~~Coupling of the Ice-sheet~~ New generic coupling adapters for ice sheet and ~~Sea-level System Model (version 4.24) with~~ subglacial hydrology ~~model CUAS-MPI~~ models (~~version~~ ISSM-preCICE Adapter 0.4, CUAS-MPI 0.1)~~using the preCICE coupling library~~

Daniel Abele[1,2,3], Thomas Kleiner[2], Yannic Fischler[4], Benjamin Uekermann[5], Gerasimos Chourdakis[5], Mathieu Morlighem[6], Achim Basermann[1], Christian Bischof[4], Hans-Joachim Bungartz[3], and Angelika Humbert[2,7]

[1]German Aerospace Center, Institute of Software Technology, Oberpfaffenhofen/Cologne, Germany
[2]Alfred-Wegener-Institut, Helmholtz-Zentrum für Polar- und Meeresforschung, Bremerhaven, Germany
[3]School of Computation, Information, and Technology, Technical University of Munich, Germany
[4]Department of Computer Science, Technical University Darmstadt, Germany
[5]Institute for Parallel and Distributed Systems, University of Stuttgart, Germany
[6]Department of Earth Sciences, Dartmouth College, Hanover, NH, USA
[7]Faculty of Geosciences, University of Bremen, Germany

Correspondence: Daniel Abele (daniel.abele@dlr.de)

Abstract. ~~Accurate earth system models must include interactions between atmosphere,~~ Adequate earth system simulations require interactions between the atmosphere, the ocean, and ~~continental~~ the ice sheets. To ~~build such models~~this end, numerical solvers that compute the evolution of the different earth system components are coupled. There are frameworks and libraries for coupling that handle the complex tasks of coordinating ~~the~~ solver execution, communicating between processes, and mapping between different meshes. This allows solvers to be developed independently without compromises on numerical methods or technology. Code reuse is improved, both over large, monolithic software systems that reimplement each coupled model as well as over ad-hoc coupling scripts.

In this work, we use the preCICE coupling library to couple the Ice-sheet and Sea-level System Model (ISSM) with the subglacial hydrology model CUAS-MPI. An adapter for each model is required ~~that passes the~~ to pass meshes and coupled variables between the model and preCICE. We ~~describe the generic, reusable adapterswe developed for both models and demonstrate their features experimentally. We also include computational performance results~~ focus mainly on the technical aspects (design, development, and use of the adapters, choice of coupling library, and large-scale performance analysis), using a synthetic setup to verify functionality and correctness.

The adapters we developed are generic, and reusable for use cases other than ice-hydrology coupling. Computational performance for the coupled system is measured on a high-performance computing cluster. ~~Coupling~~ We find that coupling with preCICE has low computational overhead and does not negatively impact scaling. ~~Therefore, the presented software facilitates~~ A comparison between unidirectional and bidirectional coupling for the synthetic ice sheet shows that the coupling captures the anticipated feedback mechanisms between the two systems. The coupled simulations are numerically stable, despite the

nonlinearities in the physical system. The generic coupling library preCICE is well suited for our use case and has advantages as well as disadvantages over earth system model specific libraries.

The new framework and code enables studies of the subglacial hydrology systems of ~~continental ice sheets~~ice sheets, as well as coupling ISSM or CUAS-MPI with other codes, such as in global earth system models or process models.

1 Introduction

Ice sheet dynamics is a gravity-driven, lubricated flow, forced by changes at ~~their~~ its boundaries, such as the ice-atmosphere interface, the ice-ocean interface and the conditions at the ice base. Beneath the ice sheet, a subglacial hydrological system exists that is formed by basal melting due to heat flux from the lithosphere and frictional heat. This hydrological system affects the ice sheet through changes in water pressure, while the ice sheet influences the subglacial system through changes in basal melt rates and ice sheet thickness. While the hydrological system ~~changes in the interior of the ice sheet~~ evolves on long time scales ~~only,~~ in the center of ice sheets, at the margins, ~~in particular in Greenland , change in the peak melt seasonin summer on short time scales. Coupled~~ particularly in Greenland during the melt season, the evolution is accelerated considerably. Both observations (Ing et al., 2024) and modeling (de Fleurian et al., 2022) have shown the high complexity of the feedback between hydrology and ice sheet. Therefore, coupled simulations are required to ~~simulate~~ study the evolution of both systems and their effect on each other. As both systems include non-linear processes, special care is required when simplifying the interactions of the system to a uniphysics problem, as also suggested by (Keyes et al., 2013). This work presents a framework for simulating interactions as a multiphysics problem and applies it to an artificial setup and to Greenland.

The Ice-Sheet and Sea-level System Model ~~(ISSM, Larour et al. (2012))~~ (ISSM, Larour et al., 2012) is a feature-rich ice sheet model. Among its capabilities, it ~~includes~~ provides a selection of subglacial hydrology models, including the Subglacial Hydrology and Kinetic, Transient Interactions ~~(SHAKTI) subglacial hydrology model (Sommers et al., 2018). SHAKTI is~~ model (SHAKTI, Sommers et al., 2018), the Glacier Drainage System (GlaDS, Werder et al., 2013), and the Double-Continuum model (DoCo, de Fleurian et al., 2016). These models are fully integrated into ISSM, using the same mesh and finite-element solvers. ~~In addition to SHAKTI, ISSM also contains an implementation of another hydrology model, Glacier Drainage System (GlaDS), which~~ GLaDS is also included in the ice sheet model Elmer/Ice (Gagliardini et al., 2013). Other ice sheet models like PISM (Khrulev et al., 2025) also have their own hydrology models. While monolithic software development can be easier (e.g., only one build process, sharing the same data structures), it also comes with disadvantages.

In this paper, we present a different, partitioned approach. ~~Fischler et al.~~ Fischler et al. (2023) recently published CUAS-MPI (subsequently referred to as CUAS), a stand-alone subglacial hydrology model. ISSM and CUAS rely on different spatial and temporal discretizations, ~~which complicates the coupling. Here, we couple ISSM and CUAS using the preCICE coupling library (Chourdakis et al., 2022). For this~~ so CUAS cannot be directly integrated into ISSM. We use an external independent coupling library, namely the general coupling library preCICE (Chourdakis et al., 2022) . To this end, we developed ~~the~~ adapters for both models that link them to preCICE. This approach has a number of advantages over integrating CUAS directly into ISSM. The models can be developed independently and make their own choices regarding discretization and numerical methods. ~~preCICE~~

is highly configurable, so the time points where data are exchanged and other properties of the coupling process can be adapted without code changes. preCICEalso provides sophisticated numerical methods that are required for accuracy, performance, and stability of the coupled system. Finally, the adapters aim to be as generic as possible to support other couplings than just ISSM-CUAS. It is therefore easy to either add components like an existing ocean circulation model or recombine the models with other ice sheet or hydrology models, which improves code reuseMoreover, coupling libraries, such as preCICE, make the setup more configurable and extendable – for instance, towards adding further components like an existing ocean circulation model or replacing existing components.

preCICE, as a general-purpose coupling library, follows an application-agnostic coupling approach. This is in contrast to specialized earth system modeling (ESM) couplers, such as OASIS3-MCT (Craig et al., 2017) or YAC (Hanke et al., 2016). Specialized couplers offer coupling methods, particularly data mapping methods, which are optimized to handle 2D fields on spherical surface meshes and are, thus, well suited to model global effects on the whole earth. preCICE, in contrast, treats meshes as clouds of scattered Cartesian points. Hocks and Uekermann (2026) compares the general data mapping of preCICE to those of specialized couplers using an ESM mapping benchmarks and concludes that preCICE is competitive. While specifically the radial-basis-function interpolation performs well for the considered smooth test functions, preCICE still lacks specialized conservative data mappings – a limitation that can, however, be overcome by bespoke pre- and postprocessing. The missing ESM specialization of preCICE is less important in non-global ESM scenarios, such as the setup of this paper, or scenarios where more flexibility is needed such as adding further models.

The functionality of preCICE goes beyond communication and data mapping, as specialized ESM couplers typically offer. Implicit coupling schemes including quasi-Newton acceleration (Mehl et al., 2016), time interpolation (Rodenberg and Uekermann, 2025), or multi-scale coupling (Desai et al., 2023) are all features that are relevant for ESM. Moreover, the generality of preCICE activates a large user and develop community, which brings further potential soft benefits: cross-domain knowledge transfer, extensive user documentation, sustainability of the software development, and integration in widely used simulation software, for example, OpenFOAM (Chourdakis et al., 2023) and FEniCS (Rodenberg et al., 2021).

Section Sect. 2 of this manuscript covers the software packages that are involved. First, we describe the coupling library preCICE and the existing solvers[1] ISSM and CUAS. Then, we present the newly developed adapters. In the next Sect.Sect. 3, we show the setup and results of a few basic experiments we performed to test the coupling and the computational performance. In Sect. 4, we discuss our findings and plans for future development.

2 Software

In the following, we describe all the software components necessary for the coupling of ISSM and CUAS. Figure Fig. 1 shows an overview of the coupling . We give short setup. We provide brief summaries of the three existing codes, with a focus on

[1]Throughout the paper, we use preCICE terminology, where *solver* refers to a complete simulation code and not rather than a single numerical routine. A solver that is part of a coupled setup is referred to as *participant*. See https://precice.org/fundamentals-terminology.html.

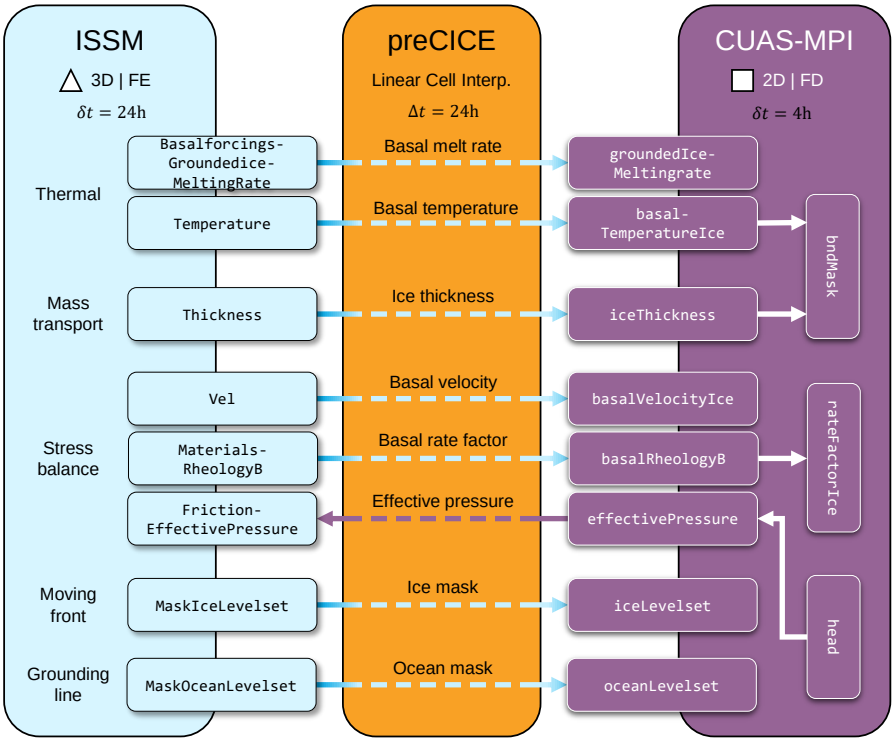

Figure 1. Overview of coupling between ISSM and CUAS-MPI. The coupling library preCICE handles communication and data mapping between the ice sheet and hydrology models. The models are free to use their own ~~mesh~~ meshes (unstructured triangular ~~,~~ or regular rectangular) and time steps. Coupled variables are listed by the name internal to the corresponding solver/adapter. ~~Some variables are not directly used by CUAS but are converted first~~ISSM uses `FrictionEffectivePressure` to compute sliding. ~~Details of~~ CUAS uses the ~~exchanged~~ variables received from ISSM as a water source and ~~how they are used by the solvers are described in~~ to update its active mask, see Sec. 2.4 for details. The ~~chosen~~ solver and coupling ~~interval of 30 d~~ configuration is ~~conservative to minimize computational overhead~~described in Sec. 3

the technical details relevant to coupling: the preCICE coupling library, the ~~ice sheet~~ ice-sheet model ISSM, and the subglacial hydrology model CUAS. The newly developed preCICE adapters for ISSM and CUAS are described in detail.

2.1 preCICE

preCICE ~~(Chourdakis et al., 2022)~~ (Precise Code Interaction Coupling Environment, Chourdakis et al., 2022) is an open source coupling library for multi-physics simulations. The library couples two or more independent parallel solvers (~~also~~ referred to as participants) and handles communication, data mapping, and coordination of the solvers, as described in the following paragraphs. Here, communication refers only to the exchange of coupling meshes and data between the solvers, not the internal

parallelism of the solvers themselves. To start a coupled simulation, the user starts each solver as usual, and each solver calls the routines for initialization and data exchange of the preCICE library. ~~This approach requires minimal changes to the source code of the solver, and all~~ All preCICE configuration options can be ~~configured at runtime . All instances of the library read a shared configuration file~~ set at runtime in a configuration file that is shared by the participating solvers, via which the respective algorithms for communication, data mapping, and time stepping are selected. The code that connects a solver with preCICE (individual lines of code, a dedicated class, or a complete standalone package) is called an adapter. An adapter calls the application programming interface of preCICE, converts between the data structures of the solver and preCICE, and steers the time evolution of the solver (i.e., adapting the time step size, or storing and reloading checkpoints, if necessary). ~~Adapters for a growing number of solvers and numerical frameworks exist.~~

2.1.1 Communication

preCICE communicates data between coupled solvers in a parallel, peer-to-peer, and point-to-point way. As back-ends, either TCP/IP or MPI can be used. ~~To establish the communication channels , coupling meshes from one solver are repartitioned on~~ Communication channels between the processes of ~~another solver during initialization using a two-level~~ both solvers are established using a highly scalable algorithm (Totounferoush et al., 2021).

