

# Exascale-ready adaptive mesh refinement and applications in Earth system modelling

ECMWF - 19th workshop on HPC in meteorology, September 2021

Johannes Holke

German Aerospace Center – Softwaretechnology | High-performance Computing



Knowledge for Tomorrow



## Disclaimer

I am a numerical mathematician, not a meteorologist.

I have been working with meteorologists/climate scientists for ~2 years now.




Knowledge for Tomorrow




# Why AMR

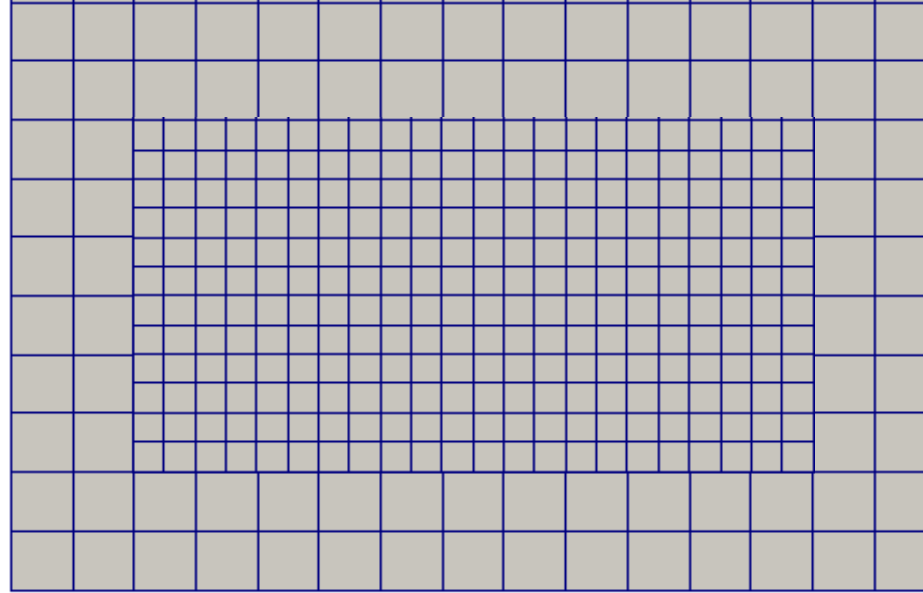
Increase resolution = increase simulation accuracy



4x/8x more DOFs and data per time step




## Common approach: Nested grids



- Increases resolution in large area of interest
- Can still use structured meshes (= implicit data and better performance)
- unflexible



# Dynamic AMR




- Dynamic AMR:
  - Refine or coarse each element **individually**
  - Change over time
- Reduce number of elements by orders of magnitude
- Enable fine scale simulations that are not possible with uniform/nested grids



