


# TOPIO IO FOR EXASCALE

## Motivation Forging a New Link

As available **computing power increases**, improvements in weather forecasting lead to a dramatic **increase in the amount of data** generated in each simulation

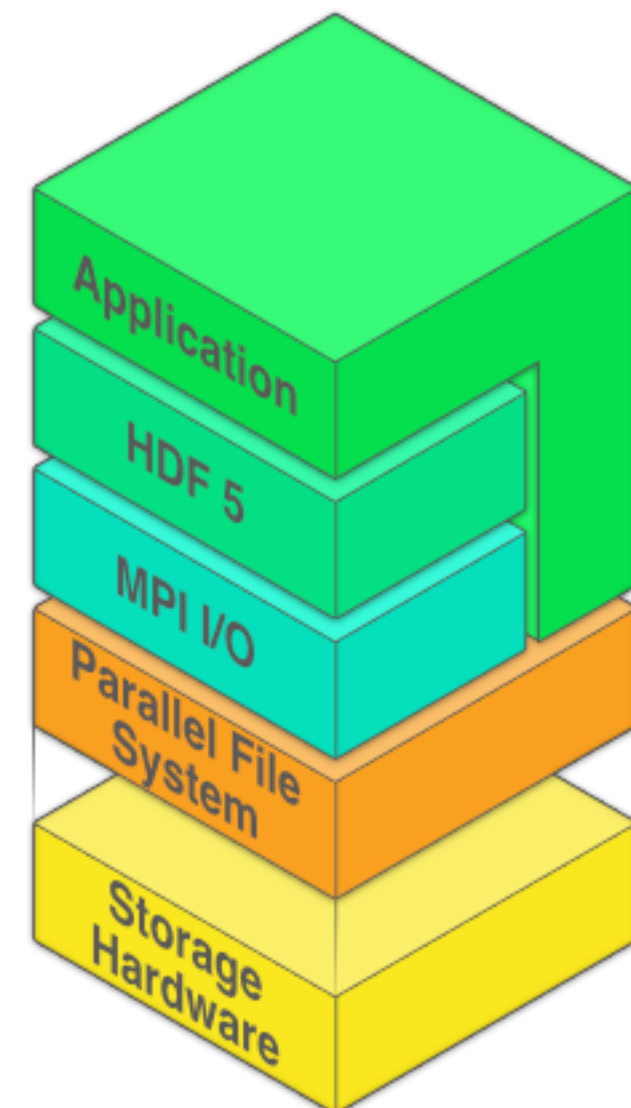
- **Current resolution** of seasonal weather forecasts is not sufficient to represent the interaction between ocean, land surface, and atmosphere with high accuracy
- Push for a global weather modeling using a horizontal **resolution of 1.5 km**
- High resolution modeling method would result in datasets that could exceed **1 petabyte in size**
- **Improving** compute resource utilisation of earth-system simulation applications (e.g. MPAS) requires a revision of their I/O approach
- Optimised **I/O** and sensible **data reduction/compression** become necessary



## I/O Twist that dial

**Optimizing the I/O activities** of an application and achieving efficient parallel data transmission are demanding tasks


- **Complex dependencies** between the layers of the I/O stack
- Each layer offers **several configuration parameters**
- Configuring these parameters **depends on various factors (e.g. striping, I/O pattern)**
- The users **may not be able to optimally tune** their applications
- In case default settings are used, this often **results in poor I/O efficiency**



## Compression What is it good for

Tests have already shown that **available compression implementations** can drastically reduce the memory requirements of large simulations

- **Lossless encoders** (e.g. Zstd): modest compression, perfect reconstruction
- **Block transform based lossy compression** (e.g. ZFP): fast, good compression and fair reconstruction
- **Discrete Wavelet Transform based lossy compression** (e.g. BigWhoop): reasonably fast, good compression and excellent reconstruction
- Compression **can and should be used** to effectively reduce the memory footprint of numerical datasets



(Original)