2.1.2 Data mapping

To handle non-matching coupling meshes, preCICE offers different methods for data mapping, including projection-based methods and kernel methods (radial-basis function interpolation) (Chourdakis et al., 2022; Schneider and Uekermann, 2025). While some projection methods require mesh connectivity information, kernel methods operate solely on point clouds. Each mapping can be configured to be either *conservative* (the total values over the interface are conserved for extensive properties, e.g., mass, forces) or *consistent* (for intensive properties, e.g., temperature, pressure). preCICE supports 2D and 3D Cartesian meshes and surface and volume coupling. In the particular case of this paper, both codes use geographically projected coordinates and can thus apply 2D Cartesian meshes. ~~However, if we wanted to couple ice-sheet codes with Earth system models, we would need to convert between both coordinate systems. This is beyond the scope of this work.~~

2.1.3 Coordination of participants

To orchestrate the simulation progress of all coupled solvers, preCICE offers different coupling schemes. On the one hand, preCICE distinguishes between serial and parallel coupling: In serial coupling, coupled solvers advance sequentially, one after the other. In parallel coupling, coupled solvers advance concurrently. In both cases, the coupled solvers synchronize and exchange data after each fixed time window. On the other hand, preCICE distinguishes between explicit and implicit coupling. In explicit coupling, each time window is only computed once. In implicit coupling, each time window is repeated, with modified exchanged values, until predefined convergence criteria are met. To this end, solvers need to go back in time, which is typically implemented by checkpointing in the adapter. Implicit coupling increases accuracy and numerical stability. The

convergence behavior can be improved with fixed-point acceleration, for instance, with quasi-Newton methods (Mehl et al., 2016). Accuracy and numerical stability can further be improved by sampling time interpolants during each time window ~~(Rüth et al., 2021)~~(Rodenberg and Uekermann, 2025).

2.1.4 ~~Alternatives~~Adapter development

~~Other coupling libraries, similar to preCICE, exist (e.g., MUI (Tang et al., 2015), OpenPALM (Duchaine et al., 2015), or DTK (Slattery, 2016)), including some that specifically target Earth system modeling (e.g., YAC (Hanke et al., 2016) or MCT (Larson et al., 2005 /cpl7 (Craig et al., 2012)). Most of them focus on handling communication and data mapping, but do not offer advanced coupling methods (e.g., time interpolation or quasi-Newton acceleration). preCICE, as a general-purpose coupling library, does, however, not offer functionality specific to Earth system modeling, such as awareness of calendars (like ISSM) or data mapping tailored to geographic or spherical coordinate systems. One of the goals of this paper is to study to what extent such a general-purpose coupling library can be used for ice-sheet modeling, to potentially benefit from the larger community and maybe more advanced numerical methods~~

Adapters for some important numerical frameworks like OpenFOAM (Chourdakis et al., 2023) have been developed and are maintained by the preCICE team. Others are provided by the community. The preCICE website maintains a list of all adapters and their development status[2].

There are different ways to implement preCICE adapters. Some are directly integrated or patched into the solver's code, for example the CalculiX-preCICE adapter (Uekermann et al., 2017). Others are developed as stand-alone software packages, either as plugins like the the OpenFOAM-preCICE adapter (Chourdakis et al., 2023) or as orchestration codes that also call the solver like the the CAMRAD II-preCICE adapter (Huang et al., 2021).

2.2 ISSM

The Ice-sheet and Sea-level System Model (ISSM) is a well-established, feature-rich code for large-scale simulations of continental ice sheets (Larour et al., 2012). Mathematical ice sheet models consist of balance equations for enthalpy, mass, and momentum and their respective boundary conditions and kinematic boundary conditions for geometry evolution. Ice sheet codes are typically structured in different modules, also referred to as cores, that either solve individual balance equations or deal with the processing of data into forcing fields. A highly versatile ice sheet code, such as ISSM, offers several options for some cores. For example, the momentum balance might solve the full-Stokes equations (FS), higher-order Blatter-Pattyn approximation (HO), shallow-shelf approximation (SSA) or the shallow ice approximation (SIA), as described in the ISSM reference (Larour et al., 2012). Several glaciological processes can to date only be described empirically, for which a code may offer various parameterizations. Calving is such an example, for which ISSM offers a multitude of parameterizations. This configurability makes ISSM a code suitable for applications from mountain glaciers up to continental-scale ice sheets but also leads to large and complex code.

[2] https://precice.org/adapters-overview.html

Most use cases of ISSM are large problems in the order of 0.1M - 10M degrees of freedom (DOF). For example, for simulating the Greenland ~~ice sheet~~ Ice Sheet in moderate resolution (e.g., G4000 in (Fischler et al., 2022) with minimal element size of approximately 4 km), the different ISSM cores compute 31.5k - 944k DOF, high resolution (e.g., G250 in (Fischler et al., 2022)) requires 1.1M - 32M DOF~~(Fischler et al., 2022)~~. Large problem sizes are computationally demanding, so the simulations need to be efficient and scale adequately. Fischler et al. (2022) investigated the performance of ISSM, showing that the code scales well and is not expected to be a significant bottleneck for the scaling of the coupled simulation.

2.2.1 Multi-physics capabilities

Ice sheets are complex systems, so even a standalone ice sheet simulation is already a multi-physics simulation (Fig. 3). In ISSM, the cores can be run individually, e.g., to get the stress balance solution only. However, more often, transient runs are conducted, where most cores are solved. The system of equations of ice sheets is not solved in a numerically monolithic way, but in a sequential (segregated) fashion. In ISSM, a typical sequence is: first, the enthalpy balance is solved, then the stress balance, and afterwards, the geometry is evolved. Each core immediately uses the results of the previous cores.

2.2.2 Mesh and solver

ISSM supports two and three-dimensional meshes. ~~Figure~~ Fig. 2 shows the mesh structure. The basic horizontal two-dimensional mesh is an unstructured triangle grid covering the horizontal computational domain, including ice-free regions. The 2D mesh is usually static, as there is limited support for adaptive mesh refinement.

The horizontal 2D mesh can be used in the SSA approximation. For HO or FS, a three-dimensional mesh is required. The 3D mesh is generated by vertically extruding the 2D mesh in multiple layers. The vertices in the top layer are aligned with the surface of the ice. The vertices in the bottom layer are ~~aligned with~~ set at the base of the ice. The vertices in the layer in between are commonly distributed equally between the top and bottom vertices, but may be unequally distributed as well. Therefore, the vertical (z) coordinate of the vertices changes in every time step as the thickness of the ice changes. The vertices are connected as truncated triangular prisms or tetrahedra.

ISSM uses the finite element method (FEM) to solve the partial differential equations (PDE) for each core. The finite element type used can be configured for most cores individually. Linear P1 elements, where nodes are placed exclusively at the vertices, is the default for many cores, but higher order elements are available.

All cores use the same mesh but generally do not have the same finite element types. The time stepping method and step size ~~is~~ are also generally identical for every core with fixed or adaptive time steps, but a few cores, e.g., ~~SHAKTI~~the hydrology core with the DoCo method (de Fleurian et al., 2016), subdivide the steps further. All cores use the same number of CPUs and the same domain decomposition for MPI parallelization. Fischler et al. (2022) showed that the cores that solve two-dimensional problems (e.g., mass transport, moving front) are significant bottlenecks for scaling as they compute fewer DOFs than the ~~three-dimensionsal~~ three-dimensional cores (e.g., thermal, stress balance). This ~~could be resolved with more flexible domain decomposition or CPU allocation~~is one of the drawbacks of integrating different models into one code that coupling libraries improve upon.

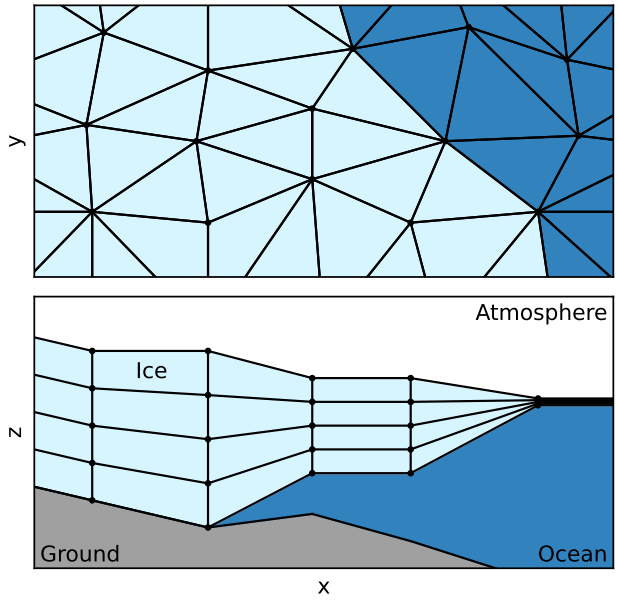

Figure 2. Schematic diagram of the ISSM mesh for 2D and 3D setups. The 2D mesh is an unstructured triangle grid covering the whole domain in the horizontal direction, even where there is no ice (top). Triangle elements are completely ice or completely ocean. The 3D mesh is generated by extruding the 2D mesh in multiple layers of triangle prism elements (bottom). The vertices in every layer have the same x and y coordinates. The z coordinate is set so that the mesh always matches the vertical extent of the ice, and it is updated when the thickness of the ice changes. In areas without ice, the mesh collapses in the vertical direction to a minimal thickness.

2.2.3 Architecture

ISSM's architecture is ~~well suited~~ well-suited for the development of a generic coupling adapter. Mesh and data access can be implemented based on abstract interfaces for different cores, mesh types, finite element types, etc. Variables are identified by runtime values (strings externally, mapped to enum values internally). With few exceptions that will be noted when describing the adapter in Sect. 2.3, the adapter does not need to include code to handle specific configurations.

2.3 The ISSM-preCICE adapter

The ISSM-preCICE adapter aims to be generic and extensible in order to support different use cases. This section explains how the features of ISSM and preCICE are handled in the implementation. The adapter is an executable that runs in place of the ISSM executable. The adapter configuration file is specified as a command-line parameter, and the command-line parameters of the ISSM executable are part of the adapter configuration file. The configuration of the adapter is done by a file in YAML format. The adapter configuration file is mostly responsible for mapping names specified in the preCICE configuration file to names expected by ISSM and defining the coupling interface. Listing 1 shows an example configuration file. The format

```yaml
precice_config_file_name: precice-config.xml
participant_name: ISSM
issm:
  root_path: some/path/to/model
  model_name: model
interfaces:
  - mesh_name: ISSM-Mesh
    patches:
      - Base # or Surface
    read_data:
      - name: effectivePressure
        solver_name: FrictionEffectivePressure
    write_data:
      - name: iceThickness
        solver_name: Thickness
      - # ...
```

Listing 1. Example adapter configuration file in YAML format. The adapter requires information about the preCICE and ISSM configurations, the coupling mesh and the names of the variables being read or written. Variable names in the preCICE configuration file are mapped to variables known to ISSM.

conforms to the adapter configuration schema defined by the preECO project[3]. Details of the entries in the file are explained in the following sections. Each section highlights limitations and missing features.

2.3.1 Architecture of the adapter

~~There are~~ Sect. 2.1.4 outlines the different ways to implement preCICE adapters. ~~Adapters can be directly integrated or patched into the solver's code (e.g., the CalculiX-preCICE adapter (Uekermann et al., 2017)), while other adapters are developed as stand-alone software packages, either as plugins (e.g., the OpenFOAM-preCICE adapter (Chourdakis et al., 2023)) or as orchestration codes that also call the solver (e.g., the CAMRAD II-preCICE adapter (Huang et al., 2021)).~~ As shown in Fig. 3, the ISSM adapter ~~follows the latter approach, structured as a~~ is implemented as a coordinating wrapper application that calls ISSM (used as a library) as well as preCICE. This approach allows for independent development and a clean architectural separation between solver and adapter, allowing, e.g., easy support for multiple ISSM versions. However, this relies on a relatively stable API of ISSM and might pose maintenance challenges in the future. Additionally, due to the architectural choice of using ISSM as a library (instead of via its command-line interface), some functionalities internal to ISSM are not available. Note that so far no changes to ISSM or its build process were necessary to support the features of the adapter.

[3] ~~https://precice.org/couple-your-code-adapter-software-engineering.html~~https://precice.org/couple-your-code-adapter-software-engineering.html

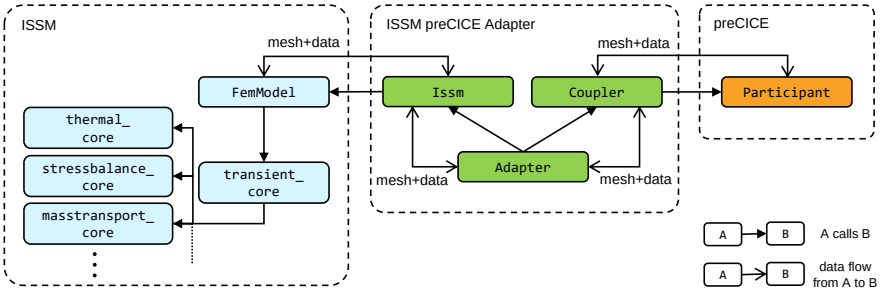

Figure 3. Structure of the ISSM-preCICE adapter. The adapter has its own time loop and calls the ISSM solver through the ISSM `FemModel` class. No interaction with the internals of ISSM (such as the constituent cores that make up the transient core) is necessary. The `Adapter` class is the entry point of the program; it coordinates the coupling and the exchange of coupled data (read and write) between ISSM and preCICE. Both ISSM and preCICE are wrapped in the adapter library to improve isolation and testability.

2.3.2 Coupling mesh

Since ISSM could potentially be coupled with many different codes, the coupling interface is configurable. The most important interfaces for coupling with the environment of the ice are the base and surface. So far, these are the only supported interfaces, see the list of limitations below for other interfaces that may be added later. In the configuration file, the user specifies which part of the mesh forms the coupling interface. The vertices of this part of the mesh are passed to preCICE. Mesh connectivity is added to support mapping schemes like linear cell interpolation[4] (Chourdakis et al., 2022). The finite element representation of ISSM variables is evaluated at the mesh vertices before writing data to preCICE.

~~**Limitations:** The current implementation of the adapter has some limitations regarding the mesh handling, all of which could be addressed in future work, as the technical requirements are already available. These limitations include:~~

Some precision of high-order finite element types is lost when evaluating variables at the mesh vertices. Instead, the nodes of the finite elements could be used. For best results, this would require multiple coupling interfaces for different finite element types, and mesh connectivity would not be available. The adaptive mesh refinement feature of ISSM is not supported~~,~~; the mesh must be static. Note that preCICE does provide facilities for mesh refinement[5], but the adapter does not yet use them.

Specialized ESM coupling support masking of irrelevant areas, e.g., masking land for ocean models. As mentioned above, preCICE does not provide this feature. Additionally, preCICE does not handle non-matching domains very well. This is not a big issue for the use case discussed here, since ISSM and CUAS are interested in approximately the same domain, and both have internal mechanisms to exclude irrelevant areas and avoid wasteful computations. But it may be beneficial for, e.g., ice-ocean coupling to only couple over floating ice.

3D (volume) coupling is ~~not supported. A 3D coupling interface~~ a possible future extension. This would allow, e.g., to couple ISSM with itself to use different meshes for different cores to optimize precision or performance or to use an external

[4]https://precice.org/configuration-mapping.html
[5]https://precice.org/couple-your-code-moving-or-changing-meshes.html

thermal solver. However, as ice is transported by the solver, the z coordinate of the vertices changes in every time step. As the interface changes significantly over time, the coupling mesh could be reset to the new locations of the interface, just like for adaptive mesh refinement above. So far, we have not tested whether the computational overhead to update the coupling mesh and corresponding ~~interpolation~~ mapping matrices in preCICE is acceptable. ~~Partial interfaces are not supported. For example, it may be beneficial for ice-ocean coupling to only couple over floating ice.~~

2.3.3 Variables

In ISSM, every state variable is identified by a unique name. These names do not follow any consistent convention, and the other coupling participant may use different names. So, it is necessary to map between names in the preCICE configuration file and ISSM names. The configuration file provides such name mappings for read and written variables.

