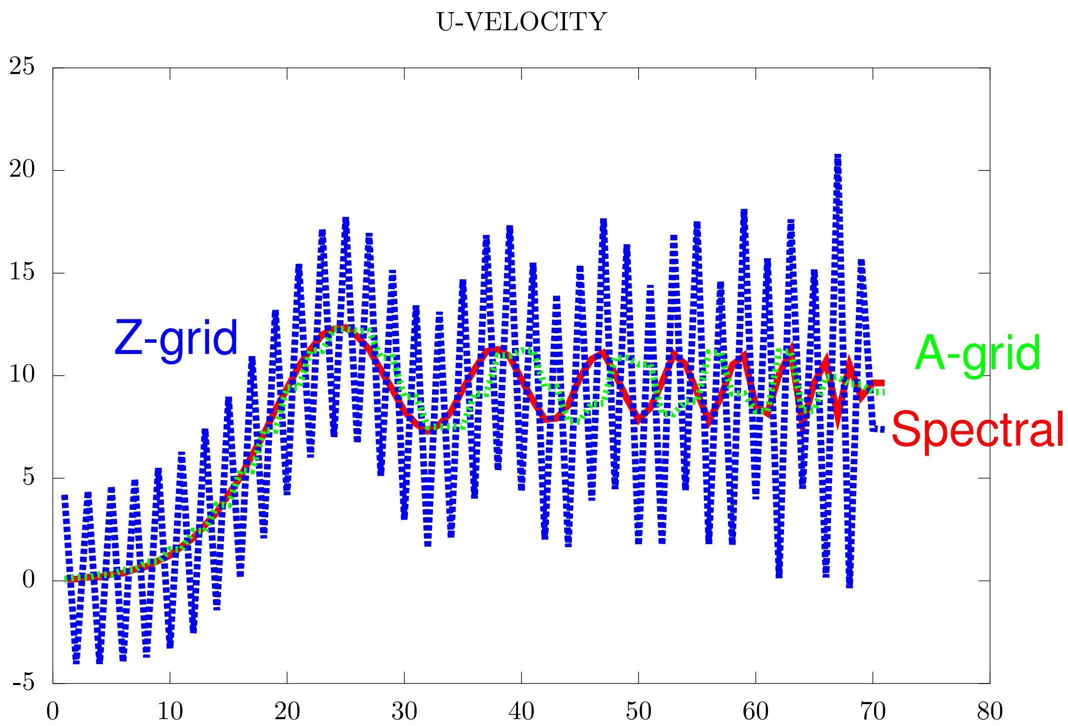

- Shallow-water toy model tests show that each grid has its own problems:
  - C-grid is staggered and requires heavy modifications
  - A-grid has bad propagation of short waves (negative group velocity)
  - Z-grid has bad projection of short waves, resulting in very noisy wind fields

