

# Veris: Fast & Efficient Sea-Ice Modeling in Python with GPU Acceleration

Jan P. Gärtner<sup>1</sup>, Martin Losch<sup>1</sup>, Suvarchal K. Cheedela<sup>1</sup>, Markus Jochum<sup>2</sup>, and Roman Nuterman<sup>2</sup>

<sup>1</sup>Alfred-Wegener-Institut, Helmholtz Zentrum für Polar- und Meeresforschung, Bremerhaven, Germany

<sup>2</sup>Niels Bohr Institute, University of Copenhagen, Copenhagen, Denmark

**Correspondence:** Jan P. Gärtner (jan.gaertner@awi.de)

**Abstract.** Climate models have traditionally been developed in Fortran, due to its long-standing use in scientific computing and its excellent computational performance. While Python offers substantial advantages in terms of code readability, maintainability, and the availability of libraries and tools, the performance gap between Python and Fortran has historically limited Python's use in large-scale climate modeling. This performance gap can be mitigated by using the JAX library, which significantly improves execution speed of Python code.

We use JAX as a backend for Veris, a new sea ice model implemented in Python. Veris builds upon the Fortran-based sea ice component of the general circulation model MITgcm. Benchmark experiments show that Veris exhibits scaling behavior with increasing process counts comparable to the Fortran reference implementation. For small CPU process counts, Veris outperforms the MITgcm, showing the great potential that JAX has for climate modeling, particularly as further improvements in inter-process communication are anticipated. When executed on a high-end GPU, a single-process Veris simulation matches the performance of the parallelized Fortran reference running on hundreds of CPU cores, but at a fraction of the energy cost. These results demonstrate the potential of Veris for large-scale HPC-based simulations. Furthermore, Veris can be coupled with the Python-based ocean model Veros to form a fully Python-based coupled sea ice-ocean model.

## Plain Language Summary

Climate simulations help us understand the Earth system and its evolution. The models used to perform these simulations are highly complex, require significant programming expertise to build and consume a lot of energy. A key component of climate models is their sea ice components. In this work, we present a sea ice model that offers an easier development process while maintaining strong performance. By using GPUs, our model can replace hundreds of CPUs, significantly reducing its energy usage.

## 1 Introduction

Climate models contain hundreds of thousands of lines of code. The programming language used in these models is typically Fortran, largely due to its legacy in scientific computing. Fortran's strong backwards compatibility allows extensive reuse of existing code in newer models written in updated versions of the language, saving significant development time compared

to rewriting legacy code in a different language. Another critical reason for Fortran's prevalence is its high computational efficiency, as it is optimized for the speed of numerical calculations (Méndez et al., 2014). Climate models require the execution of an extensive number of floating-point operations over many time steps and across large spatial grids, making computational performance particularly essential for this field.

Despite these advantages, Fortran is relatively unpopular compared to modern languages like Python, especially among scientists outside climate modeling. This can create a barrier for new researchers, who have to invest substantial time in learning an unfamiliar programming language. Additionally, using a more popular language such as Python would enable access to more libraries and tools. Fortran, though well-suited for high-performance computing (HPC) due to its optimization for numerical and array-based calculations, can be more challenging to work with than higher-level languages. Using a high-level language such as Python would not only facilitate structural modifications, such as implementing new parameterizations and simplifying debugging, but would also make HPC workflows more accessible to researchers who may lack extensive experience in low-level programming.

Although Python's popularity and readability offer advantages, these factors alone do not justify switching from Fortran to Python due to the significant performance gap between the two languages without optimizations or extensions. To overcome this gap, the Python library JAX can be used to significantly increase Python's performance for numerical modeling (Bradbury et al., 2018). JAX is a library for array-based computing with a syntax very similar to NumPy's (Harris et al., 2020), allowing NumPy code to be adapted to JAX without structural changes. JAX includes a set of function transformations like vectorization, parallelization, automatic differentiation and Just-In-Time (JIT) compilation. The JIT compilation significantly reduces the runtime of iterative calculations, making JAX well-suited for climate models as they iterate over many time steps.

In addition to these transformations, a key advantage of JAX is that it is agnostic to the type of accelerator, enabling the same code to run on both CPUs and GPUs. Due to their higher core counts, GPUs are much better suited for parallel computing than CPUs, provided that the available computational and memory resources are fully utilized. As a result, they are more energy-efficient for parallel tasks such as climate simulations, as they complete these simulations faster than CPUs. By leveraging GPUs, JAX not only enhances computational speed but also reduces the power consumption and therefore the carbon footprint of running climate models (e.g. Marquetti et al., 2018; Portegies Zwart, 2020; Li et al., 2023).

As a recent example, the Python based Versatile Ocean Simulator (Veros, Häfner et al., 2018, <https://veros.readthedocs.io/en/latest/>) is a translation of the Fortran backend of the ocean model pyOM2 (v2.1.0, Eden and Olbers, 2014; Eden, 2016). By using the Just-In-Time compilation of JAX, the performance gap to the original Fortran model was closed (Häfner et al., 2021). When using JAX as the backend, Veros can be run on both CPUs and GPUs. With a single GPU, dozens to hundreds of CPU cores can be replaced which substantially reduces the energy consumption of running Veros.

To date, Veros does not contain a sea ice model, limiting its applicability in high-latitude regions where sea ice significantly influences heat, momentum, and moisture exchanges between the ocean and atmosphere, thereby impacting the local climate. This work introduces Veris (Versatile Sea Ice Simulator), a Python-based sea ice model originally developed to be coupled to Veros to extend its range of applications. Like Veros, Veris supports both NumPy and JAX backends. When used with JAX, Veris can also run on GPU architectures, providing substantial computational advantages. In order to parallelize Veris,

the model infrastructure was subsequently redesigned, resulting in a separate, updated version of Veris that is able to leverage  
60 JAX’s parallelization capabilities on both CPU- and GPU-based systems. Owing to these architectural changes, this parallelized  
version is no longer directly compatible with Veros, but instead represents an independent, high-performance sea ice modeling  
framework optimized for large-scale simulations.

To solve the momentum equation, Veris employs the modified and adaptive Elastic-Viscous-Plastic (EVP) solver (Hunke and  
Dukowicz, 1997; Bouillon et al., 2013; Kimmritz et al., 2016). Benchmark experiments by Rasmussen et al. (2024) assessed  
65 a refactored implementation of the standard EVP solver used in the Fortran-based sea ice model CICE (Hunke et al., 2024),  
optimized for parallel execution on both CPU- and GPU-based architectures. In these experiments, execution on a single GPU  
was found to be 1.2 times faster than on 112 CPU cores. Accordingly, a significant performance increase is anticipated for  
Veris when comparing GPU-based execution to CPU-based configurations.