If ~~ISSM~~ the ISSM setup uses a 3D mesh, ~~it is possible to~~ the adapter can perform depth-averaging of a variable before writing it and extruding a variable (i.e., copying the variable values to the other layers of the mesh) after reading using the existing ISSM functionality.

~~**Limitations:**~~ ISSM accepts time series as input variables. For example, the user can set specific values for these variables at the beginning, middle, and end of the simulation. ISSM calls these "transient variables" and temporally interpolates between these values when necessary. However, ISSM does not allow overwriting the user-provided values of such transient variables. Therefore, the adapter requires coupled variables to be set up as non-transient, i.e., with one fixed value that is overwritten with the value read from preCICE during the simulation. This is a purely technical restriction and does not reduce modelling capabilities, since coupling ~~can also include~~ also models input variables that change over time.

2.3.4 Data initialization

In ISSM, most variables are not zero as the simulation starts at some point in time in a defined state. For some variables, e.g., ice thickness, zero is not even valid at all. For non-zero initial values, preCICE ~~allows~~ requires data initialization[6] by writing variables once ~~in~~ at the beginning of the simulation. The ISSM adapter assumes that such initialization is necessary for every variable it writes.

For most variables, the adapter ~~simply writes~~ can simply write the initial value specified in the ISSM setup. However, users of standalone ISSM are not required to specify true initial values for variables that are computed by ISSM before they are used. For velocity, the initial value in the setup is merely used as an initial guess for the non-linear iteration of the stress balance core. ~~To initialize velocity, the adapter runs the stress balance core during the initialization phase to get the true initial velocity. This also applies to other variablesand cores, e.g.,~~ Other variables, like rheology parameter B that is computed by the thermal core, do not need to be specified in the setup at all when the corresponding core is active.

~~**Limitations:** There is no generic way to handle the computation of initial values for all such variables. Special handling has to be added separately for each variablethat requires it~~ For selected variables, the adapter can be configured to compute the true initial value. For example, to initialize velocity, the adapter would run the stress balance core once. But this is not possible for

[6] https://precice.org/couple-your-code-initializing-coupling-data.html

every variable, at least not in a generic way, so proper initial values must be provided by the user of the adapter even when it is not required by ISSM itself.

2.3.5 Boundary conditions

ISSM can set discrete Dirichlet boundary conditions (mainly called single point constraints in ISSM) for some cores, e.g., velocity (`Vx`, `Vy`, `Vz`) in the stress balance core. The adapter has limited support to set some of these constraints by coupling. In the ISSM code, the constraints are stored differently from normal variables, so generic support for all constraints that ISSM uses is currently not possible. The association between constraints and the core that they apply to has to be hard-coded for each one. For example, the adapter has no way to find out automatically that constraints on velocity are used by the stress balance core. Further complicating the implementation is that manual MPI communication is required in the adapter to synchronize constraints on ghost vertices. For variables, no such extra communication was necessary.

~~**Limitations:**~~ Constraints can currently only be coupled for P1 finite elements. Internally to ISSM, constraints are stored per finite element node, not per mesh vertex. ~~Since~~ But the coupling mesh is defined at the vertices instead of ~~nodes, only P1 finite elements (or similar) are supported. Unlike for variables, there is no interpolation from P1 to other element types for Dirichlet constraints. Only~~ at the finite element nodes. Supporting constraints on other finite element types would require manual interpolation, whereas for normal variables, the interpolation from different finite element types is handled by ISSM automatically. Additionally, only velocity and pressure constraints of the stress balance core ~~are supported~~ have been implemented so far.

2.3.6 Time stepping

ISSM performs multiple time steps per coupling window, depending on the step size set in the ISSM setup.

~~**Limitations:** Subcycling is not used~~During a coupling window, the adapter does not perform subcycling as it is defined by preCICE, i.e., `precice::Participant::advance` ~~is not called on every ISSM time step. Therefore, time interpolation by preCICE is not enabled. This is to avoid excessive file output: ISSM unavoidably writes output at least once per call to the solver routine (in addition to fixed intervals set by the user). So we call the routine only once per coupling time window , performing multiple time steps per call, and then call~~ `advance` ~~once after the routine ends. This could only be resolved with code changes in ISSM itself.~~ it does not read or write intermediate values, only snapshots at the beginning or end. So, coupled variables are constant over one coupling window. The time interpolation features of preCICE mentioned in Sect. 2.1.3 are not currently used. This may be added in the future, but would probably require changes to the ISSM code. In the current version, the coupling window has to be set short enough to capture the dynamics of the system.

Implicit coupling is not supported. The adapter ~~currently cannot~~ does not create the necessary checkpoints of ISSM, neither in memory nor on disk. Implicit coupling is required for numerical stability in some setups (see Sect. 2.1). Our experiments show no instability with explicit coupling. This is consistent with the internal explicit coupling of ISSM cores explained in Sect.~~2.2.1.~~ 2.2.1. This is another possible future extension.

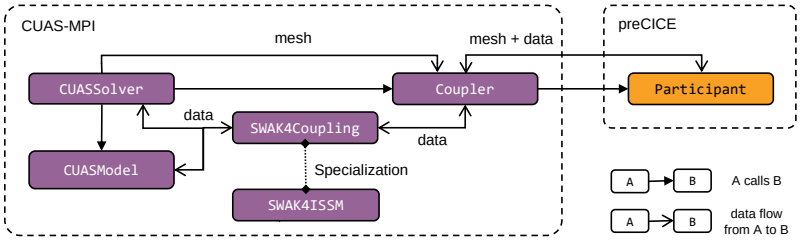

Figure 4. Structure of the CUAS-preCICE adapter. The coupling is integrated into and coordinated by the existing `CUASSolver`, but most of the coupling logic is isolated in added classes. The coupling data that is read and written flows through a utility class (specialized for coupling with ISSM) that handles transformations such as deriving ice pressure and transforming units.

2.4 CUAS and the CUAS-preCICE adapter

CUAS (Fischler et al., 2023) is the MPI parallel implementation of the Confined-Unconfined Aquifer System (CUAS) model for subglacial hydrology (Beyer et al., 2018). It employs an equivalent porous medium approach in which both a distributed and channelised system are represented by one porous layer. The model solves a vertically integrated groundwater equation (Fischler et al., 2023, Eq. 1) using effective quantities for storativity and transmissivity, which evolve over time based on parameterisations (Fischler et al., 2023, Eqns.2–4). CUAS uses a finite difference spatial approximation on a regular rectangular grid and an implicit Euler time stepping scheme. CUAS solves for the hydraulic head that is proportional to the water pressure. The effective pressure, N (ice overburden pressure minus water pressure), is a diagnostic quantity that is computed in each time step and is used in ISSM for sliding.

We added an experimental preCICE adapter to CUAS. An adapter configuration file has not been specified yet, so the adapter is specific for coupling with ISSM. The adapter currently does not support implicit coupling schemes~~or interpolations that require mesh connectivity. The CUAS-preCICE adapter~~. It is implemented within the CUAS code base and is not a standalone application. This ~~adapter implementation strategy~~ implementation approach is similar to e.g., the CalculiX-preCICE adapter (Uekermann et al., 2017) and offers a high degree of flexibility. Additionally, the code base of CUAS is much smaller and more modern than of ISSM, so maintenance is not significantly impacted by this choice.

~~Figure~~ Fig. 4 shows an overview of the module structure and data flow. preCICE is integrated directly into the existing CUAS time step iteration, enabling subcycling and time interpolation. The CUAS-preCICE adapter consists of two parts. The `Coupler` class coordinates the main coupling operations of initialization, reading and writing data, and advancing the coupling window. It exchanges data with the CUAS model and solver through the `SWAK4Coupling` interface, which applies necessary transformations to the data based on the model physics. While some transformations are generic, e.g., deriving ice pressure used by CUAS from ice thickness provided by ISSM, others are very specific to the data that the ISSM-preCICE adapter can provide, hence the `SWAK4ISSM` specialization of the interface. Some of the tasks are as simple as converting rheology parameter B from ISSM to `rateFactorIce` (A) in CUAS using $A = B^{-3}$ (assuming a Glen's flow law exponent

of $n = 3$). Others are more complex. For example, we use the ice thickness to compute the ice pressure. We further compute the pressure melting point using the ice pressure and the absolute temperature from ISSM to decide whether the ice is frozen at the bed or not, and adjust the mask (active versus inactive) in CUAS. Users can configure whether the mask is allowed to change based on the simulated temperature ~~or not~~and the temperature threshold at which it becomes active. The transformations applied to the coupling data are implemented directly in the adapter, but the motivation is similar to preCICE Actions[7].

Each time new data is available from preCICE (see Fig. 1~~or Fig. ??~~), the CUAS-preCICE adapter needs to perform several tasks, which are briefly outlined below.

- `iceThickness` is translated into ice overburden pressure using ice density.
- `groundedIceMeltingrate` is rescaled from ~~$\mathrm{m\,s^{-1}}$ ice equivalent to $\mathrm{m\,s^{-1}}$ water equivalent~~ $\mathrm{m/s}$ ice equivalent (IE) to $\mathrm{m/s}$ water equivalent (WE) using ice and water density. This is then used as a time-independent (steady) forcing for CUAS during the duration of the current coupling time window.
- Use the `iceThickness` and the steady bed elevation field from CUAS to compute a new `bndMask` using the flotation condition. The `bndMask` contains the information where we have active hydrology (warm base, grounded ice) and where boundary conditions need to be applied (e.g. floating ice or open ocean). Here we also initialize grid points that turned from ocean ~~boundary condition into~~ into ice and thus active CUAS due to grounding line advance. Grounding line retreat is also handled.
- We use `iceLevelset` to disable grounded ice areas in the CUAS domain that are not part of the ISSM domain. We do not use `oceanLevelset` in the prototype implementation of the adapter.
- `basalTemperatureIce` together with `iceThickness` is used to decide if the base is at the pressure melting point to further constrain the `bndMask` in CUAS, if needed.
- The `basalVelocityIce` is copied over without modifications.
- Finally, `rateFactorIce` (A) in CUAS is computed based on `basalRheologyB`.

Because the effective pressure is computed directly in CUAS, the adapter can ~~provided~~ provide this field for coupling without further modifications. We use `iceThickness` instead of `iceLevelset` and `oceanLevelset` provided by ISSM to define the `bndMask` in CUAS. This ensures that the grounding line (GL) in CUAS is consistent with the local bed topography. The ISSM mesh resolution at a given location may be substantially coarser than the CUAS grid resolution. Through the preCICE coupler, we only receive interpolated level-set fields representing the GL position on the coarser ISSM mesh. Consequently, the transferred GL location may not align with the grounding line implied by the higher-resolution bed topography used in CUAS.

To ~~allow for~~ enable coupled simulations, ~~other~~ the model capabilities of CUAS~~besides the adapterhave been enhanced. We reimplemented the way CUAS is handling the forcing so that different model forcings can be registered and are aggregated.~~

[7] https://precice.org/configuration-action.html

~~Ice sheet basal melt from coupling is just one source of water for CUAS. This is important because simulations of~~ , beyond the adapter, have been further enhanced. Simulations of the Greenland Ice Sheet require water not only from ~~basal melt~~the basal melt computed by ISSM, but also from surface runoff. Until now, there has been only one input field and a corresponding internal storage field for the combined water from all sources, used as model forcing. We added the ability to store multiple water sources that can be changed independently (by a time series from input files or by coupling) and added up when required. Various parameters, such as water and ice density, that were compile-time constants are now configurable at runtime to match the parameters set in ISSM.

3 Experiments

This section ~~describes~~ presents experimental verification of the ISSM-preCICE adapter for coupling ~~of ISSM to a model of subglacial hydrology. Here, only synthetic setups are used. Use in real world setups will be explored in future research. Both functionality and computational performance are demonstrated in the following sections~~ISSM to CUAS. First, we demonstrate functionality using a synthetic setup, then we analyze the performance using preexisting large-scale setups of the Greenland Ice Sheet.

3.1 Functionality

We ran simulations with a synthetic setup. The simpler geometry compared to a real-life setup allows us to verify the correctness of the implementation and study the complex behavior of the coupled system in a more controlled setting.

3.1.1 Experimental setup

~~State of the ice sheet in the CalvingMIP Thule domain at steady state. The setup is used to test coupling functionality.~~

We use the synthetic Thule geometry developed for the CalvingMIP project (Jordan, 2024). This setup is based on analytical functions for the bed elevation and ~~results in an ice sheet geometry that contains all the major parts of an ice sheet~~ yields an ice-sheet geometry that encompasses all major components of an ice-sheet model domain (grounded ice, floating ice, and open ocean). In CalvingMIP it is used to study how different ice sheet models handle calving in a very controlled setup. Instead of fixed thermal and friction conditions as in the ~~CalvingMIP project~~original definition, we enable the thermal core of ISSM with surface temperature $250\,\mathrm{K}$ and geothermal flux $0.05\,\mathrm{W\,m^{-2}}$ to compute basal melt rates to be used by CUAS and we use the default Budd friction law ($\sigma_\mathrm{b} = C^2 N u_\mathrm{b} \boldsymbol{\sigma}_b = -C^2 N^{\frac{q}{p}} |\boldsymbol{v}_\mathrm{b}|^{\frac{1}{p}-1} \boldsymbol{v}_\mathrm{b}$) with coefficient $C = 100$, ~~and~~ exponents $p = 3$ and $q = 2$, and effective pressure N supplied by CUAS. Surface mass balance was increased to $8\,\mathrm{m\,a^{-1}}$ ~~ice equivalent to get the balanced geometry shown in 5~~$8\,\mathrm{m/a\,IE}$. The goal is ~~a mix of~~ to have a balanced geometry with grounded and floating, ~~fast and slow flowing regions to see~~ fast- and slow-flowing regions to assess the effects of coupling.

The combined model (ISSM + CUAS) initialization consists of two phases. First, we run the ISSM model for ~~12500 years with 2 year time steps into a steady-state~~ 12,000 years with a time step of $0.1\,\mathrm{a}$ to reach a steady state, using an effective pressure ~~field of half~~ equal to the ice overburden pressure~~that is only derived from the ice sheet model geometry~~. In the second

phase, CUAS ~~is coupled just once to ISSM at the beginning of its own spinup to receive a consistent ice sheet model state (see Fig. ??). The CUAS spinup runs for 100 years (resolution: 2.5 km, time step: 4 h) into steady state~~ and ISSM run coupled without melt runoff forcing, exchanging data every 0.1 a (equal to the ISSM time step) for 200 a. This allows CUAS to find its own steady state while keeping ISSM state consistent with the effective pressure computed by CUAS.