# Challenges of AMR


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



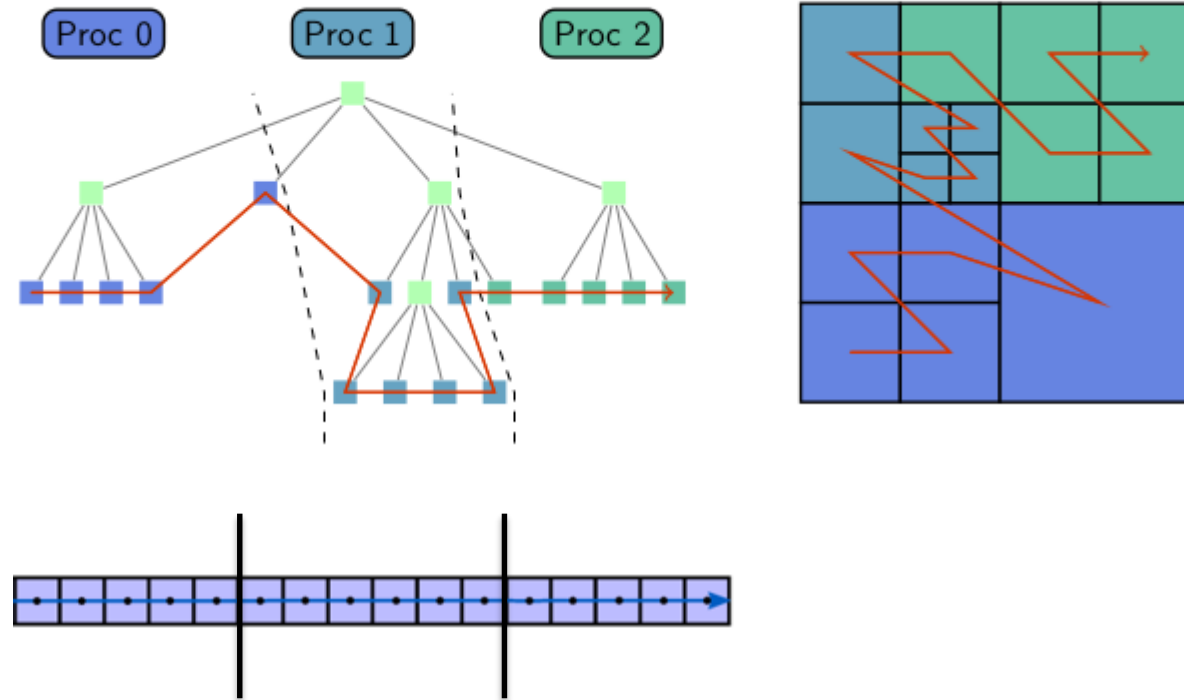
Unstructured meshes: memory usage, do not scale well, no implicit structure



# AMR data structure: Trees and SFCs





# AMR data structure: Trees and SFCs









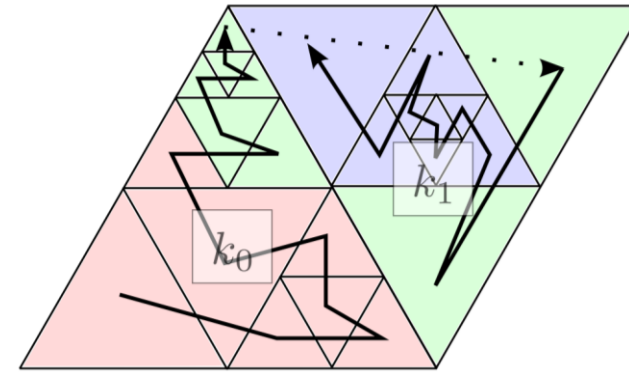

# From tree to forest



Science 329 (5995), p. 1033-1038



# From Quad to all



# t8code (“tetcode”) – AMR library






holke / t8code  
[github.com/holke/t8code](https://github.com/holke/t8code)

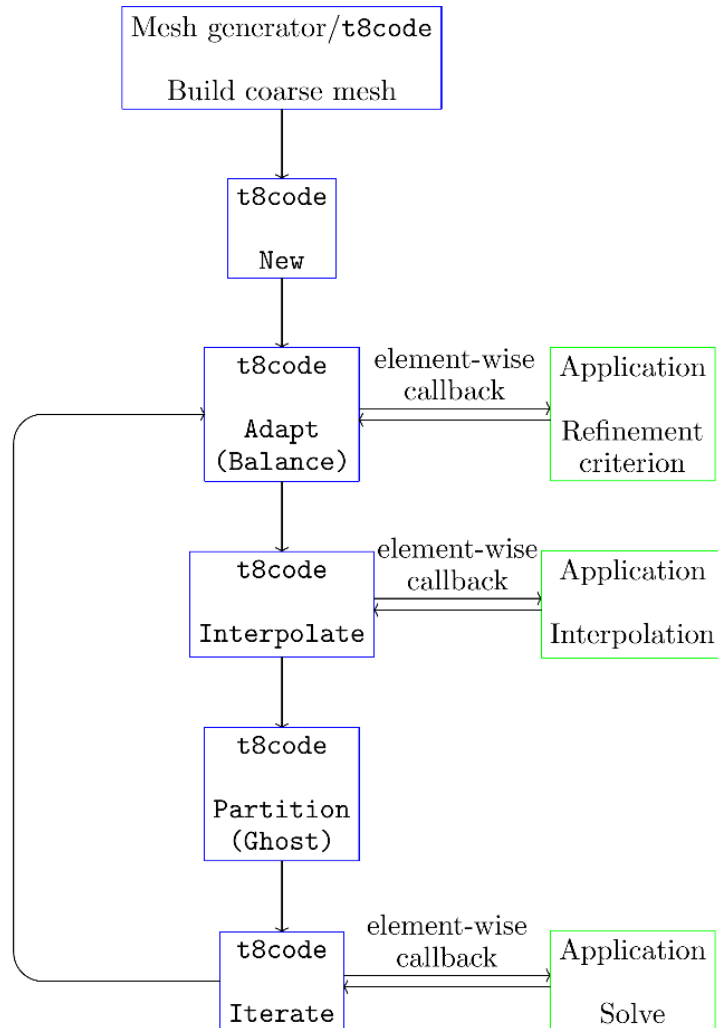
Application developer should not be concerned with mesh management

