DYNAMIC ADAPTIVE REFINEMENT
IN EARTH SYSTEM MODELLING

David Knapp, Johannes Holke, Niklas Böing

DLR

Disclaimer

This Talk is not	This Talk is

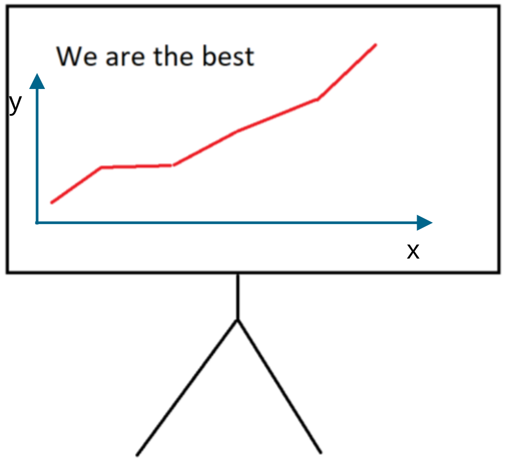

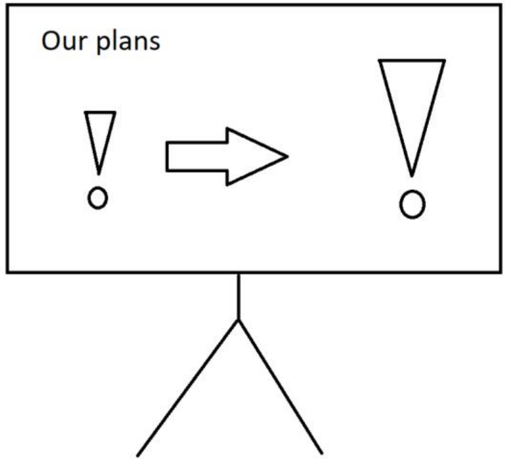

Why AMR

Increase resolution = increase simulation accuracy

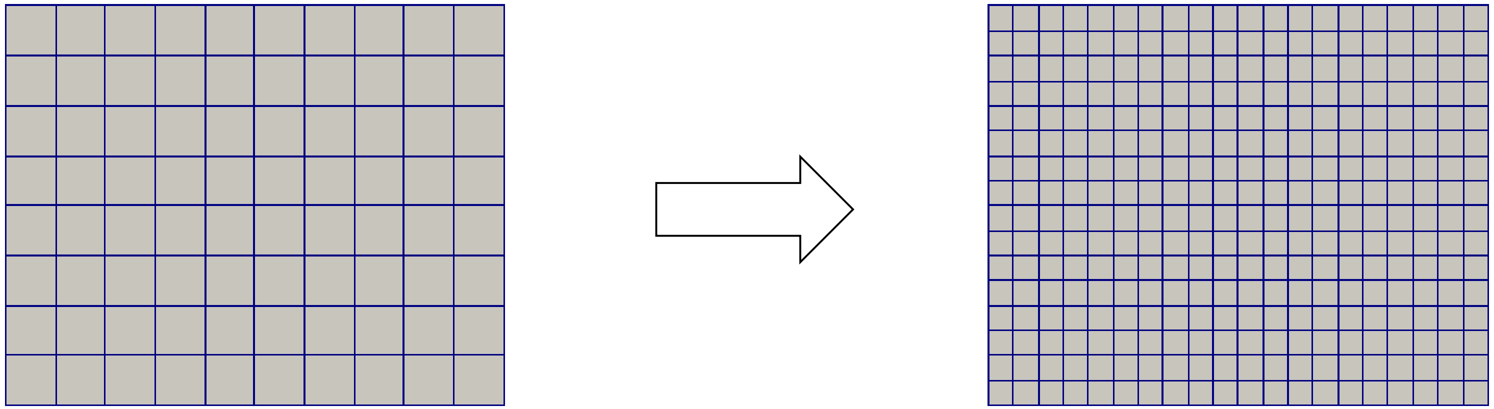

4x/8x more DOFs and data per time step
2x number of time steps (for explicit methods)
System becomes more stiff

4x - 16x computing power
memory
hard drive space

Common Approach

- Increase resolution in large area of interest
- Can still use structured meshes (implicit data (coordinates/connectivity etc.) and better performance)
- Unflexible