~~To investigate the effect of coupling, we perform two experiments. In the control simulation, the models run fully coupled for 20 years starting from the end of~~ For the actual experiment, we add forcing in the form of seasonal and spatially varying meltwater runoff (in addition to the basal meltwater received from ISSM) to CUAS to induce a dynamic system. We ran simulations with low and high runoff. The peak forcing in summer (day 210 of the year) in areas with ice surface below 500 m is 0.2 m/a WE for low runoff and 2.0 m/a WE for high runoff. Over the year, the forcing follows the shape of a Gaussian distribution with parameters $\mu = 210\,\mathrm{d}$ and $\sigma = 20\,\mathrm{d}$. Spatially, forcing is maximal for ice surface $h_s < 500\,\mathrm{m}$, zero $h_s > 1500\,\mathrm{m}$, and in between follows a monotone cosine curve, i.e., forcing increases as the surface decreases. The surface elevation h_s is taken at the end of the ~~spinup. The coupling interval is 720 h (30 d) . This is a conservative choice and the low overhead for coupling (see Sect. 3.2) shows that the frequency of coupling could be increased without significantly impacting computational performance. The time step of the solvers is still 4 h in CUAS and 1 d in ISSM. In the second experiment, we apply an additional, time-dependent water source of about $1\,\mathrm{m\,a^{-1}}$ (water equivalent) for 60 days along a straight line in one half of the domain (Fig. ??) to simulate the effect of surface melt water draining down to the base through a long crack during the summer season. We refer to this additional source of water as an anomaly (anom.) forcing. All other parameters are identical to the control simulation. As mentioned in Sect. 2.4, CUAS has been extended to aggregate water sources from coupling and from input files to enable simulations with anomaly forcing~~ spin-up, the forcing is not dynamically updated.

~~Details of the coupling configurationare~~ The basic coupling configuration, both for spin-up and the following experiment, is shown in Fig.~~??~~ 1. As described above, CUAS writes the effective pressure that is used by ISSM to determine basal friction. The basal melting rate from ISSM provides a ~~source of water~~ water source for CUAS. Basal temperature, ice thickness, and ice and ocean masks are used to update the active CUAS mask. The temperature threshold for activating the mask is set to 269.15 K. Ice velocity and rheology ~~influence the opening and closing of channels. We are using the simplest mapping scheme, nearest neighbour~~govern channel opening and closure. We use linear cell interpolation to map data between the two meshes. Details on how the coupled variables are used in CUAS are described in Sect.~~2.4~~ 2.4. Simulations are run over two years to assess changes from one year to the next. Data is exchanged daily to accurately capture rapid changes during summer. ISSM uses time steps of 1 d same as the coupling window, CUAS 4 h.

We compare a fully coupled (2-way) run with one in which coupling is performed only in one direction (1-way). In 2-way, ISSM and CUAS exchange all variables as per the coupling configuration described above. In 1-way, ISSM receives the updated effective pressure from CUAS, but CUAS does not receive updated ice geometry from ISSM. This is equivalent to offline coupling: first run CUAS standalone, then use the resulting time series as input for a standalone ISSM simulation. To ensure comparable aggregate forcing, CUAS receives the steady basal melt field from the end of the spin-up from an input file.

~~Coupling configuration between ISSM and CUAS. ISSM writes the state of the ice and reads the effective pressure to be used in the friction law. The figure was generated using the preCICE config visualizer ().~~

3.1.2 Results

~~Figures ?? and ?? show the results of the coupled runs~~ The state at the end of the ~~simulation with (anom.) or without (control) an additional water source.~~combined model initialization (spin-up) is presented in Fig. 5. The ice thickness and basal velocity are selected as the key quantities describing the ice sheet state, while effective pressure and subglacial discharge are selected to describe the state of the subglacial hydrology. The ~~hashed areas in both figures indicate the~~ dark gray areas in Fig. 5a–d (and subsequent maps) indicate the grounded parts of the domain where basal temperatures from ISSM are below the pressure melting point and ice sheet basal melt is zero (cold base) ~~. In a CUAS stand-alone simulation, these areas would be considered too cold to allow for an active hydrology~~ and hence no active hydrology exists. The ice thickness (Fig. 5a) in the warm based area varies over several thousand meters, with low thickness in region A and along the northern and southern grounding lines, as well the upstream end in region B. The effective pressure (Fig. 5b) is lowest at the grounding line and reaches magnitudes above $6\,\mathrm{MPa}$. Basal velocities (Fig. 5c) range from 10 to about $400\,\mathrm{m/a}$ with largest values along the long grounding lines in the north and south. Along those grounding lines the subglacial discharge (Fig. 5d) is highest, with a similar pattern of N, v_b and D. ~~Nevertheless, ISSM allows for sliding independent of the basal temperature and thus, also the cold-base areas are active in the CUAS setup and feed back the effective pressure to the ice sheet model.~~

~~In areas where ISSM is at the pressure melting point (warm base)~~

For a further investigation of the differences in 1- and ~~ice sheet basal melt provides a source of water for the subglacial hydrology, coupling of ISSM and CUAS results in lower effective pressure(Fig. ??a), leading to enhanced basal sliding (~~2-way coupling, we present simulations with different magnitude in seasonal forcing. Figure 6 displays the time series of total water source, Q_tot, mean effective pressure, $\bar{N}$, mean effective transmissivity magnitude , $\overline{\log(T_\mathrm{e})}$, and masal basal velocity magnitude, $\overline{\log(v_\mathrm{b})}$ showing 1-way and 2-way coupling and for each full forcing (high) and only 10% amplitude forcing (low). The imposed runoff forcing is periodic with a one-year period and attains its maximum at day 210, end of July. As model output is written every 10 days, this maximum is not sampled exactly in the second simulation year. The nearest stored time step (day 215) therefore exhibits a slightly reduced signal amplitude. The effective pressure (Fig. ~~??c) and thus higher ice discharge into the ocean. This results in a reduction of ice thickness~~ 6b) has a lower difference between 1- and 2-way coupling for low runoff forcing. In both runoff cases, the 2-way coupling has higher $\bar{N}$. The timing in $\bar{N}$ is in the low runoff forcing is slightly delayed to Q_tot. The timing of $\overline{\log(T_\mathrm{e})}$ (Fig. ~~??a) in most of those areas and grounding line retreat. The increase in subglacial discharge~~ 6c) in the low runoff forcing is similar to $\bar{N}$ and $\overline{\log(v_\mathrm{b})}$ (Fig. ~~??d)is a direct consequence of higher basal melt rates in the coupled simulations, also contributing to the reduction~~ 6d). The effect of the coupling is small for low runoff forcing in $\overline{\log(T_\mathrm{e})}$, with larger differences outside the peak runoff and higher $\overline{\log(T_\mathrm{e})}$ in 1-way. The basal velocity resembles the inverse shape of $\bar{N}$.

Comparing the results for 1-way and 2-way coupling in ice thickness ~~. In the cold-based areas of the domain, the effective pressure increases~~ difference (Fig. 7a, 8a), we find distinct differences in the regions A and B: in A the 1-way case exceeds

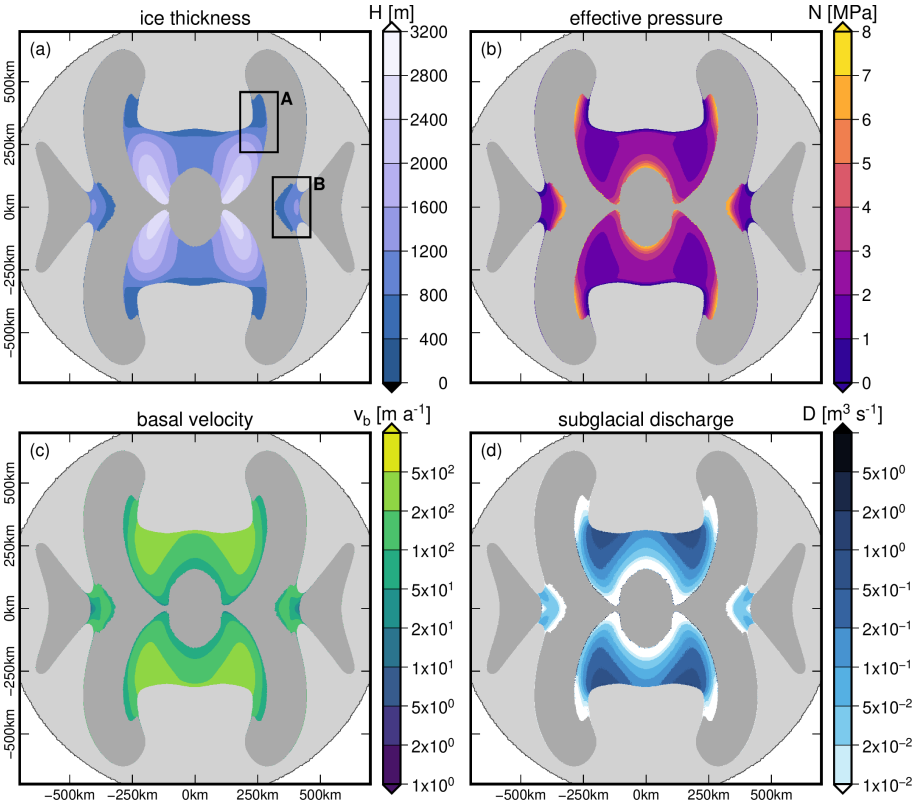

Figure 5. Model state at the end of the spin-up for the Thule domain setup. Ice-dynamical fields are shown in the left column: (a) ice thickness and (c) basal velocity. Subglacial hydrological fields are shown in the right column: (b) effective pressure and (d) subglacial discharge. Colored shading is shown only for the grounded portion of the ice sheet; floating ice shelves are indicated in light gray. Dark gray areas are inactive in CUAS. The maps are cropped to focus on the grounded ice region and therefore do not display the full computational domain. Rectangles (A, B) indicate regions of high activity referenced in the text.

$\pm 50\,\mathrm{m}$, while the 2-way case has small differences. The pattern of ice thickness difference is similar for the 1- and 2-way coupling, but the magnitude is in the 1-way case by far larger. For both, 1-way coupling and 2-way coupling, we find a reduction in effective pressure compared to the ~~uncoupled simulation, leading to lower basal velocities and thus slightly thicker ice.~~

~~Result of the fully coupled control run using the CalvingMIP setup. The panels a–d show the difference of some selected fields with respect to the initial state (end of spinup) at the end of the 20 year simulation. The gray area indicates the extent of the floating ice. The basal melt rate computed by ISSM is the only water source. The hashed areas indicate the parts of the domain where basal temperatures from ISSM are below the melting point and the basal melt rate is zero.~~

~~The simulation with additional water from the surface (anom.) is intended to test if CUAS can handle different water sources (forcing). Figure ?? illustrates that this is the case and that anomaly forcing leads to an alteration (relative to the control simulation)of the ice thickness and basal velocity. The difference is still visible at the end of the simulation, which is also the~~

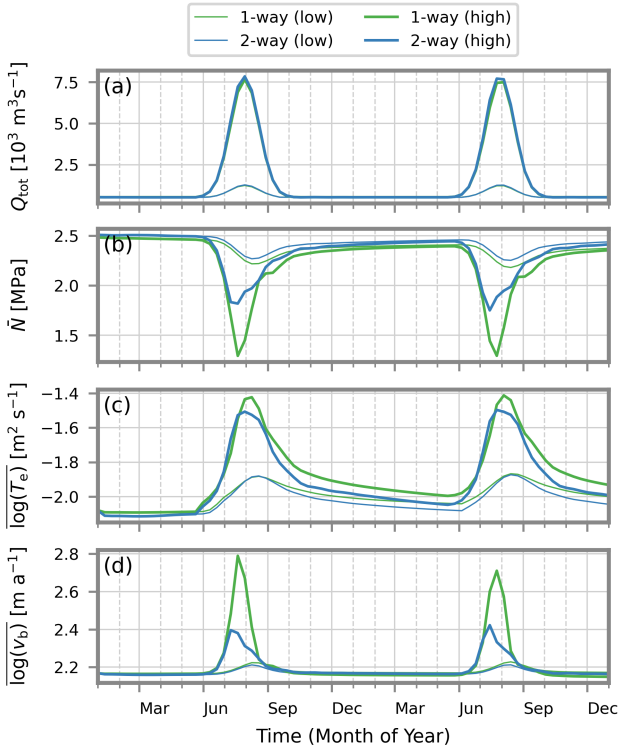

Figure 6. Time series of total water source Q_{tot}, mean effective pressure $\bar{N}$, mean effective transmissivity magnitude $\overline{\log(T_{\mathrm{e}})}$, and mean basal velocity magnitude $\overline{\log(v_{\mathrm{b}})}$ for the four simulations: one-way coupling in green, two-way coupling in blue; high-amplitude forcing shown with thick lines, low-amplitude forcing with thin lines. The means are evaluated as spatial means on the time-dependent active CUAS mask.

~~end of the year, although the anomaly is only applied during a fraction of the year~~spin-up that does not contain seasonal runoff (Fig. 7b, 8b). The basal velocity is in both cases, 1-way coupling and 2-way coupling, we have an increase in basal velocity compared to the spin-up (Fig. 7c, 8c), with an extreme increase in region A for the 1-way coupling. The discharge (Fig. 7d, 8d) is in both cases larger than in the spin-up, but the difference to the spin-up is only moderate in magnitude, except for region B in the 1-way case.

