



MAX PLANCK INSTITUTE  
OF GEOANTHROPOLOGY

**PISM**  
PARALLEL ICE SHEET MODEL



**Sprint 17**

# Challenges and results experienced during the PISM sprint

April 2025 – October 2025

**Torsten Albrecht (MPI-GEA/PIK), Wilton Jaciel Loch (DKRZ), Constantine Khrulev (UAF)**

# General information

PISM Home About Usage Download More

## PISM

### PARALLEL ICE SHEET MODEL

Get it on GitHub

$$\rho_i c_i (\partial_t T + v_\mu \partial_\mu T) = k_i \partial_{zz}^2 T + \Sigma$$

```

void TemperatureModel::update_imp(double t, double dt, const Inputs &inputs) {
    // current time does not change here
    (void) t;

    using namespace PISM;

    Logger::log(MPI_COMM_SELF, m_log-&gtget_threshhold());

    const double
        ice_density      = m_config-&gtget_number("constants.ice.density"),
        ice_c            = m_config-&gtget_number("constants.ice.specific_heat_capacity"),
        L                = m_config-&gtget_number("constants.fresh_water.latent_heat_of_fusion"),
        melting_point    = m_config-&gtget_number("constants.fresh_water.melting_point_temperature"),
        beta_cc_grad     = m_config-&gtget_number("constants.ice.beta_Clausius_Clapeyron") * ice_density * m_config-&gtget_n

    const bool allow_above_melting = m_config-&gtget_flag("energy.allow_temperature_above_melting");

    // this is bulge limit constant in K; is max amount by which ice
    // or bedrock can be lower than surface temperature
    const double bulge_max = m_config-&gtget_number("energy.enthalpy.cold_bulge_max") / ice_c;

    inputs.check();
    
```

Photo: C. Matias / Unsplash



**What**  
... is PISM?

Heard about PISM for the first time and want to know what exactly it is?

Read the overview



**How**  
... to use PISM?

Using PISM yourself and forgot about a specific option or want to learn how to use it?

Read the manual



**Where**  
... to get PISM?

Want to start using PISM yourself or just want to have a peek into the code?

Get the code

**Website**  
[pism.io](https://pism.io)

**Open-source**  
[github.com/pism](https://github.com/pism)

**Documentation**  
[pism.io/docs](https://pism.io/docs)

# General information

Co-development at PIK since 2008 (Winkelmann et al., 2011) and at UAF, Alaska (Bueler & Brown, 2009) [1,2]

Many users and contributions worldwide (see map [3]) with more than 240 peer-reviewed publications [4]

## Key research topics

- Interaction between Greenland and Antarctica via ocean, atmosphere and solid Earth
- Glacial cycle simulations and future sea-level projections
- Effect of extreme events on ice dynamics
- Tipping points and risk of tipping cascades

## Technical details

- C++ code, MPI parallelization using PETSc toolkit
- interfaces to ocean, atmosphere and Earth components
- regular grid
- NetCDF based I/O



[1] <https://www.pism.io/history>

[2] <https://www.pism.io/team>

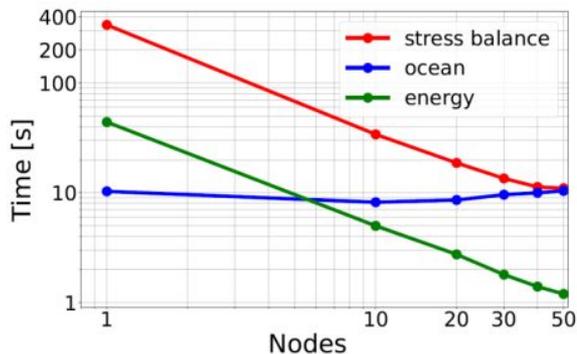
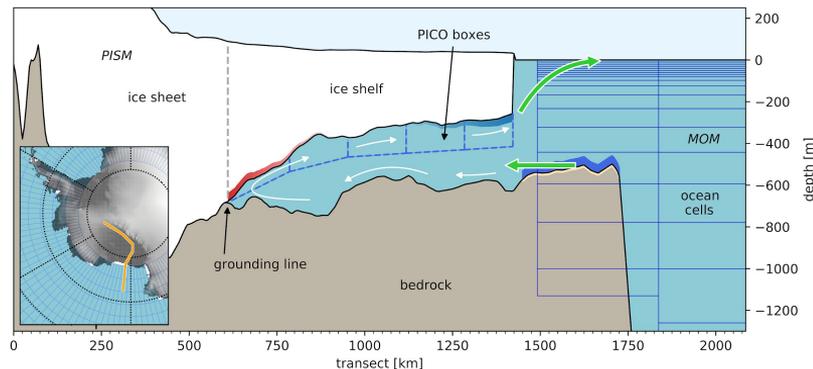
[3] <https://www.pism.io/usersmap>