The properties of Veris are described in section 2. Section 3 presents benchmark simulations to verify that Veris accurately  
70 simulates sea ice dynamics, as well as performance benchmarks comparing Veris to the Fortran reference.

## 2 Model Description

### 2.1 Basic Properties

Veris is a Fortran to Python translation of part of the sea ice component of the Massachusetts Institute of Technology general  
circulation model (MITgcm, Marshall et al., 1997; Losch et al., 2010, <https://github.com/MITgcm/MITgcm>, <https://mitgcm.org/documentation/>), originally based on the formulation by Hibler (1979). The core properties of Veris are as follows:

- The variables are stored on a two-dimensional, staggered C-grid (Arakawa and Lamb, 1977). Tracer variables, such as  
ice thickness and concentration, are stored at the centers of grid cells, while the  $u$ - and  $v$ -components of velocity are  
stored at the respective  $u$ - and  $v$ -grid points.
- Ice grows thermodynamically when the ocean loses heat to the atmosphere through a conductive heat flux across the  
80 ice. This heat loss is balanced by freezing at the ice bottom. Additionally, new ice may form in regions of open water  
when the ocean is losing heat. For the growth rate calculation, a two level thickness configuration is assumed (Hibler,  
1979). Each grid cell is assumed to contain ice of thickness  $h_{actual}$  and open water. The proportion of the grid cell that  
is covered with ice is defined as the ice concentration  $A$ . The actual ice thickness  $h_{actual}$  is derived from the mean ice  
thickness per grid cell  $h$  according to  $h_{actual} = \frac{h}{A}$ . The growth of the mean thickness over a single time step is

$$85 \quad \Delta h = \Delta t \left( f(h_{actual})A + f(0)(1 - A) + f_{ml} \right) \quad (1)$$

where  $f(h)$  is the growth rate of ice of thickness  $h$ , and  $f_{ml}$  accounts for ice melting at the base due the mixed layer  
temperature.

- The ice has no heat capacity, and the adjustment of the ice surface temperature to changes in the surrounding temperature  
occurs instantaneously (zero-layer thermodynamics, Semtner, 1976).

- 90 – Snow can accumulate on the ice cover under conditions of snowfall or freezing rain.
- Additional ice growth can occur if the weight of accumulated snow submerges the ice. When the ice-snow interface lies below the water surface, a simple mass conserving parameterization of snow-ice formation turns snow into ice until the ice surface is back at the water surface (Leppäranta, 1983).
- 95 – A viscous-plastic rheology (Hibler, 1979) is implemented using the explicit solution methods of the modified elastic-viscous-plastic solver (EVP\*, revisited EVP or mEVP, Lemieux et al., 2012; Bouillon et al., 2013; Kimmritz et al., 2015) and the adaptive elastic-viscous-plastic solver (aEVP, Kimmritz et al., 2016). For small internal ice stresses, the ice flow is viscous and the internal stress increases linearly with the strain rate. Once a certain stress threshold is reached, the ice deforms plastically, such that the internal stress does not exceed the yield criterion. Both the mEVP solver and the aEVP solver performing a specified number of subcycling iterations to update the ice stress and velocity fields at each time
- 100 step 2.3). In the aEVP solver, the relaxation parameters for the subcycling process are computed based on the grid cell size and the viscosities of ice, resulting in a faster convergence in most grid points.
- In addition to the EVP solver, a freedrift solver is implemented, which does not account for ice stress, allowing the ice to drift independently of the presence of ice in adjacent grid cells.
- Landfast ice is sea ice that remains immobile or nearly immobile and attached to a coast for a certain time period.
- 105 The ice cover may be fixed to the seafloor by down-reaching ice keels or laterally anchored to a coast. The formation of landfast ice due to the ice being anchored to the seafloor in shallow waters is parameterized using a basal drag coefficient (Lemieux et al., 2015). For landfast ice formed due to the ice being fixed to a nearby coast, a coastal drag parameterization (Liu et al., 2022) is implemented. This coastal drag is applied in conjunction with a free-slip boundary condition. Additionally, a no-slip boundary condition is also implemented.
- 110 In the Veros-compatible version of Veris, it is possible to switch between NumPy and JAX as the backend, which facilitates the development and implementation of new features. Debugging within JAX can be challenging due to its use of just-in-time compilation, which restricts access to intermediate outputs in compiled functions. By initially debugging with NumPy — where intermediate outputs and clearer error messages are available — developers can efficiently implement and debug new code before switching to JAX to take advantage of its optimized performance.

## 115 **2.2 Code Verification**

To verify the correctness of the translation from Fortran to Python, we conducted a comparative analysis of model outputs from Veris and the MITgcm. The thermodynamic and dynamic components were evaluated separately. For the thermodynamic component, a one-dimensional benchmark was employed, simulating the temporal evolution of sea ice thickness, snow thickness, and ice concentration within a single grid cell. This setup is sufficient for thermodynamic verification because thermodynamic

120 growth and melt processes depend exclusively on local state variables and atmospheric and oceanic forcing, without lateral coupling to neighboring grid cells.

The thermodynamic benchmarks covered a range of initial conditions for ice thickness, snow thickness, and ice concentration, as well as variations in atmospheric and oceanic forcing, including air temperature, ocean surface temperature, precipitation, and shortwave and longwave radiation. Simulations were performed for 365 time steps with a daily time step. Differences between Veris and the MITgcm remained on the order of  $\mathcal{O}(10^{-7})$  for all prognostic variables and exhibited random, non-accumulative behavior (Fig. A1). Small differences are expected in code porting due to a variety of reasons, including differences in compiler-dependent ordering of operations, intrinsic library implementations, and floating-point arithmetic. The bounded and non-growing nature of the differences indicates that the translated thermodynamic formulation represents a stable, non-divergent system. The deviations are therefore attributed to numerical precision limitations (Rosinski and Williamson, 1997).

For the dynamic component, a two-dimensional benchmark following Mehlmann et al. (2021) was conducted on rectangular grids with resolutions of  $128 \times 128$ ,  $256 \times 256$ ,  $512 \times 512$ , and  $1024 \times 1024$  grid cells. As in the thermodynamic tests, varying initial conditions for ice thickness, snow thickness, and ice concentration were employed, along with variations in atmospheric and oceanic forcing. In contrast to the thermodynamic component, the sea ice dynamics solver is subject to stricter stability and accuracy constraints, requiring shorter time steps. Accordingly, a time step of 10 minutes was used for 1000 time steps.