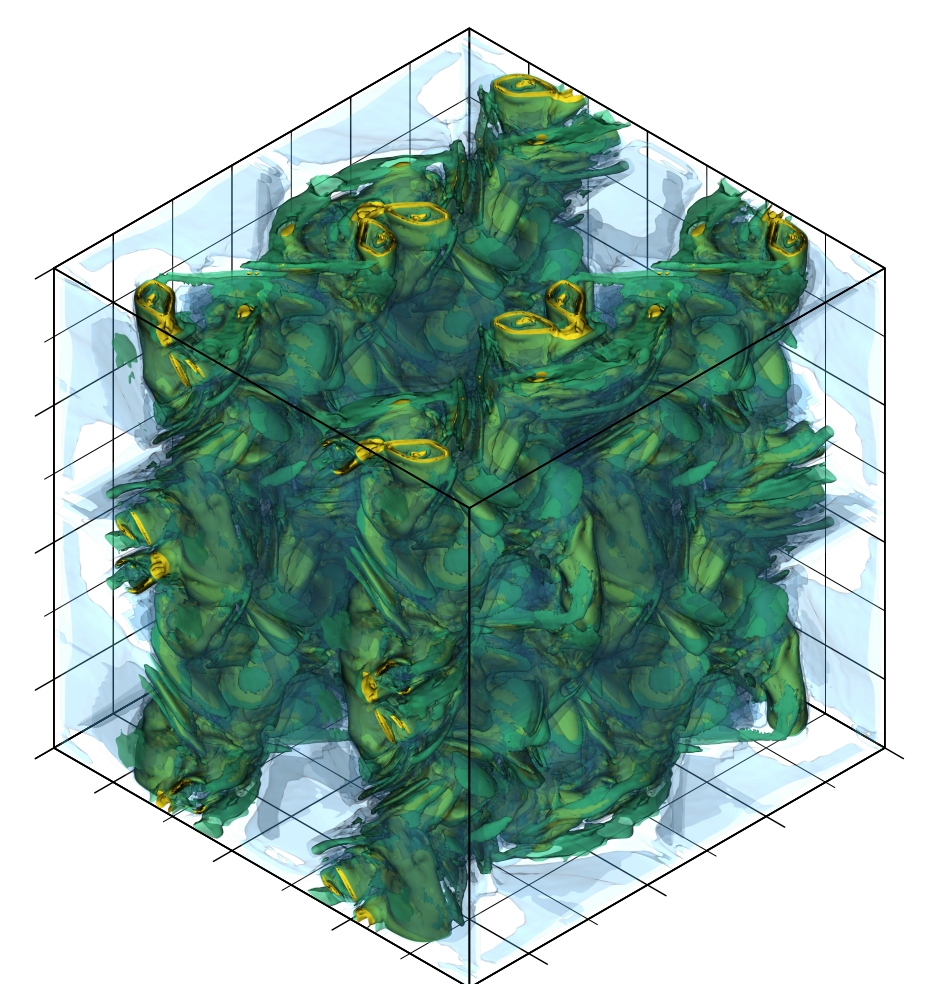
## I/O Project Goals

- **Evaluation** of the I/O performance on the **HLRS Lustre file system**
- **Optimisation** of the I/O performance on the HLRS Lustre file system **(with a simple strategy)**
- **Integration** of different optimisation strategies into the auto-tuning tool, **analytical/machine learning models**
- **Optimisation** of the I/O performance on the HLRS Lustre file system with the new optimisation methods
  - Further developments on the auto-tuning tool and on the compression libraries should then further improve performance; the evaluation is based on the MPAS model application provided
- A combination of the I/O optimisation tool and lossless compression could therefore reduce the amount of data without significantly increasing the total runtime


## Compression Project Goals

Compression approach **needs to fit** the MPAS use cases

- Analyse **reconstruction requirements** for global, high-resolution weather simulations
- Evaluate **appropriate tools** and approaches for online compression and offline storage
- **Optimisation and adaptation** of the self-developed compression library for the current use case
- **Formalise an API** for the compression methodology
- **Integrate compression methodology** into MPAS



(Compressed - 400 : 1)



## Project Target Engage

TOPIO aims to provide the HPC community with a tool that enables both auto-tuning and data compression as a **quasi-black-box** process

- The auto-tuning and compression tools will be **published** in a repository
- The software modifications applied to MPAS **will be made available** to their respective developers as well as wider HPC community
- Global **high-resolution** sub-seasonal simulations with MPAS, which will be computed on the HPC Cluster of HLRS within the framework of the project
- I/O performance of the high-resolution simulation will be **analysed and evaluated** in detail and published in scientific papers

### Project Coordinator

High-Performance Computing Center Stuttgart (HLRS)  
Prof. Dr. Michael M. Resch  
Nobelstraße 19  
D-70569 Stuttgart  
Tel: +49 711 685-87200  
E-Mail: resch@hlrs.de

### Project Volume

0.75 Mio. €  
BMBF-Funding: 0.75 Mio. € (100 %) with add. 0.15 Mio. € of so-called Project Allowance.  
Funded under the "New Methods and Technologies for Exascale High Performance Computing (SCALEXA)" funding priority.

### Project Duration

01.11.2022 until 31.12.2025

### Project Partners

- High Performance Computing Center Stuttgart (HLRS), University of Stuttgart
- Institute of Physics and Meteorology (IPM), University of Hohenheim

### Contact

Anna Mack, Patrick Vogler - HLRS  
E-Mail: anna.mack@hlrs.de, patrick.vogler@hlrs.de

Arun Kumar Dwivedi - IPM  
E-Mail: arunkumar.dwivedi@uni-hohenheim.de