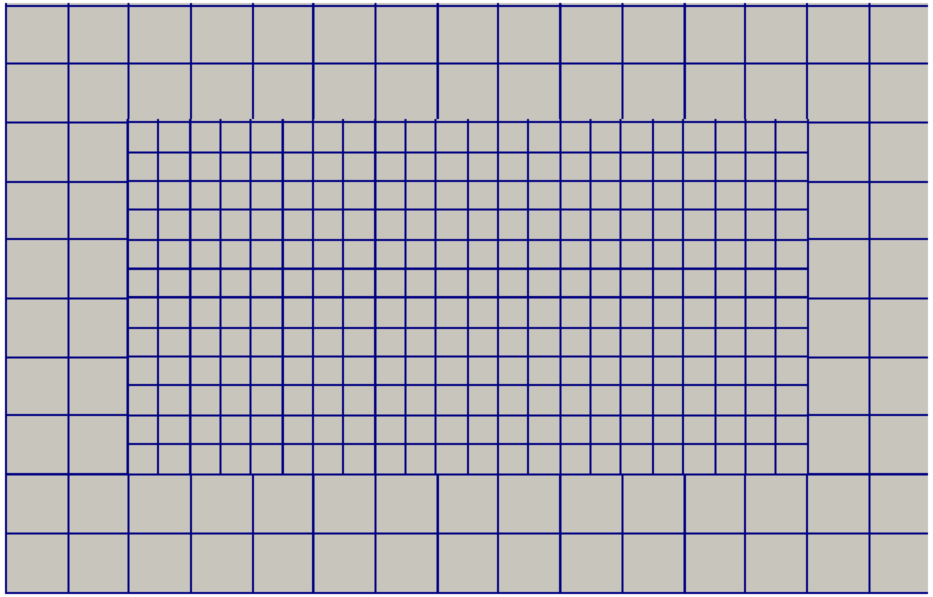

Dynamic AMR

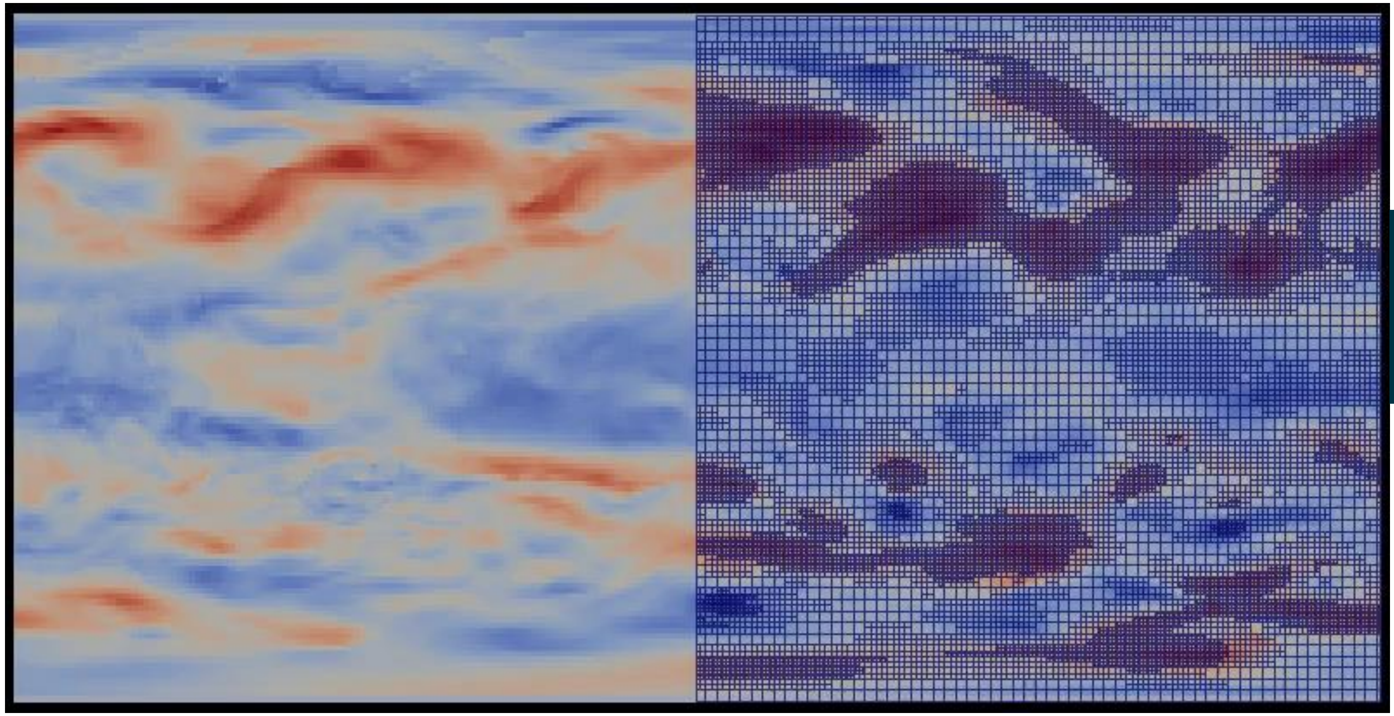

MESSy-Simulation, Wind speeds in the Earth‘s atmosphere

- Refine or coarse each element individually
- Change over time
- Concentrate computing power and memory usage on areas of interest while keeping a low error
- Reduces a number of elements by orders of magnitude
- Enables fine scale simulations that are not possible with uniform/nested grids

Challenges of AMR

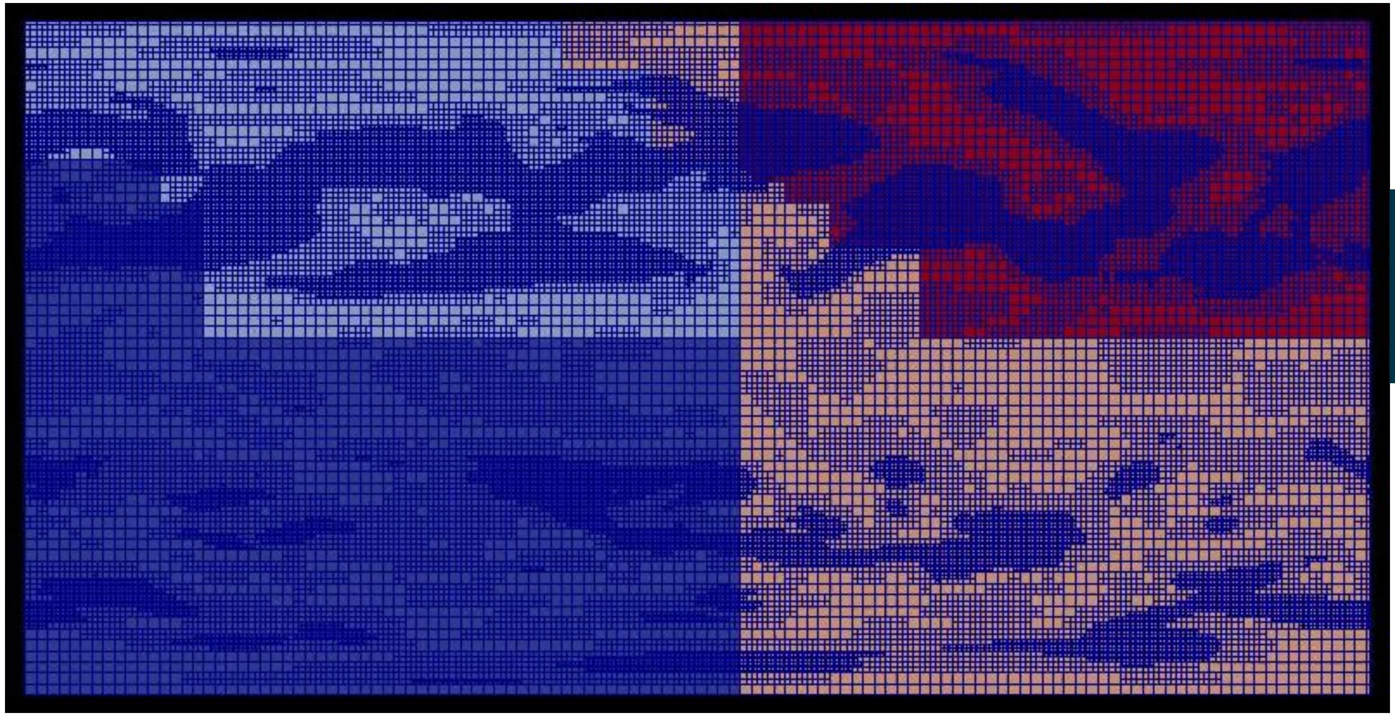

MESSy-Simulation, Wind speeds in the Earth‘s atmosphere

- Storage of mesh elements
- Load-balancing
- Ghosts
- Etc.

Unstructured meshes: memory usage, do not scale well, no implicit structure

⇨ Semi-structured meshes/space-filling curves (SFC)