The dynamical evolution of the sea ice field is substantially more sensitive to small perturbations than the thermodynamic evolution. To establish a reference for acceptable divergence, the MITgcm benchmark was repeated with perturbations on the order of  $\mathcal{O}(10^{-15})$  added to the initial conditions and forcing fields. This procedure follows Rosinski and Williamson (1997), who propose that successful code porting is achieved when the divergence between the original and translated solutions is not exceeding the growth of an initial perturbation introduced into the lowest-order bits of the original code solution. Both Veris and the MITgcm use double precision (float64) arithmetic. Even perturbations of this magnitude led to differences on the order of  $\mathcal{O}(10^{-2})$  in ice speed and  $\mathcal{O}(10^{-1})$  in ice concentration. These deviations did not accumulate further over time, displayed noise-like patterns, and were constrained primarily to regions with strong ice concentration gradients.

The differences between Veris and the MITgcm are of the same or smaller magnitude than those obtained in the perturbed MITgcm reference simulations (Fig. A2), indicating successful code porting and variations within acceptable numerical limits. In addition to the 1000-time step benchmarks, a longer integration of 3 months (12960 time steps) was performed to assess Veris’s long-term behavior (Rosinski and Williamson, 1997), which remained consistent with the MITgcm.

### 2.3 Parallelization

In the initial implementation of Veris, developed as a sea ice plug-in for Veros, the EVP solver cannot be executed in parallel when using JAX as the backend. This limitation arises because the EVP solver requires halo exchanges between neighboring partitioning regions within the solver iteration. In the original implementation, these exchanges are performed inside the EVP loop body via a `fill_overlap` routine:

```
for i in range(nEVPsteps):  
    sigma = calculate_stress(...)
```

```
155     uIce, vIce = calculate_velocity(...)

    uIce, vIce = fill_overlap(uIce, vIce)
```

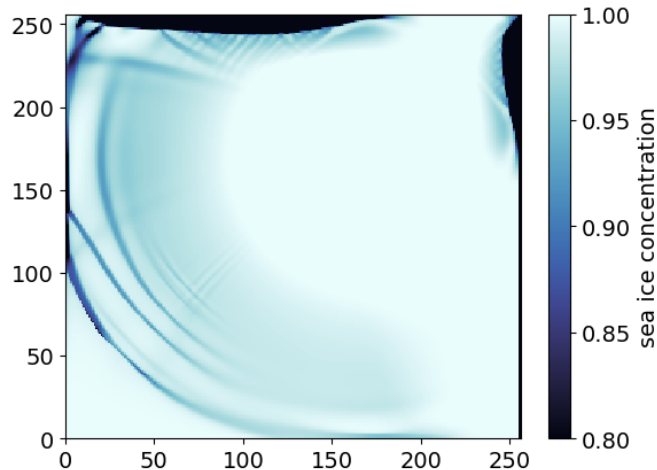
When NumPy is used as the backend, the EVP solver is implemented with a standard Python for-loop, allowing Veris to rely on the existing MPI-based communication infrastructure provided by Veros. When JAX is used, however, iterative loops are implemented using the `jax.lax.fori_loop`, which compiles the loop body without unrolling every iteration during compilation. This compilation model is not compatible with the communication mechanism employed in Veros. The communication mechanism of Veros enables MPI-based halo exchanges through a decorator that contains both the JIT compilation and the MPI communication. This decorator can only wrap functions that are not automatically compiled by JAX, and since the EVP loop is automatically compiled, this decorator-based approach cannot be applied, and the `fill_overlap` routine cannot be invoked inside the compiled loop body. Enabling MPI-based parallelism in this context would therefore require a custom Veros-compatible mechanism that supports communication inside a compiled loop while ensuring the loop body is traced and compiled only once. The development of such an interface is beyond the scope of this work.

To enable parallel execution, we instead adopted JAX’s native sharded array framework. Sharded arrays encode the distribution of array data across devices and provide point-to-point communication primitives that can be invoked within compiled functions. The model infrastructure of Veris was restructured around sharded arrays, while preserving the existing formulation of sea ice physics. This redesign resulted in a standalone, fully JAX-based version of Veris that supports scalable parallel execution on both CPU- and GPU-based architectures. All performance benchmarks presented in section 3 were conducted with this updated version of Veris, except for the Numpy-based simulations.

### 3 Results

#### 175 3.1 Simulated Sea Ice Field

To demonstrate that Veris accurately simulates sea ice dynamics and its rheology, a simulation was conducted using a forcing that is designed to induce fracturing of the sea ice cover (Fig. 1). The forcing fields include a circular anti-cyclonic ocean surface current, which grows stronger towards the boundaries, and a circular cyclonic wind field with some horizontal convergence, displaced from the grid center and is strongest in its center. These fields are taken from the sea ice benchmark experiment described by Mehlmann et al. (2021), with the exception that the wind field is stationary in this work. The quadratic model domain is enclosed by land, resulting in boundary conditions of zero velocity. The simulation was run with a simulation time of 10 days with a time step of 10 minutes. The initial conditions included a constant ice thickness of  $h = 0.3$  m and an ice concentration of  $A = 1$ . In the sea ice field, narrow features in the sea ice concentration are observed, resulting from inhomogeneities in the deformation field caused by the internal ice stress responding to the forcing fields. These linear kinematic features are discussed in Mehlmann et al. (2021) and show that Veris’s dynamics are on par with state-of-the-art sea-ice dynamics codes.



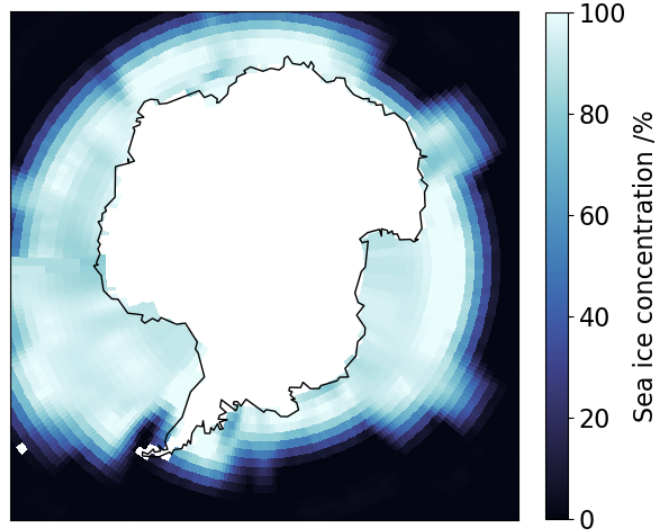
**Figure 1.** Simulated sea ice concentration generated by Veris after 1440 time steps of 10 minutes (equivalent to 10 days), driven by circular ocean surface currents and wind fields. The initial conditions included a uniform sea ice concentration of 1 and a sea ice thickness of 0.3 m. A reference simulation with the MITgcm, employing the same forcing, produced the same features (Fig. A2).

