natESM

Working-group report

Evaluating ice-sheet-model candidates for natESM

28.01.25

Summary

The working group gathered representatives of Germany's ice-sheet-modeling community to discuss potential ice-sheet-model candidates for natESM. In an online kickoff meeting (10/23) and in a breakout session at the annual community workshop (02/24) we identified key features between the two selected model types. In two sprint checks we could identify main bottlenecks in model performance which will be tackled in a separate sprint. The working-group discussion was also helpful to initiate thinking beyond institutes' and computational constraints to identify the next big scientific questions arising, related to ice-sheet dynamics as part of the Earth system.

General information

Estimated duration:	3 months
Responsible contact:	Torsten Albrecht (MPI-GEA, PIK)
WG-members:	Angelika Humbert (AWI) Thomas Kleiner (AWI) Clemens Schannwell (MPI-M) Moritz Kreuzer (PIK) Ricarda Winkelmann (MPI-GEA, PIK) Martin Rückkamp (BADW) Christian Rodehacke (AWI, DMI) Enrico Degregori (DKRZ) Wilton Loch (DKRZ) Thomas Jung (AWI) Kira Rehfeld (Uni Tübingen) Jochem Marotzke (MPI-M) Iris Ehlert (DKRZ)
KickOff:	25.10.23 (KickOff - online)
WG dissolution	28.02.24 (Community Workshop Leipzig)

Working-group objective

The objective of the working group was to explore candidates for the ice-sheet-model component to be implemented into the natESM system.

The model candidate should encompass the following natESM technical criteria:

1. Well-defined Interfaces between Earth System Components
2. Allows Simulations from Global to Local

3. Exascale-Ready
4. Scalable Workflows
5. Portability
6. Modularity
7. Data Assimilation Capacity
8. Diagnostic Capacity
9. User-Friendly and Well-Documented
10. Traceability, Reproducibility, and Version Control
11. Standardization
12. License of Useful Open-Source Type

Two ice-sheet-model candidates, both originating in the USA, used by and co-developed (in case of PISM, the PIK/MPI-ESM groups are core developers) within the German ice sheet modelers' community have been evaluated:

ISSM (Ice-sheet and Sea-level System Model) https://issm.jpl.nasa.gov

- Used for large-scale polar ice sheets Greenland/Antarctica; also for small-scale glaciers, on centennial time scales
- Scientific questions focus on processes like subglacial hydrology, calving laws, data assimilation ("inversion")
- Capability for sea-level fingerprints
- Works on unstructured grids (refinement close to grounding line)
- Primarily use and development at AWI, further developers and users at TU Darmstadt (TUD), Bayerische Akademie der Wissenschaften (BAdW), DLR
- Code parts are developed, written, and find submission in main branch
- Performance analysis existing, with good scalability up to 512 CPU cores (up to 6000 MPI-processes), with high memory consumption (up to 15GB for one million 3D mesh cells)
- Sprint check by Sergey Sukov and Jörg Benke (JSC), for benchmark of Greenland application with 1.7 mio. 3D cells, with a detailed study of the requirements of the size of RAM and the potential for GPU usage (via PETSc), scaling analysis and profiling/tracing for Levante/JUWELS (see Fig. 1)

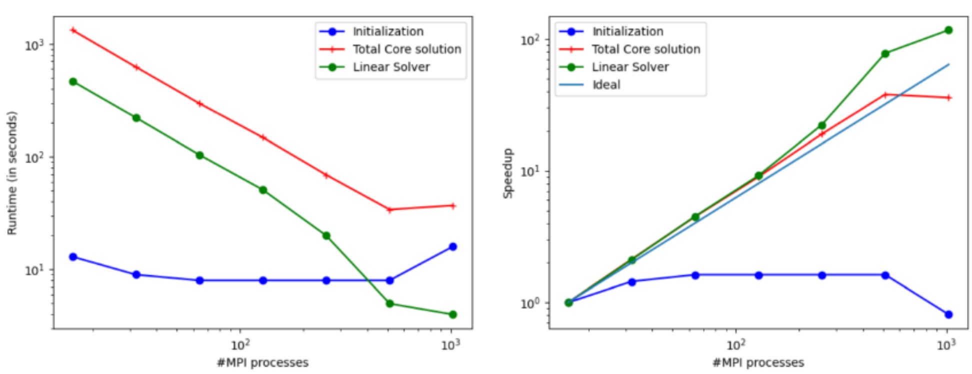

Fig 1: speed-up (a) and runtime (b) for ISSM Greenland application from natESM sprint check.