Parallel mesh and data management library

C/C++ with MPI



# The AMR simulation cycle



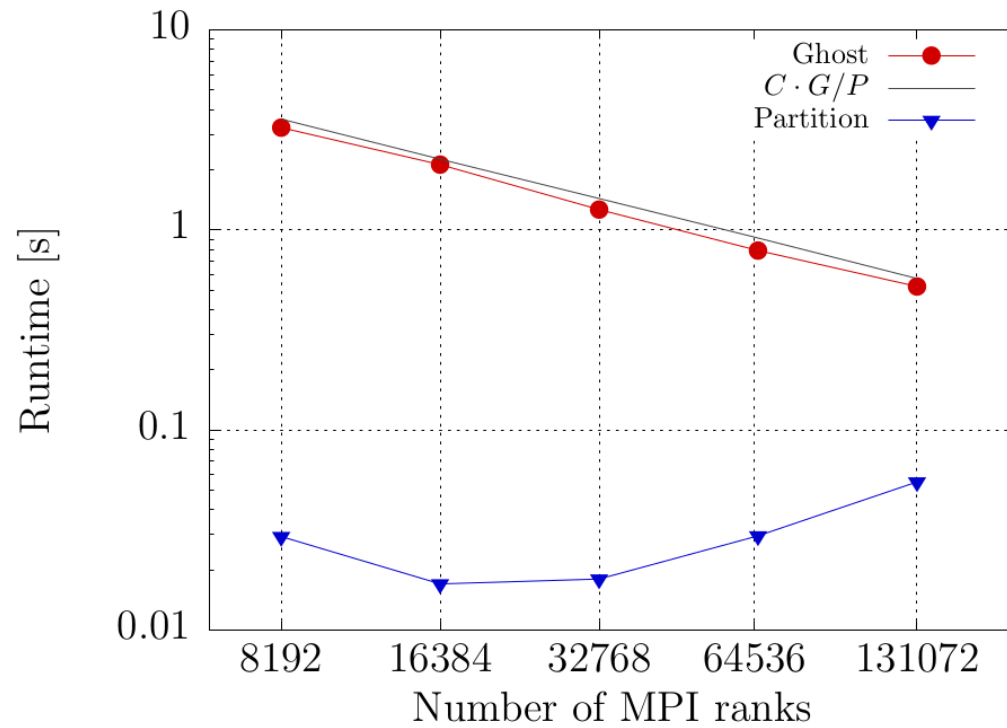
Application can freely specify how to

- Adapt the mesh
- Interpolate the data
- Solve the equation

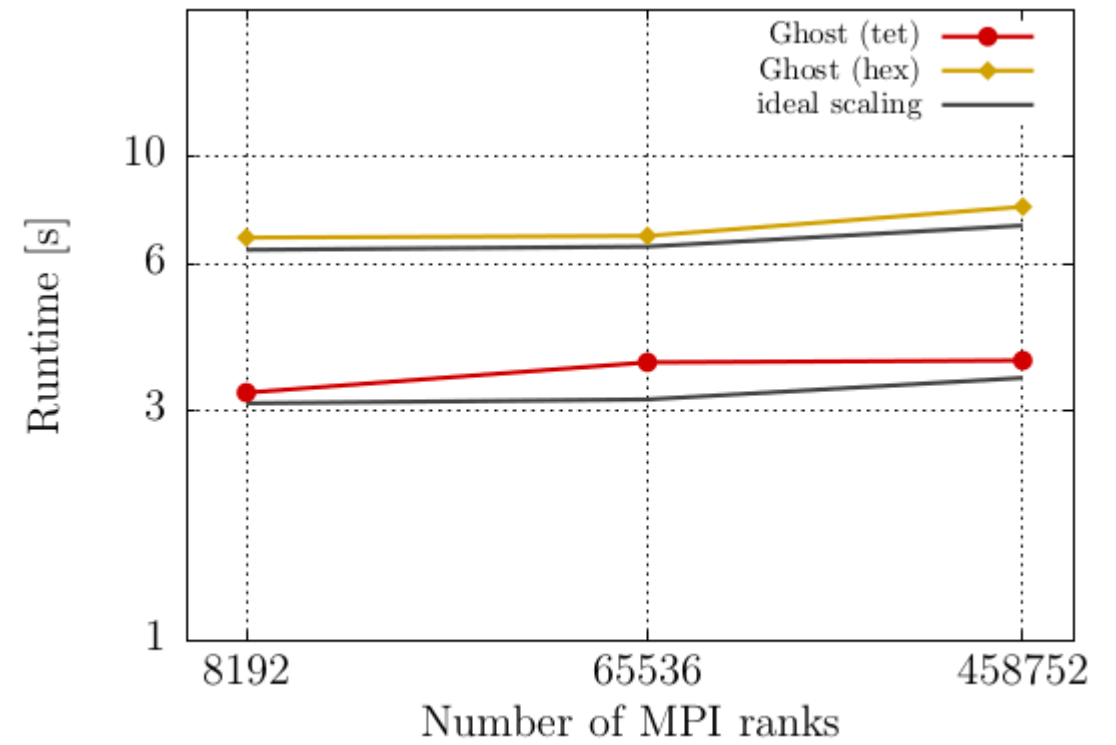


# Performance Juqueen

Strong scale 1.9 billion tets



Weak scale, up to 142 billion elements



# Performance

milestone: >1 Trillion elements on JUWELS

#processes	#Elements	#Elements/process	Ghost	Partition
98,304	1,099,511,627,776 $\approx 1.1e12$	11,184,811	1.43s	0.33s



# Two applications in ESM



[www.exaesm.de](http://www.exaesm.de)

„Explore potential of dynamic AMR in ESM“