Veris is also configured for integration with Veros to form a coupled sea ice-ocean model. Veros is run using setup files that define the configuration of the model, including specific extensions and the spatial discretization. Currently, Veris is integrated into the setup of a global simulation with a spatial resolution of  $4^\circ$ . This serves as a proof of concept, demonstrating that Veris can be successfully coupled with Veros to form a coupled sea ice-ocean model. This integration enables comprehensive simulations of sea ice dynamics within an oceanic framework. The results of a simulation using this setup are shown in Fig. 2, which displays a snapshot of simulated Antarctic sea ice concentration on July 1, based on a model run that included zero ice thickness and concentration on January 1 as initial conditions. More realistic simulations require adequate configurations of Veros, for example with grids covering the Arctic Ocean. These configurations are not yet available and creating new setups for Veros is beyond the scope of this manuscript. In the following we concentrate on the computational performance of Veris.

### 195 3.2 Benchmark Results

To evaluate the performance of Veris, a series of idealized and scalable benchmark experiments were conducted. These experiments were performed on a rectangular grid with varying process count and problem size, utilizing the scaled forcing fields described in section 3.1. The benchmarks serve two primary purposes. First, they compare the performance of Veris to the Fortran reference, assessing the impact of transitioning from Fortran to Python. Second, the benchmarks analyze the scaling behavior of Veris with increasing process count and domain size. Additionally, Veris running on a GPU is compared to the Fortran reference across various parallel CPU configurations, focusing on both time to solution and total power consumption.

During the first time step, the compilation process with JAX’s JIT compiler introduces significant overhead. To properly evaluate the model performance, the first time step is therefore discarded.



**Figure 2.** Simulated sea ice concentration field in the Antarctica generated by the Veris-Veros coupled model on July 1 in a one-year simulation using the ECMWF ERA5 reanalysis data from 1986 as forcing fields (Hersbach et al., 2020). The model was initialized with zero sea ice thickness and concentration, and the results were interpolated onto a grid with a spatial resolution of  $1^\circ$ .

The reference runs were performed with the MITgcm configured to include only the dynamical component, as the thermodynamical calculations are computationally less expensive and negligible in comparison to the dynamics. Veris was similarly run with just the dynamical component enabled. The models are compared using the aEVP solver with 120 subcycling iterations.

The benchmarks are run on two different nodes of the HPC Levante, a supercomputer of the German Climate Computing Center (DKRZ). The specifications of the different node types are listed in Tab. 1.

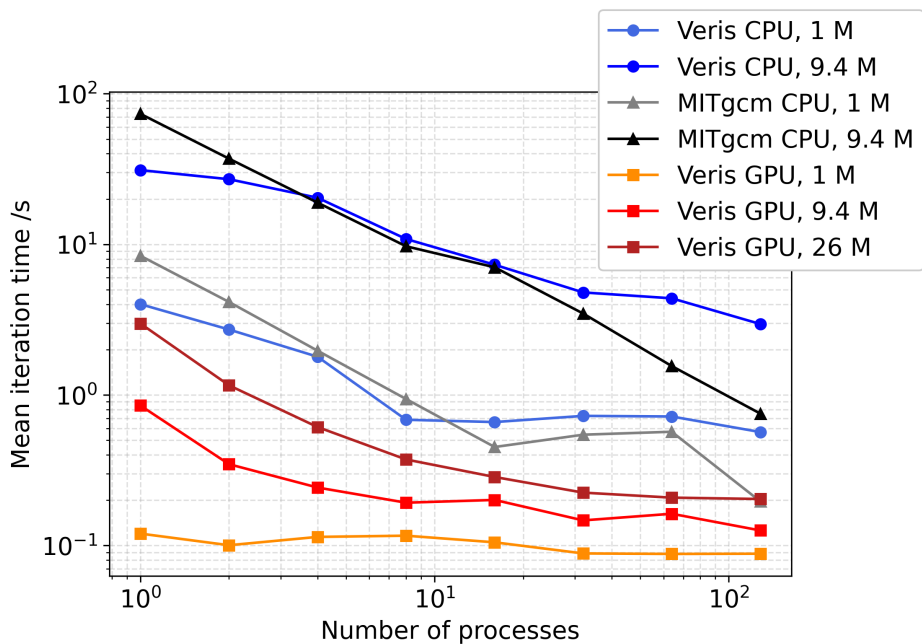
	CPU node	GPU node
CPU	2x AMD 7763 CPU (128 cores in total)	2x AMD 7763 CPU (128 cores in total)
GPU	-	4x NVIDIA A100 (80 GB memory)
Memory	512 GB	512 GB
Operating system	Red Hat Enterprise Linux 8.9	Red Hat Enterprise Linux 8.9
Software	OpenMPI 4.1.2, jaxlib 0.4.30	CUDA 13.0, jaxlib 0.4.30

**Table 1.** Specifications of the benchmark platforms.

### 3.2.1 Scaling With Number of Processes

This section analyzes the performance scaling with respect to the number of processes (Fig. 3). For the CPU-based simulations, grids with approximately  $10^6$  and  $9.4 \times 10^6$  grid cells were evaluated, hereafter referred to as the 1 M and 9.4 M grids. The 1 M

grid experiments were conducted on a single node. Both Veris and the MITgcm exhibit a similar initial increase in performance with rising core counts, a plateau in performance between 8–64 and 16–64 processes, respectively, and a renewed increase at 128 processes. For Veris, increasing the process count requires additional memory due to JAX’s compilation cache, larger communication buffers, and the overhead associated with Python objects and metadata. For the 9.4 M grid, this necessitates distributing the simulation across multiple nodes, with at most 16 processes per node. The MITgcm is able to run the 9.4 M grid on a single node; however, this configuration results in a substantial performance degradation and poorer scaling behavior. When the simulation is distributed across multiple nodes, the MITgcm outperforms Veris, achieving approximately a fourfold speedup over Veris at 128 processes (Fig. 3). When constrained to a single node, the MITgcm exhibits both worse scaling and lower absolute performance than Veris running across multiple nodes.