AMR data structure

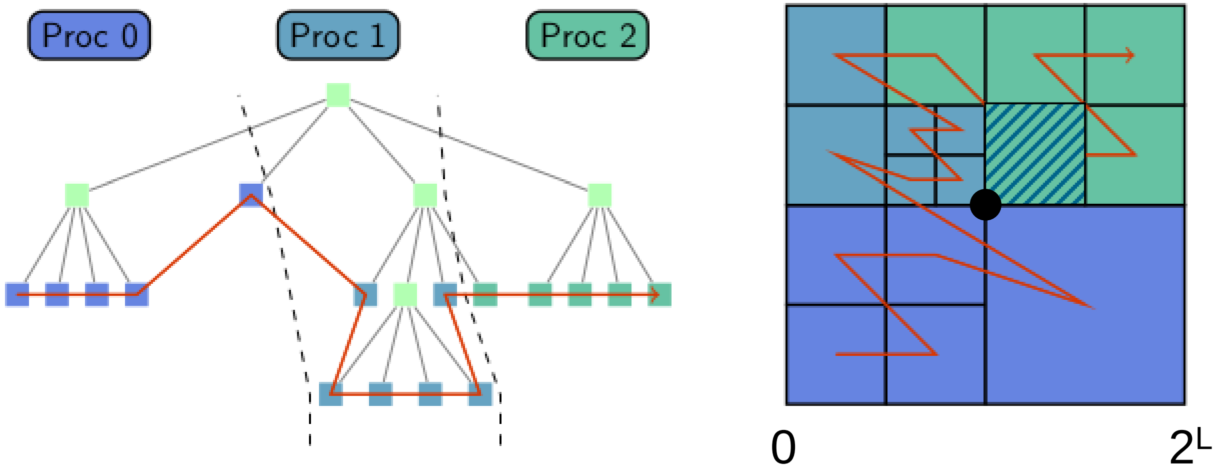

The SFCvisualized in the refinement Tree and in the associated tree

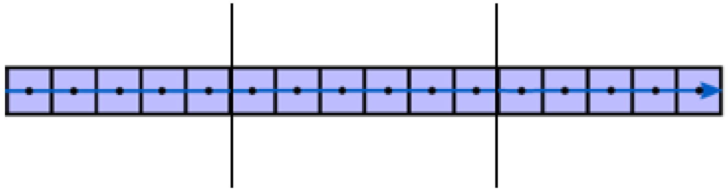

The SFC is linear in the memory.

Memory per Element:
3 (4) Integers, **x, y, (z)** and **level**

Memory efficient!