PISM (Parallel Ice Sheet Model) https://www.pism.io

- Used for large-scale polar ice sheets and regional studies, tipping-point analysis (process understanding), ensembles and sea-level projections (uncertainty quantification), easily coupled to other models, covered period in various studies: deep time (snowball Earth) to long-term future
- Focus both on long-term evolution as well as short-term dynamics (e.g., effects of extreme events, calving), facilitated by adaptive time stepping
- Free running model (dynamic spin-up accounting for memory of the past is key), but also python tools for inversion available
- Capability for global sea-level fingerprints (VILMA coupling, including global deformational, gravitational and rotational effects (GRD))
- Works on structured regular grids, simplifies parallelization for computation efficiency and code readability
- Primarily use and development at PIK (open repository at github.com/pism/pism), in close collaboration with UAF, Alaska
- Broad developer and user base around Germany (user manual and code documentation), also as part of AWI-ESM and MPI-ESM, but also internationally (ModelE, EC-EARTH, MAR coupling, NCAR) with more than 30 institutions (https://www.pism.io/usersmap)
- Performance analysis and scaling tests on various HPC systems (incl. DKRZ, HLRN) exist
- Sprint check by Enrico Degregori and Wilton Loch (DKRZ) for 4km resolution benchmark for Antarctic Ice Sheet (see Fig. 2), identified bottlenecks in MPI-communication and in serial parts of ocean boundary module, with potential speedup by asynchronous I/O

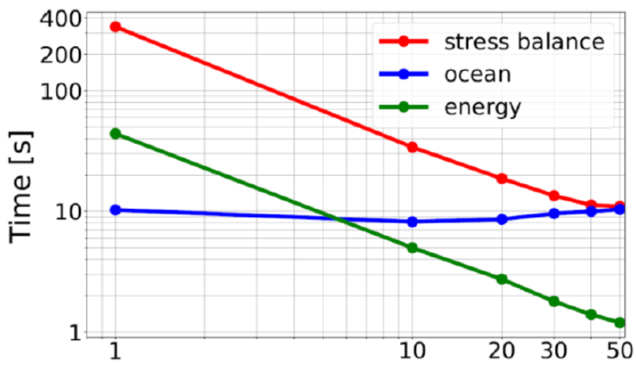

Fig. 2 Scaling analysis for PISM Antarctica application from natESM sprint check, showing elapsed time for used number of CPU nodes.

Outcomes

1. Exascale-readiness

Both model candidates use CPU parallelization. ISSM has investigated scalability in Fischler et al, 2022 (https://doi.org/10.5194/gmd-15-3753-2022) with good scalability up to 512 CPU cores. PISM has shown good scaling for high resolution applications up to several 1000 CPU cores on various HPC systems (see also performance analysis by Bueler et al., 2022, (https://doi.org/10.1017/jog.2022.113). Further optimization could be gained by better domain composition, or implicit time stepping (Bueler & Farrell, 2024, https://doi.org/10.1137/23M1594200). GPU capability may be relevant for better energy efficiency for high-resolution applications, but is tightly linked to the fundamental

PETSc toolkits. PETSc is used by both models and currently adding GPU support for NVIDIA using CUDA, and AMD and Intel using OpenCL/ViennaCL and HIP, so it might be easier in the future. But it would also require rewriting model code.

2. ESM-Coupling
PISM is integrated into several ESMs (e.g. MPI-ESM, AWI-ESM, EC-Earth, NASA/GISS ModelE, POEM, MAR). Open question: If models (atmosphere/ocean) run on GPU, is it then a problem if the ice-sheet model runs on CPU?

3. Stress balance
Mountain glaciers have different requirements to the stress-balance approximation than large ice sheets, due to the different aspect ratio. ISSM provides a variety of dynamics formulations, including a Stokes solver, PISM also provides high-order approximations, but these are rarely used, mainly for mountain glacier or regional studies. Both models use the same enthalpy-based energy model by Aschwanden et al. 2012 (https://doi.org/10.3189/2012JoG11J088).

4. Resolution
A high resolution of ice-sheet models becomes particularly important at margins, grounding lines, and mountain regions. ISSM can adjust grid resolution in key regions, PISM can use high resolution in the entire domain (order of 1km). But there are also strong biases (uncertainties) in forcing data coming from ESM to be considered. Both, ISSM and PISM can use sub-grid schemes for basal friction and ice shelf basal melt, when the grounding line motion is not sufficiently resolved. They differ in the way how calving-front motion is parameterized: ISSM used a level set method while PISM used subgrid-scale front motion (Albrecht et al., 2011, https://doi.org/10.5194/tc-5-35-2011).

5. Communication
A natESM Mattermost-Channel for the land-ice-component working group was established (https://mattermost.mpimet.mpg.de/natesm/channels/natesm-working-group-land-ice-component).

6. Conclusions
Both model candidates fulfill most of the technical criteria listed above. ISSM has a strong focus on scientific questions related to processes at the ice-sheet boundary, such as calving or subglacial hydrology, and uses assimilation to observational data in key regions. PISM has been used for process-understanding as well as in simulations both on longer timescales (e.g., long-term sea-level projections and commitment, tipping dynamics, glacial cycles) as well as shorter timescales (e.g., effects of extreme events, calving); because of its computational efficiency, PISM can be easily coupled and used to investigate stability aspects and explore multiple uncertainties by running systematic parameter ensembles. For projections, the historical spin-up can be highly relevant, which is the usual practice for a PISM spin-up procedure. In the context of the interactive coupling to the global atmosphere and ocean in natESM, new applications and challenging scientific questions arise and should be discussed within the community.

The implementation of a YAC-based output server that allows for asynchronous output will significantly improve PISM's performance for higher resolution applications. This bottleneck in Input/Output has been identified during the sprint check and will be tackled in a sprint which has been accepted for 6 months starting in April 2025 with RSE Wilton Loch.