

**NEMO Development Strategy 2023-2027**

**Version 3.0**

**4<sup>th</sup> November 2022**

**NEMO Developers Committee**

Introduction .....	3
Part I: Executive overview.....	4
1 What is NEMO ? .....	4
2 The drivers, requirements and vision for NEMO .....	6
3 How the NEMO teams work .....	11
4 Technical issues and opportunities with major impact on the code .....	14
5 Implementation Plan .....	20
Part II: Detailed Analyses .....	21
6 Team-work .....	21
7 Ocean model Kernel.....	27
8 Eddy closures .....	36
9 Surface fluxes and vertical mixing.....	50
10 Sea ice .....	57
11 Land ice / Ocean interactions .....	70
12 Tides .....	86
13 Marine Biogeochemistry and TOP Interface.....	91
14 Two-way nesting capability (AGRIF) .....	95
15 Adaptations for High Performance Computing .....	100
16 Machine Learning.....	107
17 Data Assimilation interfaces .....	124
18 Tools, Verification and Validation.....	126
19 NEMO Community engagement.....	131

## Introduction

The main objective of this Strategy is to agree on a set of proposals regarding the directions in which the NEMO code and communities should develop over the 5-year period from 2023 to 2027. These proposals are framed within a 10-year perspective and outline the initial steps to be taken during 2023 and 2024. The proposals outline both what we aim to do and the steps to get there. In this sense the document is a Roadmap rather than a Strategy, and has a more pragmatic focus than the two previous 5-year strategies. It will be used to determine the priorities of the NEMO System Team and Working Groups and of project proposals to develop NEMO. Its target audience includes the developers of NEMO, the users of NEMO, managers within the NEMO consortium and those involved in decisions on the funding of work proposed for the 2023-2027 period.

The Strategy has two parts: Part I provides a synoptic view of the development strategy and its implementation roadmap, Part II provides a more detailed analysis of the underlying challenges and opportunities.

Part I opens by describing NEMO as a code-base, a consortium and a community. It then outlines the consortium's drivers and requirements for NEMO and its development and its vision for NEMO as a coding framework. As good team-work is crucial to NEMO and complicated by the distributed nature of the NEMO community, the next chapter focuses on team-work. The following chapter describes our proposals on the main scientific and technical issues that we expect to have the greatest impact on the NEMO code over the next 5-10 years. Part I closes with a summary of the first steps that have been taken to implement the strategy and the steps we intend to take in future to progress and monitor its implementation.

Part II collects a series of more in-depth topical analyses of the issues at stake in several key areas for NEMO developments that have been considered for elaborating this roadmap. Each of the topical analyses comes with clear propositions as to the priorities and milestones over the next 5-10 years and actions to be undertaken in 2023 and 2024.

Part II is the natural successor to the NDS 2022-2026. Its chapters have been designed to:

- a) Capture the main areas where development of NEMO will be important
- b) Organise the work into working groups of people who can work efficiently together and make proposals for funding of their area
- c) Align with the existing working groups unless there are good reasons to change
- d) Make it clear who will lead on each topic
- e) Either avoid or clearly identify difficult overlaps between working groups

The order of the chapters is not so important, but has been chosen to give a natural flow from one chapter to the next.

## Part I: Executive overview

### 1 What is NEMO ?

*NEMO is an open-source codebase used for ocean dynamics, biogeochemistry and sea-ice modelling applications, including research, forecasts and climate projections. It is managed by a consortium of institutes that support and develop NEMO for their own purposes and the wider community.*

The Nucleus for European Modelling of the Ocean (NEMO<sup>1</sup>) is a geoscientific model (*the NEMO codebase*), used for a variety of applications covering research on ocean and sea-ice dynamics, operational forecasts (short-range, seasonal and decadal), re-analyses, climate projections and the preparation of ocean observing systems. The development of the NEMO codebase is supported by a group of institutions (*the NEMO consortium*) who pool resources to develop the codebase in a sustainable way. This undertaking serves not only the needs of the NEMO consortium institutions but also those of a broader group of users and interested parties (*the NEMO community*). The NEMO development strategy is intended to define priorities for developing the NEMO codebase to the benefit of the NEMO consortium and of the broader NEMO community.

**The NEMO Codebase:** NEMO is an open source geoscientific model that can be used to numerically represent the ocean and sea-ice from the global scale to sub-kilometre-scale. The NEMO codebase consists of three main components: the ocean circulation component NEMO-OCE; the sea-ice dynamics and thermodynamics component NEMO-SI; the tracer transport component with an interface for ocean biogeochemistry NEMO-TOP and the biogeochemistry component NEMO-PISCES. These physical components are described in their respective reference manuals, available on zenodo. The NEMO codebase comes with a series of Reference configurations and test cases, for users to set up new applications, and for developers to test and validate new functionalities. The NEMO codebase also includes additional capabilities such as an IO server, a two-way nesting package, a coupling interface, and pre- and post-processing tools<sup>2</sup>. The NEMO codebase is freely available under the public CeCILL licence, with official release versions issued every few years. The code development is hosted openly on a Gitlab server<sup>3</sup> and the code is distributed with a list of reference configurations with their input and set-up files.

**The NEMO Consortium:** The NEMO codebase can be traced back to the 1980s, but the contours of the NEMO platform and its community have been delineated since 2008 when the NEMO consortium was established. The NEMO consortium is currently composed of five institutions - the Euro-Mediterranean Centre for Climate Change (CMCC), the Institut National des Sciences de l'Univers at Centre national de la recherche scientifique (INSU-CNRS), Mercator Ocean International (MOI), UK Met Office (UKMO), and the National Oceanography Centre (NOC). The NEMO consortium therefore brings together operational agencies and research institutions, with complementary needs and expertise. The consortium institutions pool together resources for developing and distributing the NEMO codebase, and work together for setting up the priorities for its future developments. The consortium institutions are jointly in charge of guaranteeing the coherence, the relevance, and the

---

<sup>1</sup> See <https://www.nemo-ocean.eu/>

<sup>2</sup> See <https://www.nemo-ocean.eu/framework/components/interfaces/>

<sup>3</sup> See <https://forge.nemo-ocean.eu/nemo>

sustainability of the development of the NEMO codebase. The consortium members would welcome additional members willing to commit resources to supporting this effort.

**The NEMO Community:** NEMO development benefits a broad community of users and interested parties within and beyond the NEMO consortium institutions. The NEMO codebase is indeed used by a community of researchers in universities and academic institutions across the world, and as a component of many systems operated by governmental or international agencies. The present strategy therefore recognizes the diversity of applications of the NEMO codebase. NEMO also sits within a wider ocean-sea ice modelling community, which indirectly benefits from NEMO development, and within the European research community, which benefits from funding by the European Commission. NEMO development therefore contributes to improving ocean-sea ice models beyond NEMO itself, and also benefits from a broader community. This is why scientific publications describing new developments to NEMO and participation in model intercomparison exercises should be encouraged. Recognizing the importance of nurturing a vibrant and inclusive community of users and interested parties, the NEMO consortium is committed to providing adequate tools for this community to share expertise (as for instance through the NEMO users' forum) and eventually contribute to the NEMO development process.

## 2 The drivers, requirements and vision for NEMO

### 2.1 Drivers and requirements for NEMO and its development

#### Drivers

The strength of NEMO lies in the fact that it is a science driven model which benefits from the complementary and balanced contribution from both the academic and operational communities. The driver for NEMO is to serve the range of applications of the consortium members. For the academic community, these applications include: ocean-climate coupled projections, developing a deeper understanding of ocean and climate processes (including idealized configurations) and teaching. For the operational community, these applications encompass: NWP/ocean forecasting, seasonal and decadal prediction, climate projections, reanalyses and digital twins. Most NEMO usage requires consideration of a wide range of spatial and temporal scales and the ability to adapt NEMO to emerging technologies. The consortium members would like to continue being associated with a world-leading model that is recognised for scientific and technical excellence and is one of the models of choice for the forecasting, climate and ocean research communities.

#### Requirements

- 1. NEMO needs to be able to reliably support the consortium applications described in the drivers. NEMO must support all these applications on scales from global to regional.**
  - A priority for NEMO is that the code must be both readable and accessible.
  - The code needs to be reliable and perform optimally (including I/O), implying that code and infrastructure is operationally robust and reproducible and that the code does what it is intended to do.
  - As the code must be available at a range of resolutions, this has implications for its efficiency, the inclusion of a variety of parameterisations, the numerical schemes/discretization (such as the vertical coordinates) and it should be compatible with machine learning
  - Configurable to support a wide range of applications, with tools to build the model grid and initial fields and a clean user interface to code options and parameter settings
  - Compatible with data assimilation schemes and ensemble simulation
  - Allow tide resolving simulations requiring high frequency, high resolution and appropriate vertical coordinates
- 2. NEMO needs to be scientifically recognised as a world-leading modelling system for the applications it is intended to support.**

- Supported by documentation and peer-reviewed publications
- NEMO code should be available for reproducible open science publications
- This implies that NEMO needs to be internationally-leading in some areas and state-of-the-art in others. Priority areas should be defined dynamically by the Steering Committee (with advice from Developers Committee)

**3. NEMO needs to provide a complete and coherent representation of the physical processes essential for its target applications.**

- Surface fluxes and ocean mixed layer processes
- Internal mixing and some spurious mixing (here the work on vertical coordinated and parameterizations is essential).
- Low and high-order (quasi-) positive numerical schemes and the coupling between barotropic and baroclinic modes (i.e. time-splitting)
- Mesoscale and sub-mesoscale processes and parameterizations
- Representation of ocean under an ice shelf
- Test cases (used for validation/testing, fundamental fluid dynamics studies and teaching and training)
- Tides

**4. NEMO needs to explicitly include a biogeochemical model/interface and cryosphere model within its framework so that it can adequately serve the variety of principal applications.**

- A biogeochemistry model (PISCES)
- A state-of-the art sea-ice model
- An iceberg model

**5. NEMO needs interfaces with multidisciplinary models which meet the best practice for data exchange**

- Coupling to the atmosphere and surface waves
- A flexible coupler between the ocean, sea-ice, ice-shelf and iceberg components

- Interfaces for biogeochemistry (done by TOP), sediments, ecosystems, fishes, benthic layers and more.
- The ability to interface to hydrological models and coastal/estuary models

**6. NEMO needs downscaling and upscaling capacities to foster applications.**

- Capability for multi grid, i.e. 2-way embedded nesting, including data assimilation to allow relocatable operational systems
- Coarsening (Biogeochemistry)

**7. The NEMO team should be prepared to consider potentially maintainable options for relaxing the primitive equation hypothesis to extend the range of applications**

- Boussinesq (assimilation, climate)
- Non hydrostatic motions

**8. NEMO needs to be sufficiently flexible to run efficiently on a variety of current & emerging HPC architectures, keeping pace with technological developments in high performance computing.**

- NEMO should have a computational performance for typical high-end and low-end applications that is at least comparable with other international modelling systems
- An interface to an I/O system flexible and efficient for generating exascale datasets suited to downstream data analysis
- Specific considerations towards GPU-based architecture
- Consideration of solutions which take advantage of Machine Learning

**9. NEMO needs to be maintained and developed efficiently to deliver the agreed strategy.**

- Teams need to work together efficiently
- System team resources should be deployed efficiently including consideration of carbon footprint



- Code needs to be well-maintained and readable - following coding rules that conform with good practice, modularity, verification (unit testing, robustness, state-of-the-art software development, continuous integration)
- Code needs to be consolidated regularly ensuring compatibility of supported configuration options and removing obsolete or unsupported options
- Reproducibility and replicability (scientific assessment)
- Development methods must be agreed and shared by developers to ensure efficient distributed developments

**10. NEMO needs to appeal to a wide user base. This means that NEMO needs to have ease of use and suitable training materials to allow users to learn to configure and run NEMO to support all relevant applications.**

- An appropriate level of support must be provided to users, bearing in mind that NEMO is a consortium enterprise but wishes to encourage a user base beyond the consortium:
  - Training materials on how to set up and optimally run the model on different computers and hardwares (including virtual machines)
  - Containerisation and open source/cloud computing should be considered
  - Organized training to support user uptake
  - Postprocessing requirements need to be considered including provision of basic diagnostic tools

## 2.2 Our vision for NEMO

**Our vision for NEMO is a *reliable, sustainable, efficient* and *scalable* modelling framework whose development is driven by research and operational needs and results**

Considering the needs of the NEMO consortium and community, but also the existing broader landscape of ocean and sea-ice models, our vision is for NEMO to be a *reliable, sustainable, efficient* and *scalable* modelling framework whose development is driven by research and operational needs and results.

**Reliable codebase:** The NEMO codebase is used as a key component in a number of operational services and climate models. The predictions from these systems are used for decision making, with a

possible large impact on society. This is why the NEMO codebase should be verified<sup>4</sup> and validated<sup>5</sup> as much as possible before being integrated into these systems. It is also why the solutions implemented into the NEMO codebase should be fully described and documented alongside the code itself.

**Sustainable architecture:** The NEMO codebase is used in a wide range of applications. These applications typically choose to use different options from the NEMO modules, sometimes use different combinations of modules and are combined with other components in order to build fit-for-purpose systems. It is therefore essential for NEMO users that the NEMO codebase allows the options within and across modules to be used independently. The modules need to be orthogonal, in the sense that they do not have complicated inter-dependencies, but coherent in their design, their development and their release.

**Efficient & scalable implementation:** The applications leveraging the NEMO codebase are run on a range of different machines, from laptops to exascale HPC clusters. The architectures of these machines are increasingly heterogeneous and their future evolution is unpredictable. Recognizing the variety of applications of the NEMO codebase and the fast evolution of computer technologies, our ambition is for NEMO to be computationally efficient across a range of different applications, with acceptably scalable performance on the most high-end use-cases.

**Development driven by research and operations:** The development of the NEMO codebase is fundamentally related to scientific research. Indeed, NEMO is used as a research tool for investigating ocean/sea-ice dynamics, and its codebase reflects our understanding of these dynamics. Improvement of the representation of ocean/sea-ice dynamics in NEMO also requires a close articulation between the development of the codebase and scientific research. This is why the NEMO codebase and its development are open. Another reason why NEMO should remain a readable and well-documented code is so that students and researchers can understand and modify it depending on their needs. Research based on the numerical experiments using NEMO performed with NEMO should also be robust and easily reproducible.

---

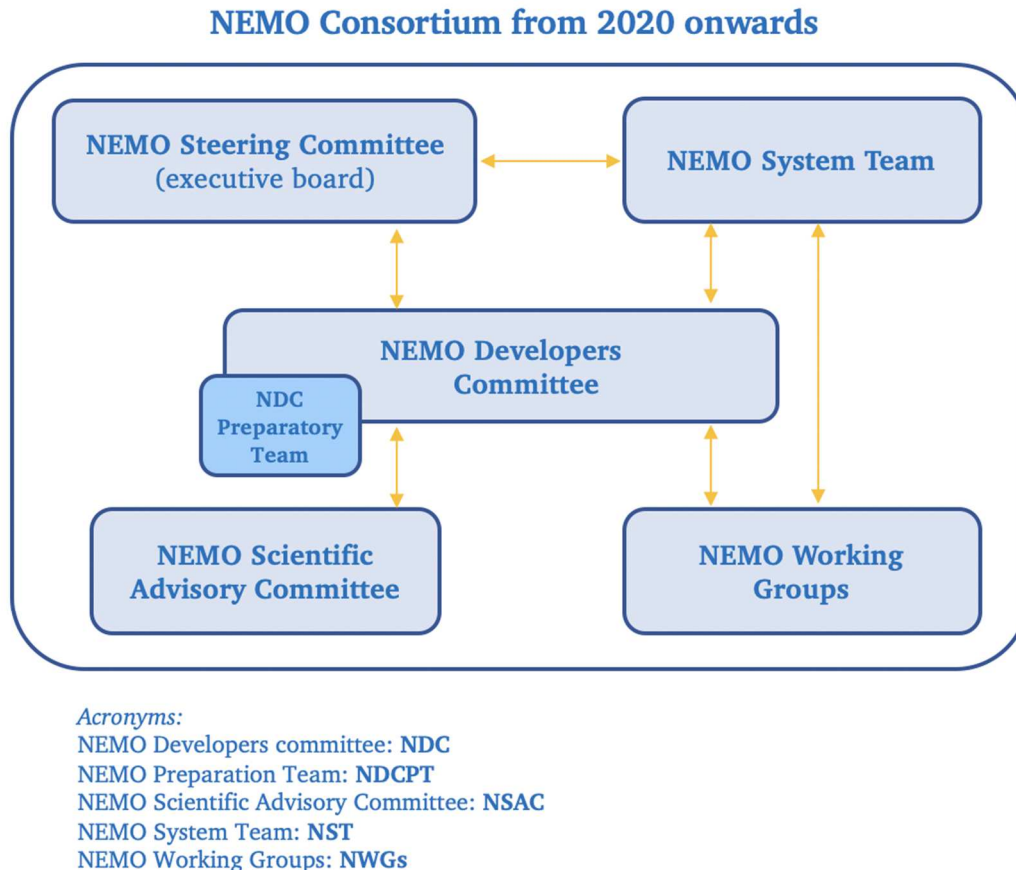
<sup>4</sup> Verification guarantees that what is actually implemented in the code is consistent with what is intended.