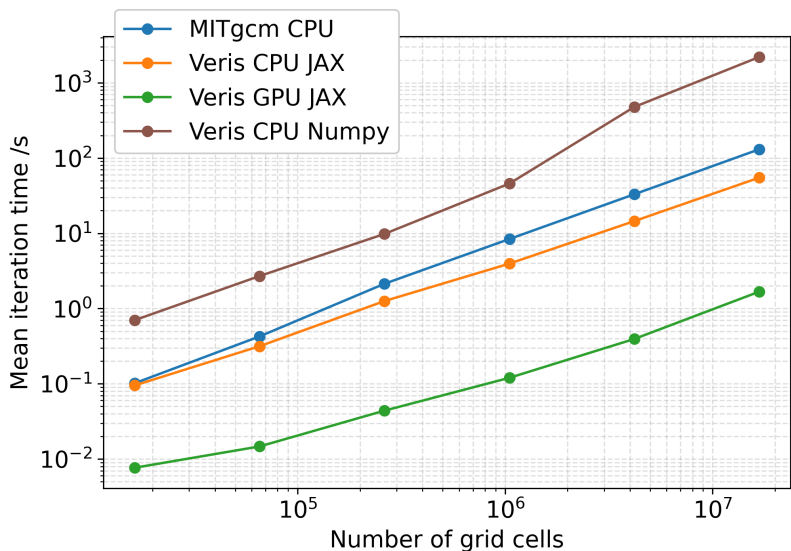
**Figure 3.** Scaling of mean iteration time with process count for Veris and the MITgcm sea ice component for varying problem sizes (1 M =  $10^6$  grid cells), calculated over 20 iterations steps.

The GPU-based simulations substantially outperform the CPU-based simulations, with a single GPU process matching the performance of the MITgcm running on hundreds of CPU cores. Increasing the number of processes introduces additional communication overhead, which is more pronounced for GPUs than for CPUs because GPUs do not share memory and therefore require explicit data movement between devices for communication. As a result, performance scaling on GPUs is strongly dependent on problem size. For the 1 M grid, increasing the GPU count from 1 to 128 yields only marginal performance gains. The 9.4 M grid exhibits improved scaling but still approaches saturation. Among the tested configurations, the 26 M grid, which contains approximately  $26 \times 10^6$  grid cells, shows the best scaling behavior, with a consistent performance increase as

the GPU count increases. However, even for this grid, only limited performance gains are observed for higher GPU counts, indicating saturation. Overall, these results demonstrate the substantial performance benefits of GPU-based execution compared to CPU-based systems. They further suggest that GPU parallelization becomes increasingly effective for larger problem sizes, highlighting the potential of Veris for very large-scale simulations on GPU-accelerated architectures.

### 3.2.2 Scaling With Domain Size

In addition to the scaling with respect to process count, we performed benchmarks to analyze the dependence of performance on domain size for single-process simulations (Fig. 4). Veris exhibits scaling behavior consistent with that of the Fortran reference implementation. Scaling in this section refers to the increase in time to solution with domain size. When using Numpy as the backend, Veris is approximately five times slower than the MITgcm for domain sizes  $10^6$  or smaller and around 15 times slower on domain sizes  $10^7$  or larger. With the JAX backend, Veris consistently outperforms the Fortran reference, with mean iteration times approximately a factor of two smaller for grids larger than  $10^6$  cells.



**Figure 4.** Benchmark results comparing the scaling of mean iteration time with grid size for Veris and the MITgcm, calculated over 20 iteration steps.

Simulations executed on a single GPU demonstrate both a substantial absolute performance increase and good scaling with increasing grid size. These results highlight the effectiveness of GPU-based execution for large domains and further emphasize the suitability of the JAX-based implementation for large-scale, high-resolution sea ice simulations.

### 3.2.3 Power Consumption

In addition to time to solution, energy consumption is a key metric in high-performance computing. Reducing energy usage not only lowers operational costs but also mitigates the carbon footprint of large-scale simulations. One promising approach for improving energy efficiency is to replace large numbers of CPU cores with a smaller number of GPUs, which can provide substantially higher performance per watt.

For the power consumption benchmark, we selected the 9.4 M grid configuration and measured the energy usage over 1000 iterations. The reference simulation was performed using a single GPU. For the MITgcm, CPU-based simulations were conducted that matched the time to solution of the single-GPU configuration, using both the minimum number of nodes and the minimum number of CPU cores (Tab. 2). In addition, a simulation utilizing all 4 GPUs available on a single GPU node was performed. The time to solution of the single-GPU run was matched by a configuration using 640 CPU cores across 5 nodes and by a configuration using 128 CPU cores across 8 nodes. Both CPU-based simulations exhibited similar total energy consumption, requiring approximately 4–4.5 times the energy of the single-GPU simulation. When all four GPUs on a GPU node were utilized, the total energy consumption was reduced by a factor of two relative to the single-GPU simulation. These results demonstrate the substantial energy-efficiency advantages of GPU-based execution for large-scale sea ice simulations.

	Time to solution	Mean power usage	Total energy usage
1 GPU	854 s	650 W	0.15 kWh
4 GPUs	251 s	1105 W	0.077 kWh
640 CPUs on 5 nodes	846 s	2845 W	0.67 kWh
128 CPUs on 8 nodes	742 s	2960 W	0.61 kWh

**Table 2.** Power consumption for Veris running on 1 and 4 GPUs (on the same GPU node) and the MITgcm sea ice component running in various CPU-based configurations. CPU here refers to a single CPU-core process. This simulation is using  $9.4 \times 10^6$  grid cells and 1000 iterations. The power usage is reported by Levante’s built-in power tracking system.

## 4 Discussion

The performance benchmarks of both the MITgcm and Veris showed a strong dependence of model performance on the underlying HPC configuration. Identical numbers of CPU-cores can yield substantially different solution times, depending on how processes are distributed across the nodes and mapped to cores within a node. This shows that optimal performance requires not only efficient model implementations, but also careful tuning of the parallel configuration to the HPC system. Such tuning involves the choice of processor grids, node layouts, and process placement strategies, and presupposes a detailed understanding of the HPC system.

For established Fortran-based model such as the MITgcm, MPI-based communication has been extensively optimized over years of development and use. In contrast, distributed execution in JAX is an active area of development, and there is con-

265 siderable potential for further improvements in communication backends and process orchestration. The results of this study  
therefore highlight both the promise of JAX for climate modeling in HPC environments and the importance of systematic  
communication analysis and optimization. The strong impact of communication overhead is also reflected in the early saturation  
observed in GPU-based simulations for smaller problem sizes. Given that Veris consistently outperforms the MITgcm at  
small process counts, it is plausible that further advances in JAX communication could enable Veris to outperform its Fortran  
270 reference in strongly parallel regimes as well.