1. Reduce I/O file size of MESSy
2. Full AMR Advection/Diffusion solver



Helmholtz-Zentrum  
Geesthacht  
Zentrum für Material- und Küstenforschung



HELMHOLTZ  
ZENTRUM FÜR  
UMWELTFORSCHUNG  
UFZ



# MESSy – lossy data compression



[messy-interface.org](http://messy-interface.org)

The highly structured  
**Modular Earth Submodel System (MESSy)**

developed by the consortium of



MAX-PLANCK-INSTITUT  
FÜR CHEMIE



Freie Universität



and supported by

Deutsches Klimarechenzentrum (DKRZ), Hamburg, Germany

MPG Rechenzentrum (RZG), Garching, Germany


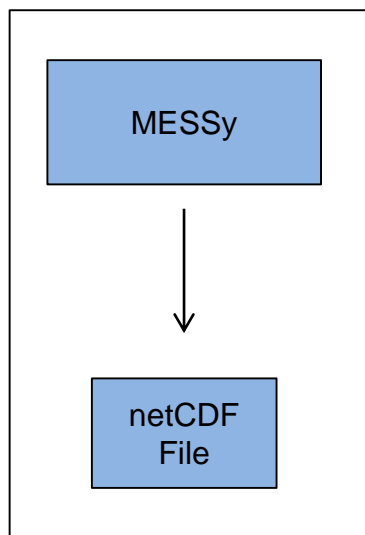
Leibniz Rechenzentrum (LRZ), Garching, Germany