<sup>5</sup> Validation is concerned with the assessment of the performance of the code.

### 3 How the NEMO teams work

**The ways of working of the NEMO teams are evolving. Adjustments to their roles and responsibilities are proposed**

Teams are groups of people with an explicitly agreed goal, purpose or objective. Each of the bodies within the NEMO consortium illustrated in figure 1 are set up and working as a team in this sense. The roles of these teams were originally articulated in the NEMO Consortium agreement.



**Figure 1. Schematic of the Consortiums bodies and their interactions**

Team-work within NEMO has evolved and should continue to evolve in response to a number of factors including the people involved and the contributions they can make, the technology available, scientific advances, restrictions on travel and sources of funding. The teams within NEMO should continue to adapt to these changing circumstances, striving always for their work to be focused and efficient and engaging and effective. More specifically we propose that in future:

- The Steering Committee agrees the tasks and resources for the NST annual work-plan, seeks funding opportunities and is responsible for delivery of staff resource to the System Team to implement the work-plan
- The Developers Committee are responsible for writing the NEMO Development Strategy and the NST annual work-plan

- The Developers Committee Preparation Team assists the Developers Committee chair in the preparation of the agenda for meetings and actions resulting from them
- The Working Groups, with assistance from NST members, propose and help to carry-out work in their areas of expertise
- The System Team maintains, verifies and, with the WGs, develops and documents NEMO to realise the Strategy.

The main actions identified in Chapter 7 of Part II are:

#### Actions relating to the NEMO System Team:

- A continuous and conscious effort should be made to value, reward and celebrate the diverse contributions to NEMO development, to match people's skills and interests with the work they are asked to do and to develop capacity where gaps are identified.
- Additional time and attention should be devoted to both training new NST members and developing the skills of more experienced developers. To address this, plans are underway to gather all existing teaching resources, update and augment where necessary, and place the organized material on a central platform for all to benefit from (Community support chapter)
- Better verification and validation of NEMO has been identified as an area for improvement. New tools for shared code development and associated automatic testing are being explored and implemented (Tools, V&V chapter)
- The process for updating the documentation of NEMO in-line with the changes introduced in new versions should be reviewed and properly followed (see the Community Engagement chapter)
- The team of NEMO officers should be better utilised to improve coordination across the NST institutions so as to ensure an equitable distribution of tasks and promote an efficient resolution of issues.

#### Actions pertaining to the Consortium:

- The articulation of the roles of the Teams should be revised in line with the way the roles have already evolved and the proposals above.
- A framework articulating the skills and contributions of most value to the NEMO Consortium should be prepared
- In order to facilitate effective communication and build interpersonal relationships between team members when meeting in person is not possible, NEMO should remain agile in employing new platforms and software that facilitate engaging virtual interactions.
- The annual Work Plan is a very important document for the NEMO community and so should undergo iterative stages of inputs and feedback from the NST, NDC, WGs and NSC. Some work is needed to ensure that a more inclusive and extensive review of the Work Plan takes place. Resourcing and prioritisation issues should be discussed with the NSC at an early stage. Furthermore, in order to meet the year-end deadline while guaranteeing a robust review procedure, it is suggested that the October to December agenda be better managed and adhered to.
- The progress and plans of the WGs have historically been reviewed once a year at NDC meetings. This approach could be improved by reviewing the WGs biennially on a rolling basis with an individual meeting for each WG.
- To make use of various funding avenues, it is suggested that more frequent communication takes place between the NSC and the NDCPT in order to identify possible financing opportunities.

### Actions relevant to the wider community:

- In order to encourage contributions to NEMO development coming from the wider scientific community, NEMO should work on improving communication regarding the structure and design of the code so as to better facilitate joint developments between NST members and external parties. (Community support chapter)
- A continuous effort needs to be made to ensure that the NEMO documentation is kept up to date and is comprehensible to a broad user base (users' guide, reference manuals, development workflows and coding rules) (Community support chapter).

## 4 Technical issues and opportunities with major impact on the code

This section outlines the proposed activities that are expected to have the largest impacts on other parts of the code system or the largest impacts on the quality, reliability or reputation of the NEMO system. The WG responsible for leading each activity is indicated in brackets in the proposal heading.

### 4.1 Discrete representation

#### 4.1.1 Flexible choice of vertical coordinate (Kernel WG)

We intend to develop a more generalised vertical coordinate (GVC) that will provide greater flexibility for users to constrain the model layer interfaces to nearly follow isopycnal surfaces in the interior (including the thermocline), to be parallel to the atmospheric interface near the surface, to nearly follow the bathymetry and to be smoothly varying. This would enable spurious diapycnal mixing in the ocean interior to be better controlled and/or flow over sills and the steering of the flow by the bathymetry to be improved. This work will focus particularly on a suite of options for the specification and calculation of the target grid and on ways to alleviate CFL restriction for vertical advection while avoiding numerical inaccuracies associated with standard implicit advection schemes. This is a substantial undertaking that will require proper planning as a work-package within a project with dedicated funding for specific tasks and individuals. We propose also to investigate further the Brinkman penalisation and multi-envelope approaches for representation of sharply varying bathymetry.

One of the kernel WGLs will look out for evidence that alternative horizontal grids have significant advantages for our applications. If/when there is clear evidence, scoping of options would become necessary.

#### 4.1.2 More effective use of horizontal resolution (AGRIF WG)

Our current strategy is to rely on AGRIF for improved horizontal or vertical resolution, with limited computational cost, in selected areas where the benefit of the highest resolution or changes in vertical coordinates is key (e.g. straits, overflow, western boundary currents...). We are also planning to use AGRIF for the coarsening of BGC components.

We propose to:

- further facilitate and promote AGRIF usage: consolidate the pre-processing tool, enhance user documentation, write a developer guide, and introduce training sessions, user/developer workshop
- use AGRIF for the coarsening of BGC components
- explore the different possibilities (mosaic of zooms versus the merge of BDY and AGRIF functionalities) to implement nests with a complex geometry (e.g. follow complex coastlines or dynamically active regions)
- adapt nesting methodology to RK3 and implement sub-step (barotropic) exchanges
- assess and improve AGRIF's computational performance (mpp optimization, compatibility with an efficient GPU usage).

The compatibility of AGRIF with icebergs, ice cavities and wetting and drying need to be discussed with other WPs.

## 4.2 Representation of high-latitude processes

### 4.2.1 Greater flexibility of the sea-ice model component (Sea-ice WG)

SI<sup>3</sup>, the NEMO sea ice component was put together over the last few years by unifying capabilities from several models previously used in the NEMO system. The resulting code is strongly modular and has a wide range of options for the representation of sea ice physics, but is not modular enough to allow prospective users to switch in/out all of the different components that they may require to tackle current research questions around sea ice.

A central activity for the NEMO sea ice strategy is to further improve the modularity of the code and the interfaces between components, both internal and external to SI<sup>3</sup>. In particular, we will enable or facilitate the coupling to other components of the Earth System (snow on sea ice, ocean waves, continental ice, biological and chemical tracers). This will widen the range of NEMO applications and leverage more contributions of the Earth System sciences community, made in the framework of ongoing or foreseen projects.

Another interface, currently in the core of the sea ice code, should be better defined and rethought — the one between vertical physics (thermodynamics, halo-dynamics, optics) and horizontal drift and deformation processes. A clean separation between the 1D and 2D processes would enable large benefits from the extensive community activities on the physical and numerical representation of sea ice drift and deformation, be they based on continuum or discrete element approaches.

### 4.2.2 Representation of land-ice interactions with the ocean (Land-ice WG)

Land ice / ocean interactions include **ice-shelves**, **glacier termini** and **icebergs** melting, as well as **surface** and **sub-glacial runoff** from the ice sheet. The scientific community modelling these processes has made rapid progress in the last decade but the subject is in its early stages and remains challenging. The 10-year roadmap for this working group is dedicated to supporting the development of ESMs that are able to model Antarctic and Greenland ice sheet / ocean interactions. The objective is to release robust and easy to use NEMO based ESMs able to simulate realistically the oceanic state and variability on the Greenland and Antarctica continental shelves, as well as the polar ice sheets states and variability over the recent and future centuries. In such ESMs, there are 2 sources of issues: missing physics and external bias. External biases are outside the scope of this document as they come from a lack of tuning, realism of the forcings and other components of ESMs. We therefore focus only on how to improve physics of the land ice ocean interactions within NEMO. In terms of major priorities, our short-term highlights focus on exploring and testing ice shelf cavity and sub-ice boundary layer parameterizations, schemes for the migration of a calving front and conservation for ice sheet coupling, improving the HPC performance of the iceberg module and the design of test cases to accompany the associated new developments.

## 4.3 Representation of oceanic processes

### 4.3.1 Resolution-dependent parameterisations of ocean macro-turbulence (Eddy Closures WG)

This new WG will develop a set of test cases and metrics to evaluate the performance of eddy closure parameterisations and help WG members to accelerate the adaptation and acceptance of closure schemes for the NEMO trunk. Driven by the need of Consortium members to target both  $\frac{1}{4}$ -deg. and 1-deg. global configurations, the primary focus will be on mesoscale eddy processes, but may also include the influence of the sub-mesoscale and surface mixed layer. Two distinct approaches have arisen in recent years, to move us beyond the era of ad-hoc application of the Gent and McWilliams, 1990 (GM) mesoscale closure: 1) the GEOMETRIC framework (Marshall et al., 2012), which specifies the required symmetry structures of an eddy forcing tensor needed to achieve desired conservation properties of the dynamics, i.e. a recipe that parameterisation schemes should adopt, rather than a solution in itself; and 2) the advent of “kinetic energy backscatter” schemes, to parameterise the inverse cascade of geophysical turbulence. At least one candidate closure scheme in each of these two categories has been identified, with one of them already undergoing implementation in NEMO trunk (Mak et al., 2018, under category 1). On the topic of scale-awareness, a resolution function has been suggested as a way to assess and adjust whether an eddy closure is acting in a region where the first baroclinic deformation radius is resolved (Hallberg, 2013). However, it is proposed that further work is needed to test its performance as a simple scaling on the forcing from new eddy closures adopted by NEMO, and additional issues may arise when parameterised processes may themselves be modified nonlinearly at certain model resolutions in a way not captured by this resolution function scaling, e.g. by partial resolution of the sub-mesoscale.

#### 4.3.2 Representation of Air-Sea interface & vertical mixing (ASI and Vertical Mixing WG)

A large amount of work has been accomplished during the last NDS period to upgrade air-wave-sea interactions & ocean mixing related parts of the NEMO code, and to provide state-of-the-art parameterisations and capacities. This is especially true concerning surface turbulent fluxes and wave-ocean interactions, which have been largely rewritten and updated to introduce new physical possibilities. Mesoscale air-sea interactions effects on the ocean have been also newly introduced during the last NDS through the development of a new parametrisation and of an atmospheric boundary layer model. Consequently, the developments proposed here for the next NDS period are less ambitious, and can be seen as a consolidation and a completion of what have been started during the previous period. In particular, the Osmosis vertical mixing scheme and the atmospheric boundary layer model developments will be completed during this new NDS period. Waves-related effects will be generalised to support more vertical mixing schemes and dynamical options. New single column test cases will be added to facilitate vertical physics model developments, intercomparison with existing models and libraries, and validation against newly available observations.

### 4.4 Technological drivers

#### 4.4.1 Support for a wide variety of HPC architectures (HPC WG)

The rapid advancement of new (pre-)exascale parallel architectures technologies requires a constant adaptation of the NEMO model to emerging architectures, finding solutions that guarantee performance portability on heterogeneous architectures and maintainability of the code even by people not expert in HPC. In the next 5 years we propose to direct the effort towards three main directions: (i) the improvement and completion of the tiling-based implementation by overcoming



the current performance limitations and the consequent optimal management of the extended halo; (ii) the gradual transition and support of parallelization towards GPU-based architectures by adopting Domain Specific Language (DSL) based solutions; (iii) lowering the time-to-solution by means of techniques based on the reduction of numerical precision while ensuring an adequate level of accuracy of the results also for long-term simulations. Finally, it is emphasized that the performance optimization activities are transversal and must be considered from the beginning each time new features are introduced in the model this requires that HPC-WG works in close collaboration with all other working groups.

#### 4.4.2 Support for and opportunities from Machine Learning (MLMU WG)

The combination of machine learning with scientific computing is an active area of research which could eventually improve geoscientific models and their integration into broader numerical systems, such as climate models and operational forecasting systems.

These emerging approaches could help quantify and reduce systematic model errors, improve the representation of unresolved processes, reduce the numerical cost of model simulations, improve the representation of uncertainty propagation and allow better leverage of observations in model design.

But, while the potential of ML for improving ocean/sea-ice models is large, the *Technological Readiness Level* of these applications is still relatively low. Yet, given the possible implications on some of the NEMO consortium requirements, we think that the NEMO consortium should be proactive on these topics.

We therefore propose to set up a new NEMO working group (WG) focusing on “*machine learning and model uncertainties*” in order to coordinate and accelerate development in this area. The ambition of this working group is to gather a group of interested parties for fostering innovation and exchanges of expertise on these interdisciplinary topics.

In parallel, we propose to leverage resources from several recently funded research projects for moving forward on an explicit list of actions for the next 5 years. These actions will overall aim at :

- (A1) introducing several ML-based components into NEMO,
- (A2) improving the representation of model uncertainties in NEMO,
- (A3) investigating the potential of differentiable emulation.

The roadmap we propose will be implemented in two successive phases with several key actions to be completed before 2025. Their successful implementation will require strong coordination with other NEMO working groups and with the NEMO ST. Community and capacity building activities will also be pivotal to our success in this area.

Additionally, we propose that the *working group on machine learning and model uncertainties* acts as a link with the data assimilation community which leverages NEMO in operational systems. We propose to monitor new developments in this field and to identify key needs that would help users to run NEMO within data assimilation frameworks.

## 4.5 Sustainability

### 4.5.1 Efficient & effective verification & validation (Tools and V&V WG)

Tools to build NEMO and check whether updates to the NEMO trunk change the results are well established. This group will develop such essential NEMO support tools in line with the priorities identified in the NEMO V&V roadmap<sup>6</sup>. The plans are:

- **Within 2 years:** To have the system team fully proficient with all aspects of a GitLab-based development environment. Including: the use of GitLab runners for automated tasks and continuous integration and regular updating of Wiki and web-based support material
- **Within 5 years:** To have Unit-testing capabilities in most code areas. Frequent, automated testing of the code base. Full support for exascale and heterogeneous computing environments (e.g. GPU co-processors) via Domain Specific Language pre-processing tools.
- **Within 10 years:** To have code testing carried out by AI-enabled agents

### 4.5.2 Verification and extension of the TOP interface

Marine biogeochemistry within the NEMO framework currently accounts for a variety of coupled system, being directly embedded within the core code (e.g., PISCES and MEDUSA) or exploiting the TOP interface to access the transport drivers (e.g., BFM, ERSEM, BAHMBI).

It is indeed a priority to ensure a reliable development of the TOP interface to sustain the coupling of built-in and non-legacy biogeochemical models over the long term, by retaining the modular structure built in previous years and extending it with complementary features. It is here foreseen the setup of a dedicated test configuration to verify the consistency of all data handlers and processes inherited from the NEMO core. Such a testbed will evaluate the correct simulation of passive tracers' dynamics due to physical schemes addressing the long-term need to ensure a resilient and reliable TOP interface.

The evolution of the interface toward biogeochemical processes in the coming years will foster the orthogonality between physical processes and oceanic tracers' dynamics to provide generalised schemes to resolve vertical dynamics of sinking particles and the handling of optical properties within the water column. In particular, a synergic action with the HPC WG will address the improvement of computationally expensive tasks (e.g., transport of tracers) and evaluate potential of tailored optimizations for biogeochemical dynamics.

Under the same long-term perspective, the TOP interface will be extended to face the future biogeochemical complexity beyond the marine pelagic environment. The main opportunity is represented by the development of a coherent and comprehensive framework to handle all marine ecosystem components (sea-ice, pelagic, benthic). This system will account for dedicated sub-modules and it will provide access for shared physical drivers of the NEMO core, along with interfaces to resolve the boundary exchanges.

---

<sup>6</sup> [forge.ipsl.jussieu.fr/nemo/attachment/wiki/WorkingGroups/Verification/nemo-validation-and-verification-roadmap\\_draft\\_1.0\\_Nov\\_2020.pdf](https://forge.ipsl.jussieu.fr/nemo/attachment/wiki/WorkingGroups/Verification/nemo-validation-and-verification-roadmap_draft_1.0_Nov_2020.pdf)

#### 4.5.3 NEMO Community (Community Engagement WG)

A thriving user base is fundamental to delivering the development strategy. Furthermore, sustainability of the framework requires a certain level of investment in the user community. Being an open source framework, NEMO has a number of procedures already in place, but there are gaps and there is also a potential to present community support material in a more practical and accessible way. In particular, the existing material is difficult for new users of NEMO to grasp, particularly if they are new to modelling.