3.2 Performance

The coupled simulation should run with minimal computational overhead and use computing resources efficiently. To demonstrate this, we use a large-scale setup of the Greenland Ice Sheet ~~. As described below, the coupling is simplified to avoid numerical problems, but should have representative performance characteristics . Realistic coupling would require more careful construction of compatible setups and is beyond the scope of this paper. This will be explored in future work.~~ as the simple setup used in Sect. 3.1 is not complex enough to have performance characteristics that are representative of real uses cases,

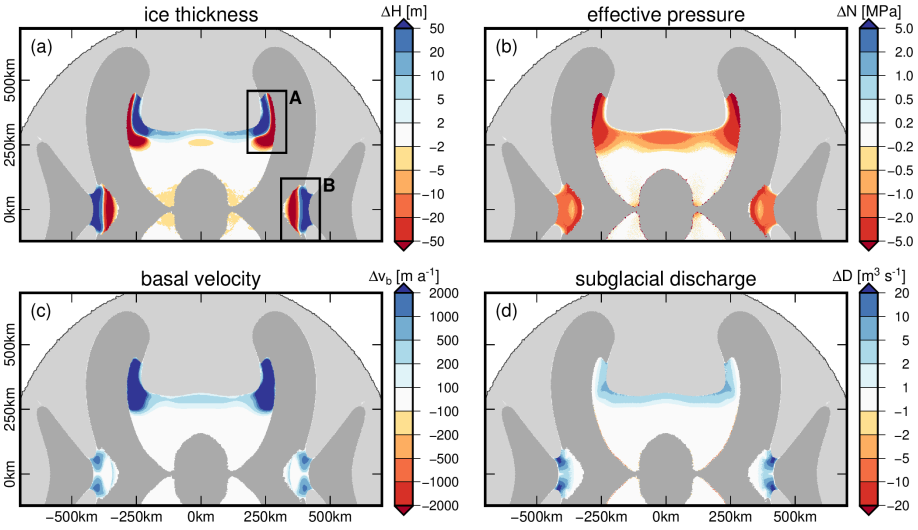

Figure 7. Result of ~~a fully~~ the one-way coupled ~~run~~ simulations using the ~~CalvingMIP~~ Thule setup~~with an additional seasonal water source, comparable to surface melt in the summer in Greenland~~. The panels a–d show the ~~difference~~ differences of ~~some selected fields~~ the summer state (day 215 in the second year) with respect to the ~~control simulation. Both at the~~ end of the ~~20-year simulation~~spin-up. The ~~hashed areas indicate~~ panel layout is identical to Fig. 5, with ice-dynamical fields in the ~~parts of~~ left column (a,c) and subglacial hydrological fields in the right column (b,d). Differences are evaluated on the common grounded domain~~where basal temperatures from ISSM~~ . The maps are ~~below~~ cropped to the ~~melting point and~~ northern part, the ~~basal melt rate~~ missing southern part is ~~zero~~symmetrical. Rectangles (A, B) indicate regions of high activity referenced in the text.

which usually have much more complexity. We have analyzed the performance of initialization and data mapping as well as scaling with the number of processes.

3.2.1 Experimental setup

The setup for ISSM is mostly the same as G1000 in Fischler et al. (2022). The only modification we made for this paper is to use the default ISSM Budd friction law ~~$\boldsymbol{\sigma}_b = -C^2 N^{\frac{q}{p}} |\boldsymbol{u}_b|^{\frac{1}{p}-1} \boldsymbol{u}_b$~~ $\boldsymbol{\sigma}_b = -C^2 N^{\frac{q}{p}} |\boldsymbol{v}_b|^{\frac{1}{p}-1} \boldsymbol{v}_b$ with coefficient $C = 12.94$, exponents $p = 3$ and $q = 2$, and effective pressure N provided by CUAS. The ~~spinup~~ spin-up process is the same as in Fischler et al. (2022). During the ~~spinup~~spin-up, the effective pressure is $N = N_{\text{opc}} = \rho_{\text{ice}} g H - \max(0, \rho_{\text{water}} g(-z_b))$ with z_b the height of the base above sea level as in Wolovick et al. (2023).

~~The setup for CUAS~~ The setup set the lower bound for effective pressure 35 % of ice overburden pressure. This parameter is likely to affect the stability of the coupled system, but investigating its effect in detail is beyond the scope of this work. The resolution of the mesh is around 0.7 km at the shear margins, 10 km at slow-moving ice in the interior, and up to 100 km at open ocean.

The CUAS setup is taken from Fischler et al. (2023). We use ~~G500~~G600, which has ~~approximately the same resolutionas~~ a uniform grid with 600 m resolution, similar to the minimum element size of the ISSM ~~setup. Ice~~ mesh. We use daily surface

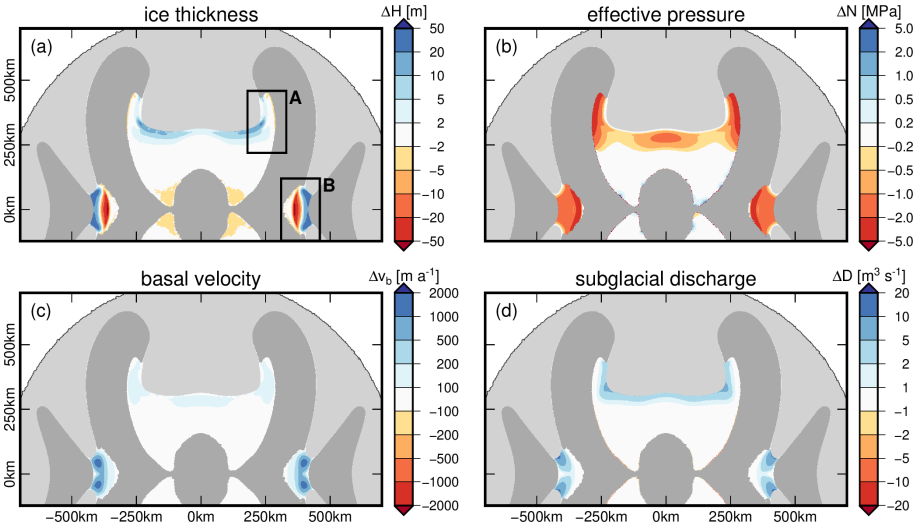

Figure 8. Result of the fully coupled (two-way) simulations using the Thule setup. The panels a–d show the differences of the summer state (day 215 in the second year) with respect to the end of the spin-up. The panel layout is identical to Fig. 5, with ice-dynamical fields in the left column (a,c) and subglacial hydrological fields in the right column (b,d). Differences are evaluated on the common grounded domain. The maps are cropped to the northern part, the missing southern part is symmetrical. Rectangles (A, B) indicate regions of high activity referenced in the text.

runoff data (R) from the regional climate model RACMO (Noël et al., 2019) for the year 2019 in addition to the ice sheet basal melt ~~from ISSM is the only water source for CUAS. No basal temperature constraint is applied to the CUAS mask (`bndMask`) to allow for active subglacial hydrology everywhere under the ice. This is needed because ISSM allows for basal sliding independent of basal temperature and thus requires getting the effective pressure everywhere. The layer thickness~~ (M) from ISSM to impose seasonality to the CUAS water source as:

$$Q(x,y,t) = M(x,y,t) + 0.9R(x,y,t), \tag{1}$$

where we only consider 90% of the surface runoff to enter the subglacial system. The percentage is arbitrary, but ensures strong seasonality for the coupled simulations. In the year 2019, particularly high melt was measured for the Greenland Ice Sheet (Tedesco and Fettweis, 2020; Sasgen et al., 2020). The aquifer layer thickness in CUAS was chosen equal to 1 m, ~~and no transition between confined and unconfined. The~~ with yield storativity $S_y = 10^{-2}$ ~~.~~ and minimal transmissivity $T_{\min} = 10^{-14}\,\mathrm{m^2/s}$. Subglacial channel creep opening/closure is parameterized using an ice flow rate factor of $A = 6.8 \times 10^{-24}\,\mathrm{Pa^{-3}\,s^{-1}}$ (Werder et al., 2013). All remaining parameters are unchanged from those reported in Fischler et al. (2023). For defining the initial active mask in CUAS, a threshold of minimum 10 m of ice thickness was chosen. The temperature threshold for activating the mask during simulation is set to 269.15 K. Similar to ISSM, the initial hydraulic head of CUAS is derived from N_{opc}.

~~As experiments, we ran both serial and parallel coupling and increase the number of processes used. We used the results of (Fischler et al., 2022) and (Fischler et al., 2023) to find ranges of process counts where we can reasonably expect the coupled~~

~~system to scale efficiently. Apart from the coupling scheme, the coupling~~ The preCICE coupling configuration is the same as in Sect.~~3.1~~ 3.1 with the exception of using different mapping methods for Sect. 3.2.4 and different coupling schemes for Sect. 3.2.3. Simulations run for 730 days.

~~For serial coupling, both participants run on the same nodes/CPUs and alternate in using all available processes. This minimizes idle time of CPUs. However, this may not work on every system: it may be necessary to allocate separate nodes for each participant to avoid deadlocks in communication.~~

~~For parallel coupling, we started with a distribution of two ISSM processes for each CUAS process ($n_{\mathrm{ISSM}} : m_{\mathrm{CUAS}} = 2.0$). Early testing suggested this would lead to approximately equal execution times for each solver and therefore minimize wait times. Based on the results of these experiments, as the solvers exhibit different scaling characteristics, we also tried other distributions of processes for a few selected total process counts.Note that this flexible allocation of CPUs would not be possible without significant development effort if CUAS was integrated into ISSM directly.~~

~~We ran the coupled setup for 18 coupling windows of 30 simulated dayseach, averaging runtimes over five runs. The remaining experiments with parallel coupling and all experiments with serial coupling and are run only once, as randomness in execution times is accounted for by the long runtime and averaging over coupling windows.~~

~~We measure the execution time required for each coupling window and how much time is used by each solver and preCICE, respectively.~~ For the measurements, we are using the profiling utility integrated into preCICE[8]. ~~In order to achieve representative results,~~ Runtime of initialization is measured at the beginning. Where we report aggregate runtime of solvers or data mapping in the following sections, only the final 365 days are included. This ensures that the analysis includes a full year of seasonal changes but excludes the noisy early-coupling windows in which the states of ISSM and CUAS are not yet aligned. In real use, depending on which parameters are changed between runs, this relaxation phase could be skipped entirely by using restart files. Runtime measurements do not include writing output, as I/O execution times can swing wildly and unpredictably. Adding moderate amounts of data output to both participants should not, on average, significantly impact the analysis. Every experiment is repeated three times, and the results are averaged to reduce variance.

~~Note that serial and~~ In parallel coupling schemes~~also give different numerical results. Serial coupling is generally more accurate as results from~~ , every participant has its own exclusive resources. CPUs are assigned according to a ratio $\alpha = n_{\mathrm{ISSM}}/n_{\mathrm{CUAS}}$ where n_{S} is the number of CPUs assigned to solver S. Care is taken to assign packed CPU blocks to participants, i.e., processes used by one participant are ~~immediately used by the other participant in the same coupling time window. However, careful investigation of simulation results will be part of future work. Here, we only made sure that both participants use comparable resolution and precision, with a relative solverresidual of $\mathrm{rtol} = 10^{-6}$.~~ as local to the nodes as possible, since each participant communicates more internally than with other participants.

~~The~~ In serial coupling schemes, we conducted experiments with both shared and exclusive resources. Shared resources are used by both solvers in turn. Exclusive resources are used by only one solver and are idle while the other solver computes. That means except during initialization, half of the CPUs will be almost completely idle. No energy is wasted, but the cluster resources are nonetheless reserved for the entire run. With shared resources, both solvers always use all allocated CPUs, as it is

[8] ~~https://precice.org/tooling-performance-analysis.html~~https://precice.org/tooling-performance-analysis.html

generally recommended to use all allocated cluster resources unless the region of flat or negative scaling is reached. For better comparisons, we also allocate the same number of CPUs to both solvers when using exclusive resources.

Where we give CPU or process counts below, they have the following meanings for the coupling and resource allocation schemes:

- Parallel: total number of processes allocated to both solvers, each solver using a fixed subset.
- Serial (shared): total number of processes allocated, used by both solvers in turn.
- Serial (exclusive): maximum number of processes allocated to either solver. Since the processes of the other participant are mostly idle, this results in the most relevant comparison regarding the required resources and it allows us to measure the effect of resource contention. As both solvers use the same number of processes, the number of processes given in plots is half the number of processes allocated.

All experiments are performed on the Albedo cluster of the Alfred Wegener Institute. The compute nodes of the cluster are equipped with 256 GB of RAM and two AMD Rome Epyc 7702 CPUs for a total of 128 CPU cores per node and are connected by 100 Gb InfiniBand network. The solvers and dependencies are built with GCC version 12.1.0 and OpenMPI version 4.1.3. ~~Care is taken to assign packed blocks of CPUs to participants, i.e., processes used by one participant are as local to nodes as possible, since each participant communicates more internally than with the other participants.~~

3.2.2 ~~Results~~Initialization

~~Timeline of two representative experiments with parallel coupling scheme and different numbers of processes (given in parentheses for each solver). Significant times spent in *Advance* indicate wait times where one participant finishes a coupling window before the other. In the first few coupling windows, ISSM's stress balance solver struggles to converge with effective pressure from CUAS. With few processes, wait times are balanced beyond the first coupling windows. With increasing numbers of processes, CUAS needs to wait more often. Initialization time of solvers and preCICE is negligible relative to the simulation.~~ In earth system models, initialization of the coupled setup is often a significant part of the runtime, see for example the runtime analysis of the OASIS3-MCT coupling library in Craig et al. (2017). For this scaling analysis, we use nearest neighbor mapping exclusively. Comparison of mapping methods follows in Sect. 3.2.4.

~~Timeline of two representative experiments with serial coupling scheme and different numbers of processes (given in parentheses for each solver). Like in **??**, a few coupling windows are required for stable solve times. During initialization, both participants compete for the shared CPUs, significantly increasing the time required relative to the simulation.~~

~~Figures **??** and **??** show the basic time line of coupling experiments with different numbers of processes and coupling schemes. The four basic components of coupling are: Initializing the solver, including loading input data and preparing initial coupling data and mesh for preCICE. Initializing preCICE, where the coupling framework builds its internal data structures for mapping and communication. This also includes wait times if one solver is initialized faster than the other. Running the solver. For legibility we do not display single time steps with preCICE calls in between, but all time steps of one coupling~~

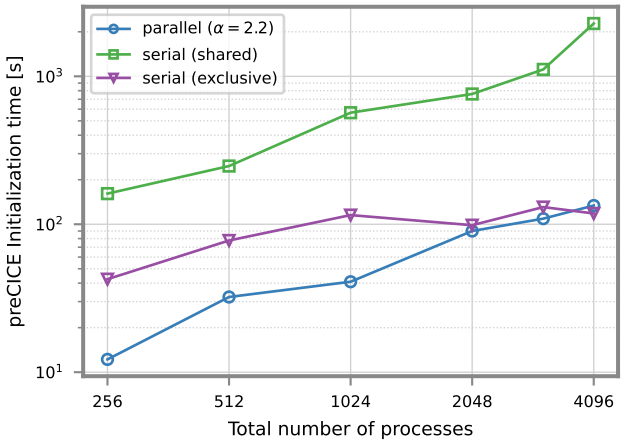

Figure 9. Time to initialize preCICE in the Greenland setup for different coupling and/or resource allocation schemes and increasing numbers of processes. Initialization includes partitioning of the coupling mesh and computation of nearest neighbor mapping weights for both solvers.

~~window in aggregate. No relevant information is lost by that since the ISSM adapter calls preCICE only once per coupling window anyway and for CUAS~~ Fig. 9 shows the time required to initialize preCICE. This includes both partition of the domain and computation of mapping weights. For comparison, Fig. 10 and 9 show the time required to initialize the solvers. Note that ISSM and CUAS initialize at the same time, so the actual time required is the ~~overhead of calling preCICE in the middle of a coupling window is negligible (on the order of 10 ms per time step.) Advancing to the next coupling window and sending the data between participants. This mainly consists of wait times if one solver finishes the computation of one coupling window before the other.~~ maximum of both. While the data is quite noisy due to a low number of repetitions and high variance of I/O, general trends can be identified.