RMSE w.r.t. spectral

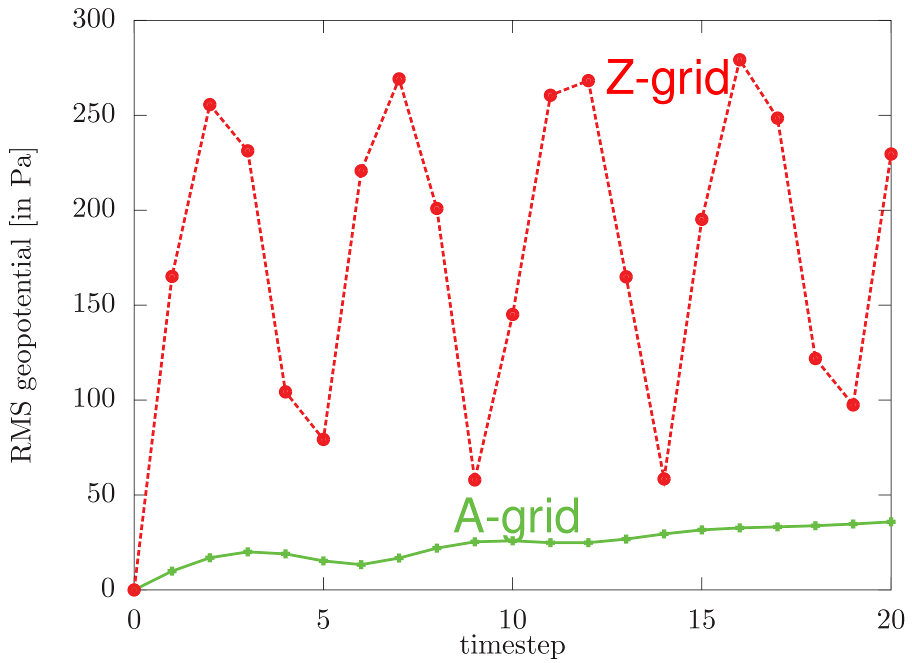

- Shallow-water toy model tests show that each grid has its own problems:
  - C-grid is staggered and requires heavy modifications
  - A-grid has bad propagation of short waves (negative group velocity)
  - Z-grid has bad projection of short waves, resulting in very noisy wind fields

RMSE w.r.t. spectral

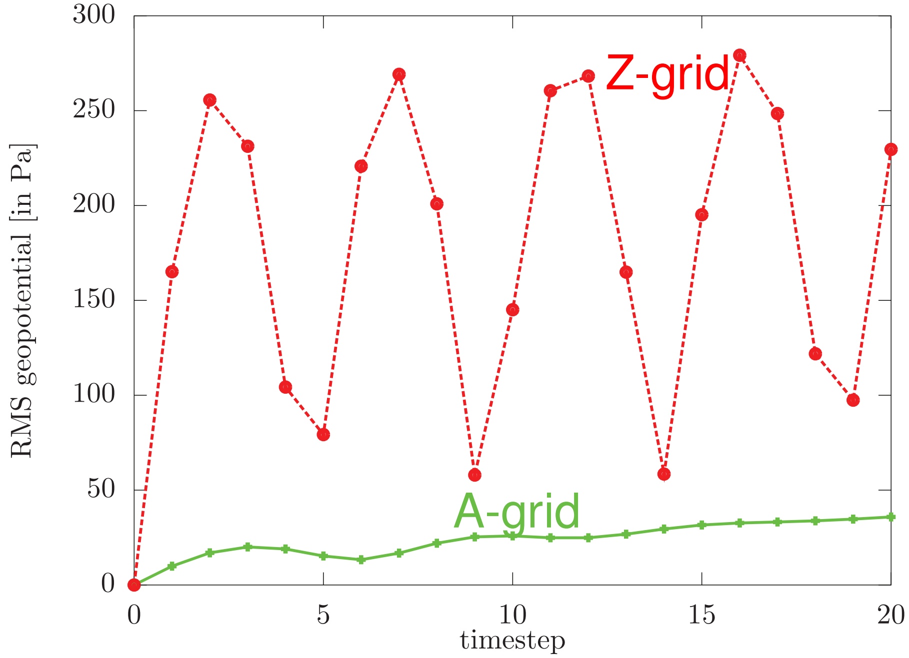


- But how representative are these toy-model results for a real atmospheric model?

- It's possible to emulate a gridpoint model with our spectral ALADIN model.

  The trick is to modify the spectral response of differential operators, e.g.

| | Spectral | | A-grid |
|---|---|---|---|
| $\frac{\partial \hat{\psi}_k}{\partial x}$ | $ik\hat{\psi}_k$ | $\longrightarrow$ | $i\frac{\sin k\Delta x}{\Delta x}\hat{\psi}_k$ |

- (of course, this doesn't allow to test the scalability.)