It is advantageous to facilitate the development of the User Community, as it will draw more users, who will include or grow into new developers and future leaders.

Principal challenges include (A) developing introductory and new-user material to NEMO (lowering the barrier to entry) and (B) incentivising contributions from experts (ascribing appropriate recognition to delivering helpful documentation, appropriate training and effective mentoring).

Specific plans are detailed in the Community Engagement Chapter, in Part II.

## 5 Implementation Plan

Our initial idea for an implementation plan was centred on a Gantt chart for each WG summarising the key milestones in the implementation of their main aims and objectives similar to that in the Land-ice chapter. Uncertainties about the resources and funding that will be available, however, make it difficult to establish reliable/meaningful Gantt charts of this sort at this stage.

Instead we have started to construct tables of the tasks and resources that each WG need to undertake and the resources that these tasks will require. The content and format of these tables<sup>[1]</sup> was agreed with the WGLs and the members of the NSC and draft tables have been constructed.

In addition, each of the WGLs was consulted about the main challenges/needs of their WG and summaries of these issues for each WG have been discussed with the NSC. The issues include obtaining suitable sources of funding for the different types of work, improving documentation and guidance for users and developers, and improving the support for NEMO System Team members. The NSC and NDC will work together to develop a plan for prioritising and addressing these issues. An initial plan will be developed by the end of April 2023 and its implementation will be monitored by the NSC and the NDC.

An Annual Work Plan for the NST is written by the NDC. Clearly these plans will be an important component of the implementation of the Strategy. Progress against the Plan at the end of each year will be monitored by both the NDC and the NSC, and the Plan for the next year will be agreed with the NSC.

---

[1] <https://docs.google.com/document/d/1O56JKHU-FeBX26487ygDjHmEiID8XAoFuzvImRJ00Tc/edit>

## Part II: Detailed Analyses

### 6 Team-work

#### 6.1 Executive Summary

The NEMO teams and their ways of working have evolved significantly over the last 5 years. This chapter discusses how the teams could further evolve and proposes

- some improvements to the process for developing the NST annual work-plan
- that a framework articulating the skills and contributions of most value to the NEMO Consortium underpins a continuous and conscious effort to reward NST and WG members
- that the NSC hold short, more frequent, meetings to identify potential avenues for funding
- that the team of NEMO officers be better utilised to improve coordination across the NST institutions so as to ensure an equitable distribution of tasks and promote an efficient resolution of issues
- that the roles and responsibilities of the Teams are revised in line with the way the roles have already evolved and the proposals outlined in this chapter.

The proposed action points contained in this section are the same as those listed in chapter 3.

#### 6.2 Introduction

Section 1 below considers issues that are relevant to more than one team whilst section 2 focuses on specific teams and how they could adapt to better meet their objectives. This chapter is designed to provide context and add detail to the action items for team-working summarised in chapter 3 of Part 1.

#### 6.3 Issues not specific to a single team

##### 6.3.1 Engagement with the wider scientific community

The NEMO teams need to entrain expertise and man-power from the wider scientific community in order to be efficient and effective in ensuring that NEMO remains a state-of-the-art model framework. It is the role of the WGs and their leaders to gather scientific know-how and inputs from the ocean modelling community (e.g. not only NEMO Consortium members) in order to make sure the relevant drivers and priorities are identified and to make it easier for scientists who wish to contribute, to engage with NEMO. The WGs are “open” precisely for this reason.

Some challenges have presented themselves regarding the inclusion of new developments into the NEMO reference code. An important driver in NEMO development is to ensure modularity and readability. As a result, a new idea can quite easily be developed and the results can be published without the author including them within the official code. Having this new idea included in the reference code, however, requires an understanding of NEMO development workflow, code design

and standards and, at present, a collaboration with the NST is usually required. This often results in a significant amount of work both by NST members and external collaborators. We should seek to improve the communication of the salient points of the code design to make this step more efficient in future.

**Action:** In order to encourage contributions to NEMO development coming from the wider scientific community, NEMO should work on improving communication regarding the structure and design of the code so as to better facilitate joint developments between NST members and external parties.

Some developments to the NEMO code can be made by informal collaborations. But many developments need to be supported by specific funding. This Strategy is intended to determine the top priorities for NEMO development and hence to underpin proposals for future activities and funded projects. Work of this sort by external scientific developers is expected to be most efficient when it is undertaken jointly with NST members.

### 6.3.2 Supporting NEMO communities

The priorities for supporting each of the NEMO communities are analysed in chapter 20.

**Action:** Additional time and attention should be devoted to both training new NST members and developing the skills of more experienced developers. To address this, plans are underway to gather all existing teaching resources, update and augment where necessary, and place the organised material on a central platform for all to benefit from.

### 6.3.3 Technologies supporting remote team working

Active collaboration requires frequent contact, through exchange of messages or conversations. The recent evolution of technologies to support communication and code development have greatly improved the effectiveness of all the NEMO teams. These technologies offer teams opportunities to adapt the way they work to become more effective.

Methodologies and tools for shared code development and verification have evolved considerably over the last 10 years. NEMO has recently produced a roadmap for improvement of its verification and validation processes and transitioned from svn to the more modern and potentially powerful Gitlab environment. We are exploring the options available with Gitlab to do various things (e.g. communications on tickets, issues, but much much more) (see Tools and V&V chapter).

**Action:** Better verification and validation of NEMO has been identified as an area for improvement. New tools for shared code development and associated automatic testing are being explored and implemented.

Because there is no travel-time involved in online meetings, they can be much more focused, shorter and more frequent than in-person meetings. The invitations to meetings can then be tailored to the topics discussed so that everyone at the meeting is properly engaged in the discussion. Sharing of screens facilitates joint technical work, as well as presentations, and makes it possible, for example, to merge code frequently without meeting in person.

At conferences and face-to-face team meetings, individuals are able to share news and opinions informally and build personal relationships that are vitally important for good team-work. Some on-

line tools, such as gather.town, facilitate somewhat similar interactions. We need to exploit these tools better and recognise that, at least at the moment, there is real value in occasional face-to-face meetings.

**Action:** In order to facilitate effective communication and build interpersonal relationships between team members when meeting in person is not possible, NEMO should remain agile in employing new platforms and software that facilitate engaging virtual interactions.

#### 6.3.4 Cross-team working

The WGs are intended to enable teams with common interests to interact effectively. However some issues and challenges need to be considered by members from more than one WG. These issues have been highlighted in the chapters of Part II written by individual WGs. The NDC is responsible for prioritising these issues, and new ones, as they arise. The mechanism we propose for doing this is for the NDC to maintain a list of these issues, to agree which of them are the highest priority for scoping, and to set up small teams to do that and report back with proposals for what to do. This Strategy document could be updated with the main proposals.

The development of the annual Work Plan for the NST requires input from the NST, the NDC and the NSC. The NDC builds the first draft of the annual Work Plan receiving information from the NST. The NDC then encourages contributions and comments from the WGs on this plan in an iterative process to refine the focus. It is suggested that the agenda of the process during the last 3 months of the year is well organised in order to facilitate these exchanges and meet the year-end deadline.

**Action:** The annual Work Plan is a very important document for the NEMO community and so should undergo iterative stages of inputs and feedback from the NST, NDC and WGs. Some work is needed to ensure that a more inclusive and extensive review of the Work Plan takes place. Resourcing and prioritisation issues should be discussed with the NSC at an early stage. Furthermore, in order to meet the year-end deadline while guaranteeing a robust review procedure, it is suggested that the October to December agenda be better managed and adhered to.

#### 6.3.5 Funding and proposals

Development of proposals for funding also requires cross-team working. As mentioned in Section 2.1.1, the most efficient way for NEMO developments to have scientific fidelity is for the development to happen in tandem with a research project. These projects often need national or European funding. To support this collaboration, the NDC is expected to write letters of support for projects that are well aligned with the NDS. One of the main motivations for the NDS is therefore to provide a well considered list of priorities that can form the basis for funding proposals. The NSC should work with the NDCPT to identify opportunities for funding proposals. The NDCPT should then work with the WGLs to develop proposals for funding.

**Action:** To make use of various funding avenues, it is suggested that more frequent communication takes place between the NSC and the NDCPT in order to identify possible financing opportunities.

#### 6.3.6 Recognition of diverse inputs

Developing an established numerical model requires a multitude of skill sets (e.g. improving its documentation, verification, technical design, scientific formulation, or coordinating activities). No

one person possesses all the necessary expertise (or the time) and thus a collaborative effort is essential. The work is very demanding and is not always given the recognition it deserves. So, the skill sets of value to NEMO should be articulated and used to inform staff assessments, a conscious, continuous effort should be made to value, reward and celebrate contributions, to match people's skills and interests with the work they are asked to do, and to develop skills where there are gaps. This is relevant to all the NEMO teams – how it is tackled will depend on the team.

**Action:** A framework articulating the skills and contributions of most value to the NEMO Consortium should be prepared

**Action:** A continuous and conscious effort should be made to value, reward and celebrate the diverse contributions to NEMO development, to match people's skills and interests with the work they are asked to do and to develop capacity where gaps are identified.

**Action:** Set up an annual Gurvan Madec medal to recognise and reward outstanding contributions to NEMO development.

### 6.3.7 Roles of Teams

Previous documents have noted that the roles of the teams include the following points:

- The Steering Committee verifies the achievements and validates the work plan;
- The Working Groups (NWGs) are responsible for the organisation of major axes for development;
- The Systems Team (NST) carries out the code development and implements the strategy;
- The Developers' Committee (NDC) supplies a bridge for communication between the Working Groups and Systems Team by conceiving and approving the annual work plans;
- The Scientific Advisory Committee (NSAC) provides an external source of feedback for strategy and proposed advances.

Team-work within NEMO has evolved and should continue to evolve in response to a number of factors including the people involved and the contributions they can make, the technology available, scientific advances, restrictions on travel and sources of funding. The teams within NEMO should continue to adapt to these changing circumstances, striving always for their work to be effective and engaging.

For clarity, the articulation of the roles of the Teams should be revised without delay. The roles have either already evolved or we have recommended above that they do evolve in the following directions:

- The Steering Committee agrees on the resources for the NST annual work-plan, seeks funding opportunities and is responsible for delivery of staff resource to the System Team to implement the work-plan
- The Developers Committee are responsible for writing the NEMO Development Strategy and the NST annual work-plan
- The Working Groups, with assistance from NST members, propose and help to carry-out work in their areas of expertise
- The Systems Team maintains, verifies and, with the WGs, develops NEMO to realise the Strategy
- The Developers Committee are responsible for writing the NEMO Development Strategy and the NST annual work-plan.



**Action:** The articulation of the roles of the Teams should be revised in line with the way the roles have already evolved and the proposals above.

## 6.4 Improving the effectiveness of particular teams

### 6.4.1 NEMO System Team (NST)

The NST meets on a regular basis (every 3 weeks) by videoconference. It is the place to discuss on-going developments, any difficulties, answers to users' questions when needed, bugs, fixes and more. Meetings usually take less than one hour. They are chaired in turn by NEMO Officers which allows specific questions to be raised by one of the local group members. As a number of developments take place in parallel, it is highly beneficial that all members are present at these meetings in order to stay current and identify points of inter-dependence and connectivity. Finally, these meetings provide the opportunity for work to be fairly shared between developers and tasks to be assigned.

Aside from these short and regular meetings, the NST meets in person once-a-year for a 3-day "Merge Party" at the location of one of the Consortium members. Aside from the technical advances achieved, this in-person meeting provides the time for people to get to know one another better so as to feel at-ease raising issues and concerns during open discussions by video-conference during the rest of the year. This annual face-to-face meeting helps the NST to function optimally.

There is also a team constituted by the NEMO Officers that meets occasionally. This team could operate as a team of managers. If it met regularly, this team could prepare the agenda for the NST meetings, make sure that timely decisions are being made at NST meetings, coordinate the finalisation of releases, seek to ensure that actions are equitably distributed between NST members and discuss difficult issues. It could also work with the NDCPT on the preparation of the annual NST work-plan.

**Action:** The team of NEMO officers should be better utilised to improve coordination across the NST institutions so as to ensure an equitable distribution of tasks and promote an efficient resolution of issues.

### 6.4.2 NEMO Steering Committee (NSC)

The members of the NSC are responsible for determining the staff resources committed to the NST work-plan within their organisation. The NSC formally agrees the work-plan and resources to be made available for the NST each year. It also considers opportunities for future funding and succession planning. At present it meets once a year (usually in January or February).

The NSC needs to identify or help to create opportunities for funding of the proposals within this Strategy. In order to do this the NSC should meet more frequently. We propose that the NSC holds short quarterly virtual meetings with the NDCPT members whilst funding is being actively sought.

The NSC should consider the skill sets that are required for development of NEMO and the formulation of criteria by which contributions to NEMO can be better evaluated and recognised. It should also review its succession planning within the next couple of years.

### 6.4.3 NEMO Developers Committee (NDC)

The NDC is responsible for the development and implementation of the NDS. The WGs are coordinated by and report to it. It meets at least twice a year and more frequently when required.

To ensure the effectiveness of the NDC, an NDC Preparation Team analyses key issues, and prepares the agenda for the NDC meetings. The NDCPT has worked well: its discussions are effective at generating ideas, it has enabled open discussion of difficult issues and it has improved the preparation and conduct of the NDC meetings. The scope and membership of the NDC has also been rationalised recently. The NDC took the lead in setting up the NEMO SAC.

Over the last couple of years, the progress and plans of the WGs have been reviewed at NDC meetings. This approach could be improved by reviewing the WGs biennially on a rolling basis with an individual meeting for each WG.

**Action:** The progress and plans of the WGs have historically been reviewed once a year at NDC meetings. This approach could be improved by reviewing the WGs biennially on a rolling basis with an individual meeting for each WG.

Improvements to the process for agreeing the annual NST work-plan and a proposal for handling issues that fall between WGs has been proposed in the section on cross-team working.

#### 6.4.4 Working Groups (WGs).

The size of the membership of the WGs varies considerably. Particularly if the WG is larger than 5-6 people, specific developments will usually involve sub-groups of 2-5 people. At least one member of each sub-group should also be an NST member. These sub-groups should report back to the main group (this does not need to be a big deal - just occasional presentations on progress). The WG leaders (WGLs) should periodically ensure that the WG members agree on the group's objectives and how they will work (focus of the meetings, frequency etc). Information should be available on the publicly available WG Gitlab web page.

## 7 Ocean model Kernel

WGLs: Mike Bell, Florian Lemarié and Gurvan Madec

### 7.1 Executive Summary

Over the last 5 years, NEMO has been substantively re-organised in order to implement an RK3 time-step scheme. We propose to focus in the next 5 years on implementation of a more flexible generalised vertical co-ordinate (GVC) than the existing  $\tilde{z}$  scheme. The aim is to provide greater flexibility for users to constrain the model layer interfaces to nearly follow isopycnal surfaces in the interior (including the thermocline), to be parallel to the atmospheric interface near the surface, to nearly follow the bathymetry and to be smoothly varying. This would enable spurious diapycnal mixing in the ocean interior to be controlled and/or flow over sills and the steering of the flow by the bathymetry to be improved. This work will focus particularly on a suite of options for the specification and calculation of the target grid and on ways to alleviate CFL restriction for vertical advection while avoiding numerical inaccuracies associated with standard implicit advection schemes. This is a substantial undertaking that will require proper planning as a work-package within a project with dedicated funding for specific tasks and individuals. We propose also to investigate further the Brinkman penalisation approach and the Multi-Envelope method (possibly including a combination of the two) for the representation of sharply varying bathymetry, the options for relaxation of the Boussinesq approximation and to work more closely with the HPC WG to pursue the co-design of algorithms suited to modern high performance computers. We also propose to continue to rely on AGRIF to enable users to vary the horizontal resolution within the model domain.

### 7.2 Introduction

#### 7.2.1 Assessment of the 2018-2022 Development strategy

The main issues discussed in the 2018-2022 strategy were

1. the vertical grid
2. the horizontal grid
3. the time-stepping algorithm

The focus in that period has been on the introduction of a Runge-Kutta (RK) time-stepping algorithm (through the EC funded IMMERSE project). A significant change to the time-indexing of the model fields (designed by Gurvan Madec and implemented by Andrew Coward and Dave Storkey) was made to enable more flexibility in the time-step scheme and analyses were made (by Florian Lemarié and Nicolas Ducouso) of the stability of options for efficient time-stepping of the external mode (Ducouso et al 2021). Good progress has been made (by Sybille Techene and Gurvan Madec with Jerome Chanut & Andrew Coward) in implementing a 3rd-order RK scheme and it is expected to become available in vn 4.2.1. Further opportunities to refine the methods used, in particular to explore compensated space-time schemes to allow more accurate and/or efficient calculations, are not expected to be intrusive. Florian Lemarié and Gurvan Madec are best placed to supervise this work.