[4] <https://www.pism.io/publications>

# Sprint check

carried out by Enrico and Wilton (both DKRZ)

- scalability sprint preparing PISM for future high-resolution simulations (e.g. natESM integration)
  - Requirements for multi-million grid points
  - GPU compatibility requires in-depth analysis of PETSc performance



- Performance limitations identified for CPU:
  - limited scaling on CPU
  - I/O bottlenecks at high resolution
  - Flat strong scaling in PICO ocean component

PICO: Reese et al. (TC, 2018), Kreuzer et al., (GMD, 2021)

# Challenges - Asynchronous Output Server

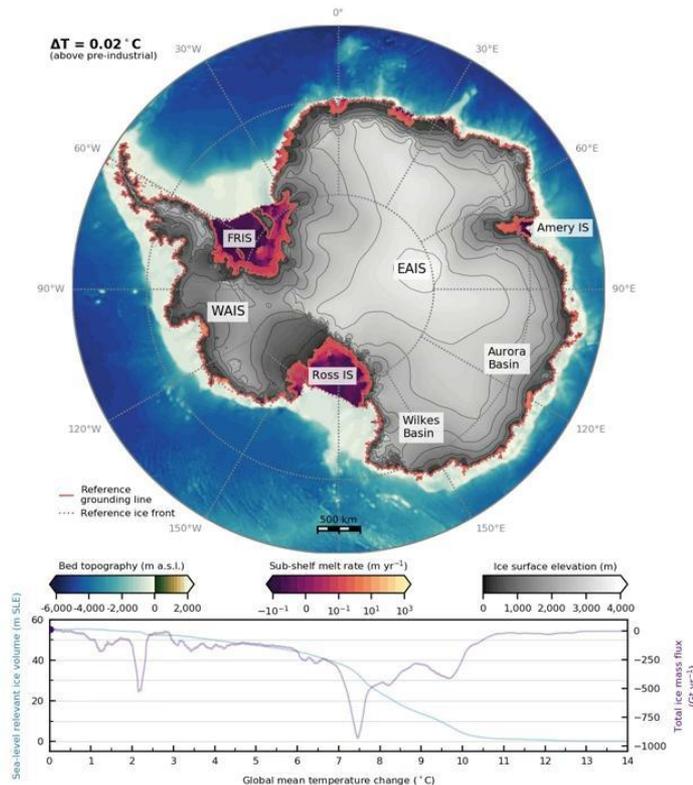
**Problem:** Synchronous I/O Bottleneck

Simulation pauses during output of different file classes:

- timeseries (1D, high time resolution)
- extra (2D or 3D diagnostics)
- snapshots/backups (PISM restartable 3D)
  - extended grid for bed deformation

Output time grows linearly with simulation length

At high resolution → severe performance penalty, as I/O dominates runtime in long simulations.



Tipping in Antarctica: Garbe et al. (Nature, 2020), Winkelmann et al., (NCC, 2026)

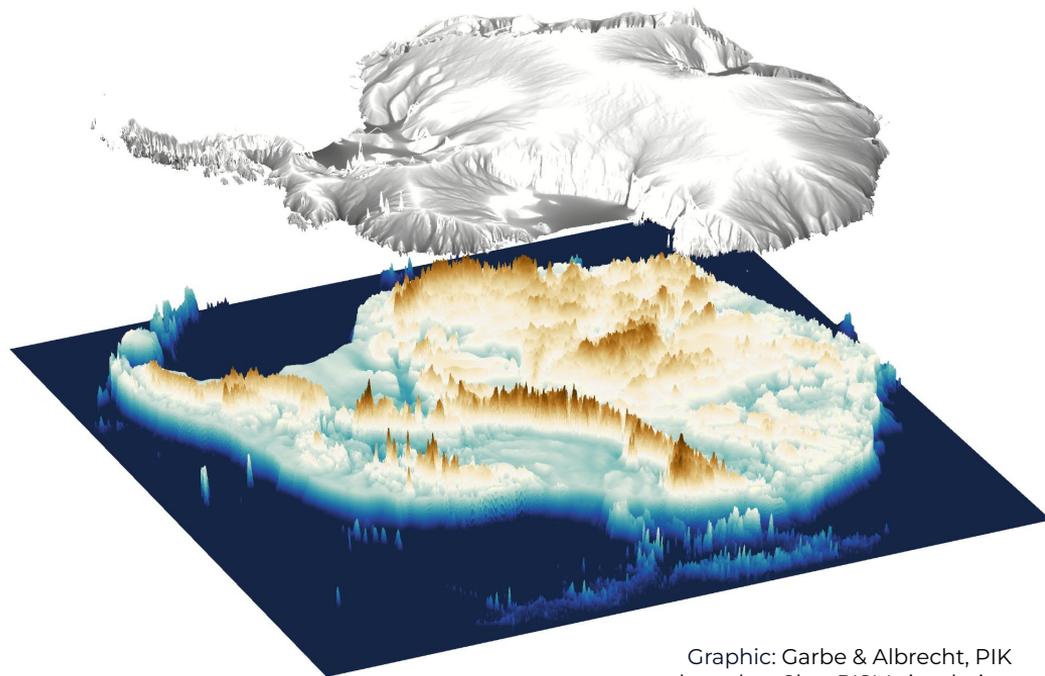
# Challenges - Asynchronous Output Server