The benchmarks further reveal a significant performance benefit when transitioning from CPU-only systems to GPU-  
accelerated architectures, underscoring the growing importance of GPUs for climate modeling. This is particularly advan-  
tageous in the context of JAX, which is accelerator-agnostic and allows newly developed model components to run on GPUs  
without additional code modifications.

275 The GPU-based simulations exhibit the expected scaling behavior, with improved efficiency for larger domains, consistent  
with the high communication overhead associated with multi-GPU execution. This emphasizes the suitability of multi-GPU  
configurations for high-resolution modeling. Moreover, the power consumption measurements demonstrate that GPU-based  
Veris simulations can achieve comparable or superior performance at a fraction of the energy cost of CPU-based MITgcm  
runs, showing the potential of GPU-accelerated Python models to reduce both the operational cost and the carbon footprint of  
280 climate simulations.

The sensitivity of model performance to the chosen HPC configuration is less pronounced for GPU-based simulations than  
for CPU-based simulations. This is partly attributable to the fact that JAX is currently better optimized for GPU execution than  
for distributed CPU-based simulations. These results suggest that JAX holds further promise for CPU-based climate modeling,  
given the ongoing development of JAX and its distributed runtime, which may lead to further performance improvements in  
285 the future.

## 5 Summary and Conclusion

By translating the sea ice component of the MITgcm into Python, we developed Veris, a fully Python-based sea ice model.  
Veris can be coupled to the Python-based ocean model Veros, forming a coupled ocean-sea ice model implemented entirely in  
Python. This version supports parallel execution with Numpy and single-process execution with JAX on both CPU and GPU ar-  
290 chitectures. Although this version does not support distributed parallelism for JAX-based simulations, the benchmarks already  
demonstrated the potential JAX has for large-scale simulations, given the substantial performance gains of GPU execution.

To enable scalable parallelism, the Veris infrastructure was subsequently redesigned around JAX's sharded array framework,  
allowing distributed execution on both CPU- and GPU-based systems. The benchmarks presented in this study demonstrate the  
great potential of JAX for climate modeling. At small process counts, Veris outperforms the MITgcm. At larger process counts,  
295 communication overhead becomes the dominant performance factor, and the MITgcm runs faster than Veris. This sensitivity  
on the communication framework was also evident in the dependence of performance on specific HPC configurations. Whereas  
MPI-based communication in the MITgcm benefits from long-standing optimizations, communication in JAX-based distributed

systems is still evolving. The superior performance of Veris at low process counts indicates strong potential for large-scale applications, particularly as JAX's distributed runtime continues to be developed and optimized.

300 GPU execution yields particularly strong performance benefits. A single-GPU run of Veris matches the performance of the MITgcm running on hundreds of CPU cores across multiple nodes, while consuming substantially less energy. These results demonstrate not only the competitiveness of Python-based, JAX-accelerated models in terms of time to solution, but also their potential to significantly improve the energy efficiency of climate simulations.

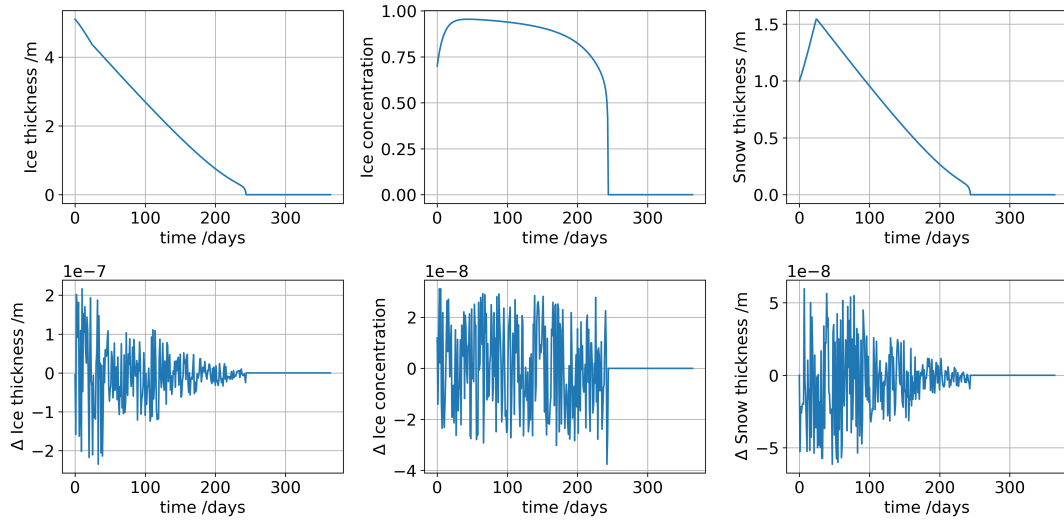
In summary, Veris demonstrates that Python, when combined with JAX, can support high-performance, scalable, and energy-  
305 efficient sea ice modeling. These findings suggest that JAX-based climate models are a promising pathway towards modern HPC workflows, with the prospect of further performance gains as communication backends and distributed execution frameworks continue to improve.

*Code availability.* The initial version of Veris is available under <https://github.com/team-ocean/veris>. It can be used as a plug-in for Veros, available under <https://github.com/team-ocean/veros>, or run as a standalone model using the scripts available under [https://github.com/jpgaertner/veris\\_minimum\\_working\\_example](https://github.com/jpgaertner/veris_minimum_working_example). The updated version of Veris for parallel execution using JAX is available under <https://github.com/jpgaertner/veris/tree/jax-only>. The scripts for running Veris in this configuration and for creating the benchmarks presented in this study are available under [https://github.com/jpgaertner/veris\\_minimum\\_working\\_example/tree/jax\\_halo\\_exchange](https://github.com/jpgaertner/veris_minimum_working_example/tree/jax_halo_exchange).  
310