Relatively little progress has been made on the vertical coordinate. However improvements to the  $z$ -tilde coordinate have been made by Jerome Chanut and assessments of the dependence of diapycnal mixing on various schemes (including  $z$ -tilde coordinates) and parameter settings in realistic global configurations have been made (Megann et al. 2020, 2022). A number of schemes for calculation of horizontal pressure forces in  $s$ -coordinates have also been tested within the NEMO framework (Bell & Young 2021) and the scoping study envisaged in the strategy has made useful progress over the last 12 months. The Multi-Envelope approach to vertical coordinate (Bruciaferri et al. 2018) has been

proven to be a robust and generalized method to improve the model's accuracy and flow-bathymetry interactions, both in regional (e.g., [Bruciaferri et al. 2020](#), [Wise et al. 2021](#), [Bruciaferri et al. 2022a](#)) and global configurations ([Bruciaferri et al. 2023](#)). On the same front, the Brinkmann penalisation method ([Kevlahan et al. 2015](#), [Debreu et al. 2020](#), [Debreu et al. 2022](#)) has also started to emerge as an elegant and powerful approach for representing small-scale variations in bathymetry (and potentially embedded sea-ice and ice shelves). Such an approach is relevant for a z-coordinate model as it eliminates some artefacts due to the step-like geometry.

A “lesson learnt” here is that it is probably unrealistic to expect to work on more than one major aspect of the code formulation within a 5-year period.

The previous strategy noted what were then relatively new formulations for hexagonal and triangular horizontal grids using either finite element or finite volume methods. The MPAS and FESOM teams are exploring these approaches using finite volumes. We decided in the previous strategy to focus our effort on variable horizontal grids on the AGRIF two-way nesting tool. We propose to continue with that strategy for the next 5-year period. As a group we don't have expertise in hexagonal or triangular grids and the impacts of making such changes on other aspects of the code would be very wide ranging. Our view is that code using such grids would probably constitute another model. A smoother representation of the “side” boundaries could be achieved instead by (nearly) terrain-following coordinates (e.g., via the Multi-Envelope and/or Brinkmann penalisation methods). There may be merit in using the alternative meshes for estuarine models, but even that is not clear ([Nudds et al 2020](#)), and it is not the main focus of the NEMO consortium members. We propose that a scoping study reviews the evidence and options on this issue for the NEMO community towards the end of this 5-year period and do not discuss alternative horizontal grids in more detail here.

### 7.2.2 Priorities for the 2023-2027 Development strategy

The outstanding issues on which we would like to make progress in the next 5 years are

1. Generalised vertical coordinates with vertical ALE (V-ALE) algorithm
2. Better representation of bathymetry-flow interaction processes
3. Green computing: improving the energy efficiency of NEMO

These topics are considered in turn below.

## 7.3 Vertical grids and representation of bathymetry

### 7.3.1 Generalised vertical coordinates (VLR vs V-ALE vs Quasi-Eulerian)

The expected benefits of generalised vertical coordinates from a physical viewpoint are: the reduction of spurious diapycnal mixing in the ocean interior (particularly in climate simulations); improvement of the flow over sills; and improvement of the steering of the flow by the bathymetry. For example, there are some suggestions that the rapid degradation of Gulf Stream separation in the first phases of spin-up depends on representation of the bathymetry ([Ezer 2016](#); [Schoonover et al. 2017](#)). Another illustration of the detrimental effect of the step-like representation of the bathymetry is that the discretisation of the Coriolis term gives rise to a false representation of the Coriolis “force” on gyre circulations ([Styles et al. 2021](#)). It may not be possible to achieve all of these goals with one choice of vertical co-ordinate. So the aim is to give users greater flexibility to specify target coordinates that are well suited to their applications.

A very educational description of vertical Lagrangian remapping (VLR), arbitrary Lagrangian Eulerian (V-ALE) and quasi-Eulerian schemes is provided in [Griffies et al. \(2020\)](#).

The VLR method uses a Lagrangian approach to advection in the vertical combined with an Eulerian approach in the horizontal. This requires a directional splitting. In standard implementations of directional splitting the order of integration is permuted at each time-step to reduce splitting errors (e.g. Strang splitting). Since there is no such permutation in VLR implementations it can be at most first-order accurate and cross-derivatives terms can lead to stability issues depending on the time-stepping (Lemarié et al. 2020). Within the ALE framework these problems can be avoided because standard 3D advection can be used and stability issues associated with vertical advection could be handled via the Shchepetkin 2015 approach. More general V-ALE formulations in which the movement of the target grid is not purely Lagrangian appear to be attractive.

The specification of the target grid is a major challenge: one can specify isopycnals as the target grid or penalise departures from isopycnal slopes and grid smoothness as in Hofmeister et al (2010) or Gibson (2019). Alistair Adcroft has noted that it is very hard to find a globally satisfactory set of isopycnal target surfaces. More generally, the difficulty is that the choice of a satisfactory target grid corresponds to the resolution of an optimization problem jointly under 1D (e.g. sufficient resolution in boundary layers, avoid vanishing layers, etc) and 3D (e.g. regularity of vertical levels) constraints. International collaboration on a shared software tool for implementation and testing of schemes for specifying target grids would be very valuable. We propose to investigate the feasibility and mechanisms to realise this.

The scoping study mentioned earlier, has been focused on the above issues. The most active members of this group are Jerome Chanut, Laurent Debreu, Gurvan Madec, Knut Klingbeil, Florian Lemarié and Andrew Shao.

A strictly monotonic advection scheme, such as the Piecewise Parabolic Method (Colella & Woodward 1984) will be required to advect the thickness of very thin cells, such as occur in very strongly stratified water, even though we do not intend to allow mass-less cells. The choice of advection scheme of course needs to be consistent with the time-stepping scheme. Schemes to calculate the horizontal pressure forces specifically designed for the presence of sloping layer geometries, nonlinear equation-of-state and non-uniform vertical stratification profiles are also required (Shchepetkin & McWilliams, 2003; Adcroft et al., 2008; Engwirda et al., 2017, Bell & Young 2021). Mike Bell & Diego Bruciaferri will work on this. One issue is to decide whether options treating tracer values as cell-mean values (which would make the pressure forces less smooth) are essential.

It is highly desirable that new eddy closure parametrisations make relatively weak assumptions about the vertical coordinate. The impact of the GVC on the eddy closures will need careful evaluation. Coordination between the kernel and Eddy Closure WGs could be achieved through the 2 WGLs, Florian Lemarié & Andrew Shao.

The GVC code will need to be developed in parallel with the existing  $z^*$  and quasi-vanishing sigma coordinates, probably as a generalisation of the  $\tilde{z}$  co-ordinate. The arrays representing the coordinate system are quite generic, so we believe that they will not need further development. The costs of a GVC formulation could be significant so it may be desirable to keep the existing formulations as alternative options for the long-term. The current NEMO implementation of the  $\tilde{z}$  (V-ALE) coordinate has been recently reported in Appendix A of Megann et al (2022) and will be the starting point for future developments.

It is difficult to determine the size of each of the tasks outlined above but it is almost certainly substantial. Calculation of the pressure force is in principle the simplest of the tasks and that has grown into quite a complex and time-consuming study. The tasks involve literature reviews,

decisions on formulations, documentation of algorithms, design for implementation, their implementation with detailed checking, and testing in simplified then real-world configurations. The process is iterative. The members of the GVC scoping study and the team that implemented the RK3 scheme have the skills to do the work. Some others (e.g. Diego Bruciaferri and Chris Subich) have the skills and aptitude to contribute. But most of these people will not be able to devote significant time to the task unless they are specifically funded to do so. The task needs to be planned as a work-package within a project. Florian Lemarié is willing and best placed to coordinate this activity. Planning this work-package and finding funding for it is our first priority.

Some of the results obtained by Alex Megann (in [Megann et al 2020, 2022](#)) with  $z$ -tilde coordinates are somewhat puzzling: increasing the viscosity, using  $z$ -tilde and 4<sup>th</sup> order rather than 2<sup>nd</sup> order fct horizontal advection each reduce diapycnal mixing but only by about 10-15%; 4<sup>th</sup> order vertical advection has relatively little impact. Some caution is needed interpreting these results as they were obtained with the leapfrog scheme and a very specific choice of numerical options (e.g. very low viscosity values see Holmes et al 2021). Also there is not a clear consensus on how to quantitatively diagnose diapycnal mixing. These results should not stand in the way of developing more generic vertical coordinates but suggest that the effectiveness of schemes in reducing diapycnal mixing will need thorough evaluation.

### 7.3.2 Multi-Envelope $s$ -coordinates to improve accuracy of terrain-following models

With the Multi-Envelope method, computational surfaces are curved and adjusted to multiple arbitrarily defined surfaces (aka envelopes ) rather than following geopotential levels or the actual bathymetry. This allows one to define model levels that can be optimised for the prevailing physics (Bruciaferri et al. 2018). This approach has been successfully tested in several regional NEMO configurations against various type of vertical coordinates (e.g.,  $z$ -partial steps, vanishing-quasi sigma or hybrid  $s$ - $z$ ), and has been proven to be a flexible and yet robust and generalised method to improve the accuracy of regional models including shelf and open-ocean regimes (e.g., [Bruciaferri et al. 2020](#), [Wise et al. 2021](#), [Bruciaferri et al. 2022a](#)). In addition, the Multi-Envelope method has been lately combined with the idea of [Colombo 2018](#) to implement localised-Multi-Envelope  $s$ -coordinates in global configurations. This new approach has been successfully applied to improve the representation of the Nordic-seas overflows in a  $\frac{1}{4}$  degree NEMO global configuration ([Bruciaferri et al. 2023](#)). The same methodology is being currently tested to assess the sensitivity of the Western North Atlantic circulation to the vertical coordinate system in global ocean models (preliminary results for a  $\frac{1}{4}$  degree configurations can be found in [Bruciaferri et al. 2022b](#)). The Multi-Envelope approach could be also useful with AGRIF when changing the vertical coordinates between the parent and the child models.

The code includes three Fortran90 modules to be added to the DOMAINcfg tool and few python modules to generate the envelopes and the localisation areas. Up to know a branch of the NEMO DOMAINcfg tool has been created and all the relevant code has been included. Some future work by D. Bruciaferri is needed in order to include the modifications of this branch in the NEMO trunk.

### 7.3.3 Brinkman volume penalization for representing complex geometry

Brinkman penalisation is a method to implicitly enforce boundary conditions for complicated or moving geometries through the addition of specific source terms to the continuous dynamical equations. With this method the solid boundaries are treated as a porous medium whose representation depends on permeability and porosity parameters. In [Debreu et al, 2020](#) it is shown

that the total energy of the penalised primitive equations cannot increase (stability) and that constants are preserved (consistency). Moreover, at a discrete level, the method does not introduce any new stability constraints. There are several possibilities to choose the Brinkman penalization parameters (permeability and porosity). A possibility is to consider a smooth (envelope) terrain following (generalised) coordinates with more rapidly varying bathymetry (lost during the smoothing procedure) represented using porosity and permeability settings. The approach also allows to eliminate the step-like detrimental effects associated with z-coordinates. It should however be noted that a lot of effort is required to specify the penalization parameters in a systematic way for a high resolution, large area model.

The methodology has been implemented in the Croco ocean model in both idealised and realistic settings and is under development in idealised settings in NEMO (PhD work of A. Nasser). In particular, [Debreu et al. 2022](#) show improved Gulf Stream separation at 0.25° resolution. There is also potential to represent ice-shelves with penalisation and embedding sea-ice in the ocean by using a time-varying penalisation. Since the penalisation is applied to the continuous equations written for generalised vertical coordinates, this approach is virtually compatible with any type of vertical coordinate. The Brinkman penalisation and V-ALE work will be able to proceed in parallel.

#### 7.4 Energy efficiency

Nowadays, given the increase of energy costs and the need to adopt environmentally friendly practices, model developers must deal with a new technological paradigm to keep under control the carbon footprint of numerical simulations. The (time/energy/cost)-to-solution for a given effective resolution is an increasingly important metric to evaluate a given numerical code (e.g. [Kalinnik et al, 2021](#)). Besides the software environment, important drivers affecting the energy-to-solution are directly related to the dynamical kernel and include the time-integration strategy as well as the dissipative/dispersive properties of numerical schemes. The environmental constraint on oceanic dynamical kernels thus necessitates to re-assess existing algorithmic strategies ([Mengaldo et al., 2019](#)) and to keep a technological watch on hardware evolution.

From this perspective the evolution from leapfrog to RK time-stepping goes in the right direction: the effective resolution is improved, the code runs ~50% faster and there are no tuning parameters like the one associated to the Robert-Asselin filter. Several initiatives around the NEMO kernel could be envisioned to help for a responsible use of the code:

→ More systematic evaluation and documentation of the cost vs benefits of the particular options available in the code (in particular for vertical physics and advection schemes). The more expensive schemes can serve as a reference.

→ Closer collaboration with the HPC WG to determine which algorithmic choices fit the best on given architectures.

→ Thinking in terms of effective resolution vs computational cost, following [Sanderson \(1998\)](#) it is beneficial (and even optimal) to use third or fourth order schemes (however such study would deserve to be updated).

#### 7.5 Additional issues

The following points have been raised during the preliminary discussions on the development strategy. The first two points (relaxing the Boussinesq and/or the hydrostatic assumptions) are mentioned because there are frequent questions about their feasibility and the extent of the associated developments.

### 7.5.1 Hydrostatic Boussinesq/Non-Boussinesq options

Simulations of tides are sensitive to the mean density of the water and could be slightly improved by relaxing the Boussinesq approximation. It is often considered that non-Boussinesq (hydrostatic) equations could be easily used by virtue of the isomorphism between the z-coordinate Boussinesq and pressure-coordinate non-Boussinesq systems (DeSzoeke & Samelson, 2002). However such isomorphism exists only under the rigid-lid assumption and does not directly apply to modern oceanic models based on a mode-splitting algorithm with a prognostic free-surface. Another possibility, following Greatbatch et al., 2001 is to replace the discretized vertical volume factor  $\rho_0\Delta z$  by  $\rho\Delta z$ . Relaxing the Boussinesq assumption also requires the use of full dynamic pressure in a fully compressible, realistic EOS (acoustic waves being still excluded via the hydrostatic assumption). Generally speaking, the Boussinesq assumption removes several interdependencies within model components (e.g. it filters out conversion between internal and kinetic/potential energy) and it is not as straightforward as it seems to relax it (Shchepetkin & McWilliams, 2011). As a first step, a thorough understanding of Shchepetkin & McWilliams 2011 and of the non-Boussinesq formulation in MOM6 seem necessary.

### 7.5.2 Non-hydrostatic option

The implementation of a non-hydrostatic option within an existing hydrostatic ocean model has profound implications for the code: either the use of an external library to solve efficiently the corresponding 3D elliptic boundary value problem in the Boussinesq (incompressible) case or the integration of a (very) fast acoustic mode in the non-Boussinesq (“pseudo” compressible) case. The “pseudo” compressible approach is implemented in the CROCO model (Auclair et al., 2018) which uses an additional super-fast level of time-step splitting for the 3D acoustic mode and is very efficient on massively parallel architectures. Such a fast 3D mode allows to integrate other terms raising stability/accuracy issues, e.g. bottom friction, vertical advection, non-traditional Coriolis terms. Overall this is a major undertaking in terms of development and the cost vs benefits ratio given the typical NEMO user requirements is unclear. Moreover, for current operational applications non-hydrostatic effects are of secondary importance. In case it turns out to be possible to incorporate a non-hydrostatic option in a modular, maintainable, way with little impact on code performance we should scope out the ideas with the aim of incorporating them as key elements in the next 5-year strategy.

In the short term, the possibility of moving to an AGRIF mother-child interface with the CROCO code could be considered as a potentially viable alternative to developing an internal NEMO NH capability. Another possibility instead of resolving explicitly NH effects would be to parameterise them. There are ongoing initiatives to use machine learning techniques to do so.

### 7.5.3 Physics/dynamics coupling

The consistency between the kernel & parametrisations particularly for the transfer of energy between resolved and sub-grid scale forms should be borne in mind (e.g. Burchard 2002; Marsaleix et al. 2008, Eden et al 2014). TKE and GLS do this whilst OSMOSIS does not. A broader reflection on physics/dynamics coupling is provided in Gross et al. (2018).

### 7.5.4 Unintrusive developments

An improved scheme for momentum advection TVD, monotonicity-preserving or WENO (as well as for tracers) is desirable to regularize the velocity field with an expected reduction of numerical diapycnal mixing.



The Shao et al (2020) isopycnal diffusion scheme can simply slot in as an alternative option.

A form of biharmonic GM based on the Greatbatch & Lamb (1990) vertical mixing of momentum formulation is being developed. Again this should be relatively unintrusive.

Only the vector invariant form for surface wave - mean flow interactions has been coded up (Couvelard et al., 2020). Implementing the flux form given by Couvelard et al (2020) is relatively straightforward (unintrusive). The Bennis et al 2011 test case should be implemented.

## 7.6 References

Adcroft, A., R. Hallberg and M. Harrison (2008). A finite volume discretization of the pressure gradient force using analytic integration. *Ocean Modelling* 22, 106-113.