- Adiabatic run, $\Delta x = 7\,\text{km}$, quadratic trunctation, no DFI

RMS w.r.t. spectral

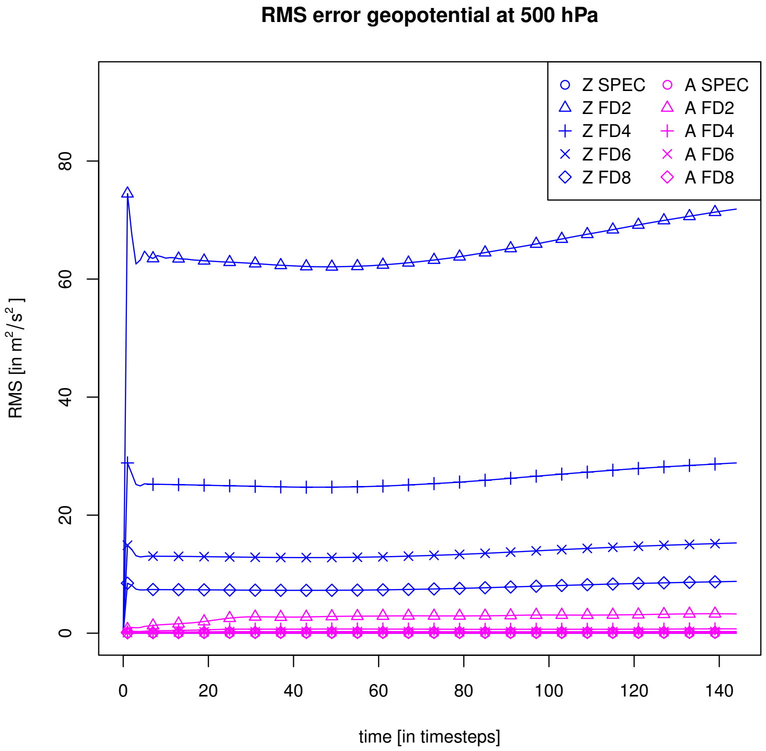


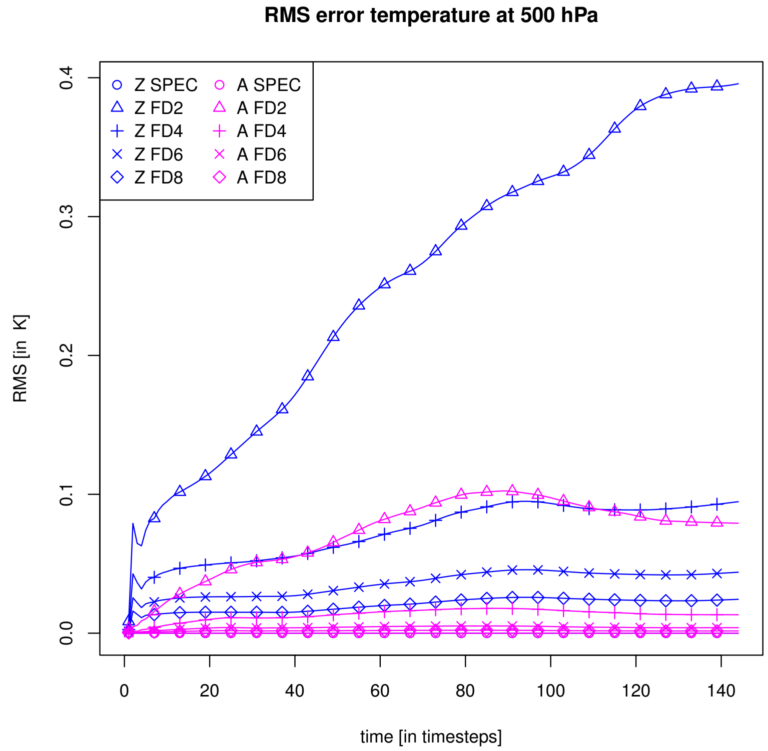

- Adiabatic run, $\Delta x = 7$ km, quadratic trunctation, no DFI

500 hPa geopotential after 1 timestep

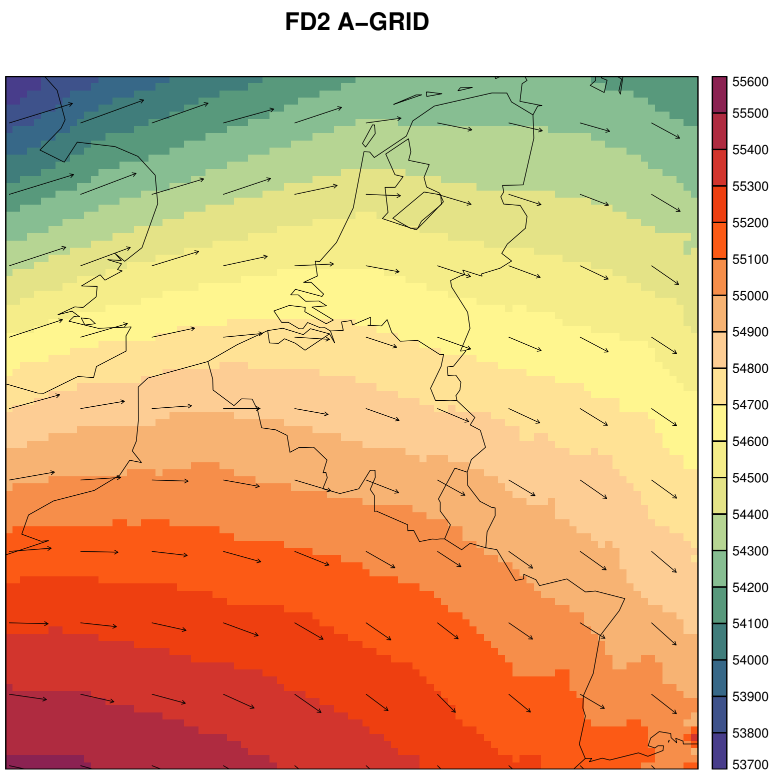


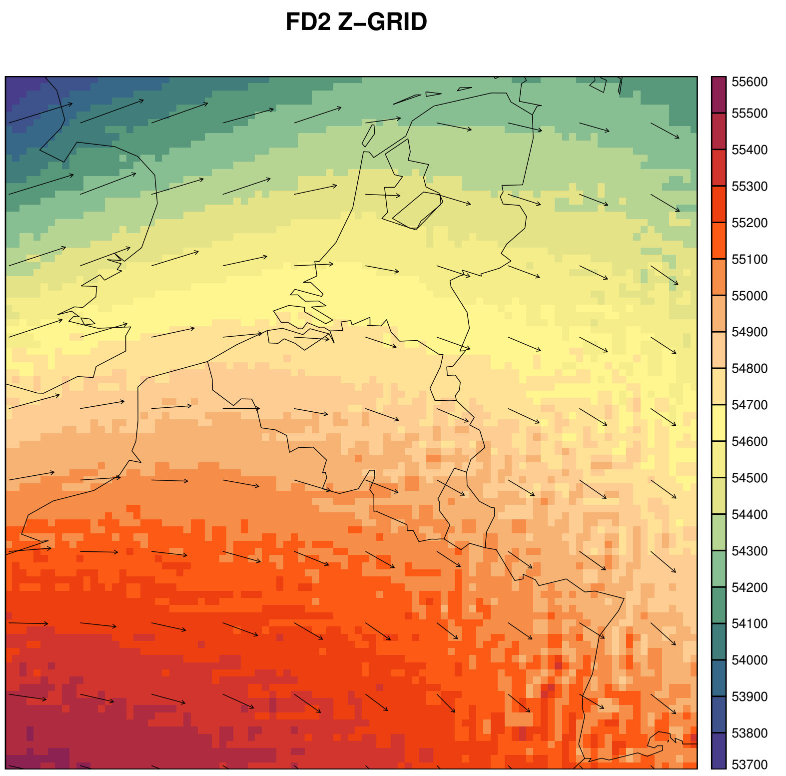

- Adiabatic run, $\Delta x = 7\,$km, quadratic trunctation, no DFI

500 hPa geopotential after 12h