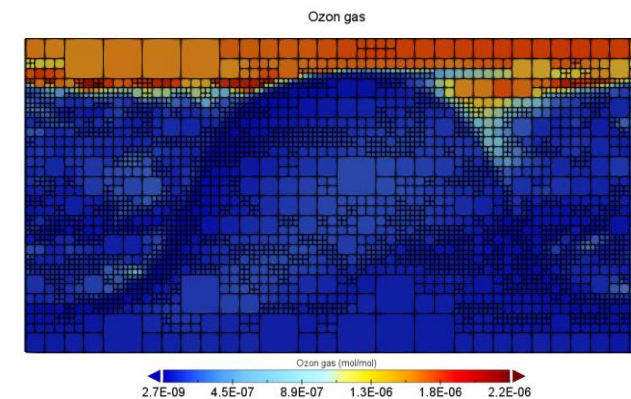
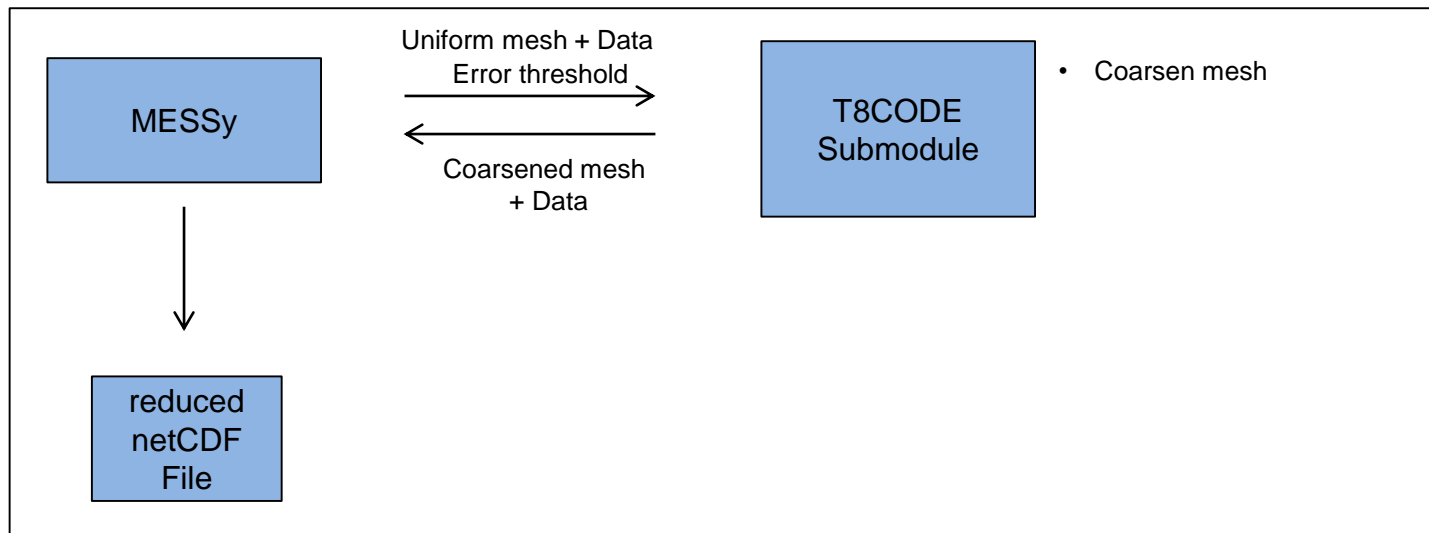
# MESSy – lossy data compression