Auclair F., L. Bordois, Y. Dossmann, T. Duhaut, A. Paci, C. Ulses and C. Nguyen (2018): A non-hydrostatic non-Boussinesq algorithm for free-surface ocean modelling, *Ocean Modelling*, Volume 132

Bell M. J. and A. Young (2021): Accurate calculation of pressure forces on cells defined by steeply sloping coordinates. Immerse Deliverable D3.3.

Bennis A. C., F. Ardhuin and F Dumas (2011): On the coupling of wave and three-dimensional circulation models: Choice of theoretical framework, practical implementation and adiabatic tests. *Ocean Modelling* Volume 40, Issues 3–4, Pages 260–272 <http://dx.doi.org/10.1016/j.ocemod.2011.09.003>

Bruciaferri, D., Shapiro, G.I. and Wobus, F. (2018) A multi-envelope vertical coordinate system for numerical ocean modelling. *Ocean Dynamics*, Volume 68(10), Pages 1239-1258, <https://doi.org/10.1007/s10236-018-1189-x>

Bruciaferri, D., Shapiro, G., Stanichny, S., Zatsepin, A., Ezer, T., Wobus, F., Francis, X., Hilton, D. The development of a 3D computational mesh to improve the representation of dynamic processes: the Black Sea test case, *Ocean Modelling*, 146 (2020), 101534, <https://doi.org/10.1016/j.ocemod.2019.101534>

Bruciaferri, D., Tonani, M., Ascione, I., Al Senafi, F., O'Dea, E., Hewitt, H. T., and Saulter, A. GULF18, a high-resolution NEMO-based tidal ocean model of the Arabian/Persian Gulf, *Geosci. Model Dev. Discuss.* [preprint], <https://doi.org/10.5194/gmd-2022-189>, in review, 2022a.

Bruciaferri, D., Hewitt, H.T., Bell, M.J., Guiavarc'h, C., Storkey, D., Roberts, M.J. and Jackson, L. Sensitivity of the Western North Atlantic circulation to the vertical coordinate system in global ocean models. 'Whither the Gulf Stream' Workshop - Woods Hole, MA, USA, June 2022b (<https://usclivar.org/sites/default/files/2022/posters/Diego-Bruciaferri-Poster.pdf>).

Bruciaferri, D., Guiavarc'h, C., Hewitt, H.T., Harle, J., Almansi, M., Mathiot, P. Improving the overflows representation in global ocean models: the Nordic seas test case - in preparation for *Journal of Advances in Modelling Earth Systems* - 2023

Burchard, H. (2002): Energy-conserving discretisation of turbulent shear and buoyancy production. *Ocean Modelling* - 4. 347-361.

Colella P., and P. R. Woodward (1984): The piecewise-parabolic method (PPM) for gas dynamical simulations. *J. Comp. Phys.* 54, 1: 174–201.

Colombo, P. (2018) - Modelling dense water flows through sills in large scale realistic ocean models : demonstrating the potential of a hybrid geopotential/terrain-following vertical coordinate <http://www.theses.fr/2018GREAU017#>

Couvelard, X., F. Lemarié, G. Samson, J.-L. Redelsperger, F. Ardhuin, R. Benshila and G. Madec (2020): Development of a two-way-coupled ocean–wave model: assessment on a global NEMO(v3.6)–WW3(v6.02) coupled configuration, *Geosci. Model Dev.*, 13, 3067–3090,

Debreu, L., N. Kevlahan and P. Marchesiello (2020): Brinkman volume penalization for bathymetry in three-dimensional ocean models. *Ocean Modelling*, 145, doi:10.1016/j.ocemod.2019.101530.

Debreu L., N. Kevlahan, P. Marchesiello (2022). Gulf Stream separation through Brinkman penalization. <https://hal.archives-ouvertes.fr/hal-03593222/>

de Szoeke, R. A. and R.M. Samelson: The duality between the Boussinesq and non-Boussinesq hydrostatic equations of motion. *J. Phys. Oceanogr.* 32, 2194–2203 (2002).

- Ducouso N., F. Lemarié, L. Debreu and G. Madec (2021): Stability and accuracy of Runge–Kutta-based split-explicit time-stepping algorithms for free-surface ocean models. *Journal of Advances in Modeling Earth Systems* (in review)
- Eden, C., Czeschel, L., and Olbers, D. (2014). Toward energetically consistent ocean models. *J. Phys. Oceanogr.* 44, 3160–3184. doi: 10.1175/JPO-D-13-0260.1
- Engwirda D., M. Kelley, and J. Marshall (2017): High-order accurate finite-volume formulations for the pressure gradient force in layered ocean models, *Ocean Modelling*, 116.
- Ezer, T. (2016): Revisiting the problem of the Gulf Stream separation: On the representation of topography in ocean models with different types of vertical grids. *Ocean Modell.*,104,15–27.
- Gibson A.: An adaptive vertical coordinate for idealised and global ocean modelling, *PhD thesis, Australian National University*, (2019)
- Greatbatch, R. J. and K.G. Lamb (1990): On Parameterizing Vertical Mixing of Momentum in Non-eddy Resolving Ocean Models, *Journal of Physical Oceanography*, 20(10), 1634-1637
- Greatbatch, R. J., Y. Lu and Y. Cai (2001) : Relaxing the Boussinesq Approximation in Ocean Circulation Models, *Journal of Atmospheric and Oceanic Technology*, 18(11), 1911-1923
- Griffies, S., A. Adcroft and R. Hallberg (2020): A Primer on the Vertical Lagrangian-Remap Method in Ocean Models Based on Finite Volume Generalized Vertical Coordinates. *Journal of Advances in Modeling Earth Systems*. 12.
- Gross, et al. (2018): Physics–Dynamics Coupling in Weather, Climate, and Earth System Models: Challenges and Recent Progress, *Monthly Weather Review*, 146(11), 3505-3544
- Hofmeister, R., H. Burchard and J.-M. Beckers: Non-uniform adaptive vertical grids for 3D numerical ocean models. *Ocean Modelling*, 33. 70-86. (2010)
- Holmes, R. M., J. D. Zika, S. M. Griffies, A. McC. Hogg, A. E. Kiss and M. H England (2021): The Geography of Numerical Mixing in a Suite of Global Ocean Models. *Journal of Advances in Modeling Earth Systems*, 13, e2020MS002333. <https://doi.org/10.1029/2020MS002333>
- Kalinnik, N., R. Kiesel, T. Rauber, M. Richter and G. Rünger (2021): A performance- and energy-oriented extended tuning process for time-step-based scientific applications. *The Journal of Supercomputing*.
- Kevlahan, N. K. -R. and Dubos, T. and Aechtner, M. (2015): *Adaptive wavelet simulation of global ocean dynamics using a new Brinkman volume penalization*. *Geosci. Model Dev.*, 8, 12, 3891–3909.
- Lemarié et al. (2020): *D3.1 - Algorithm specification and code design for 2LTS scheme*, IMMERSE Deliverable
- Marsaleix, P., F. Auclair, J.-W. Floor, M. Herrmann, C. Estournel, et al. (2008): Energy conservation issues in sigma-coordinate free-surface ocean models. *Ocean Modelling*, 20 (1), pp.61-89
- Megann, A. and D. Storkey (2021): Exploring viscosity space in an eddy-permitting global ocean model: Is viscosity a useful control for numerical mixing? *Journal of Advances in Modeling Earth Systems*, 13
- Megann, A., J. Chanut and D. Storkey (2022): Assessment of the  $z\sim$  time-filtered Arbitrary Lagrangian-Eulerian coordinate in a global eddy-permitting ocean model. Submitted to JAMES.
- Mengaldo, G., A. Wyszogrodzki, M. Diamantakis, S.J. Lock, F. Giraldo and N. Wedi (2019): Current and emerging time-integration strategies in global numerical weather and climate prediction. *Arch. Comput. Method E*. 26, 663-684
- Nudds S. and co-authors. (2020): Evaluation of Structured and Unstructured Models for Application in Operational Ocean Forecasting in Nearshore Waters *J. Mar. Sci. Eng.* 2020, 8, 484; doi:10.3390/jmse8070484.
- Sanderson, B. G. (1998). Order and Resolution for Computational Ocean Dynamics, *Journal of Physical Oceanography*, 28(6), 1271-1286
- Schoonover, J., Dewar, W. K., Wienders, N., & Deremble, B. (2017). Local Sensitivities of the Gulf Stream Separation, *Journal of Physical Oceanography*, 47(2), 353-373
- Shao, A., A. Adcroft, R. Hallberg and S. Griffies (2020): A General-Coordinate, Nonlocal Neutral Diffusion Operator. *Journal of Advances in Modeling Earth Systems*. 12.

Shchepetkin, A.F. and J.C. McWilliams (2003): A method for computing horizontal pressure-gradient force in an oceanic model with a nonaligned vertical coordinate. *J. Geophys. Res.* 108, C3, 30390.

Shchepetkin A.F. and J. C. McWilliams (2011): Accurate Boussinesq oceanic modeling with a practical, “Stiffened” Equation of State, *Ocean Modelling*, Volume 38, Issues 1-2

Shchepetkin A.F. (2015): An adaptive, Courant-number-dependent implicit scheme for vertical advection in oceanic modeling, *Ocean Modelling*, Volume 91, 2015, Pages 38-69.

Styles, A., M. Bell, D. Marshall, D. Storkey (2021): Spurious forces can dominate the vorticity budget of ocean gyres on the C-grid. *JAMES*, 14, e2021MS002884, doi:10.1029/2021MS002884.

Wise, A., Harle, J., Bruciaferri, D., O'Dea, E., Polton, J. The effect of vertical coordinates on the accuracy of a shelf sea model. *Ocean Modelling*, 170, 101935. <https://doi.org/10.1016/j.ocemod.2021.101935>.

## 8 Eddy closures

### 8.1 Executive Summary

The Eddy Closure Working Group (ECWG) recognises that many Consortium members are targeting both eddy-permitting 1/4-degree and 1-degree global simulations for their applications. The latter will likely continue to be used for Earth system and paleoclimate modelling whereas the former will be used for coupled climate modelling and to provide boundary conditions for regional models. To ensure that NEMO remains an appropriate model for these applications, the ECWG has identified a dual need to bring up to date the existing eddy closures (parameterisations) in NEMO, which have not evolved significantly for many years, and to facilitate the continued development of future eddy closures, including scale-awareness.

The WG recognises that facilitating the sharing of new eddy closure schemes throughout the broader community of NEMO users (outside the NEMO Consortium) is crucial for NEMO's use as a proving ground for novel research as well as integrating them into current use cases. We thus propose the following structure, to encourage community development efforts that allow NEMO's suite of eddy closures to be state-of-the-art without requiring undue burden on the Systems Team. The ECWG proposes to create a code-sharing framework, making use of sandboxing, which is complementary to the NEMO Gitlab framework, and facilitates development, testing, and evaluation of eddy closures. In particular, the ECWG will provide a suite of test cases which the community may use to evaluate eddy closures. This suite will range in complexity from idealised simulations to realistic global configurations. Additionally, the ECWG will develop metrics which may be used as (idealised and realistic), indicators of eddy parameterisation skill. These metrics will target important processes and known biases which are common to most non-eddy resolving simulations, (e.g. ensemble variability, AMOC variability). The availability of standardised test cases and metrics will facilitate easy comparison between eddy parameterisations. An online collaborative space will offer users a forum to share and test source code outside of the official NEMO trunk which the WG sees as crucial to entraining new developments into the NEMO ecosystem. From there, the Working Group may recommend that internal NEMO resources be used to elevate specific parameterisations based on proven scientific benefit and impact on the overall codebase. This represents a departure from the typical NEMO development process, but is designed to capture the extant enthusiasm of the broader community. Via this community forum, the ECWG will be reaching out to the NEMO community to provide two global configurations based on eORCA025 and eORCA1. These will not be tested as part of the System Team's work, but some output and metrics will be provided to the community serving as a baseline configuration as a comparison for new parameterisations.

Another aim of the working group is to recommend specific eddy closures for adoption into the NEMO trunk. The set of closures currently included in the NEMO trunk is quite limited compared to other widely-used ocean models. For example, two distinct approaches have arisen in recent years to move us beyond the era of ad-hoc application of the Gent and McWilliams, 1990 (GM) mesoscale closure: 1) the GEOMETRIC framework (Marshall et al., 2012), which specifies the required symmetry structures of an eddy forcing tensor needed to achieve desired conservation properties of the dynamics, i.e. a recipe that parameterisation schemes should adopt, rather than a solution in itself; and 2) the advent of "kinetic energy backscatter" schemes, to parameterise the missing sub- and near-gridscale eddy effects in the inverse cascade of geophysical turbulence. At least one candidate closure scheme in each of these two categories has been identified. The Mak et al. (2018) scheme is based on GEOMETRIC and parameterises eddy buoyancy fluxes via a GM eddy transfer coefficient, which is constrained by a prognostic equation for the evolution of sub-gridscale eddy energy. This has shown great promise and is being implemented in the NEMO trunk. The WG recommends continued support for this effort and future efforts based on these approaches. There

are a few very promising avenues to explore for KE backscatter, including an example already implemented in an idealised configuration of NEMO, described later.

The issue of scale-aware eddy closures, which adapt depending on the relative ability of the model grid to resolve the relevant dynamics, is present to some extent in some closures, but is a desirable effect to seek to include. A relatively simple way to include this effect for mesoscale eddy closures is to apply the Hallberg (2013) resolution function. The ECWG recommends this approach if scale-awareness is not already contained and it is a high priority for implementation and testing.

The ECWG also recognises the ongoing importance of diffusive, down-gradient eddy closures. Some of these include new approaches to determine eddy diffusivities (e.g. Groeskamp et al. (2020) for isoneutral mixing) and alternate formulations of the parameterisations themselves (i.e. Ferrari et al. (2010) for Gent-McWilliams-style eddy transport or Roberts and Marshall (1998) for a recipe for biharmonic GM). The ECWG aims to compile and maintain a list of eddy parameterisations for inclusion in the NEMO trunk which is likely to include many of those discussed above.

## 8.2 Introduction

### 8.2.1 Assessment of the 2018-2022 Development Strategy

The previous strategy identified as a key oceanic parameterised physical process, “the representation of the impact of balanced turbulence (mesoscale and submesoscale eddies) on larger scale/lower frequency flows [...] this field of research is very active [...] with several new approaches [...] (scale-aware closures, stochastic closures, energy backscatter, approaches from LES)”. It also proposed the formation of an Eddy Closure Working Group (ECWG), which first met in September 2021.

### 8.2.2 Working Group members

The current membership of this working group covers a very broad range of interests and expertise, including observational analysis (O), theory (T), modelling (M), deterministic (D), stochastic (S) and alternative (A) approaches to eddy parameterisation. It also includes members of the Systems Team (ST).

**Chairs:** Andrew Shao (O,M,D,A), Chris Wilson (O, T, M, D, S)

**System Team Members:** Jérôme Chanut (M,D), Andrew Coward (M)

**General:** Scott Bachman (T,M,D), Ekaterina Bagaeva, Pavel Berloff (T,M,D,S,A), Casimir de Lavergne, Julie Deshayes, Carsten Eden, Helene Hewitt, Andy Hogg, Malte Jansen (T,M,D), Stephan Juricke (M,D,S), Milan Kloewer (M,D), Julien Le Sommer, Till Kuhlbrodt (O,T,M,D), Julian Mak (T,M,D), David Marshall (T,M,D), Pavel Perezhogin (M,D,S), Zofia Stanley (S), Dave Storkey, Andrea Storto (S, O), Anne-Marie Treguier, Takaya Uchida (T, M, D), Robin Waldman (T, M, D), Stephanie Waterman

### 8.2.3 Goals and objectives of the working group

The primary goal of this working group is to determine which eddy closures will best meet the needs of the community of NEMO users. In particular, scale-aware parameterisations of meso- and submeso-scale turbulence are crucial to the continued use of NEMO. While some NEMO applications are able to resolve mesoscale eddies, it remains computationally infeasible to run global simulations used for climate and Earth System modelling at eddy-resolving resolutions. For instance, resolving the mesoscale on the coastal shelves and in the Arctic requires O(1km) resolution. This

computational limitation, and hence the need for meso- and submeso-scale parameterisations, is likely to remain for the foreseeable future.

Eddy-permitting (1/4-degree) and laminar (1-degree) simulations of the kind used for global operational weather forecasting and climate simulations are the primary resolutions where the working group has identified a need for development in the NEMO codebase. Some well-established eddy closures that are useful for the laminar regime (e.g. Gent-McWilliams and Smagorinsky viscosity) already exist within NEMO. This working group draws on research from the active modelling and parameterisation community in its recommendations for further parameterisations and new theoretical frameworks which target the eddy-permitting regime as well.

The objectives of this working group are driven by desired scientific applications of current and future NEMO users. To this end, the working group aims to facilitate a structured environment where the community can develop, test, evaluate, and experiment with new parameterisations. In pursuit of this aim, the working group has identified two concrete, attainable outcomes which will provide large benefits to the community.