From tree to forest

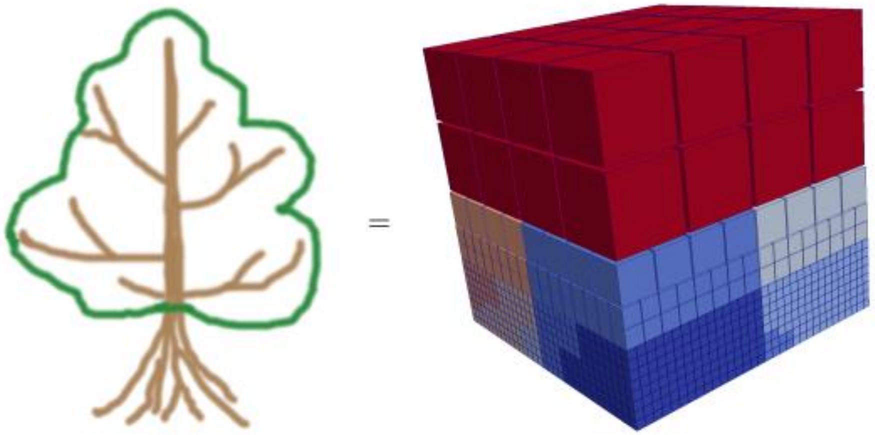

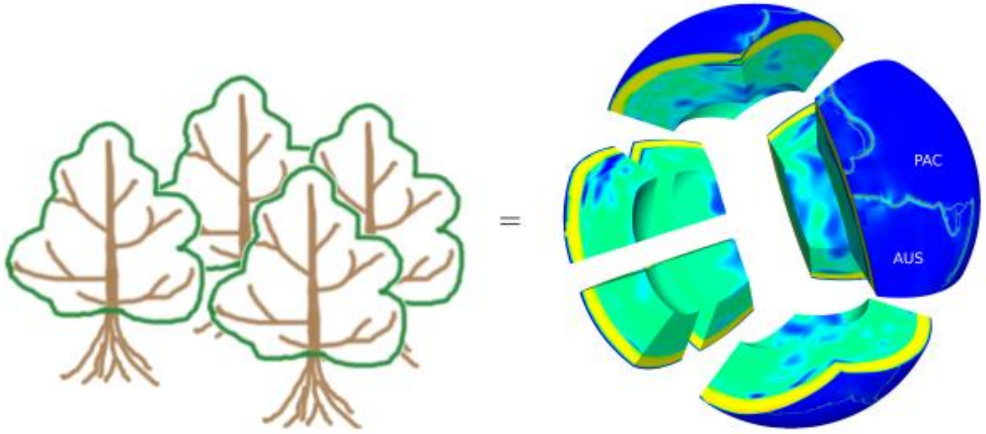

Science 329 (5995), p. 1033-1038

Images by Carsten Burstedde

From Quad to hybrid 3D

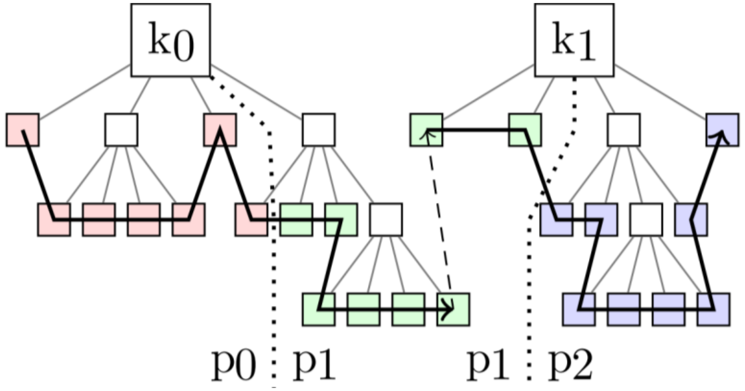

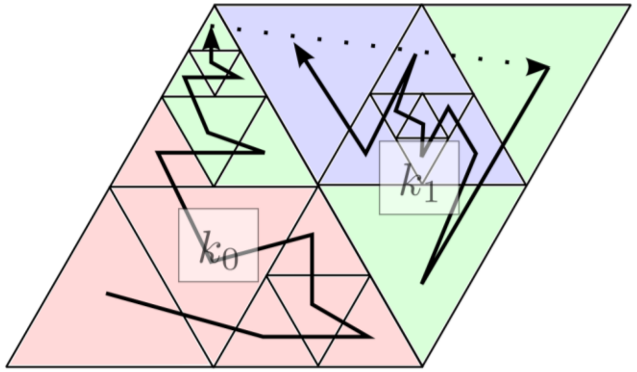

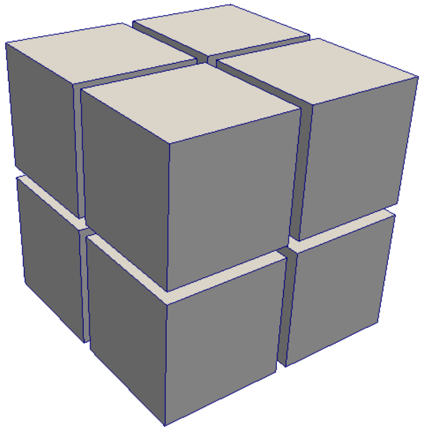

Hexahedra

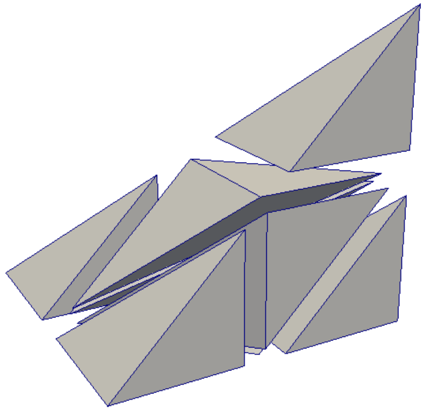

Tetrahedra

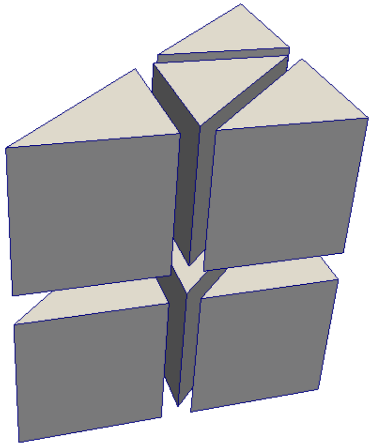

Prisms

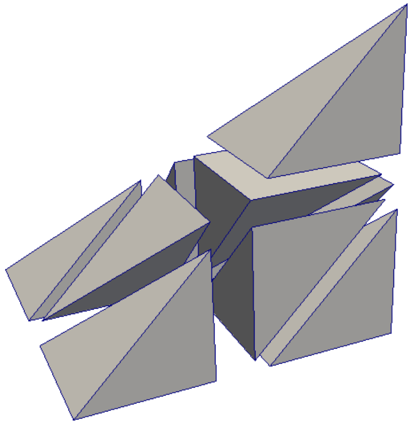

Pyramids

T8code („tetcode“)

DLR-AMR / t8code Public

DLR

- Parallel management of **adaptive meshes** and **data**
- Refine/Coarsen, Load-balancing, interpolating, ghost elements, etc…
- Tree-based/semi-structured with **space-filling curves**
- Supports all standard element shapes in 1D, 2D, 3D
- Scales up to **1 mio. MPI** ranks (with >90% efficiency)
- C/C++ and MPI
- Collaborations with: MESSy, JSC, University of Cologne, University of Bonn, …