~~During initialization of solvers and preCICE in parallel coupling, participants have dedicated compute and~~

ISSM initialization time grows linearly with the number of processes. CUAS is basically constant, trending slightly downward. Solver initialization is dominated by CUAS at low CPU counts. The larger mesh means higher data input requirements. ISSM catches up between 1024 and 2048 processes due to less optimal parallel I/O ~~resources. In serial coupling, both participants compete for the same resources, leading to much increased initialization times . Initialization, i.e., everything that happens before the participants or the first participant in a serial coupling , also does not scale well. Execution times do not change much with moreprocesses. These effects may not be representative of how other clusters handle resource contention, specifically~~ accesses.

In the range of CPU counts tested, preCICE initialization times also grow linearly with the number of processes, as more communication is necessary to partition the meshes and compute weights. Initializing the coupler and initializing the solvers take similar time except when sharing resources.

Sharing resources during initialization in serial coupling has a strong negative effect. Naive estimation would suggest the required time doubles for shared resources, since the same amount of work is being done by half the CPUs. However,

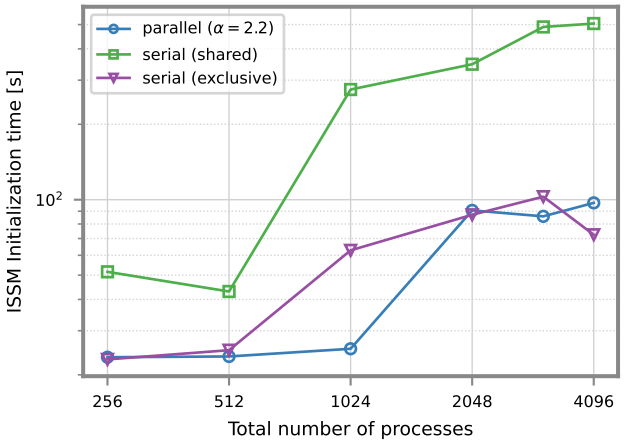

Figure 10. Time to initialize the ISSM Greenland setup for different coupling and/or resource allocation schemes and increasing numbers of processes. Initialization consists mostly of loading of input data from files.

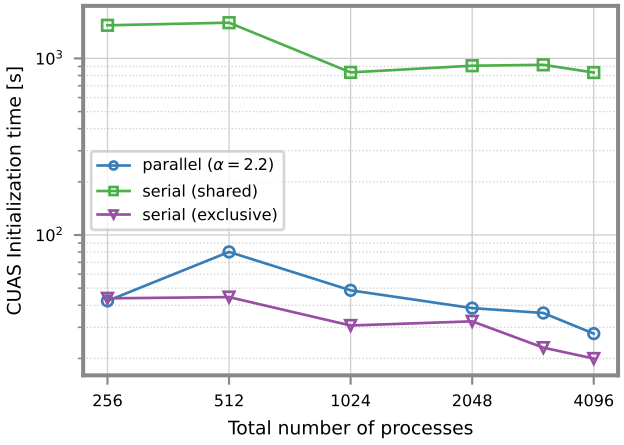

Figure 11. Time to initialize the CUAS Greenland setup for different coupling and/or resource allocation schemes and increasing numbers of processes. Initialization consists mostly of loading of input data from files.

scheduling conflicts can increase times by an order of magnitude or more. CUAS is especially badly affected, probably due to contention of I/O ~~.~~ resources. Parallel and serial exclusive require approximately the same amount of time. The increased effort required to compute and communicate the weights of a more partitioned mesh is counterbalanced by the increased resources.

~~The first few coupling windows are clearly dominated by the execution time of ISSM. The stress balance solver struggles to converge, often hitting the maximum number of iterations. The effective pressure delivered by CUAS is too different from the effective pressure used during to set up the ISSM model. To some degree this is unavoidable, but it can probably be reduced with careful construction of compatible solver setups.After a few coupling windows, execution times stabilize.~~

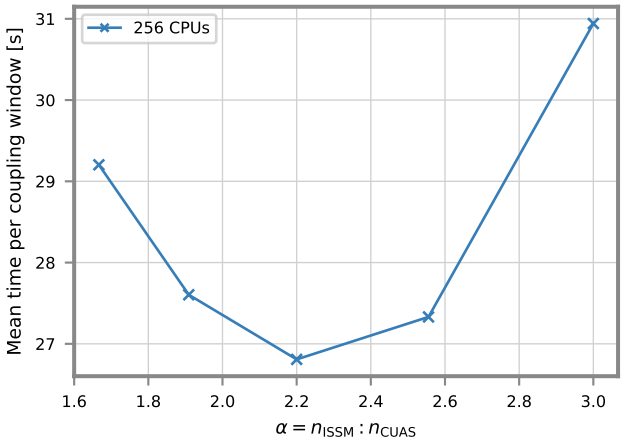

Figure 12. Average time to compute one coupling window with parallel coupling scheme in the Greenland setup for different distributions of 256 total CPUs.

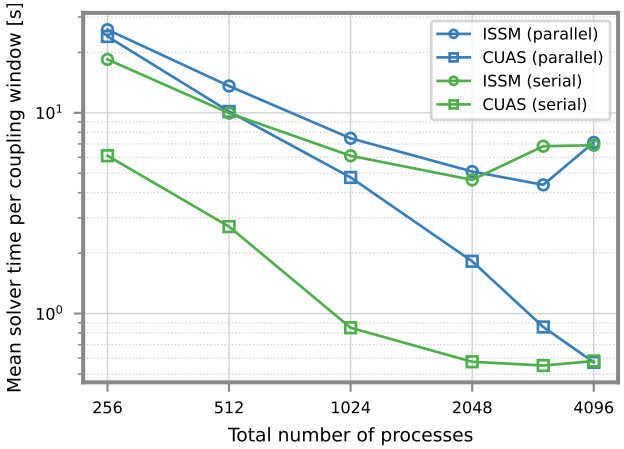

Figure 13. Average execution time required by the CUAS and ISSM solver during one coupling window for CUAS and ISSM in the Greenland setup for increasing numbers of MPI processes.

3.2.3 Simulation

~~The following analysis ignores the initialization phase and the first six coupling windows and only includes the remaining 12~~

For parallel coupling, we first needed to determine the optimal distribution of available CPUs among the solvers. This is the ratio of CPUs used by ISSM and CUAS where both solvers take approximately the same time for one coupling window. Fig. 12 shows the times required to compute a single coupling window with different distributions of 256 total CPUs. The best result is at $\alpha = 2.2$. If both participants scale equally, this would be the ideal distribution for larger numbers of CPUs as well.

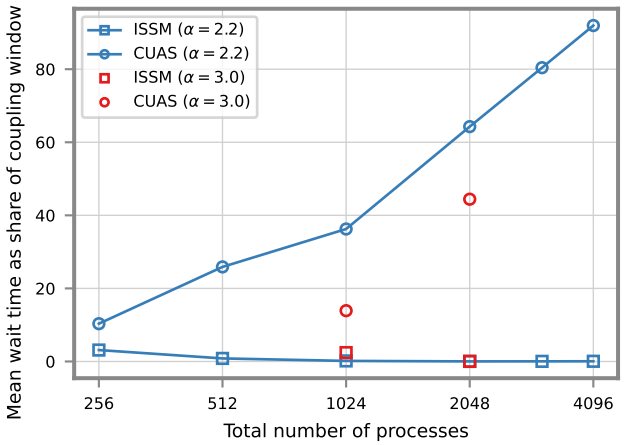

Figure 14. Average wait time relative to execution time of one coupling window for CUAS and ISSM in the Greenland setup in a parallel coupling scheme for increasing numbers of MPI processes and two different distributions of processes to participants.

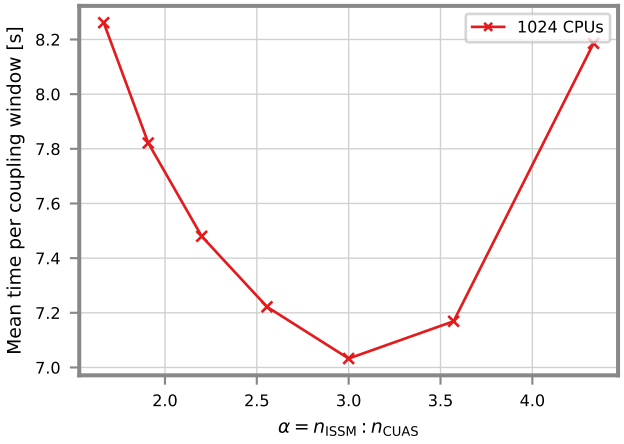

Figure 15. Average time to compute one coupling window with parallel coupling scheme in the Greenland setup for different distributions of 1024 total CPUs.

However, as Fig. 13 displays, ISSM does not scale as well as CUAS. This is consistent with the scaling analyses of standalone ISSM (Fischler et al., 2022) and CUAS (Fischler et al., 2023) as well, even though the uncoupled and coupled solvers are not directly comparable. Accordingly, with increasing numbers of processes, the duration that CUAS waits for ISSM increases as shown in Fig. 14. At 256 processes, where the distribution of $\alpha = 2.2$ had the best result, the wait times are close. Note that even here, the wait time is not zero, as the computational effort of each solver varies with the seasons, and both solvers wait in some coupling windows. ~~This covers 360 simulated days. Simulations usually run for longer, so any one-time effects at the beginning of the simulation do not significantly impact the overall runtime .~~

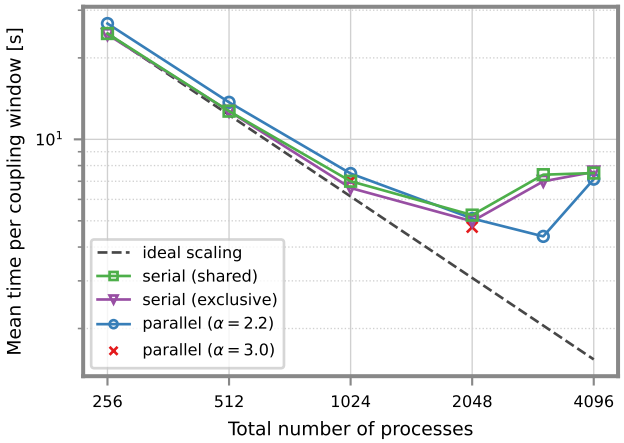

Figure 16. Average execution time for one coupling window (includes solver~~and~~ , communication, and data mapping) using a parallel or serial coupling scheme in the Greenland setup for ~~different~~ increasing numbers of MPI processes. Parallel coupling is tested with different distributions of processes to participants, serial coupling with exclusive or shared resources.

These results suggest that more than 256 CPUs should be distributed differently. Fig. 15 shows that for 1024 processes, the ideal distribution would be $\alpha = 3.0$. Fig. 14 shows reduced wait times for this distribution over $\alpha = 2.2$. Note that when running the simulation once at any moderate distribution α and measuring the average solver runtime t_X, the equation $\alpha_{opt} = t_{ISSM}/t_{CUAS} * \alpha$ gives a decent approximation of the optimal distribution. For example, with $\alpha = 2.2$ and the measured solver times $t_{ISSM} = 7.5\,s$ and $t_{CUAS} = 4.8\,s$, the estimated optimal distribution is $\alpha_{opt} = 3.4$, which is close to the minimum in Fig. 15.

~~Figure 16 shows the comparison of scaling between the serial or parallel coupling scheme. In the range where both parallel and serial coupling exhibit almost perfect strong~~ With these preparations, we ran scaling experiments of the coupled system. Fig. 16 shows a comparison of average run times for a coupling window for parallel and serial coupling schemes with different resource allocations.

In the lower range of processes, where both serial and parallel show almost ideal scaling, serial coupling is ~~faster since it is not possible to entirely eliminate wait times in parallel coupling due to uneven execution times of the solvers~~slightly faster. There is an unexpected bend in the measured times for serial coupling. The same bend is observed in the ISSM solver execution times, but not in CUAS. We were unable to identify the root cause or any difference in the simulation that could explain it, the numbers of solver iterations are similar between all experiments. A deeper analysis of the ISSM solver is beyond the scope of this work. It is possible that a technical aspect of the cluster interferes in the simulation with the most processes.

Exclusive resources for serial coupling give a small advantage over shared resources, but probably insignificant compared to the difference in initialization times discussed in Sect. 3.2.2. Scaling is basically identical. The changed distribution of CPUs in parallel coupling also gives a small advantage over the original distribution.

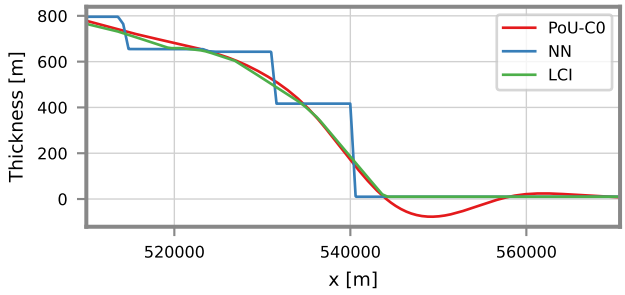

Figure 17. Results of mapping ice thickness from ISSM to CUAS mesh in the Greenland setup using different mapping methods along a horizontal section that crosses the ice boundary at $y = -1277863.4\,\mathrm{m}$. Partition-of-unity radial basis functions (PoU-C0) overshoots and produces unphysical values at the ice boundary. Linear cell interpolation (LCI) is equivalent to the P1 finite elements used by ISSM. Nearest neighbor (NN) is constant in regions around ISSM vertices.