The first outcome is to provide a suite of test cases which the community may use to evaluate eddy parameterisations. The working group will choose test cases which range from idealised scenarios (some of which may have analytical solutions or are otherwise verifiable) to realistic global simulations. These cases will be carefully chosen to cover a wide range of research interests. A researcher developing a new parameterisation may pick a test case which captures the phenomenon or phenomena they are interested in, add their bit of code, run the model, and compare to the published baseline. The established test cases will not be exhaustive and researchers will still be encouraged to set up their own baseline runs when beneficial. However, establishing a common suite of test cases will save an enormous amount of time for researchers who no longer have to set up and validate, or ask a collaborator for, model runs which serve only as comparisons. Further, researchers will have a centralised place to look for test cases. It will also save computational time as each test case need only be run once in its baseline configuration.

The second outcome identified by this working group is to provide a set of eddy-related metrics which may be used to evaluate eddy parameterisations. The metrics target important processes and known biases which are common to most non-eddy resolving global simulations. As with the test cases, the metrics will not be a complete list, but rather a starting point for researchers to understand if their parameterisations are likely to improve key aspects of the model. Standardising these metrics and making them easily accessible will save researcher time in that each researcher need not develop their own measure of, for example, AMOC variability. Together with the established test cases, established eddy metrics will also facilitate comparisons between different parameterisations. Two researchers who both use the same test case and eddy metrics will easily be able to discuss advantages and disadvantages to each approach, perhaps generating new ideas for additional improvement.

Further, this working group suggests that these test cases and eddy-related metrics be shared in a code-sharing framework which will facilitate and easily ingest community contributions. This code-sharing framework would be complementary to the core NEMO framework. Parameterisations which are assessed using the test cases and metrics described above will be more easily evaluated by the working group. The working group will make recommendations about implementations of particular scientific merit which may be incorporated into the NEMO codebase. In addition, the test cases will provide the core of a code-testing subsystem which will ensure that newly implemented parameterisations are compatible with NEMO releases.

Finally, this working group will make specific recommendations about existing parameterisations which could be incorporated into the NEMO trunk. These recommendations consider scientific impact as well as their maturity of development, documentation of any testing that has been done, and potential difficulty of implementation and/or impact on the NEMO codebase.

#### 8.2.4 Synergistic activities within the broader oceanographic community

The eddy closure problem is a difficult one, and as with most difficult problems, benefits from fruitful collaborations. Such collaborations provide both intellectual and practical resources needed to tackle the theoretical and computational tasks involved in eddy closure research. Therefore, part of the role of the working group includes providing a forum to support the building and sustaining of collaborations, encouraging members to work together to identify funding opportunities. The membership includes those who are early in their careers as well as senior members to encourage two-way knowledge transfer.

Significant coordinated efforts are underway in the international community, particularly the Eddy Energy Climate Process Team in the United States and the Energy Transfers in the Atmosphere and Ocean (TRR181) group in Germany, to both suggest new eddy parameterisations and also develop idealised test cases. This working group differs from these research and development efforts as it offers suggestions for NEMO specifically. As such, the working group will continue to evaluate the state of these efforts and determine what can be adapted for use in NEMO. The working group intentionally includes members of these two consortia to ensure that activities within the working group are complementary. Of primary interest is to evaluate which of the metrics and test cases identified by these groups would be useful to include in NEMO.

### 8.3 Development of test cases to evaluate eddy parameterisations

The working group recognises the need for test cases to provide a consistent benchmark to compare new parameterisations that might be developed within the community (see Section 4 for a discussion of metrics). This WG is thus proposing to establish a set of test cases that span both physical and computational complexity to provide a useful development and testing environment for researchers. This can largely be split into two categories, described below.

#### 8.3.1 Idealized simulations

Simple simulations have the benefit that their behaviours are well understood and are usually computationally cheap enough that most scientists can run on a small amount of computational resources. Relevant test cases would be a baroclinic channel to test the spinup and spindown of baroclinic turbulence and the 'double-gyre' model already included with NEMO. The NEMO GYRE model can be run at multiple resolutions, is easy to configure with a wide range of forcing regimes, and by default is coupled to a biogeochemical (BGC) model, and calculations have been reported over the past decade at least (e.g. Lévy et al., 2010 - up to 1/54 degree; Couespel et al., 2021 - with BGC). Additionally, the Eddy Energy CPT has developed a 'Neverworld 2' test case, which is a model of intermediate realism that has analogous features to the Atlantic basin and Southern Ocean (Marques et al., 2022). Other idealised cases may be considered if warranted by feedback from the broader community. These and other published test cases will be reviewed and recommended as benchmarks for simple tests of eddy parameterisations. As described below in the discussion of eddy diagnostics and metrics, there is an aim to engage with the community, ideally via discussions leading to a collaborative paper, to define shared testing and evaluation methods, and idealised simulations also fit into this approach. Putting a new eddy closure through a set of idealised simulations and metrics, including some which may not have been in the minds of the original developer of the closure, is a good approach to 'pressure testing' the closure in various dynamical

regimes. There needs to be a mechanism (likely agreed with the community engagement WG) to share and disseminate these test cases especially those which require large amounts of input data. If peer-reviewed publications are not appropriate for any of these cases, we will recommend that datasets/results are published to a citable, DOI-issuing archive, such as Zenodo.

### 8.3.2 Realistic simulations

Global simulations are the most relevant applications to the climate and weather domains. An informal poll of the WG membership has identified that the major modelling centres and academic researchers are targeting both eddy-permitting ( $\frac{1}{4}$ -degree) and laminar (1-degree) ocean simulations in their next generation coupled models. The eddy-permitting case is particularly challenging because the mesoscale is partially resolved, resulting in a mix of requirements where some eddy effects must be parameterised, but where traditional closures can either double-count the effect of eddies or suppress resolved eddies and jets.

Currently NEMO 'officially' only supports one realistic, global configuration - a 2-degree model driven by CORE forcing, which is not really relevant to CMIP6 onwards, where 1-degree and higher is more typically used. This coarse resolution has the benefit of being able to be run on a small number of cores and so appropriate for testing and code validation. The inherent diffusiveness associated with this grid size limits the testing to diffusion-based eddy closure schemes and is inappropriate for 'scale-aware' parameterisations.

Outside of the official repository, a number of groups share 1-degree and  $\frac{1}{4}$  degree configurations. These require significant high-performance computing resources to run multi-decadal or ensemble simulations but will remain the 'gold standard'. A goal of this working group is thus to recommend 1-degree and  $\frac{1}{4}$ -degree configurations with a recommendation for particular choices for existing parameterisations within NEMO (e.g. Gent-McWilliams, isoneutral mixing, and a particular viscosity). We will also investigate ways of providing a standardised model output; nominally we will aim to use a reference configuration based on an existing CMIP5/6 contribution with published OMIP data. Full namelists and input data would need to be shared - we will aim to use NEMO Gitlab for this, but may need to augment with large input files archived elsewhere, e.g. Zenodo. These configurations will serve as the baseline for eddy-related metrics to evaluate new parameterisations. Due to the computational expense of a global  $\frac{1}{4}$ -degree model, the ECWG will investigate the feasibility of providing a regional configuration (e.g. an Atlantic Sector or Southern Ocean model) that may be of sufficient complexity to recreate the main dynamical features (e.g. jets, boundary currents, etc.) that are of broad relevance.

## 8.4 Diagnostics and metrics for assessment of eddy parameterisations

Assessing the skill of eddy closures on a non-eddy-resolving model is crucial for being able to decide whether to focus attention and resources on one particular scheme over another, or whether a particular eddy closure would improve model skill more for one particular question than another. In contrast to diagnostics which are quantities that enable validation or comparison of one simulation against another, here we define metrics as diagnostics which may further be verified against observations or theory. Following initial discussion within the WG, we have come up with a shortlist of eddy-related diagnostics and metrics, i.e. characteristic properties of the ocean state that we care about (often large-scale climate, but not solely) and which we believe are affected by the eddy processes not being correctly simulated. However, there are many other groups and individuals working on the eddy closure problem, both outside of this WG and also using models other than NEMO, and we may benefit from interacting with them. The Scientific Advisory Group was



supportive of our goal to devise a set of shared eddy-related metrics with other major groups (such as the US and German consortia and other modelling centres), but recommended that we pursue a peer-reviewed, collaborative paper on this topic - not an easy goal, but a worthwhile one. Ideally, we would like to achieve such a goal within the first two years but, if a paper proves impossible, we should at least have explored the topic with a discussion and agreed and defined (both in model-agnostic and NEMO-specific terms) the diagnostics and metrics. A challenge here would simply be to get buy-in from the other modelling groups and to convince them of the potential benefit, but we will seek to do so, first by putting forward a summary of our position and motivation, before setting up meetings to take the next steps.

It is expected that a subset of our shortlist would be on any agreed intersection of preferred diagnostics and metrics across the international modelling community. We would therefore focus on developing and applying the subset, for intercomparison, but may retain one or more outside the subset if they would be useful within the NEMO community.

The definitions of the metrics as "model-agnostic", e.g. transports across a full-depth section between two lon-lat points, are not only intended to be independent of basic differences in model grids, but to also permit observational dataset diagnostics and comparison.

Other key points for further discussion with the community are:

1. The role of averaging operators for defining the "eddy" component of the ocean state - smoothing in space/time vs ensemble averaging. Specifically, should parametrizations target spatiotemporal variability or intrinsic variability of the eddies (although not necessarily mutually exclusive)?
2. How to include the component of chaotic intrinsic variability in assessing skill using metrics - are ensemble experiments always needed?
3. The potential resource burden on some researchers, e.g. ensemble simulations for eORCA025 might be too large a computation, although a single realisation may not be; and some energy-based diagnostics may be informative even without an ensemble approach.

Here is the shortlist of eddy-related diagnostics and metrics proposed by this WG, to be refined as proposed above:

Note that there may be inter-dependencies between these individual diagnostics and metrics. Ideally, an eddy closure would, without much tuning, be able to improve the model skill simultaneously.

- Representation of the intensity and shape of jets (e.g. Gulf Stream extension and Azores Current) and of the extension of mesoscale-active regions (In non-eddy-resolving models, these are often too weak and too broad, with a limited extension)
- Energy spectrum and transfers of energy between resolved, parameterised, and turbulent reservoirs
- Climate variability and ensemble spread (often underestimated and under-dispersive)

- Southern Ocean eddy saturation and eddy compensation (which may be important for transient heat and carbon transports and budgets) in non-eddy models.
- Realism of the circulation on the Antarctic and Greenland coastal shelves, as a boundary condition for dynamic ice-sheet models
- Mean state and variability of AMOC, ACC transport and global stratification (mean and variability are often biased)
- Ocean heat uptake particularly for projecting changes under global warming (often biased)

The ECWG will work with the Tools, Verification and Validation group to develop tools to define and diagnose these metrics and provide a standardised analysis package that the community can run.

## 8.5 Cultivating a community of developers

### 8.5.1 Use of Gitlab and sandboxing

NEMO is increasingly being used and developed outside of the traditional set of NEMO Consortium members. Eddy closures have been a particularly popular research avenue for such members. Moving forward, it would only benefit NEMO development to make it easier for these developers to share source code, diagnostics, metrics and test cases that are not necessarily yet officially supported by the NEMO System Team at the working group level. This serves the dual purpose of establishing NEMO as a common framework to develop such schemes while also mobilising developments across the entire community. The recent migration of NEMO from the SVN-based forge to Gitlab further provides an opportunity for such collaboration and to provide recommendations for natural entry points for researchers interested in developing eddy closures

This working group proposes to establish a community-oriented package of eddy closures to be made available either on the NEMO-managed Gitlab, or another public platform. The barriers for inclusion in this space will be by default fairly low, in contrast to the ‘formal’ process by which developments in other aspects of NEMO are included into the trunk. This space will be a supervised playground where researchers may work on creative solutions to the eddy closure problem with tools to test and compare between emerging parameterisations. In the eddy parameterisation community, with its numerous, concurrent research activities all targeting different processes and biases, it is crucial that developments from the community be shared and incorporated in a bottom-up development style. For this to be successful, the eddy closure code space must be accessible. This working group sees a great need for effective training on how to work within the NEMO development workflow. This includes documenting best practices for communication, both within the eddy closure code space and between this space and the main NEMO-managed Gitlab. We find that no matter how good the documentation, development in a new environment is only possible with the help of a knowledgeable mentor. This is especially true for students and early career scientists who are just getting familiar with collaborative coding practices. Therefore, this working group recommends that we actively facilitate connections between new community members and experts in the NEMO codebase who are willing to serve as mentors. This will encourage new voices and new ideas, while ensuring that NEMO standards are upheld.

An additional part of this activity will be to have a repository where users can access the common set of test cases (Section 3) and a common set of scripts that implement best practices to calculate diagnostics and metrics (Section 4).

The ambition is that this will simplify the systematic intercomparison of closures that are not yet available into NEMO, therefore simplifying the experiments/workflow for an intercomparison paper. Within NEMO Gitlab, the aim is to set up a 'sandbox' for this WG to enable the sharing and interactions described above.

The sandbox should probably be a collection of repos more than a single repos (so that they can be maintained independently, especially if they started from different NEMO versions). All of them should follow a similar structure (and build procedure). The sandbox should only store files that have been modified wrt the NEMO main codebase (in order to avoid the risk of a “shadow NEMO version”). The sandbox (or rather each of its repos) should point toward one and only one NEMO tag. The build process could use the `makenemo -u` procedure which allows the use of user-provided code in the compilation process. An example of a similar type of NEMO sandboxing (but with Github, rather than Gitlab) has been used by NOC for several years and may provide a helpful basis to build upon. In NOC's example, the makenemo script is adapted to download large files (that can't be stored on Github/Gitlab) from external repositories.

The sandbox could also store the WG-relevant test cases / configurations and metrics/diagnostics. The structure of the sandbox should probably be agreed by (or elaborated with) the NEMO Tools Verification & Validation WG.

The creation of a new repos should be as smooth as possible for users and without too much requirement for “human” operation. Documentation for the test cases can be included as README files etc on the repository, eg. as here [https://github.com/JMMP-Group/AMM7\\_surge](https://github.com/JMMP-Group/AMM7_surge). We need a “champion” to take over the responsibility of proposing a first structure and ideally they should probably be the person leading the discussions/paper on eddy-related diagnostics and metrics.

This WG has identified some potential issues and open questions to be addressed in the first 2 years:

- maintenance : do we need someone to be in charge of phasing the code available in the sandbox with the evolution of the NEMO reference codebase ?
- support : should someone be in charge of providing support on the code stored in the sandbox ?
- phasing : we will need to find a way to phase to different repos from the sandbox from time to time so that they start from the same NEMO tag
- we need a mechanism for discussing the (non-local) large scale code changes that may be required for implementing specific closures (or simply bug fixing the reference NEMO codebase to which the sandbox is pointing)
  - A good way in git is to create a pull request for a branch/set of changes which provides a forum for discussion.
- should the sandbox be open and visible to anyone or should it be restricted to the members of the WG ?
- How do we share test case results/data - is Zenodo or another archive server suitable?

#### 8.5.2 Effective interaction with the Tools, Verification and Validation WG, and with the wider academic research community

The NEMO Eddy Closure Working Group can both benefit from and contribute to developments from the grassroots level (e.g. students and independent researchers outside the NEMO Consortium), right up to those of international climate centres (e.g. GFDL/MOM6) and large, funded consortia (e.g. Eddy Energy Climate Process Team in the United States and the Energy Transfers in the Atmosphere and Ocean (TRR181) group in Germany). Within the NEMO Consortium, there are several Working Groups which are closely related, including the Machine Learning WG and the Tools, Verification and Validation WG. In this chapter, the Eddy Closure WG has already identified two major outcomes to benefit the community: a) to provide a suite of test cases; and b) to provide a set of eddy-related metrics to evaluate eddy parameterisations. It is intended that a) and b) will be used together to form a Parameterisation Evaluation Protocol (PEP).

Following feedback from the Scientific Advisory Committee, a key aspect to address is the ease and speed at which external investigators might be able to adopt a codebase (for a test configuration, such as eORCA025, eORCA1, NeverWorld2), modify them, assess their performance and, if skillful, work with a NEMO development "buddy" to finalise improvements to the NEMO trunk. The typical timescales to assess performance skill could range from 6 months to 3 years, for example.

As has been advocated in several chapters, the Eddy Closure WG will also utilise the new Gitlab site for sharing (both internally and externally) code, documentation and other information, including test configurations (and ideally an auto-setup script for each configuration, to allow easy installation and running on a range of architectures), relevant literature references and summary results for comparison. The target timescale for spinning this up on Gitlab is within 2 years.

Also, within 2 years, we will draw upon our WG members who are also involved in the major international eddy closure academic research consortia in the US (Bachman, Jansen) and Germany (Juricke, Eden, Bagaeva) to discuss with both groups the following issues: a) can we agree a set of test cases that would benefit our shared development interests, to enable intercomparison? b) can we agree on a basic set of eddy-related diagnostics and metrics, with the aim of writing a collaborative, peer-reviewed paper? - this would hopefully allow better intercomparison and community involvement (and was encouraged by the Scientific Advisory Committee).