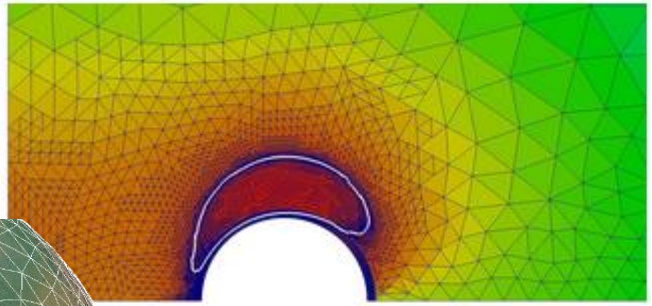

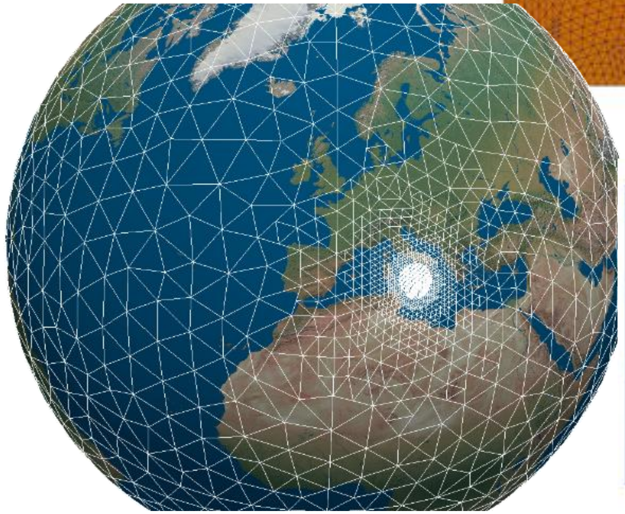

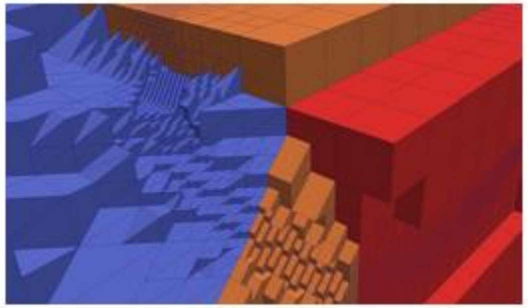

Performance – mesh management

Juqueen

Strong scale 1.9 billion tets

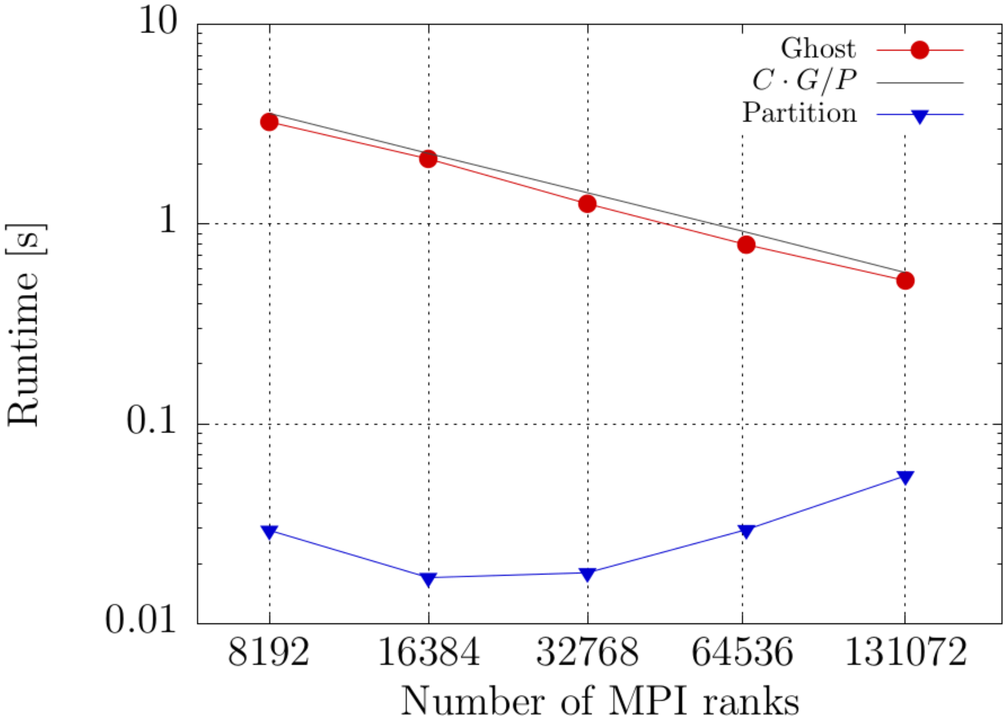

Weak scale, up to 142 billion elements

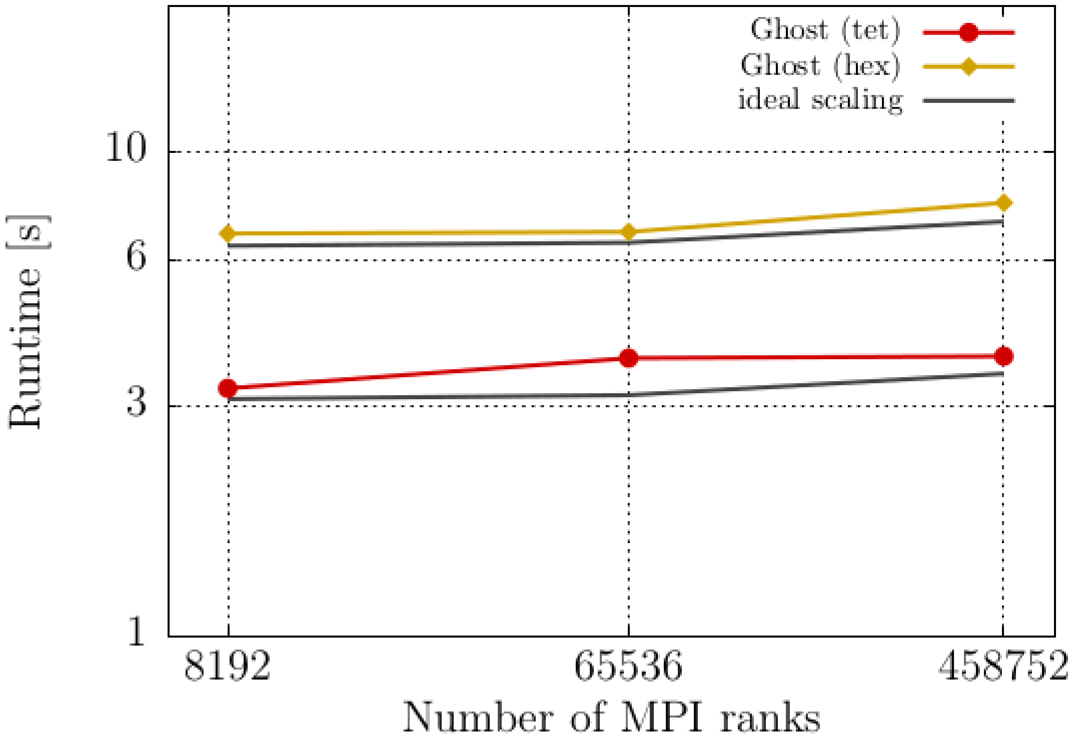

Source: An Optimized, Parallel Computation of the Ghost Layer for Adaptive Hybrid Forest Meshes, J. Holke, D. Knapp, C. Burstedde, 2019

t8code

- Designed as third-party library tool, not an application/simulation
- Controlled via Callbacks („Don‘t call us, we‘ll call you“)
- Originally developed for numerical simulation
- Now also used in post-processing, data compression, visualization

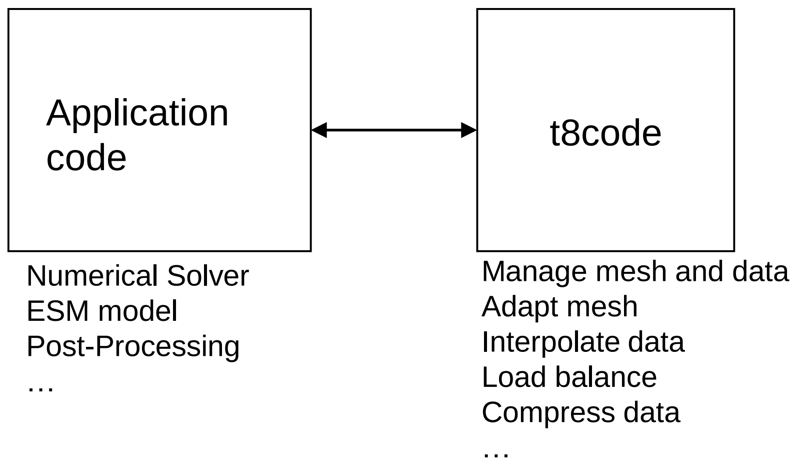

PilotLab ExaESM

www.exaesm.de

- PL-ExaESM was a Helmholtz incubator project from 2019-2022
- “explores specific concepts to enable exascale readiness of Earth System models and associated work flows in Earth System science.”
- Our research: Explore the potential of AMR in ESM
  - Advection-Diffusion DG solver with AMR
  - Data compression for MESSy as post-processing