~~Average execution time required by the CUAS and ISSM solver during one coupling window for CUAS and ISSM for different numbers of MPI processes.~~

3.2.4 Mapping

~~Average communication time relative to execution time of one coupling window for CUAS and ISSM in a parallel coupling scheme for different numbers of MPI processes and two different distributions of processes to participants.~~

~~Figure 13 shows the time required by the solvers to finish one~~ preCICE offers a choice of different methods for mapping data between meshes. The methods differ in the order of approximation and computational cost. We ran experiments with a selection of three methods of different order:

- nearest neighbor (NN), a first order projection method
- linear cell interpolation (LCI), a second order projection method
- partition-of-unity radial basis functions (PoU), a kernel based method designed for large scale mappings

We tested all mapping methods with a parallel coupling scheme on 1024 total CPUs, distributed $2.2 : 1$ to ISSM and CUAS and measured the initialization time and mapping time per coupling window. ~~CUAS scales very well in the entire tested range , whereas ISSM only scales close to perfect strong scaling up to around 768 processes. As shown~~

The PoU method and basis functions must be parameterized[9]. Our setup is particularly challenging for the current, relatively recent implementation of PoU in preCICE. As described in Sect. 3.2.1, the ISSM mesh has a wide range of element sizes. However, preCICE currently permits only a single global parameter each for the basis function radius and PoU cluster size. In addition, the fields of an ice sheet model include discontinuities at the calving front and grounding line. Even with the best parameterization we found (C0 compact polynomial basis functions, radius $200\,\mathrm{km}$ for mapping from ISSM to CUAS,

[9] https://precice.org/configuration-mapping.html

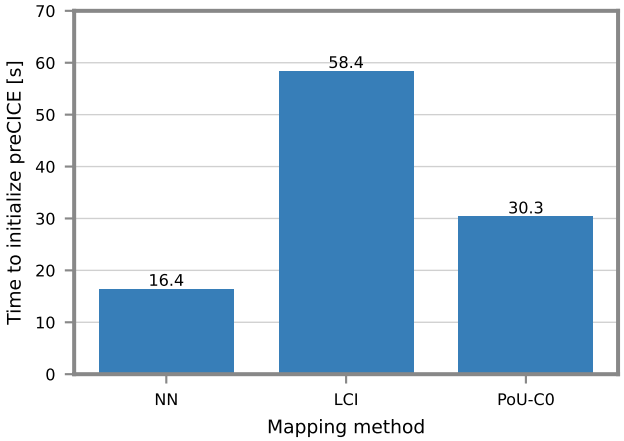

Figure 18. Time required to initialize preCICE for both coupling participants in the Greenland setup for different mapping methods: nearest neighbor (NN), linear cell interpolation (LCI) and partition-of-unity radial basis functions with C0 compact polynomial basis (PoU-C0). Both participants use the same mapping method, mapping is computed by the reading participant.

10 km from CUAS to ISSM, default values for the other parameters) there is overshooting, e.g., negative ice thickness values all around the ice boundary as in Fig. ~~14, this leads to increasing wait times for CUAS in parallel coupling. In both cases, execution time of ISSM is the limiting factor.~~ 17. Performance measurements are included here with the hope that improvements in preCICE and/or the setup will be implemented. The other methods did not show such problems, detailed analysis of numerical errors is beyond the scope of this work but is covered in Chourdakis et al. (2022), Schneider and Uekermann (2025), and Hocks and Uekermann (2026).

~~An attempt was made to improve execution times of parallel coupling by minimizing wait times by assigning more processes to ISSM. For 192 processes, a ratio of 2.0 ISSM processes per CUAS process balances wait times well. For 768 total processes, a ratio of 2.7 works better~~

Fig. 18 shows the total runtime required to initialize preCICE for both participants. This time includes both communication, mesh partitioning, and initializing data mappings. For all methods, initialization of preCICE is around the same order as initialization of the solvers (see Fig. 10 and 11). LCI is about twice as expensive as PoU and almost four times as expensive as NN.

Fig. 19 shows the runtime required to map the data in each coupling window. The cost difference between the mapping directions (ISSM to CUAS or CUAS to ISSM) can be attributed to the number of fields that must be mapped. NN and LCI are very close in cost, whereas PoU is significantly more expensive. However, for all mapping methods, the runtime of a coupling window (see Fig. ~~14.) However, this leads to only slightly lower execution times. Giving more resources to ISSMwhen it is at its scaling limit does not seem to significantly improve performance. At 1536 processes, lowering the wait times by redistributing processes actually has a slightly negative effect. These are the experiments that were repeated five times to increase confidence in the results. There is no large effect in either direction.~~ 16) is dominated by the solvers themselves. NN was measured at

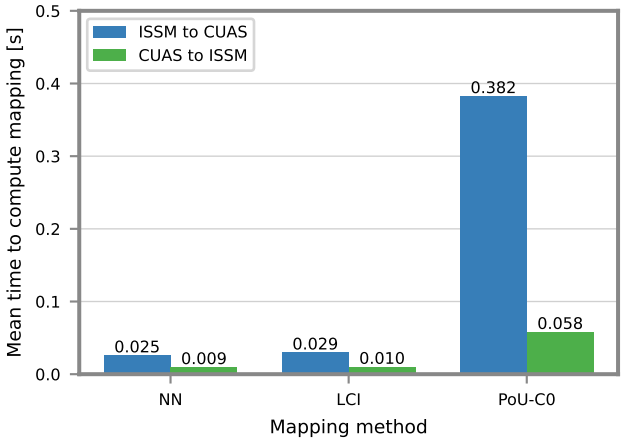

Figure 19. Time required to compute data mapping per coupling window in the Greenland setup for different mapping methods: nearest neighbor (NN), linear cell interpolation (LCI) and partition-of-unity radial basis functions with C0 compact polynomial basis (PoU-C0). Mapping is computed by the reading participant.

approximately 9 ms from CUAS to ISSM, 25 ms from ISSM to CUAS, or ~ 0.1 % and ~ 0.3 % of the total time of a coupling window. LCI was measured only slightly slower, at 10 ms (~ 0.1 %) and 29 ms (~ 0.4 %). Note that the difference is too small to be reliably measured with this setup. The cost of PoU on the other hand is not as negligible, with 58 ms (~ 0.7 %) and 380 ms (~ 6 %). The detailed profile shows significant imbalance in the distribution of work, probably due to the uneven mesh resolution. The slowest processes work up to four times longer than the fastest.

4 Discussion

~~We have presented the technical aspects of a new framework for coupling a subglacial hydrology model to the ice sheet model ISSM~~The new coupling framework presented in this paper is a promising approach for the development of coupled ice sheet and earth system models. Compared to SHAKTI (Sommers et al., 2018), DoCo (de Fleurian et al., 2016), and other hydrology models directly integrated into ISSM, external coupling ~~with preCICE allows greater~~ (with preCICE or another library, see discussion below) allows more choice of implementation and ~~in the~~ numerical treatment. ~~CUAS and ISSM use different meshes, time steps, and numerical methods.~~ Development of the models can progress independently and neither model is restricted by the choices of the other. The ~~setup for~~ effort to set up the coupling is minimal. Coupling scheme, data mapping, and coupling time window are easy to adapt ~~.~~and extend to additional participants.

~~Our performance experiments show that preCICE coupling does not negatively affect scaling and has minimal computational overhead. Therefore, our choice of a coupling window size of 30 d is excessively conservative and coupling on the same time scales as the ISSM time step is feasible. While not a comprehensive review of all possible setups, the scaling results gives future users a basis for how to run the coupling efficiently. Parallel coupling widens the range of CPU counts that can be used~~

efficiently. On the other hand, imbalance in solver run times adds overhead. Therefore, serial coupling is faster for lower CPU counts. This imbalance can not be significantly reduced by redistributing CPUs.

This approach offers many opportunities for the development of ice sheet and earth system models. Many independent models can be more easily recombined to try different setups than tightly integrated models. Rather than reimplementing the CUAS model, only a simple reusable preCICE adapter needs to be developed to add CUAS's capabilities to another ice sheet model . It also would be straightforward to add an ocean circulation model to the experiments presented here. We plan to use preCICE to couple ISSM with a new code for capturing calving dynamics in future work.

The ISSM-preCICE adapteraims to be generic, but still has limitations that need ISSM-preCICE adapter provides the basic features necessary for most use cases. It significantly reduces the effort required to couple other codes to ISSM, but it will almost certainly become necessary to allow the adapter to be customized more, e.g., to add logic specific to each coupling setup such as deriving model variables as we already had to do in the CUAS adapter. Additionally, we listed some general limitations to be resolved in future work as described in Sect.2.3Sect.Accuracy 2.3. For example, accuracy is lost because the coupling interface is based on the mesh vertices instead of finite element nodes. And , and it is not yet possible to couple the full three-dimensional volume of the ice sheet. Some features of ISSM and preCICE are not supported at this stage. , among them dynamic adaptive meshes and time interpolation. It was not necessary to modify the code of ISSM to implement the adapter, but future extensions may require such changes.

The CUAS-preCICE adapter is an early prototype and does not yet follow most of the guidelines for preCICE adapters. Specific adaptations had to be made to CUAS to support the coupling with ISSM. These can be more generalized to open the adapter for different use cases. In addition, the adapter needs to be cleanly integrated into the code, as CUAS-MPI itself aims for a high degree of software quality.

preCICE has We have demonstrated the stability and functionality of the coupled system in Sect. 3.1. Within both 1-way and 2-way coupling, the fields for N, v_b, T_e and H are consistent, indicating that the hydrological system is well set up. Our results reveal that the differences between the physical fields in 1- and 2-way coupling are significantly larger with high seasonal forcing. The larger the runoff, the more important is a multi-physics approach. The more realistic feedback is found by using 2-way coupling, as the two systems adapt to a joint state with reasonable physical fields in both systems. The time series (Fig. 6) exemplifies that the winter state in the runoff forced case has a higher $\overline{\log(T_\mathrm{e})}$ in winter than without seasonal runoff, which is reasonable. The lower $\overline{\log(T_\mathrm{e})}$ in the 2-way coupling corresponds to lower $\bar{N}$ in the high runoff forcing and vice versa, as N also influences the evolution of T_e. To summarize, the behavior of the ice and hydrology systems are reasonable, the simulated feedback is expected, but the drastic difference in 2-way coupling was not expected to this extent. As this paper focuses on the technical aspects of the coupling, we have not quantified the numerical accuracy of the coupled system nor fully demonstrated its ability to represent real-world cases. We have also not fully explored all preCICE features. In particular, informed choices need to be made regarding the optimal coupling scheme and the data mapping method.

Our performance experiments in Sect. 3.2 show that preCICE coupling does not negatively affect scaling during the simulation. Data mapping has negligible computational overhead, enabling coupling windows at least as short as the ISSM time step.

preCICE time interpolation can be used in the future to improve temporal resolution if necessary, as shortening the ISSM time step would be expensive. Adaptive coupling windows would also be an option, but are not currently supported by preCICE.

Serial coupling is around 5 % faster than parallel coupling up to 1024 processes. There is tentative evidence that parallel coupling widens the range of CPU counts that can be used efficiently. On the other hand, an imbalance in scaling of the solvers limits scaling of the coupled system. Reducing the imbalance by redistributing CPUs does not significantly improve overall execution times because it involves reassigning CPUs from well-scaling CUAS to worse-scaling ISSM. If there are many runs with the same number of CPUs and similar setups, it may be worth searching for the optimal distribution. Otherwise, an approximate solution is sufficient and probably uses fewer resources overall. It is also technically easier to allocate a distribution where every solver is assigned entire cluster nodes.

The cost of coupling and data mapping is more significant during initialization than during the actual simulation. Partitioning the coupling mesh and computing the mapping weights takes about the same amount of time as initializing the solvers themselves. The required time depends on the chosen mapping method. Linear cell interpolation, a second order method, is twice as expensive as first-order nearest neighbor interpolation. The partition-of-unity method shows promising performance for second-order (or higher) mapping, particularly the initialization times, but algorithmic improvements and special handling of discontinuous fields (e.g., specially adapted meshes) are required for it to be suited for our setup.

The performance results give future users a basis for how to run the coupling efficiently. The relative solver and coupling execution times depend on too many factors (meshes, parameters, tolerances, etc.) to cover here. But estimates can be used to translate our results to other setups. Unlike the computation during simulation, initialization takes more time with added processes. Therefore, the length of the simulation, and with it the balance of initialization and simulation time, need to be considered when choosing which how to allocate resources.

preCICE has so far not been widely used in the earth system modeling community. While few (if any) ready-to-use adapters relevant to earth system modeling exist at the moment, developing such adapters is a current opportunity. Besides ESM-specific codes, models in frameworks like OpenFOAM, Elmer, or FEniCS/FEniCSx, for which preCICE adapters already exist, can be coupled with minimal effort. This makes these frameworks a good choice for the development of new models. For example, the ice sheet code Elmer/Ice could easily be coupled to CUAS-MPI using the adapters presented in this work. ~~It is an open question whether a generic coupling library like preCICE has significant downsides over a~~ As mentioned in the introduction, Elmer/Ice already includes a different hydrology model, but being able to compare different approaches is highly valuable as demonstrated by the various model intercomparison projects.

Without implementing comparable adapters for the same models, it is not possible to make a strong determination whether to prefer preCICE over a coupling library like YAC (Hanke et al., 2016) ~~that is~~ or OASIS3-MCT (Craig et al., 2017) that are specialized on earth system models~~, at least in global setups. To answer this question, we are in the process of evaluating the data mapping between global meshes using preCICE in a Cartesian coordinate system.~~

~~This paper was focused on the technical aspects of the coupling and has neither quantified the numerical accuracy of the coupled system nor fully demonstrated its capabilities to represent real world cases. We have also not fully explored all preCICE features. In particular, we will need to make qualified choices for the coupling scheme and data mapping method~~. We believe

preCICE is at least on par regarding the basic functionality. The software is mature. Basic adapters can be developed quickly, but of course a generic adapter for a large model like ISSM still requires significant effort to build and maintain. The adapters that we presented here as well as the other existing adapters show that preCICE is compatible with any solver architecture. The regridding benchmark by Hocks and Uekermann (2026) already demonstrated that numerical performance is equivalent. The results in Sect. 3.2 are in line with those reported for other libraries (Craig et al., 2017; Hanke et al., 2016), but a full benchmark – maybe an extension to the aforementioned regridding benchmark – with identical meshes would be necessary to get an accurate picture. The features of preCICE (mapping methods, cartesian coordinates) are well suited for the use case discussed here. However, we have identified several missing features (spherical coordinate systems, masking, specialized mapping methods) that may be highly relevant for other ESM applications. Further development is probably needed to satisfy the requirements of the community. The main advantage of preCICE due to its generic nature is the potentially much larger community, collaborating in the development of adapters and the coupling library itself.

5 Conclusions

In this paper, we presented the software for coupling the ice sheet model ISSM to the subglacial hydrology model ~~CUAS. The coupling is~~ CUAS-MPI. The main work has been to develop the adapters for the models for the coupling library preCICE. The ISSM adapter is generic and supports other use cases such as ice-ocean coupling, but adaptations will probably be necessary in some cases. The CUAS-MPI adapter is still a prototype and specific to coupling with ISSM. Future development will focus on generalizing the adapter and better integrating it into the CUAS-MPI code.

The coupling is easy to use, adaptable, and extensible due to the generic coupling library~~preCICE~~. We have ~~given performanceresults that show low computational overhead for communication and data mapping . The parallel scaling results will inform~~ demonstrated its functionality in a synthetic setup to verify its correctness and stability. We have also analyzed different aspects of the system's performance, including initialization times, scaling, and mapping methods. The system scales well with the number of processes and the overhead for coupling is low. These experiments can inform the efficient use of the software in the future.