Simulation Hundreds of chemical species result in very large output files



# MESSy – lossy data compression

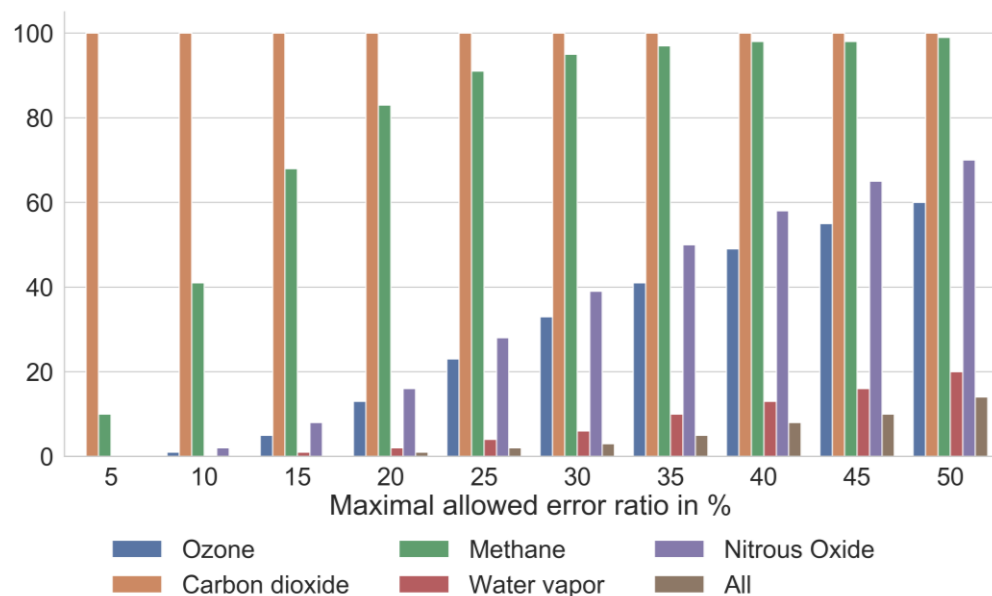
- Prototype/proof of concept



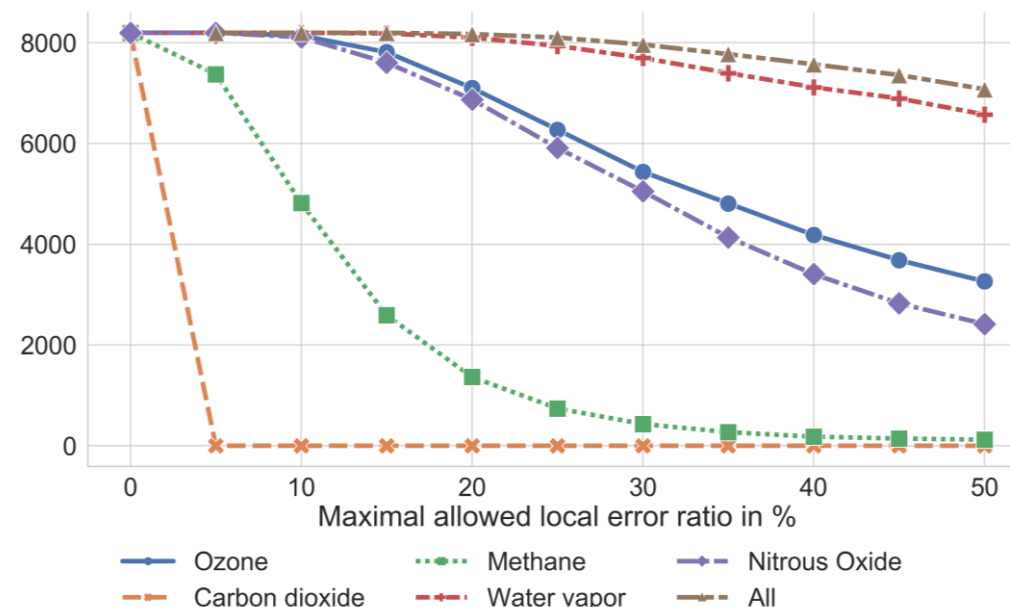
# First results

RC1-base-07<sup>1</sup> ECHAM, 128 x 64 x 90

File size reduction in %



Number of mesh elements



*Lossy data compression for atmospheric chemistry using adaptive mesh coarsening*