Overhead of AMR

Scaling on JUWELS

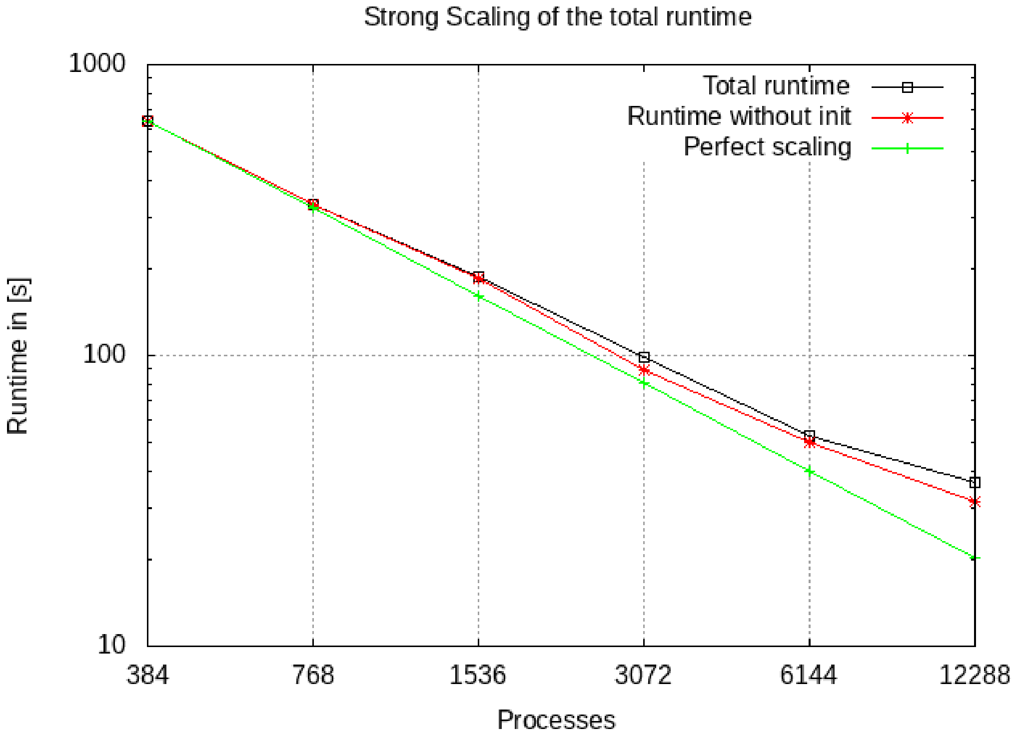

Solver

AMR

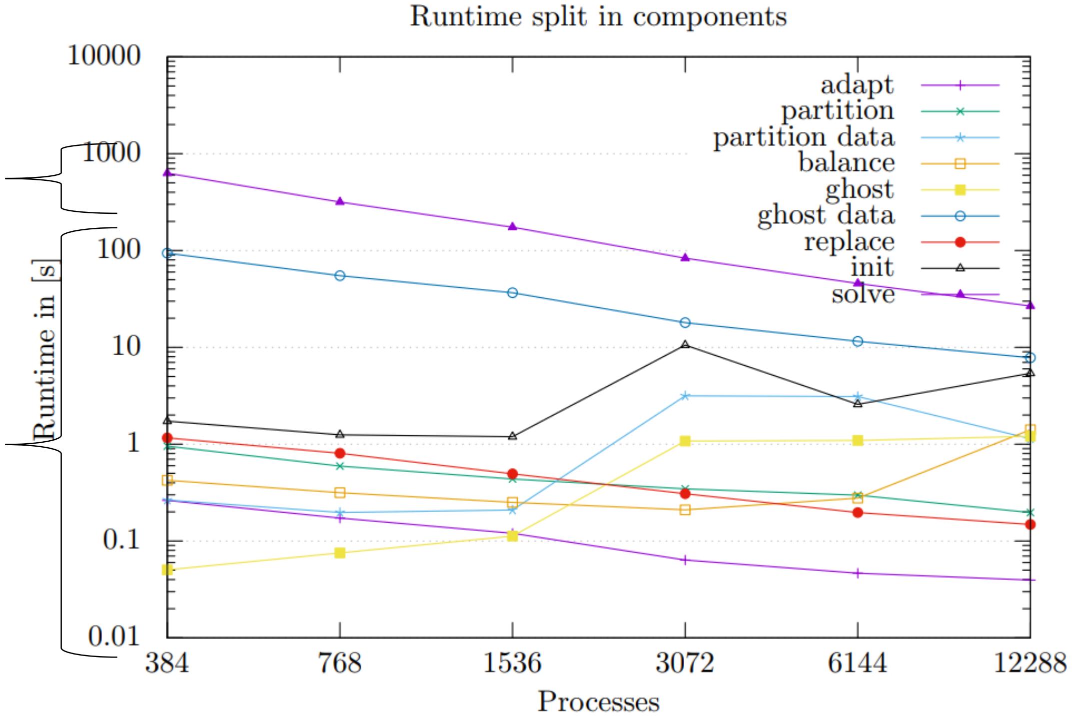

Advection/Diffusion with AMR, matrix-free, hex-mesh

The Local Discontinuous Galerkin Method for the Advection-Diffusion Equation on adaptive meshes
Master‘s Thesis by Lukas Dreyer at Uni Bonn

Adaptive vs. non-adaptive

Tests on JUWELS

	Runtime	Error	#DOFs
Uniform 3D	7057s	1.3e-3	16.777.216
Adaptive 3D	561s	1.5e-3	~1.920.000

12.6x speedup
8.7x less DOFs

Outlook: t8code + Trixi

Trixi

- Simulations on line/quad/hex/simplex meshes
- DG-methods
- High-order accuracy in space and time
- Hierarchical quadtrees/octrees with AMR
- Written in Julia

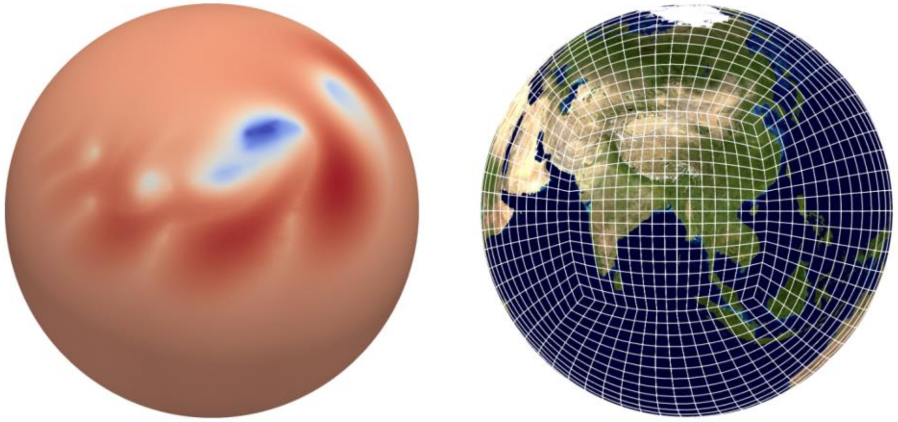

Figure 6.5: Surface pressure (hPa) at day 10 (left) of our most highly resolved simulation and the corresponding mesh (right), projected on a picture of the Earth [45], from the same perspective. The mesh has $16 \times 16 \times 8$ elements per cube face, a polynomial degree of $N = N_{\text{geo}} = 7$, and was stretched with a stretching factor $c = 2$. The same polynomial degrees have been used for the geometry approximation and the DGSEM.

Discontinuous Galerkin Methods for Atmospheric Simulations on Hierachical Meshes in Julia, Erik Faulhaber

Outlook: t8code + Trixi

Trixi	t8code + Trixi
• Simulations on line/quad/hex/simplex meshes • DG-methods • High-order accuracy in space and time • Hierarchical quadtrees/octrees with AMR • Written in Julia	• Extend AMR to hybrid meshes with our tree-based approach • Focus: Prisms • Written in C/C++