~~The new preCICE adapters~~ We were able to at least qualitatively compare preCICE to libraries specialized on earth system modeling. We found preCICE to be competitive in all aspects, at least for the current use case, but closer analysis will be required to give a definitive answer. We identified possible future improvements for preCICE to better serve this community. We also provided arguments that either approach is superior in the long run to integrating different models into the same monolithic code.

The new coupling will facilitate studies of the interaction between continental ice sheets and the hydrology systems underneath. ~~The~~ We will also use the generic ISSM-preCICE adapter ~~can also be used~~ in other setups. For example, we are developing a new solver for capturing the ice sheet calving fronts that can be coupled with ISSM to improve upon its existing moving front core. The use of preCICE to integrate ISSM into a global earth system model will be evaluated. Finally, the adapters can be extended to lift some of the limitations described in this paper and open even more use cases.

~~Multiphysics problems involving ice sheets have always been solved as split uniphysics problems to date. Keyes et al. (2013) suggests a 'coupled until proven decoupled' strategy for multiphysics problems. The coupling we presented here enables such an investigation.~~

We hope that the software we developed enables researchers to implement and test new ice sheet model capabilities more quickly. In general, researchers should consider preCICE coupling when developing new models or extending existing ones.

Code and data availability. The current version of the ISSM-preCICE adapter is available at https://git.rwth-aachen.de/terrabyte-dnn2sim/issm-precice. Version 0.4.0 of the ISSM-preCICE adapter used in this paper is available at https://doi.org/10.5281/zenodo.18846020 (Abele and Humbert, 2026). The current version of CUAS-MPI is available at https://github.com/tudasc/CUAS-MPI. Version 0.1 of CUAS-MPI with added preCICE adapter used in this paper is available at https://doi.org/10.5281/zenodo.18846076 (Fischler et al., 2026). Input data, scripts to run the experiments and produce the plots for all the simulations presented in this paper as well as results of performance measurements are available at https://doi.org/10.5281/zenodo.18846105 (Abele et al., 2026)

Author contributions. DA developed the ISSM-preCICE adapter. YF and TK developed CUAS-MPI and the CUAS-preCICE adapter with contributions by DA. TK and DA ran the experiments to test functionality. DA ran the experiments to measure performance. AH supervised the project. CB supervised YF, HJB and AB supervised DA. MM provided support in the use of ISSM. BU and GC provided support in the use of preCICE. DA prepared the manuscript with significant contributions by TK, AH, BU, GC, and MM. All authors commented on and approved all parts of the manuscript.

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

Acknowledgements. Part of this work was funded by HELMHOLTZ IMAGING, a platform of the Helmholtz Information & Data Science Incubator [grant number: ZT-I-PF-4-026].

We further thankfully acknowledge funding by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany's Excellence Strategy EXC 2075 – 390740016 and under project number 528693298, and the support by the Stuttgart Center for Simulation Science (SimTech).

The authors gratefully acknowledge the computing time provided to them on the high-performance computer Lichtenberg 2 at the NHR Center NHR4CES at TU Darmstadt under grant p0020118. NHR4CES is funded by the Federal Ministry of ~~Education and Research~~Research, Technology and Space, and the state governments participating on the basis of the resolutions of the GWK for national high performance computing at universities.

References

Abele, D. and Humbert, A.: ISSM-preCICE adapter, https://doi.org/10.5281/zenodo.18846020, Zenodo [code], 2026.

Abele, D., Kleiner, T., Fischler, Y., and Humbert, A.: Coupling ISSM and CUAS-MPI: example cases, https://doi.org/10.5281/zenodo.18846020, Zenodo [data set], 2026.

Beyer, S., Kleiner, T., Aizinger, V., Rückamp, M., and Humbert, A.: A confined–unconfined aquifer model for subglacial hydrology and its application to the Northeast Greenland Ice Stream, The Cryosphere, 12, 3931–3947, https://doi.org/10.5194/tc-12-3931-2018, 2018.

Chourdakis, G., Davis, K., Rodenberg, B., Schulte, M., Simonis, F., Uekermann, B., Abrams, G., Bungartz, H., Cheung Yau, L., Desai, I., Eder, K., Hertrich, R., Lindner, F., Rusch, A., Sashko, D., Schneider, D., Totounferoush, A., Volland, D., Vollmer, P., and Koseomur, O.: preCICE v2: A sustainable and user-friendly coupling library [version 2; peer review: 2 approved], Open Research Europe, 2, https://doi.org/10.12688/openreseurope.14445.2, 2022.

Chourdakis, G., Schneider, D., and Uekermann, B.: OpenFOAM-preCICE: Coupling OpenFOAM with external solvers for multi-physics simulations, OpenFOAM® Journal, 3, 1–25, https://doi.org/10.51560/ofj.v3.88, 2023.

Craig, A., Valcke, S., and Coquart, L.: Development and performance of a new version of the OASIS coupler, OASIS3-MCT_3.0, Geoscientific Model Development, 10, 3297–3308, https://doi.org/10.5194/gmd-10-3297-2017, 2017.

Craig, A. P., Vertenstein, M., and Jacob, R.: A new flexible coupler for earth system modeling developed for CCSM4 and CESM1, The International Journal of High Performance Computing Applications, 26, 31–42, https://doi.org/10.1177/1094342011428141, 2012.

de Fleurian, B., Morlighem, M., Seroussi, H., Rignot, E., van den Broeke, M. R., Kuipers Munneke, P., Mouginot, J., Smeets, P. C. J. P., and Tedstone, A. J.: A modeling study of the effect of runoff variability on the effective pressure beneath Russell Glacier, West Greenland, Journal of Geophysical Research: Earth Surface, 121, 1834–1848, https://doi.org/https://doi.org/10.1002/2016JF003842, 2016.

de Fleurian, B., Davy, R., and Langebroek, P. M.: Impact of runoff temporal distribution on ice dynamics, The Cryosphere, 16, 2265–2283, https://doi.org/10.5194/tc-16-2265-2022, 2022.

Desai, I., Scheurer, E., Bringedal, C., and Uekermann, B.: Micro Manager: a Python package for adaptive and flexible two-scale coupling, Journal of Open Source Software, 8, 5842, https://doi.org/10.21105/joss.05842, 2023.

Duchaine, F., Jauré, S., Poitou, D., Quémerais, E., Staffelbach, G., Morel, T., and Gicquel, L.: Analysis of high performance conjugate heat transfer with the openpalm coupler, Computational Science & Discovery, 8, 015 003, https://doi.org/10.1088/1749-4699/8/1/015003, 2015.

Fischler, Y., Rückamp, M., Bischof, C., Aizinger, V., Morlighem, M., and Humbert, A.: A scalability study of the Ice-sheet and Sea-level System Model (ISSM, version 4.18), Geoscientific Model Development, 15, 3753–3771, https://doi.org/10.5194/gmd-15-3753-2022, 2022.

Fischler, Y., Kleiner, T., Bischof, C., Schmiedel, J., Sayag, R., Emunds, R., Oestreich, L. F., and Humbert, A.: A parallel implementation of the confined–unconfined aquifer system model for subglacial hydrology: design, verification, and performance analysis (CUAS-MPI v0.1.0), Geoscientific Model Development, 16, 5305–5322, https://doi.org/10.5194/gmd-16-5305-2023, 2023.

Fischler, Y., Kleiner, T., Abele, D., and Humbert, A.: CUAS-MPI with adapter for the preCICE coupling library, https://doi.org/10.5281/zenodo.18846076, Zenodo [code], 2026.

Gagliardini, O., Zwinger, T., Gillet-Chaulet, F., Durand, G., Favier, L., de Fleurian, B., Greve, R., Malinen, M., Martín, C., Råback, P., Ruokolainen, J., Sacchettini, M., Schäfer, M., Seddik, H., and Thies, J.: Capabilities and performance of Elmer/Ice, a new-generation ice sheet model, Geoscientific Model Development, 6, 1299–1318, https://doi.org/10.5194/gmd-6-1299-2013, 2013.

Hanke, M., Redler, R., Holfeld, T., and Yastremsky, M.: YAC 1.2.0: new aspects for coupling software in Earth system modelling, Geoscientific Model Development, 9, 2755–2769, https://doi.org/10.5194/gmd-9-2755-2016, 2016.

Hocks, A. and Uekermann, B.: Evaluation of preCICE (version 3.3.0) in an Earth System Model Regridding Benchmark, EGUsphere, 2026, 1–19, https://doi.org/10.5194/egusphere-2025-5618, 2026.

Huang, Q., Abdelmoula, A., Chourdakis, G., Rauleder, J., and Uekermann, B.: CFD/CSD coupling for an isolated rotor using preCICE, 14th World Congress of Computational Mechanics and ECCOMAS Congress, https://doi.org/0.23967/wccm-eccomas.2020.081, 2021.

Ing, R. N., Nienow, P. W., Sole, A. J., Tedstone, A. J., and Mankoff, K. D.: Minimal Impact of Late-Season Melt Events on Greenland Ice Sheet Annual Motion, Geophysical Research Letters, 51, e2023GL106 520, https://doi.org/https://doi.org/10.1029/2023GL106520, e2023GL106520 2023GL106520, 2024.

Jordan, J. R.: Calving Model Intercomparison (CalvingMIP) wiki, https://github.com/JRowanJordan/CalvingMIP/wiki, [Accessed 2025-02-26], 2024.

Keyes, D. E., McInnes, L. C., Woodward, C., Gropp, W., Myra, E., Pernice, M., Bell, J., Brown, J., Clo, A., Connors, J., et al.: Multiphysics simulations: Challenges and opportunities, The International Journal of High Performance Computing Applications, 27, 4–83, 2013.

Khrulev, C., Aschwanden, A., Bueler, E., Brown, J., Maxwell, D., Albrecht, T., Reese, R., Mengel, M., Martin, M., Winkelmann, R., Zeitz, M., Levermann, A., Feldmann, J., Garbe, J., Haseloff, M., Seguinot, J., Hinck, S., Kleiner, T., Fischer, E., Damsgaard, A., Lingle, C., van Pelt, W., Ziemen, F., Shemonski, N., Mankoff, K., Kennedy, J., Blum, K., Habermann, M., DellaGiustina, D., Hock, R., Kreuzer, M., Degregori, E., and Schoell, S.: Parallel Ice Sheet Model (PISM), https://doi.org/10.5281/zenodo.14991122, Zenodo [code], 2025.

Larour, E., Seroussi, H., Morlighem, M., and Rignot, E.: Continental scale, high order, high spatial resolution, ice sheet modeling using the Ice Sheet System Model (ISSM), Journal of Geophysical Research: Earth Surface, 117, https://doi.org/10.1029/2011JF002140, 2012.

Larson, J., Jacob, R., and Ong, E.: The Model Coupling Toolkit: A New Fortran90 Toolkit for Building Multiphysics Parallel Coupled Models, The International Journal of High Performance Computing Applications, 19, 277–292, https://doi.org/10.1177/1094342005056115, 2005.

Mehl, M., Uekermann, B., Bijl, H., Blom, D., Gatzhammer, B., and van Zuijlen, A.: Parallel coupling numerics for partitioned fluid–structure interaction simulations, Computers & Mathematics with Applications, 71, 869–891, https://doi.org/10.1016/j.camwa.2015.12.025, 2016.

Noël, B., van de Berg, W. J., Lhermitte, S., and van den Broeke, M. R.: Rapid ablation zone expansion amplifies north Greenland mass loss, Science Advances, 5, https://doi.org/10.1126/sciadv.aaw0123, 2019.

Rodenberg, B. and Uekermann, B.: A waveform iteration implementation for black-box multi-rate higher-order coupling, https://doi.org/10.48550/arXiv.2511.07616, 2025.

Rodenberg, B., Desai, I., Hertrich, R., Jaust, A., and Uekermann, B.: FEniCS–preCICE: Coupling FEniCS to other simulation software, SoftwareX, 16, 100 807, https://doi.org/10.1016/j.softx.2021.100807, 2021.

Rüth, B., Uekermann, B., Mehl, M., Birken, P., Monge, A., and Bungartz, H.-J.: Quasi-Newton waveform iteration for partitioned surface-coupled multiphysics applications, International Journal for Numerical Methods in Engineering, 122, 5236–5257, https://doi.org/10.1002/nme.6443, 2021.

Sasgen, I., Wouters, B., Gardner, A. S., King, M. D., Tedesco, M., Landerer, F. W., Dahle, C., Save, H., and Fettweis, X.: Return to rapid ice loss in Greenland and record loss in 2019 detected by the GRACE-FO satellites, Communications Earth & Environment, 1, https://doi.org/10.1038/s43247-020-0010-1, 2020.

Schneider, D. and Uekermann, B.: Efficient Partition-of-Unity Radial-Basis-Function Interpolation for Coupled Problems, SIAM Journal on Scientific Computing, 47, B558–B582, https://doi.org/10.1137/24M1663843, 2025.

Slattery, S. R.: Mesh-free data transfer algorithms for partitioned multiphysics problems: Conservation, accuracy, and parallelism, Journal of Computational Physics, 307, 164–188, https://doi.org/10.1016/j.jcp.2015.11.055, 2016.

Sommers, A., Rajaram, H., and Morlighem, M.: SHAKTI: Subglacial Hydrology and Kinetic, Transient Interactions v1.0, Geoscientific Model Development, 11, 2955–2974, https://doi.org/10.5194/gmd-11-2955-2018, 2018.

Tang, Y.-H., Kudo, S., Bian, X., Li, Z., and Karniadakis, G. E.: Multiscale universal interface: a concurrent framework for coupling heterogeneous solvers, Journal of Computational Physics, 297, 13–31, 2015.

Tedesco, M. and Fettweis, X.: Unprecedented atmospheric conditions (1948 –2019) drive the 2019 exceptional melting season over the Greenland ice sheet, The Cryosphere, 14, 1209–1223, https://doi.org/10.5194/tc-14-1209-2020, 2020.

Totounferoush, A., Simonis, F., Uekermann, B., and Schulte, M.: Efficient and Scalable Initialization of Partitioned Coupled Simulations with preCICE, Algorithms, 14, https://doi.org/10.3390/a14060166, 2021.

Uekermann, B., Bungartz, H.-J., Cheung Yau, L., Chourdakis, G., and Rusch, A.: Official preCICE adapters for standard open-source solvers, pp. 210 – 213, GACM Colloquium on Computational Mechanics for Young Scientists from Academia and Industry, https://doi.org/10.18419/opus-9334, 2017.

Werder, M. A., Hewitt, I. J., Schoof, C. G., and Flowers, G. E.: Modeling channelized and distributed subglacial drainage in two dimensions, Journal of Geophysical Research: Earth Surface, 118, 2140–2158, https://doi.org/https://doi.org/10.1002/jgrf.20146, 2013.

Wolovick, M., Humbert, A., Kleiner, T., and Rückamp, M.: Regularization and L-curves in ice sheet inverse models: a case study in the Filchner–Ronne catchment, The Cryosphere, 17, 5027–5060, 2023.