## Goals:

- Decouple simulation from file writing
- Preserve existing output structure
- Allow runtime switch (sync vs. async)
- Keep flexibility for future coupling

## Tools:

- Evaluated: ADIOS2, YAXT and YAC
- Chosen: YAC
  - Already integrated in PISM
  - Multi-language support
  - Metadata & time management
  - Flexible design
  - Yet, YAC not designed as output library



Graphic: Garbe & Albrecht, PIK  
based on 2km PISM simulation

# Major Technical Challenges - Asynchronous Output Server

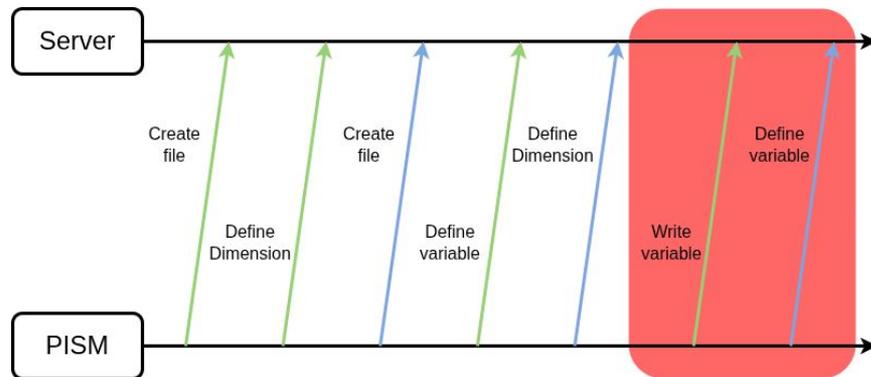
## Grid information availability:

- Grid only exists inside PISM at runtime
- Not guaranteed in external files
- Needed to transfer grid metadata via MPI intercommunicator

## Interleaved Output Operations (core architectural conflict):

- PISM:
  - Arbitrary order of file operations
  - Multiple files handled interleaved
- YAC:
  - Strict phase-based design
  - Definitions phase
  - Data exchange phase

➔ Deadlock risk



*One possible order of file operations performed on PISM.  
Different arrow colors represent different file classes and the red region represents the conflict for the output server.*

# Results - Solution for Asynchronous Output Server

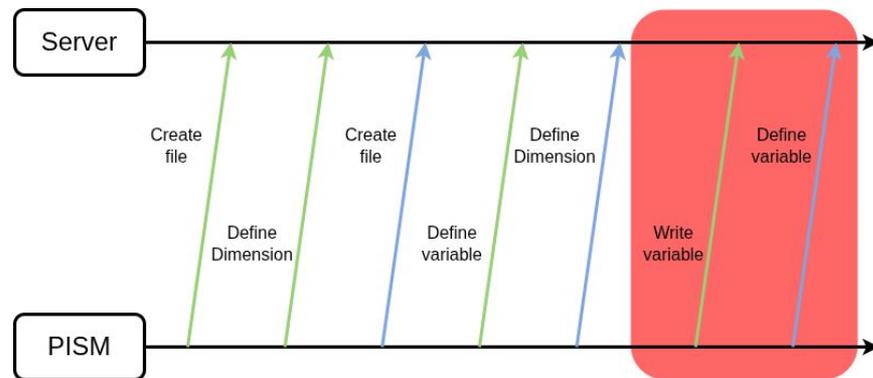
## Key developments:

- ✓ MPI intercommunicator between PISM & server
- ✓ Action-based client-server communication
- ✓ Redesign of PISM output interface
- ✓ Pre-definition of variables to avoid YAC deadlocks

## Result:

- Flexible action-driven server
- Snapshot support fully working

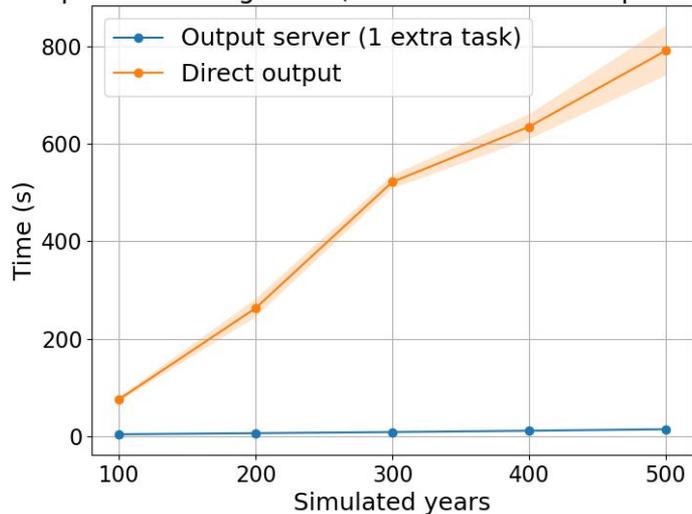
➔ Infrastructure now reusable beyond this sprint.