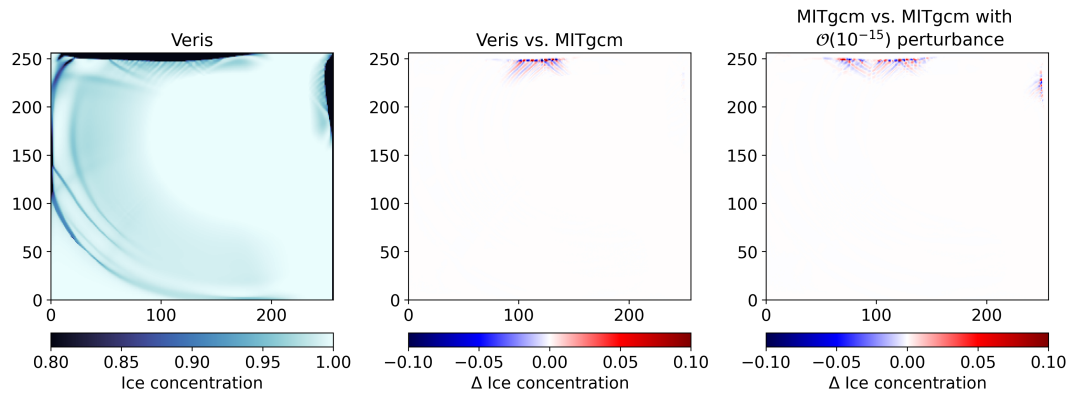
The MITgcm is available under <https://github.com/MITgcm/MITgcm.git>. The configurations used to create the thermodynamic and dynamic benchmarks are available under [https://github.com/mjlosch/MITgcm/tree/seaice\\_test\\_1d](https://github.com/mjlosch/MITgcm/tree/seaice_test_1d) and [https://github.com/mjlosch/seaice\\_benchmark.git](https://github.com/mjlosch/seaice_benchmark.git).  
315

The benchmark results presented in this manuscript and the plotting scripts are available under <https://doi.org/10.5281/zenodo.18620150>.

## Appendix A: Code Verification, Comparison of Simulated Sea Ice between Veris and the MITgcm



**Figure A1.** Example benchmark simulation used for verifying the thermodynamic component of Veris. Top row: Thermodynamic evolution of ice thickness, ice concentration, and snow thickness simulated by Veris. Bottom row: Comparison of the simulated sea ice between Veris and the MITgcm employing the same initial and forcing conditions. This figure is representative of the results of different thermodynamic benchmark comparisons using different initial and forcing conditions.



**Figure A2.** Example benchmark simulation used for verifying the dynamic component of Veris. Left panel: Sea ice concentration simulated by Veris after 1000 10-minute time steps using forcing fields designed to induce ice fracturing. Middle panel: Comparison of the simulated sea ice between Veris and the MITgcm employing the same initial and forcing conditions. Right panel: As a reference for acceptable divergence, the difference between the MITgcm and a rerun of the MITgcm with perturbed initial and forcing conditions is shown. This figure is representative of the results of different dynamic benchmark comparisons using different grid sizes and initial and forcing conditions.

*Author contributions.* JPG wrote the python code with ML and SKC, carried out the benchmark experiments with ML, SKC and RN, coupled the model with Veros with RN and MJ, and ran the global simulations with ML. JPG wrote the manuscript with help of ML, MJ, RN.

320 *Competing interests.* The authors declare that they have no conflict of interest.

*Acknowledgements.* This paper builds upon the Master's thesis of Gärtner (2023), during which Veris was developed.

We express our gratitude to Dion Häfner for his assistance in coupling Veris with Veros.

## References

- Arakawa, A. and Lamb, V. R.: Computational Design of the Basic Dynamical Processes of the UCLA General Circulation Model, in: 325 General Circulation Models of the Atmosphere, vol. 17 of *Methods in Computational Physics: Advances in Research and Applications*, pp. 173–265, Elsevier, <https://doi.org/10.1016/B978-0-12-460817-7.50009-4>, 1977.
- Bouillon, S., Fichefet, T., Legat, V., and Madec, G.: The Elastic–Viscous–Plastic Method Revisited, *Ocean Modelling*, 71, 2–12, <https://doi.org/10.1016/j.ocemod.2013.05.013>, 2013.
- Bradbury, J., Frostig, R., Hawkins, P., Johnson, M. J., Leary, C., Maclaurin, D., Nedula, G., Paszke, A., VanderPlas, J., Wanderman-Milne, 330 S., and Zhang, Q.: JAX: composable transformations of Python+NumPy programs, GitHub, <http://github.com/jax-ml/jax>, 2018.
- Eden, C.: Closing the energy cycle in an ocean model, *Ocean Modelling*, 101, 30–42, <https://doi.org/10.1016/j.ocemod.2016.02.005>, 2016.
- Eden, C. and Olbers, D.: An energy compartment model for propagation, nonlinear interaction, and dissipation of internal gravity waves, *Journal of Physical Oceanography*, 44, 2093–2106, <https://doi.org/10.1175/JPO-D-13-0224.1>, 2014.
- Gärtner, J.: Sea Ice Modeling in Python, Zenodo, <https://doi.org/10.5281/ZENODO.11474409>, 335 <https://zenodo.org/doi/10.5281/zenodo.11474409>, 2023.
- Häfner, D., Jacobsen, R. L., Eden, C., Kristensen, M. R., Jochum, M., Nuterman, R., and Vinter, B.: Veros v0.1—A Fast and Versatile Ocean Simulator in Pure Python, *Geoscientific Model Development*, 11, 3299–3312, <https://doi.org/10.5194/gmd-11-3299-2018>, 2018.
- Häfner, D., Nuterman, R., and Jochum, M.: Fast, Cheap, and Turbulent—Global Ocean Modeling with GPU Acceleration in Python, *Journal of Advances in Modeling Earth Systems*, 13, e2021MS002717, <https://doi.org/10.1029/2021MS002717>, 2021.
- 340 Harris, C. R., Millman, K. J., van der Walt, S. J., Gommers, R., Virtanen, P., Cournapeau, D., Wieser, E., Taylor, J., Berg, S., Smith, N. J., Kern, R., Picus, M., Hoyer, S., van Kerkwijk, M. H., Brett, M., Haldane, A., del Río, J. F., Wiebe, M., Peterson, P., Gérard-Marchant, P., Sheppard, K., Reddy, T., Weckesser, W., Abbasi, H., Gohlke, C., and Oliphant, T. E.: Array Programming with NumPy, *Nature*, 585, 357–362, <https://doi.org/10.1038/s41586-020-2649-2>, 2020.
- Hersbach, H., Bell, B., Berrisford, P., Hirahara, S., Horányi, A., Muñoz-Sabater, J., Nicolas, J., Peubey, C., Radu, R., Schepers, D., et al.: The 345 ERA5 Global Reanalysis, *Quarterly Journal of the Royal Meteorological Society*, 146, 1999–2049, <https://doi.org/10.1002/qj.3803>, 2020.
- Hibler, W. D.: A Dynamic Thermodynamic Sea Ice Model, *Journal of physical oceanography*, 9, 815–846, [https://doi.org/10.1175/1520-0485\(1979\)009%3C0815:ADTSIM%3E2.0.CO;2](https://doi.org/10.1175/1520-0485(1979)009%3C0815:ADTSIM%3E2.0.CO;2), 1979.
- Hunke, E., Allard, R., Bailey, D. A., Blain, P., Craig, A., Dupont, F., DuVivier, A., Grumbine, R., Hebert, D., Holland, M., Jeffery, N., Lemieux, J.-F., Osinski, R., Poulsen, J., Steketee, A., Rasmussen, T., Ribergaard, M., Roach, L., Roberts, A., Turner, M., Winton, M., and 350 Worthen, D.: CICE-Consortium/CICE: CICE Version 6.5.1, Zenodo, <https://doi.org/10.5281/zenodo.11223920>, 2024.
- Hunke, E. C. and Dukowicz, J. K.: An Elastic–Viscous–Plastic Model for Sea Ice Dynamics, *Journal of physical oceanography*, 27, 1849–1867, [https://doi.org/10.1175/1520-0485\(1997\)027%3C1849:AEVPMF%3E2.0.CO;2](https://doi.org/10.1175/1520-0485(1997)027%3C1849:AEVPMF%3E2.0.CO;2), 1997.
- Kimmritz, M., Danilov, S., and Losch, M.: On the convergence of the modified elastic–viscous–plastic method for solving the sea ice momentum equation, *Journal of Computational Physics*, 296, 90–100, 2015.
- 355 Kimmritz, M., Danilov, S., and Losch, M.: The Adaptive EVP Method for Solving the Sea Ice Momentum Equation, *Ocean Modelling*, 101, 59–67, <https://doi.org/10.1016/j.ocemod.2016.03.004>, 2016.
- Lemieux, J.-F., Knoll, D. A., Tremblay, B., Holland, D. M., and Losch, M.: A comparison of the Jacobian-free Newton–Krylov method and the EVP model for solving the sea ice momentum equation with a viscous-plastic formulation: A serial algorithm study, *Journal of Computational Physics*, 231, 5926–5944, 2012.