For both the test cases and the eddy-related metrics, successful interaction with the Tools, Verification and Validation WG is important, as is ongoing interaction whenever any promising eddy closure generates new code changes which need to be assessed before including in the NEMO Trunk. There is a question about how/whether the tools developed by this WG would include diagnosis of eddy-related metrics, e.g. perhaps for Drake Passage volume transport or regional KE or a measure of jet sharpness. Draft diagnostic code could be prepared by the Eddy Closure WG first. Other routine verification and validation such as restartability, reproducibility, etc. will also be relevant for the development process of new eddy closures. There are two members (Coward and Le Sommer) in both the Eddy Closure WG and the Tools, Verification and Validation WG. Within the first 2 years, we will agree upon the overlap of any shared goals and specific tasks and responsibilities, as well as refining a process for users to assess eddy-related metrics (to be documented on the Github site) and, if relevant, any aspects of the V&V related to code adoption into the NEMO Trunk.

A remaining challenge is how to foster and improve the involvement of students and independent researchers, who are often the ones who develop the next generation of eddy closures. One challenge here is likely to be how quickly and easily a model codebase can be understood and set up, complete with a relevant test case. Another challenge is the general awareness of the existence of the NEMO framework and example applications (e.g. compare with <http://mitgcm.org>). We will liaise with the NEMO Community Engagement WG and follow its initial recommendations. Some of these (Gitlab documentation, easily setup test configurations, testing skill) are already planned, but we will also contribute to the Discourse group discussions to support users and will publish diagnosed test case results of eddy-related metrics, as well as code configurations (with a DOI). In addition, we will continue to work with the Community Engagement WG over the next 5 years to adapt to any changing drivers, develop tutorials, etc. Some demonstrations at major conferences might also be helpful to attract interested researchers. In particular, to increase accessibility it is important to show that NEMO and the chosen test cases may be easily adopted as a versatile platform from which to easily explore dynamics and try out new eddy closures, rather than solely showcasing the capability for a niche application, e.g. global 1/36 deg. even if that would be impressive. A further important issue is the need for "pull-through" from eddy closures proposed in the literature through all the stages to reach the NEMO trunk, if skillful. It is perceived that students or other researchers who publish viable closures may find the implementation and testing more time-consuming or less interesting and perhaps do not see it as their job (or don't have funding) to see the concept through to wider application. More direct interaction (communication/partnership) with the developers of eddy closures, from quite early stages, is therefore important too, so that the NEMO Consortium can encourage and assist, either directly or by helping to gain access to external funding.

### 8.5.3 Effective interaction with the Machine Learning WG

Both NEMO WGs share several members in common (le Sommer, Perezhugin, Shao, Storto, Wilson), so there will be ongoing collaboration. Three main topics of relevance to machine learning and mesoscale closure approaches in NEMO have been identified and should be further explored:

- Topic 1: machine learning-based (tuning of) mesoscale eddy closures. Machine learning approaches could be used to tune more automatically existing parameterisations given a more explicit cost function. Also, as suggested by Zanna and Bolton GRL 2020, ML could serve to discover novel formulations for mesoscale closures.
- Topic 2: deep emulation of ocean mesoscale eddies in NEMO. Deep emulation of geophysical models has recently been proposed (e.g. Doury et al CD 2022). By drastically reducing computational costs, they allow for a wider exploration of parameter ranges and uncertainty. Such an approach could address specific modules of ocean models, such as eddy closures.
- Topic 3: sampling of uncertainty with design of ensembles / stochastic approaches to mesoscale dynamics using ML. Mesoscale dynamics are stochastic in essence, and as such they generate an internal self-emerging oceanic variability. Stochastic approaches to mesoscale eddy closures have already been proposed (e.g. Zanna et al OM 2017). ML could assist the formulation of such approaches.

In terms of organisation we formulate a few proposals to ensure that 1. knowledge is shared between both activities in NEMO; 2. interactions are fostered between the two working groups; and 3. some planning emerges around ML-based mesoscale eddy closures within the next few years.

Namely:

- Make sure that there is some overlap in the membership of the two groups
- Organise regular meetups between the chairs of each working-group for catching up on the ongoing activities in each group
- Organise joint scientific meetings on ml-based subgrid closures, what about a joint seminar series ?
- Make sure that the sandboxing mechanism is adapted to share ml-based closures too.
- Identify a specific ongoing project at the interface between the two groups (eg : a on-going PhD or postdoc project)

## 8.6 Recommending parameterisations for NEMO

Lastly, this working group will recommend existing parameterisations and identify promising research directions that may inform high-level design decisions within NEMO. Due to the limitations available by the System Team and to maintain a reasonable size codebase, this WG takes into account the following considerations when recommending to the NEMO Developers Committee:

- Is the barrier to implementation sufficiently low to allow for testing and incorporation within NEMO within the timeframe of the strategic plan?
- What are the fundamental ‘needs’ vs. ‘wants’ for the immediate applications being pursued within the Consortium?
- Can the skill of the parameterisation be evaluated against the scientific objectives?
- With the push from the Kernel Working Group to implement new vertical coordinates, the implementation of parameterisations will need to be evaluated for appropriateness in generalised vertical coordinates

### 8.6.1 Initial recommendations for eddy closures to develop further for NEMO

The first two schemes have received significant support from the System Team already. Their inclusion in this document is to further recommend their inclusion into NEMO. Each of the schemes below has demonstrated improvements in large-scale climate metrics in preliminary applications with NEMO.

- Prognostic eddy energy-based scheme to generate a modified GM eddy transfer coefficient, following the GEOMETRIC framework (Marshall et al., 2012; Mak et al., 2017, 2018).
- Biharmonic operator for Gent-McWilliams (Gent and McWilliams, 1990; Roberts and Marshall, 1998) to allow explicit, harmonic GM eddy stirring to be reduced in magnitude.
- Energy backscatter scheme, most likely either one based on GM+E scheme of Bachman, 2019, which optimally combines with the QG Leith scale-dependent viscosity (Bachman, 2019; Pearson et al., 2017; Fox-Kemper et al, 2011); or a scheme based on Jansen et al., 2015a,b - see also Juricke, 2020; Perezhugin, 2020.

- QG Leith scale-dependent viscosity (Leith, 1996; also see above) - to enable reduction in the typically excessive gridscale viscous damping of KE.

The interest in scale-aware eddy parameterisations arise from the increasing push for climate models to increase spatial resolution (e.g. from 1 to ¼ degree, Hewitt et al., 2020, 2022), and a desire to reduce the number of parameterisation frameworks and/or tuning parameters in eddy permitting models, which span over a range of horizontal resolutions for a given nominal grid size (e.g. from 5 to 25km in the ¼ degree global configuration).

For backscatter approaches, it is proposed that the focus be on the GM+E approach of Bachman (2019): from a theoretical point of view, GM+E follows the GEOMETRIC formulation (Marshall et al., 2012, Mak et al, 2022), and since the GM-version of GEOMETRIC is already being considered in NEMO, there is consistency of parameterisation via a single theoretical framework; GM+E utilises a Leith-type eddy viscosity (e.g. Fox-Kemper et al, 2011), which is being recommended here as a high priority implementation; GM+E has been tested in idealised models (MITgcm) as well as in realistic models MOM6 (NCAR version, communication from Scott Bachman). Thus, the concrete actions recommended are:

- implement a Leith-type eddy viscosity (can leverage the NEMO Smagorinsky implementation and be made scale-aware),
- implement GM+E with some interfacing with the GM version of GEOMETRIC (the latter available as an option in the NEMO LDF modules)
- provide test cases highlighting the effect on simulations utilising a combination of the parameterisation and/or parameters

Suggested actions that would be desirable are:

- to have as few system parameters relating to the parameterisations as feasible, and ideally favour parameters that are physically informed and/or can be constrained
- document within the code, as part of the manual, or otherwise, some suggested choice of (and known sensitivities to) system parameters relating to the parameterisations
- consolidating the existing GM-based and isoneutral diffusion schemes via one consistent framework, if possible and/or beneficial (cf. Smith & Marshall, 2009)
- investigate whether scale-aware approaches of adiabatic diffusive closures are in fact possible and/or beneficial (e.g. activation approaches from Hallberg, 2013; field splitting-type approaches suggested by Greatbatch et al., 2004)

## 8.7 References

Bachman, S. D., Fox-Kemper, B., and Pearson, B., (2017), A scale-aware subgrid model for quasi-geostrophic turbulence, *J. Geophys. Res. Oceans*, 122, 1529– 1554.

Couespel, D., M. Lévy, and L. Bopp, (2021), Oceanic primary production decline halved in eddy-resolving simulations of global warming. *Biogeosciences*, **18**, 4321–4349, <https://doi.org/10.5194/bg-18-4321-2021>.

Doury, A., *et al.* (2022), Regional climate model emulator based on deep learning: concept and first evaluation of a novel hybrid downscaling approach. *Clim Dyn.* doi:[10.1007/s00382-022-06343-9](https://doi.org/10.1007/s00382-022-06343-9).

- Ferrari, R., *et al.* (2010). A boundary-value problem for the parameterized mesoscale eddy transport. *Ocean Modelling* **32**, 143–156 (2010).
- Fox-Kemper, B. *et al.* (2011), Parameterization of mixed layer eddies. III: Implementation and impact in global ocean climate simulations. *Ocean Modelling* **39**, 61–78.
- Gent, P. R., and J. C. McWilliams (1990), Isopycnal mixing in ocean circulation models. *J. Phys. Oceanogr.*, **20**, 150–155.
- Greatbatch, R. J. & Lamb, K. G. On Parameterizing Vertical Mixing of Momentum in Non-eddy Resolving Ocean Models. *Journal of Physical Oceanography* **20**, 1634–1637 (1990).
- Greatbatch, R. J. *et al.* The semi-prognostic method. *Cont. Shelf Res.* **24**, 2149–2165 (2004).
- Groeskamp, S., *et al.* (2020). Full-Depth Global Estimates of Ocean Mesoscale Eddy Mixing From Observations and Theory. *Geophysical Research Letters* **47**, e2020GL089425.
- Hallberg, R., (2013), Using a resolution function to regulate parameterizations of oceanic mesoscale eddy effects. *Ocean Model.* **72**, 92–103.
- Hewitt, H. T. *et al.* (2020), Resolving and Parameterising the Ocean Mesoscale in Earth System Models. *Curr Clim Change Rep* **6**, 137–152.
- Hewitt, H., Fox-Kemper, B., Pearson, B., Roberts, M. & Klocke, D. The small scales of the ocean may hold the key to surprises. *Nat. Clim. Chang.* **12**, 496–499 (2022).
- Jansen, M. F., Adcroft, A. J., Hallberg, R. & Held, I. M., (2015a), Parameterization of eddy fluxes based on a mesoscale energy budget. *Ocean Model.* **92**, 28–41.
- Jansen, M. F., Held, I. M., Adcroft, A. & Hallberg, R., (2015b), Energy budget-based backscatter in an eddy permitting primitive equation model. *Ocean Model.* **94**, 15–26.
- Juricke, S. *et al.* A Kinematic Kinetic Energy Backscatter Parametrization, (2020), From Implementation to Global Ocean Simulations. *J. Adv. Model. Earth Syst.* **12**, e2020MS002175.
- Klöwer, M., Jansen, M. F., Claus, M., Greatbatch, R. J. & Thomsen, S., (2018), Energy budget-based backscatter in a shallow water model of a double gyre basin. *Ocean Model.* **132**, 1–11.
- Leith, C. E., (1996), Stochastic models of chaotic systems, *Physica D.*, **98**, 481-491.
- Lévy, M. *et al.*, (2010), Modifications of gyre circulation by sub-mesoscale physics. *Ocean Model.*, **34**, 1–15, <https://doi.org/10.1016/j.ocemod.2010.04.001>.
- Mak, J., Marshall, D. P., Maddison, J. R. & Bachman, S. D., (2017), Emergent eddy saturation from an energy constrained eddy parameterisation. *Ocean Model* **112**, 125–138.
- Mak, J., *et al.* (2018) Implementation of a Geometrically Informed and Energetically Constrained Mesoscale Eddy Parameterization in an Ocean Circulation Model. *J. Phys. Oceanogr.* **48**, 2363–2382.
- Mak, J., *et al.* (2022), Acute Sensitivity of Global Ocean Circulation and Heat Content to Eddy Energy Dissipation Timescale. *Geophysical Research Letters* **49**, e2021GL097259.



Marshall, D. P., Maddison, J. R. & Berloff, P. S. (2012). A Framework for Parameterizing Eddy Potential Vorticity Fluxes. *J Phys Ocean.* 42, 539–557.

Marques, G. et al., (2022), Neverworld2: an idealized model hierarchy to investigate ocean mesoscale eddies across resolutions. *Earth and Space Science Open Archive*.  
<https://doi.org/10.1002/essoar.10511043.1>

Pearson, B. et al. (2017), Evaluation of scale-aware subgrid mesoscale eddy models in a global eddy-rich model. *Ocean Modelling* **115**, 42–58.

Perezhogin, P. A., (2020), Testing of kinetic energy backscatter parameterizations in the NEMO ocean model. *Russian Journal of Numerical Analysis and Mathematical Modelling*, **35**, 69–82, <https://doi.org/10.1515/rnam-2020-0006>.

Roberts, M. J., and Marshall, D. P., (1998), Do we require adiabatic dissipation schemes in eddy-resolving ocean models? *J. Phys. Oceanogr.*, 28, 2050-2063.

Smith, K. S. & Marshall, J. Evidence for Enhanced Eddy Mixing at Mid Depth in the Southern Ocean. *J. Phys. Oceanogr.* **39**, 50–69 (2009).

Zanna, L. & Bolton, T. (2020). Data-Driven Equation Discovery of Ocean Mesoscale Closures. *Geophysical Research Letters* **47**, e2020GL088376.

Zanna, L., et al. (2017), Scale-aware deterministic and stochastic parametrizations of eddy-mean flow interaction. *Ocean Model.* **111**, 66–80 (2017).

## 9 Surface fluxes and vertical mixing

WGL: Guillaume Samson

### 9.1 Executive Summary

A large amount of work has been accomplished during the last NDS period to upgrade AWSI & ocean mixing related parts of Nemo code, and to provide state-of-the-art parameterisations and capacities. This is especially true concerning surface turbulent fluxes and wave-ocean interactions, which have been largely rewritten and updated to introduce new physical possibilities. Mesoscale air-sea interactions effects on the ocean have been also introduced during the last NDS through the development of a new parametrisation and of an atmospheric boundary layer model. Consequently, the developments proposed here for the next NDS period are less ambitious, and can be seen as a consolidation and a completion of what has been accomplished during the previous period. In particular, the Osmosis vertical mixing scheme and the atmospheric boundary layer model developments will be completed during this new NDS period.

### 9.2 Context

A specific chapter concerning air-wave-sea interactions (AWSI) and ocean mixing parameterisations has been introduced for the first time in the previous NDS document for the 2018-2022 period. A working group focused on wave-ocean coupling led by Emanuela Clementi (INGV) has also been active over the 2013-2017 period. It allowed to organize, gather and coordinate several independent waves-related developments and resulted in the proposed waves-related actions for the last NDS. But it has not been active since yet. The initiation of the new NDS has been the opportunity to setup a new WG focused on this specific chapter to review the progress made and the related plans for next period. It is important to note that this WG is fairly new and met only once. Consequently, it has not achieved yet the maturity to produce a well-structured chapter, compared to other long-standing NEMO-related topics. But despite its relative youth, a significant amount of developments have been done during the previous NDS period; their status and continuation are summarized hereafter.

### 9.3 Priorities and timescales

The main priority of this WG is to ensure that NEMO proposes a state-of-the-art representation of AWSI and ocean mixing processes for oceanic applications ranging from kilometeric to climatic scale. Practically, it was organised in the last NDS into 4 main objectives:

- more accurate estimation of turbulent air-sea fluxes through the use of more sophisticated bulk algorithms/parameterizations to compute exchange coefficients in use in bulk formulae (i.e. consideration of wave information, skin temperature, salinity, etc.)
- include representation of relevant wave-current interaction processes
- more realistic representation of the coupling between the atmospheric boundary layer (ABL) and the ocean surface boundary layer (OSBL); e.g. the action of the ABL on the OSBL, and the response of the ABL to the evolving prognostic ocean surface properties, must be better accounted for.
- better representation of subgrid-scale vertical physics through improved turbulence closures in the OSBL and below.

It is difficult to give a clear timescale for each development presented hereafter, but they are small-to-medium size and incremental developments, mostly independent from one another, and will not significantly impact other parts of the NEMO code (ie outside the SBC, ZDF and ABL directories). Consequently, it should be feasible to implement them during the next five years. Note also that some developments related to this chapter may not be identified yet because of the recent creation of the associated WG or because they will come from community-driven developments.

The general philosophy is to introduce new physical parametrisations, or to update existing ones, or to add an explicit representation of small-scale processes which have an impact on ocean large-scale evolution; noting that some applications such as coastal modelling require specific processes. The priorities given to these possible evolutions will mostly depend on community advances, user needs and man power available. On a 10-year timescale, we expect that ocean simulations including online-coupling with wave models and with the ABL model (when full atmospheric simulations are not needed) or with wave and atmosphere models will become very common. Consequently, we must prepare NEMO to totally support these kind of configurations, including all the related necessary processes during the two next NDS periods.