*One possible order of file operations performed on PISM. Different arrow colors represent different file classes and the red region represents the conflict for the output server.*

# Results - Performance for Asynchronous Output Server

Snapshots Writing Time (8km Grid - 128 Compute Tasks)



*Total snapshot writing time for the original synchronous output and the asynchronous output server with different simulation lengths.*

## Async Output Performance

Experiment - 8km Antarctica (760 x 760):

- Snapshots every 5 simulated years
- Single node + 1 server task
- Averaged over 5 runs

### Result:

- ~98% reduction in snapshot writing time
- Benefit increases with simulation length

# Challenges - PICO ocean module performance

**Technical Challenge:** Original Eikonal solver:

- Approximate discrete algorithm
- Many unnecessary loops & checks
- Inefficient domain traversal

**Implemented:**

- BFS-based algorithm
- Efficient domain updates

**Results:**

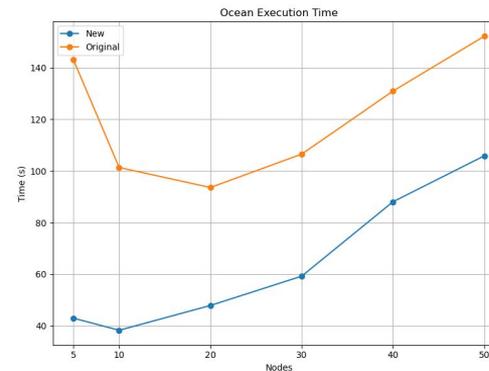
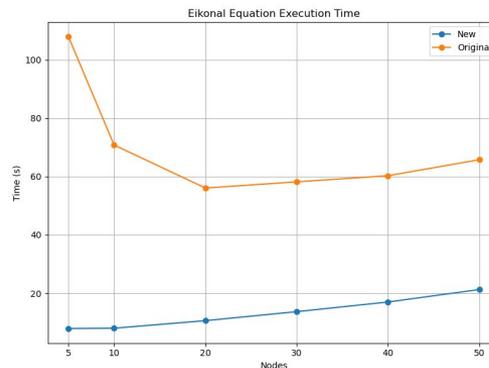
Eikonal solver:

- 80–98% runtime reduction

Whole ocean component:

- 62–69% runtime reduction

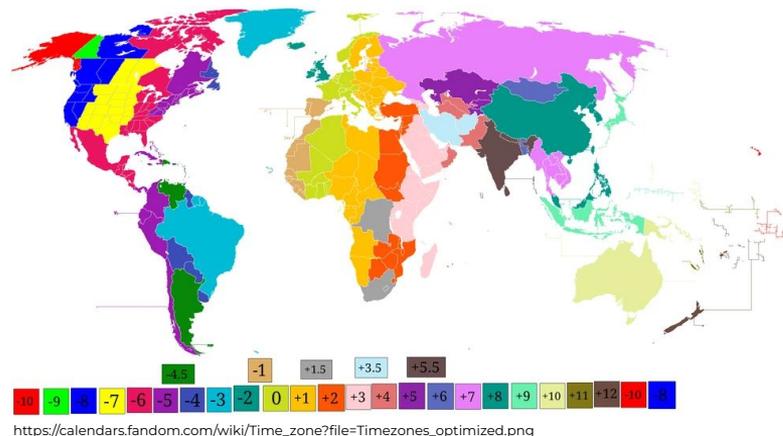
➔ Significant improvement with moderate implementation effort.



*Eikonal equation solution time (top) and total PICO execution time (bottom) for original and newly implemented algorithms.*

# Non-technical Challenges

- Time-zone differences (Germany–Alaska)
- Dependency on external feedback, but limited availability of maintaining institution



## Mitigation:

- definition of multiple, partially independent goals helped to reduce the impact of such interruptions
- RSE can continue work on alternative tasks whenever progress on the primary objective

➔ Important lesson for future collaborations

## Key takeaways

Asynchronous output reduces snapshot writing time by ~98%

Ocean melt module runtime reduced by ~60–70%

Flexible client–server architecture established

Sprint demonstrated importance of communication reliability