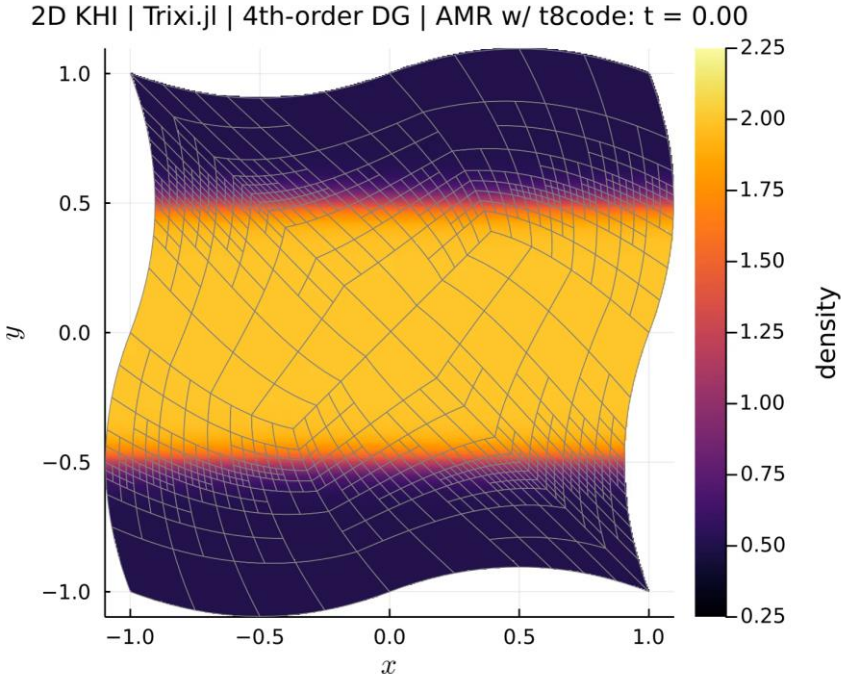

- Combine the best of each software
- In Progress: Julia Interface for t8code

Prisms in t8code

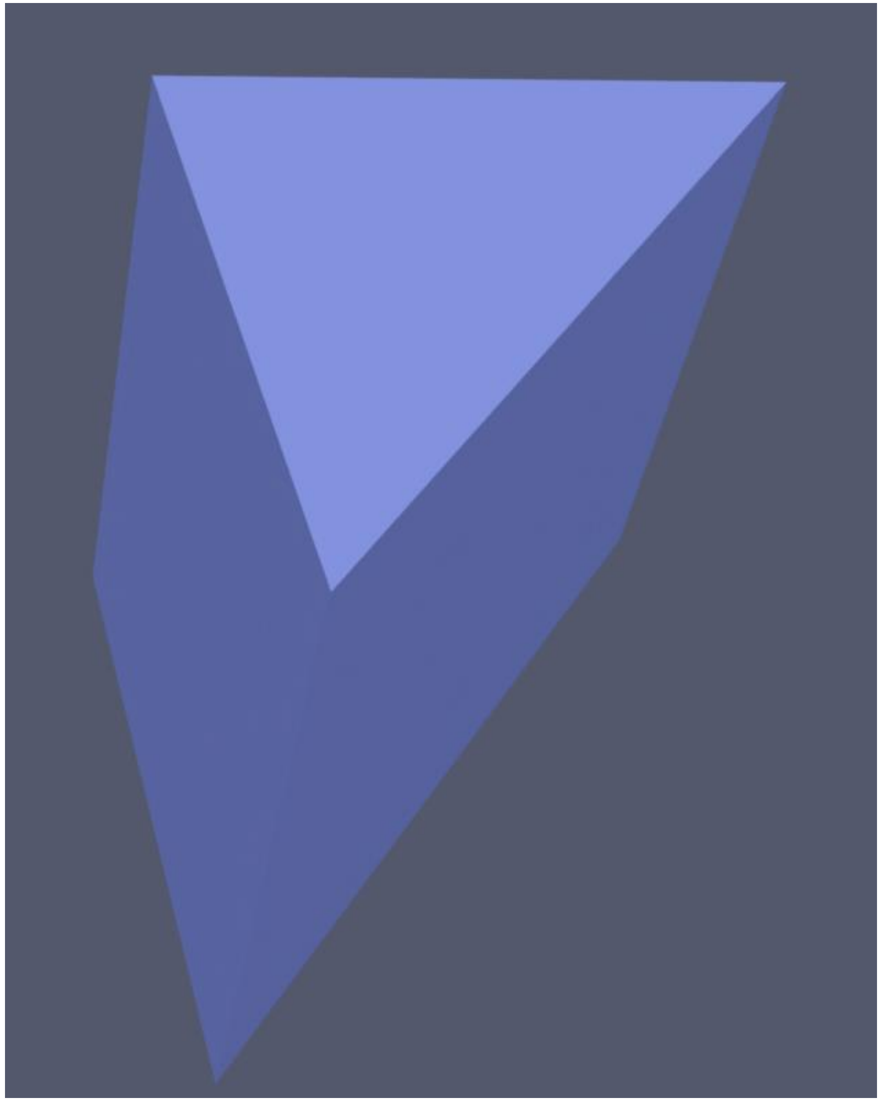

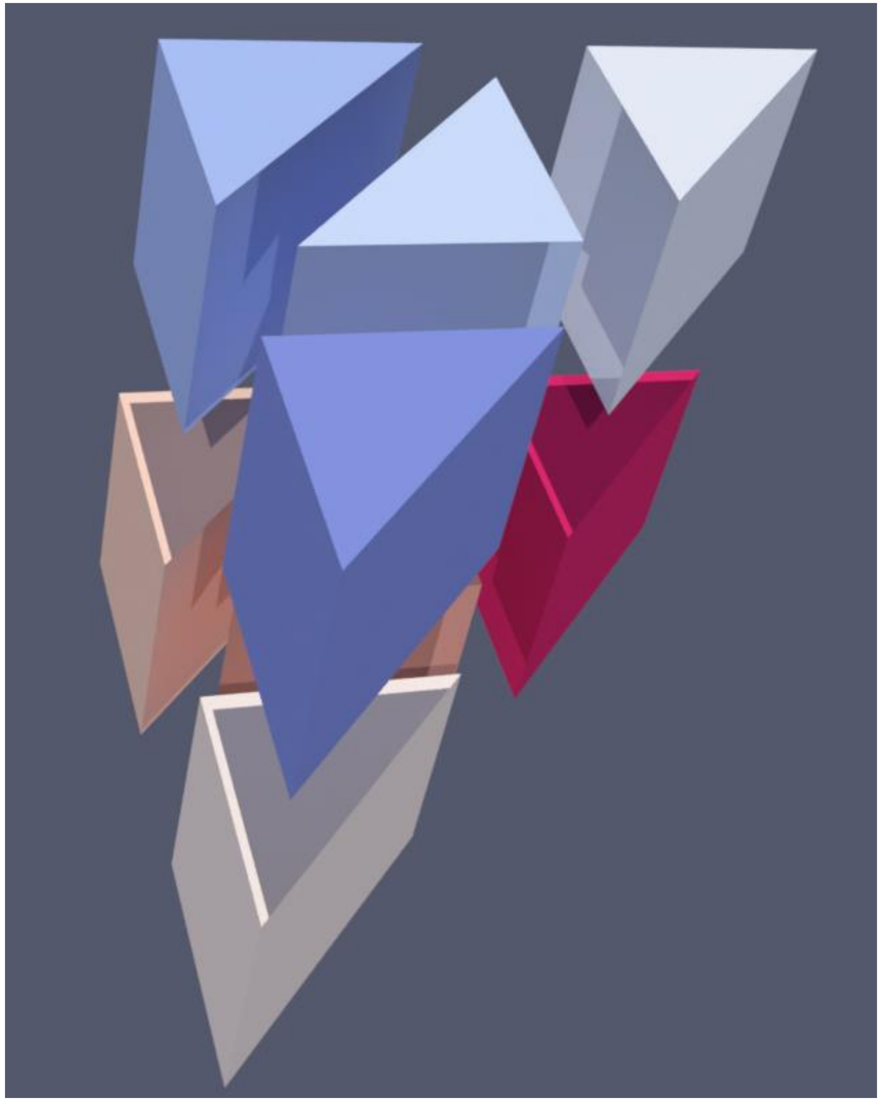

Data Compression with AMR

- AMR in **computation:**

  Concentrate **computing power** and **memory usage** on areas of interest while keeping a **low error**

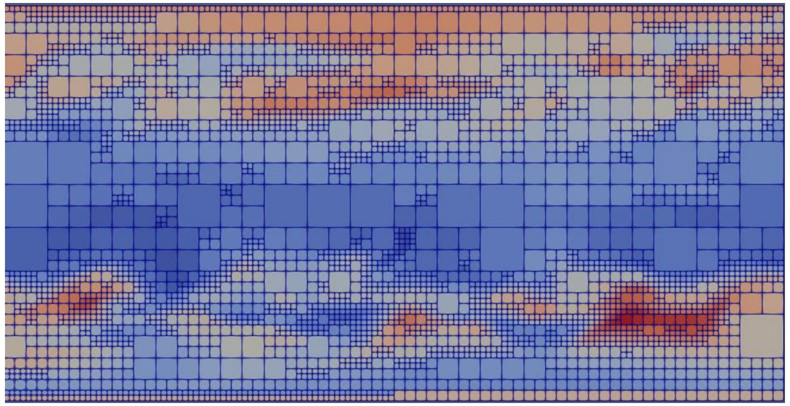

- AMR in **data compression:**

  Concentrate **data** and **disk usage** on areas of interest while keeping a **low error**