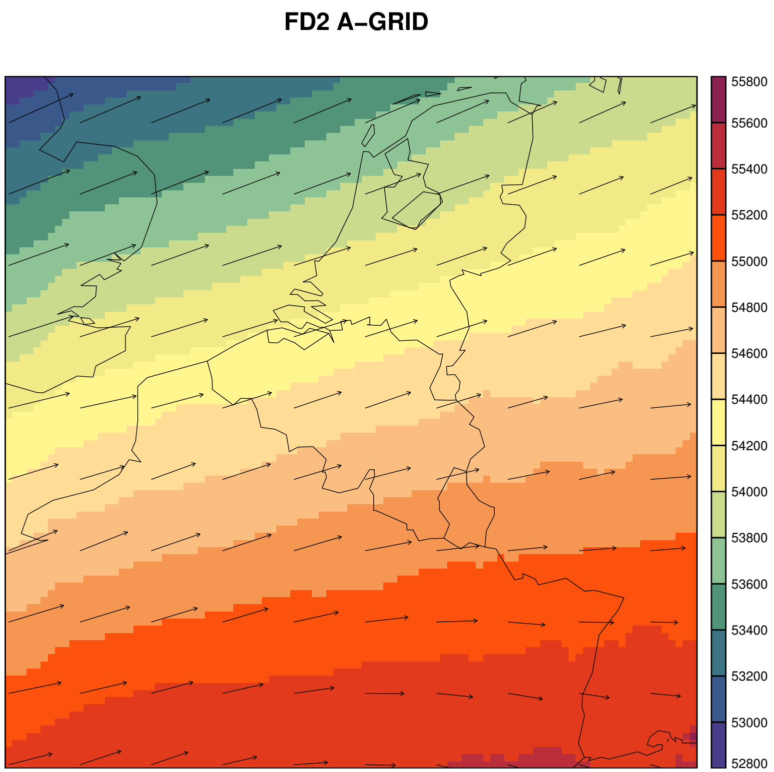


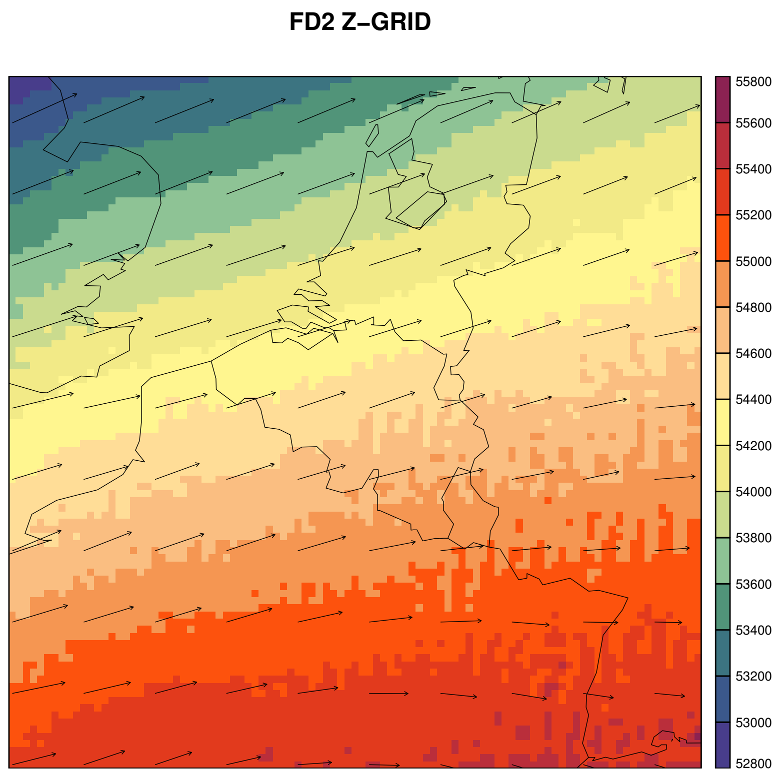

# Some results with ALADIN

- Adiabatic run, $\Delta x = 7\,$km, quadratic trunctation, no DFI

- Run with ALARO physics

RMS w.r.t. spectral

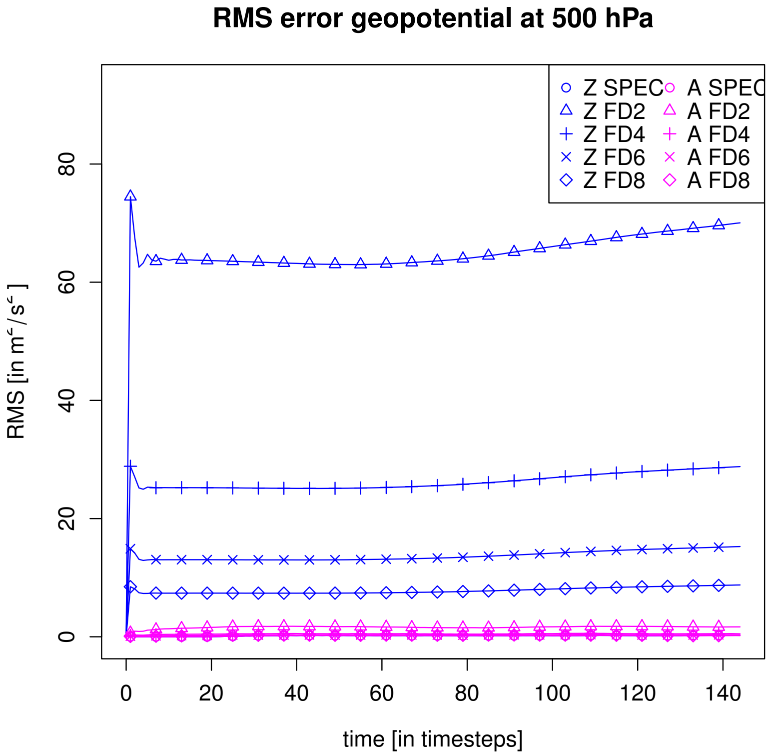


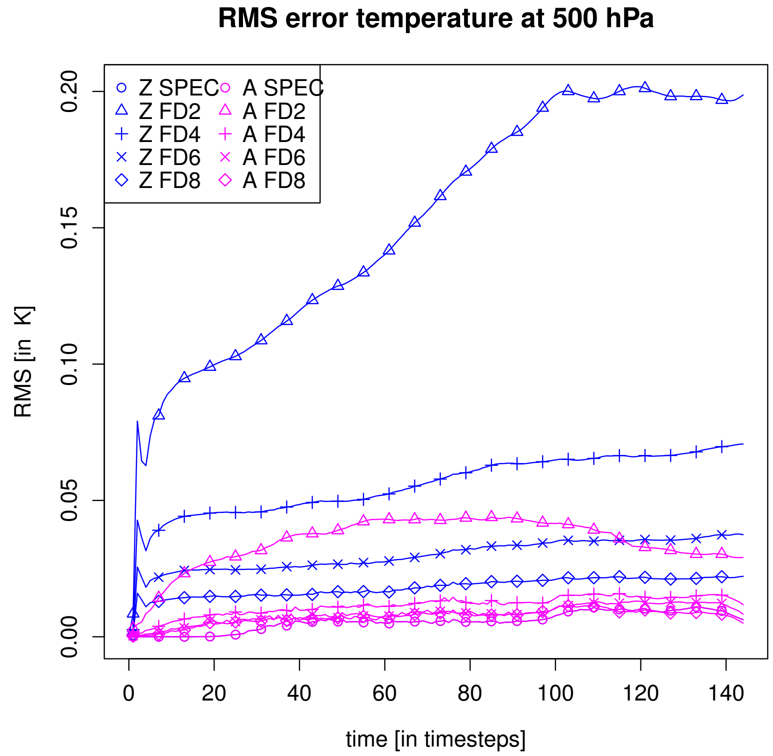

# Some results with ALADIN



- The footprint of the faulty projection in the Z-grid discretization is clearly visible!
- The fields remain noisy during the entire forecast, so a simple DFI will not solve the problem.
- Why is this so bad? It turns out that short scales do matter, even in a 3D model with diffusion and quadratic spectral truncation.

# Some results with ALADIN



- The footprint of the faulty projection in the Z-grid discretization is clearly visible!
- The fields remain noisy during the entire forecast, so a simple DFI will not solve the problem.
- Why is this so bad? It turns out that short scales do matter, even in a 3D model with diffusion and quadratic spectral truncation.
- On the sunny side: the A-grid scheme performs quite okay, especially the higher-order variants.

So we have a giant (IFS/ARPEGE/ALADIN/HARMONIE), who is restrained by a grid.

**Question**: How to break him free?

So we have a giant (IFS/ARPEGE/ALADIN/HARMONIE), who is restrained by a grid.

**Question**: How to break him free?

**Answer**: you cut him into pieces.

# The way out



So we have a giant (IFS/ARPEGE/ALADIN/HARMONIE), who is restrained by a grid.