- 360 Lemieux, J.-F., Tremblay, L. B., Dupont, F., Plante, M., Smith, G. C., and Dumont, D.: A Basal Stress Parameterization for Modeling Landfast Ice, *Journal of Geophysical Research: Oceans*, 120, 3157–3173, <https://doi.org/10.1002/2014JC010678>, 2015.
- Leppäranta, M.: A growth model for black ice, snow ice and snow thickness in subarctic basins, *Hydrology Research*, 14, 59–70, <https://doi.org/10.2166/nh.1983.0006>, 1983.
- Li, B., Basu Roy, R., Wang, D., Samsi, S., Gadepally, V., and Tiwari, D.: Toward Sustainable HPC: Carbon Footprint Estimation and Environmental Implications of HPC Systems, in: *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*, pp. 1–15, <https://doi.org/10.1145/3581784.3607035>, 2023.
- 365 Liu, Y., Losch, M., Hutter, N., and Mu, L.: A New Parameterization of Coastal Drag to Simulate Landfast Ice in Deep Marginal Seas in the Arctic, *Journal of Geophysical Research: Oceans*, 127, e2022JC018413, <https://doi.org/10.1029/2022JC018413>, 2022.
- Losch, M., Menemenlis, D., Campin, J.-M., Heimbach, P., and Hill, C.: On the Formulation of Sea-Ice Models. Part 1: Effects of Different Solver Implementations and Parameterizations, *Ocean Modelling*, 33, 129–144, <https://doi.org/10.1016/j.ocemod.2009.12.008>, 2010.
- 370 Marquetti, I., Rodrigues, J., and Desai, S. S.: Ecological Impact of Green Computing using Graphical Processing Units in Molecular Dynamics Simulations, *International Journal of Green Computing (IJGC)*, 9, 35–48, <https://doi.org/10.4018/IJGC.2018010103>, 2018.
- Marshall, J., Adcroft, A., Hill, C., Perelman, L., and Heisey, C.: A Finite-Volume, Incompressible Navier Stokes Model for Studies of the Ocean on Parallel Computers, *Journal of Geophysical Research: Oceans*, 102, 5753–5766, <https://doi.org/10.1029/96JC02775>, 1997.
- 375 Mehlmann, C., Danilov, S., Losch, M., Lemieux, J.-F., Hutter, N., Richter, T., Blain, P., Hunke, E., and Korn, P.: Simulating Linear Kinematic Features in Viscous-Plastic Sea Ice Models on Quadrilateral and Triangular Grids with Different Variable Staggering, *Journal of Advances in Modeling Earth Systems*, 13, e2021MS002523, <https://doi.org/10.1029/2021MS002523>, 2021.
- Méndez, M., Tinetti, F. G., and Overbey, J. L.: Climate models: challenges for Fortran development tools, in: *2014 Second International Workshop on Software Engineering for High Performance Computing in Computational Science and Engineering*, pp. 6–12, IEEE, <https://doi.org/10.1109/SE-HPCCSE.2014.7>, 2014.
- 380 Portegies Zwart, S.: The Ecological Impact of High-Performance Computing in Astrophysics, *Nature Astronomy*, 4, 819–822, <https://doi.org/10.1038/s41550-020-1208-y>, 2020.
- Rasmussen, T. A. S., Poulsen, J., Ribergaard, M. H., Sasanka, R., Craig, A. P., Hunke, E. C., and Rethmeier, S.: Refactoring the elastic–viscous–plastic solver from the sea ice model CICE v6. 5.1 for improved performance, *Geoscientific Model Development*, 17, 6529–6544, <https://doi.org/10.5194/gmd-17-6529-2024>, 2024.
- 385 Rosinski, J. M. and Williamson, D. L.: The accumulation of rounding errors and port validation for global atmospheric models, *SIAM Journal on Scientific Computing*, 18, 552–564, 1997.
- Semtner, A. J.: A Model for the Thermodynamic Growth of Sea Ice in Numerical Investigations of Climate, *Journal of Physical Oceanography*, 6, 379–389, [https://doi.org/10.1175/1520-0485\(1976\)006%3C0379:AMFTTG%3E2.0.CO;2](https://doi.org/10.1175/1520-0485(1976)006%3C0379:AMFTTG%3E2.0.CO;2), 1976.