Master's Thesis by Luca Spataro at TU Munich

<sup>1</sup>Earth system chemistry integrated modelling (ESCIMO) with the modular earth submodel system (MESSy) version 2.51, Jöckel et. al., 2016




# Results

The first results are promising.  
Currently Coarsening over all z-levels simultaneously:




Much better compression rates expected when handling  
z-levels independently:



# T8DG Advection/Diffusion




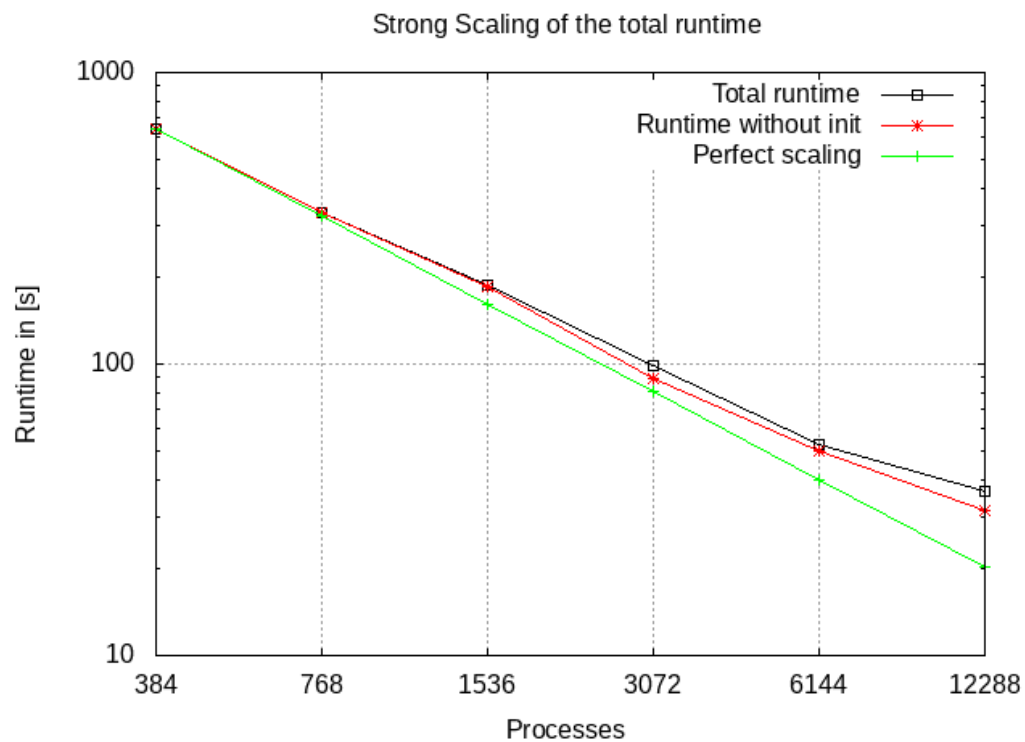
- Motivating example: Volcanic ash distribution
- High-order discontinuous Galerkin
- Matrix free
- Geometry support
- Implicit
- Multigrid



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



# Scaling on JUWELS



# Adaptive vs. non-adaptive

	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

# Conclusion/Outlook

- AMR can significantly reduce computing time and memory/disk usage
- AMR is efficient and scales (for all element shapes)
- You should not do AMR yourself
- Stronger Coupling MESSy + t8code plannend
- Ongoing Coupling to MPTRAC from JSC
- Suitable adaptation criteria in ESM?





# Thank you

## Special thanks to

Luca Spataro, Lukas Dreyer, Niklas Böing, David Knapp

Patrick Jöckel, Kerstin Hartung

Lars Hoffmann, Olaf Stein

Carsten Burstedde

Gregor Gassner

Michael Bader