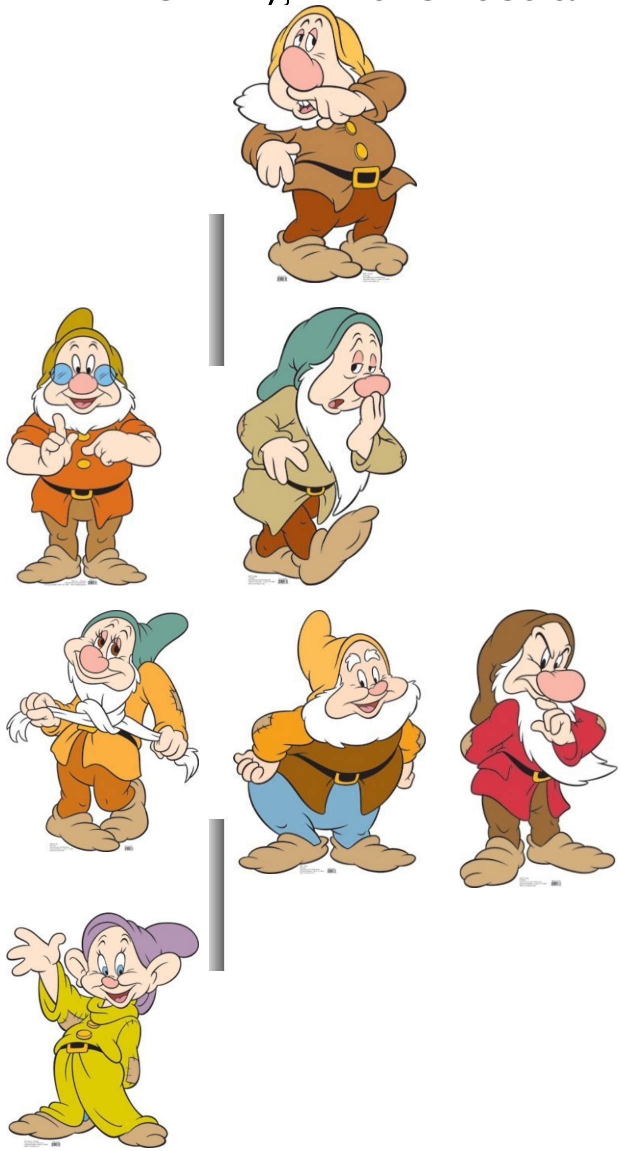

**Question**: How to break him free?

**Answer**: you cut him into pieces.

# The ESCAPE project



- ESCAPE (Energy-efficient Scalable Algorithms for Weather Prediction at Exascale) is a H2020 project, coordinated by ECMWF, and involving HIRLAM and ALADIN members, HPC hardware manufacturers, universities and supercomputing centers.
- The core of ESCAPE is the identification of fundamental algorithm building blocks. These are the so-called NWP dwarfs (an idea based on the Berkeley Dwarfs)
- So far, 4 dwarfs have been defined:
  - Global spectral transform
  - (2D) elliptic solver with iterative method
  - Cloud scheme (column physics)
  - Bifourier spectral transform
- Candidates for more dwarfs:
  - Advection schemes (semi-Lagrangian or Eulerian)
  - 3D Helmholz sparse solver
  - Time integration schemes
  - Radiation scheme

# The ESCAPE project



- The concept of the NWP dwarfs offers the modularity that is necessary to cleanly compare alternatives.
- Further steps in the ESCAPE project are
  - Code adaptation: accelerator directives, interoperate with a domain-specific language
  - Hybrid computing: optimize dwarfs for accelerators
  - Benchmarking and diagnostics: time-to-solution, energy consumption, determine NWP benchmark references.
- The common basis for the dwarfs is the ATLAS library, which supports structured and unstructured grids.

# Conclusions



- Be careful when relying on 'common knowledge'
  - 'A-grid has problem at the shortest scales'
  - 'Formulation in vorticity/divergence (Z-grid) solves this'
- Every method has its drawbacks; only clean testing can bring answers
- Clean testing means modularity
- Modularity (dwarfs) is the central idea of the ESCAPE project

# Conclusions

- Be careful when relying on 'common knowledge'
  - 'A-grid has problem at the shortest scales'
  - 'Formulation in vorticity/divergence (Z-grid) solves this'
- Every method has its drawbacks; only clean testing can bring answers
- Clean testing means modularity
- Modularity (dwarfs) is the central idea of the ESCAPE project
- Big challenges ahead:
  - Quoting Enda O'Brien:
    *How to build a giant from several dwarfs?*

  - Don't forget there's also a scientific impact of switching methods!

# Thank you !

## And many thanks and congratulations to Steven!