- **Lossy Data Compression** by coarsening simulation data with AMR

Data Compression with AMR

- Developing a flexible AMR Lossy Compression Tool

- Expecting higher compression ratios for…

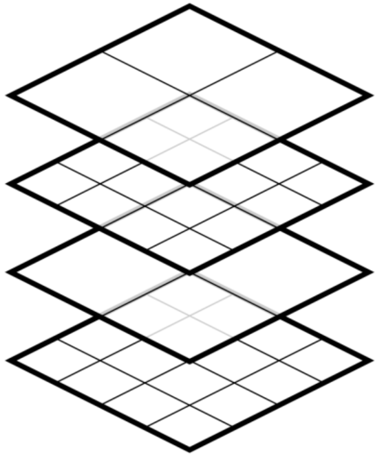

Coarsening each „atmospheric layer“ independently

Coarsening chemical species alone or grouped

Coarsening the atmospheric data on its 3D domain

Thank you

https://github.com/DLR-AMR/t8code

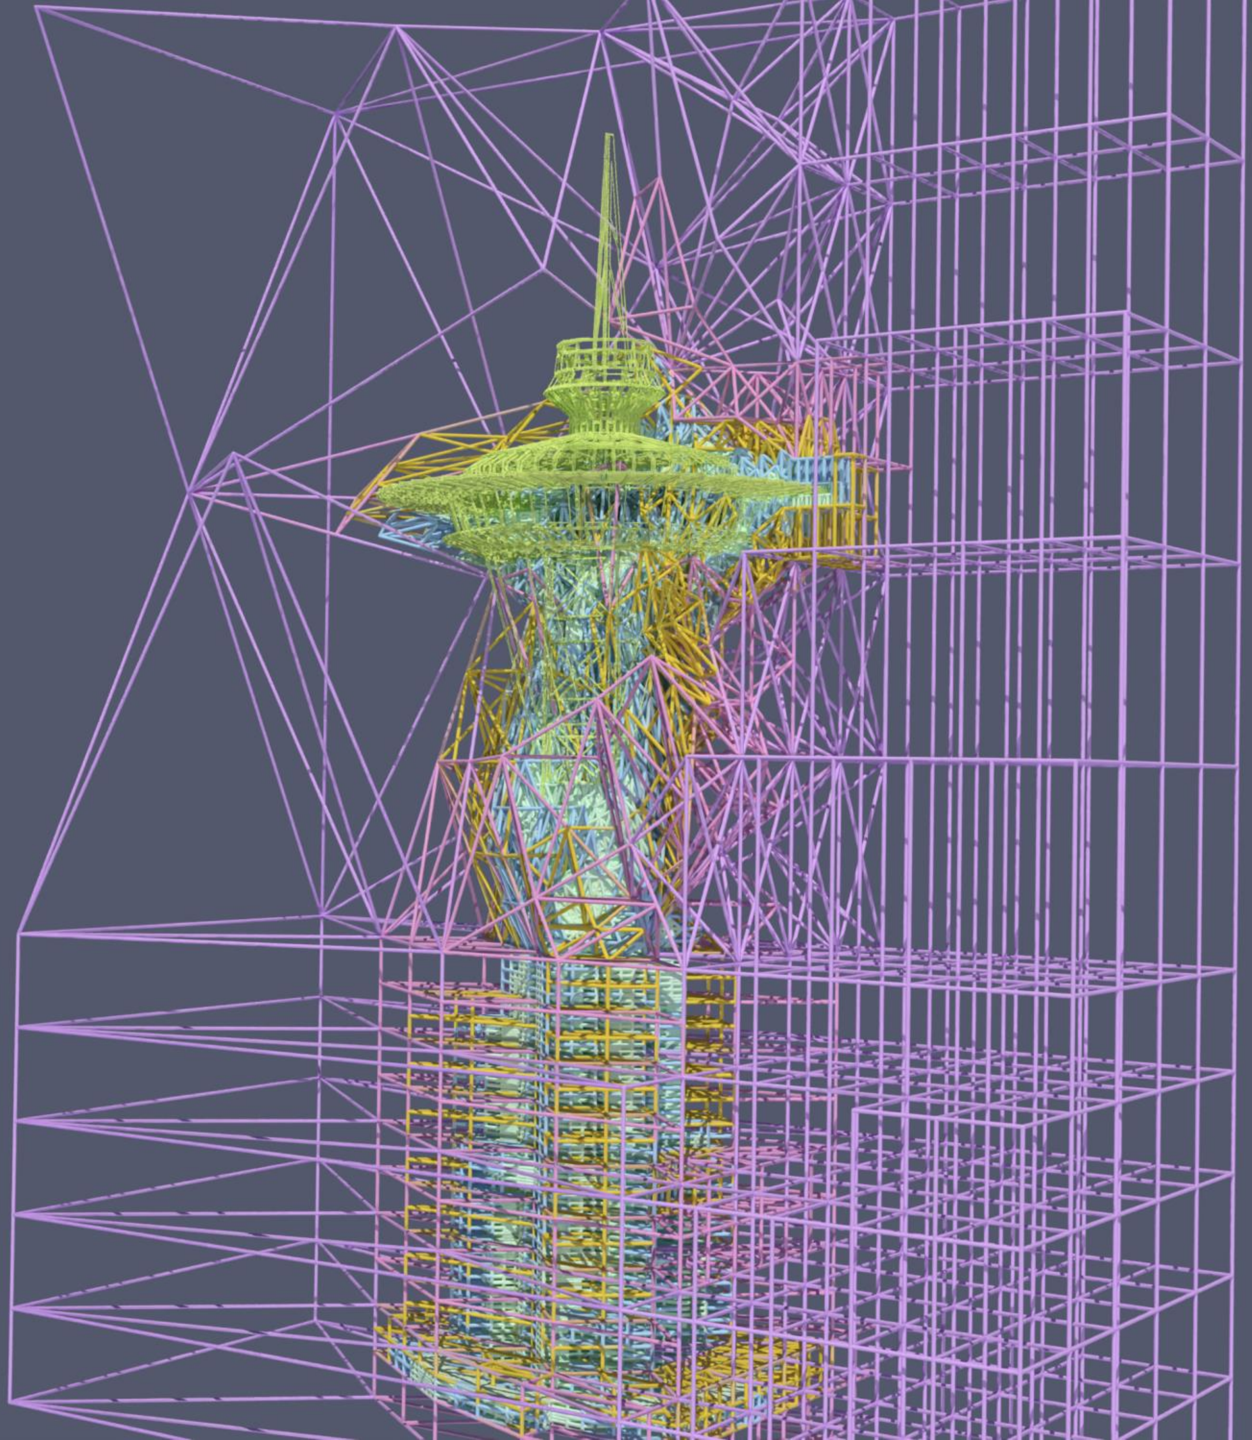

Question:

- What are good criterions to adapt?
  - -> Gradients/concentration etc.
  - -> In a new project(ADAPTEX) Error Estimator.
    Multiskala-methods, which are directly coupled with the DG-Method can give a criterion