#### 9.4 People involved by topic (alphabetical order)

- Turbulent air-sea/ice fluxes : Clement Rousset, Guillaume Samson, Laurent Brodeau
- Waves-ocean interactions : Aimie Moulin, Emanuela Clementi
- Mesoscale ocean-atmosphere coupling : Guillaume Samson, Sébastien Masson
- Vertical mixing: Alan Grant, George Nurser, Casimir de Lavergne

#### 9.5 Turbulent air-sea/ice fluxes

Concerning turbulent air-sea fluxes (TASF), only 3 outdated bulk formulae (CORE from Large & Yeager 2004; CLIO from Goosse et al. 1999 and MFS from Castellari et al. 1998) were available in NEMO 3.6. A major rewriting has been undertaken by Laurent Brodeau to introduce the “AeroBulk” library (<https://brodeau.github.io/aerobulk/>) in NEMO (Brodeau et al. 2017). Within this framework, it is now possible to estimate the TASF in an efficient and unified way by sharing common thermodynamical functions between the bulk parametrisations. It is also easier to introduce new bulk algorithms following a standardized implementation.

Four state-of-the-art (and CORE for retrocompatibility) algorithms are currently available: COARE v3.0 & v3.6 (Fairall et al. 2003, 2018), ECMWF (IFS cy40) and ANDREAS (Andreas et al. 2015). COARE and ECMWF schemes also provide a cool-skin/warm-layers parametrisation to accurately capture the SST diurnal cycle without requiring a high vertical resolution to represent it. This set of bulk algorithms and parametrisations provides a nice framework to explore the sensitivity of oceanic processes to the uncertainties associated with TASF. In this context, the MFS bulk scheme will be reintroduced in the Aerobulk framework to propose an additional choice. Gryanik et al. (2021) also recently proposed non-iterative formulations of momentum and heat transfer coefficients for stable conditions. This approach (which can be seen as an equivalent of the GLS framework for bulk schemes) will be explored and evaluated during the next NDS period. Finally, in order to improve the Charnock parameter estimation which underlies the computation of transfer coefficients, a new formulation including a wave-age dependency proposed by Sauvage et al. 2020 will be introduced in the Aerobulk library.

Another major caveat of the current generation of bulk algorithms is the lack of representation of subgrid variability in TASF. It has also been shown that introducing effects of unresolved processes on TASF, such as convective wind gustiness or using stochastic approaches, improves their estimation and can have a significant impact on the ocean variability (Williams 2012, Berner et al. 2017, Blein et al. 2022). Consequently, even if no related development within NEMO has been identified for now, a careful review of the existing literature will be undertaken to follow this subject and identify which new sub-grid-scale process parameterizations are relevant for TASF and should be included in the existing bulk algorithms. New promising community-driven developments related to this topic will also be encouraged to integrate NEMO in a sustainable way.

Concerning turbulent air-ice fluxes, it was only possible to use a unique and constant transfer coefficient in NEMO 3.6. During the last NDS, 3 sea-ice dependent bulk parametrisations have been introduced following Andreas 2005, Lupkes 2012 and Lupkes & Gryanik 2015. These schemes allow to have distinct transfer coefficients for heat and momentum, but also to represent some subgrid effects on TASF such as melt ponds and ice floes. They also propose sea-ice specific stability functions well adapted to stable environments. No new air-ice TASF related development has been identified yet for this NDS period. However, in link with the SI WG, a bulk mixed layer scheme equivalent to the one included in CICE will be developed in order to represent sea-ice – ocean interactions without using the full ocean model. This will greatly facilitate and speedup sea-ice related developments.

Several 1D test cases have also been implemented during the last NDS to easily evaluate TASF and their effect on the water column (or in the presence of sea-ice) using atmospheric and oceanic observations or reanalysis. Three configurations are currently available: Papa station located in the North-East Pacific, Lion buoy in the Mediterranean Sea and a virtual station including sea-ice located north of Greenland. These 3 test cases have proven very useful to compare and to validate new developments related to TASF and vertical physics. New 1D test cases will be implemented during the next NDS period to easily evaluate NEMO vertical physics in various oceanic environment, such as the Tropics (Voltaire et al. 2022), and the OSMOSIS observational site located in the North Atlantic. As discussed during the SAC evaluation meeting, the 1D test case from Li et al. 2019 will also be implemented in NEMO to get a first overview of NEMO mixing schemes behaviours and sensitivities.

## 9.6 Wave-ocean interactions

A first version of the different processes and of the coupling interface with wave models has been implemented in NEMO 3.6. In particular, it was possible to use the neutral drag coefficient from the wave model to compute momentum transfer to the ocean and the surface boundary condition has been adapted to waves effects. Wave-enhanced vertical turbulence has been added following Qiao (2010) and the vertical profile of the Stokes drift has been introduced following Breivik et al. 2014. The Stokes-Coriolis force has been added to the momentum equation and the Stokes drift advection in the tracer equation. During the last NDS period, these initial developments have been improved and completed in order to represent wave-ocean interaction in a more coherent way. In particular, a new generic coupling interface between NEMO and wave models has been implemented, which allows to exchange more fields between the models and to easily add new ones if needed. Wave-induced terms in NEMO equations have also been updated and completed to better represent Stokes-Coriolis, vortex force and wave-induced pressure contributions in the momentum equation following Bennis et al. 2011, as well as a new momentum boundary condition. Only the “vector invariant” form of these terms has been implemented for now. A new Stokes drift velocity profile based on a finite volume approach has also been introduced to avoid limitations of the finite difference approach of Breivik et

al. 2014, 2016. Finally, different wave effects on the TKE vertical mixing scheme have been added (shear production term, wave-breaking surface energy injection, length scales dependency to wave surface roughness, new Langmuir turbulence parameterization), giving a physical basis to avoid using the adhoc “ $\epsilon$ ” mixing enhancement. All these developments presented in Couvelard et al. 2020 led to significant improvements in terms of MLD and SST.

During the next NDS period, a development effort will be made toward coastal applications and shallow-water environments. In these conditions, new wave-related developments are needed to correctly take into account interactions with bathymetry. In particular, a new wave-induced bottom pressure term will be introduced and the Stokes drift velocity profile will be adapted to correctly deal with bottom friction. However, we keep in mind that NEMO targets coastal kilometric applications, so smaller-scale waves-related effects will not be considered. The actual wave-effects implementation in NEMO equations will also be adapted to the “flux form” formalism, which will provide more flexibility to use it with other available dynamical schemes. GLS vertical mixing scheme will also be modified to take into account wave-induced mixing following what has been done with the TKE scheme in Couvelard et al. 2020. Approaches proposed by Kantha & Clayson 2004 or Harcourt 2013 will be tested. Note that waves-sea ice interactions are out of the scope of this chapter and will be treated in the sea-ice chapter. Finally, the coupling interface between NEMO and OASIS has continuously increased as new NEMO versions come out, with more and more coupled models supported, and more granularity in terms of variables which can be exchanged. Hence, its readability and maintenance have become difficult. This is why a cleaning and splitting of the coupling interface is considered during the next NDS period. It would facilitate future developments specific to coupled models.

## 9.7 Mesoscale feedbacks between the OSBL and the ABL

With the increase in resolution of NEMO applications, it becomes more and more important to consider air-sea interactions at play at mesoscale, for several reasons. An important one is that they constitute a significant and physical source of kinetic energy damping in the ocean. Until recently (NEMO 4.0), it was possible to consider the surface current speed in the wind stress computation only by coupling an atmospheric model to NEMO (Jullien et al. 2020), or by tuning the “ $rn\_vfac$ ” parameter. However, both approaches have significant shortcomings. This is why two alternative choices have been developed during the last NDS period.

The first one is a parameterization of the surface current effect on the wind stress following Renault et al. 2020. This scheme uses a linear relationship between the wind forcing and the surface stress from an observation-based statistical regression to compute a correction coefficient that mimics the dynamical coupling. This scheme uses standard 10m-wind forcing to be activated. The other one is a simplified atmospheric boundary layer model (ABL1d) following Lemarié et al. 2020. ABL1d is a single-column model with a vertical discretization of the lower atmosphere. It allows an explicit representation of both dynamical and thermodynamical coupling between OSBL and ABL by solving the air temperature, humidity and wind evolutions through vertical mixing and sea-surface boundary conditions. Contrary to the Renault et al. parameterization, ABL1d needs additional external fields (3D atmospheric forcings) in order to relax the model toward the atmospheric forcing.

These two new options can successfully and efficiently replace a classic atmospheric model or the  $rn\_vfac$  parameter to represent mesoscale air-sea interactions and their effect on the ocean. However, some important features are still missing in the ABL model to fully represent these coupled processes. Consequently, the ABL model development will continue during the next NDS period with a focus on the last missing important processes. First, the air pressure adjustment process which can

be locally important will be added in addition to vertical mixing processes. Then, horizontal advection of momentum and tracers will be implemented to decrease the model relaxation toward the atmospheric forcing and to improve the model solution. Dealing with coastal lateral boundary conditions will require special care and could be handled using the BDY routines already available in NEMO in the longer term.

Finally, the last proposed development is related to the ABL model, but also more generally to the way atmospheric forcings are read by NEMO and to the HPC WG. It is directly related to new atmospheric forcing datasets such as ERA5 and JRA55-DO, which have strongly increased their spatial and temporal resolutions, as well as their data volume. This is even more the case with the ABL model which needs 3D atmospheric forcings. For now, these data are read sequentially by NEMO, which can slowdown the code up to 50%. We propose here to use XIOS to read these atmospheric forcings in an asynchronous way, and hence to speed up substantially the NEMO code.

## 9.8 Parametrisation of mixing in the OSBL

The NEMO code already offers a large panel of vertical mixing parametrisations and options, including historical and state-of-the art schemes: constant values, Richardson number, TKE, GLS (including 4 different turbulent models) and OSMOSIS (OSM). Waves and sea-ice effect on turbulence can also be taken into account for these 3 last schemes. Convection processes can also be parametrised using a non-penetrative convective adjustment, enhanced vertical diffusion or a new eddy-diffusivity mass-flux (EDMF) scheme from Giordani et al. 2021. Consequently, the available choice for users is already very rich, and due to the plurality of NEMO applications, it is quite difficult to remove some of these parametrisations or options from the code. Even simple mixing schemes such as constant diffusion and Richardson number are still used by some institutions and can be useful for debugging purposes or idealized test case validation. Consequently, we do not intend to suppress or simplify the vertical mixing routines in the near future.

Concerning interfacing external community codes such as GOTM or CV-Mix with NEMO such as in Li et al. 2021, we agree that it would greatly simplify NEMO code and that efforts could be put and shared on the libraries itself. Thanks to SAC comments, it has been decided that an interface between NEMO and CVMix will be developed in order to offer an alternative to current mixing schemes, but more importantly to open NEMO code to new communities and to facilitate intercomparison exercises.

Concerning the next NDS period, only one new parametrisation to represent mixing induced by near-inertial waves will be developed. The other developments will be mostly dedicated to improving or completing the existing ones. OSM scheme development will continue (including optimizations and code cleaning), Langmuir-circulation induced mixing will be improved in TKE scheme and EDMF scheme will be generalised to momentum and passive tracers and compared to LES simulations to improve its calibration. This LES-comparison methodology could be applied to other vertical mixing schemes in order to verify and compare their accuracy in the longer term. Finally, the various 1D test cases described in the TASF section will also be very useful to validate and compare the existing and new vertical mixing related developments.

## 9.9 Parametrisation of mixing in the interior

Vertical mixing in the ocean interior is largely fuelled by breaking internal tides. NEMO has long employed a background diffusivity, together with bottom-intensified tidal mixing following Simmons

et al. (2004) and elevated mixing in the Indonesian seas following Koch-Larrouy et al. (2007). This approach does not explicitly account for internal tides dissipating far from generation sites, nor does it ensure energetic consistency. The scheme of de Lavergne et al. (2020) remedies both caveats and has been recently incorporated in the trunk. Only marginal improvements to this scheme are thus envisioned for the forthcoming NDS period.

Wind-generated inertial oscillations are also a significant source of mixing in the ocean interior. There are two distinct effects of these oscillations: (1) they create shear at the base of the mixed layer, thus causing episodic shear-driven mixing that deepens the mixed layer; (2) they cause vertical oscillations of the mixed-layer base that radiate near-inertial internal waves in the interior. Effect (1) is currently parameterised with the ad-hoc “ $\epsilon$ ” scheme of NEMO; it can have a substantial impact in coupled climate simulations (Jochum et al. 2013). Effect (2) is currently not accounted for. Alford (2020a,b) recently mapped energy sources (1) and (2) using observations and theory. These maps may provide guidance for the design of a new energy-constrained parameterisation of both effects. A robust and versatile parameterisation would require the energy sources to be dependent on the (resolved) wind forcing. The possibility of parameterising energy sources (1) and (2) as a function of 10m winds and simulated surface currents, building on Jochum et al. (2013) and Alford (2020a,b), will be explored during this NDS period.

Double-diffusive processes are an additional minor source of mixing in the interior. Their impact is currently parameterised by elevating the vertical diffusivity when the simulated stratification is thought to be favourable to salt fingering or diffusive convection (Merryfield et al. 1999). This parameterisation has a modest impact on circulation but can have large regional impacts on passive tracers and biogeochemistry. We plan to review the literature to identify more recent and better constrained parameterisations that may deserve replacing the scheme of Merryfield et al. (1999).

## 9.10 References

- Large, W. G., and S. Yeager, 2004: Diurnal to decadal global forcing for ocean and sea-ice models: the data sets and flux climatologies. NCAR Technical Note, NCAR/TN-460+STR, CGD Division of the National Center for Atmospheric Research
- Goosse, H., E. Deleersnijder, T. Fichefet, and M. England, 1999: Sensitivity of a global coupled ocean-sea ice model to the parameterization of vertical mixing. *J. Geophys. Res.*, 104, 13,681-13,695
- Castellari, S., N. Pinardi, and K. Leaman, 1998: A model study of air-sea interactions in the mediterranean sea. *J. Mar. Sys.*, 18, 89-114
- Brodeau, L., B. Barnier, S. Gulev, and C. Woods, 2016: Climatologically significant effects of some approximations in the bulk parameterizations of turbulent air-sea fluxes. *J. Phys. Oceanogr.*, 47 (1), 5–28, 10.1175/JPO-D-16-0169.1
- Fairall, C. W., E. F. Bradley, J. S. Godfrey, G. A. Wick, J. B. Edson, and G. S. Young, 1996: Cool-skin and warm-layer effects on sea surface temperature. *Journal of Geo-physical Research: Oceans*, 101 (C1), 1295–1308
- Fairall, C. W., E. F. Bradley, J. E. Hare, A. A. Grachev, and J. B. Edson, 2003: Bulk parameterization of air-sea fluxes: Updates and verification for the coare algorithm. *Journal of Climate*, 16 (4), 571–591
- Gryanik, V. M., Lüpkes, C., Sidorenko, D., & Grachev, A. (2021). A universal approach for the non-iterative parametrization of near-surface turbulent fluxes in climate and weather prediction models. *Journal of Advances in Modeling Earth Systems*, 13
- Williams, P. D. (2012), Climatic impacts of stochastic fluctuations in air–sea fluxes, *Geophys. Res. Lett.*, 39, L10705, doi:10.1029/2012GL051813
- Berner, J., Achatz, U., Batte, L., Bengtsson, L., De La Camara, A., Christensen, H. M., ... & Yano, J. I. (2017). Stochastic parameterization: Toward a new view of weather and climate models. *Bulletin of the American Meteorological Society*, 98(3), 565-588
- Blein, S., Roehrig, R., Voltaire, A., & Faure, G. (2020). Meso-scale contribution to air-sea turbulent fluxes at GCM scale. *Quarterly Journal of the Royal Meteorological Society*, 146(730), 2466–2495.

- Lüpkes, C. and V. M. Gryanik, 2015: A stability-dependent parametrization of transfer coefficients for momentum and heat over polar sea ice to be used in climate models. *Journal of Geophysical Research*, 120 (2), 552–581
- Lüpkes, C., V. M. Gryanik, J. Hartmann, and E. L. Andreas, 2012: A parametrization, based on sea ice morphology, of the neutral atmospheric drag coefficients for weather prediction and climate models. *Journal of Geophysical Research Atmospheres*, 117 (13), 1–18
- Voldoire, A, Roehrig, R & Giordani, H & Waldman, R & Bouin, Marie-Nöelle. (2022). Assessment of the sea surface temperature diurnal cycle in CNRM-CM6-1 based on its 1D coupled configuration. *GMD*, 15. 3347-3370. 10.5194/gmd-15-3347-2022
- Qiao, F., Y. Yuan, T. Ezer, C. Xia, Y. Yang, X. Lu, and Z. Song, 2010: A three-dimensional surface wave–ocean circulation coupled model and its initial testing. *Ocean Dynamics*, 60 (5), 1339–1335
- Breivik, Ø., J.-R. Bidlot, and P. A. E. M. Janssen, 2016: A Stokes drift approximation based on the Phillips spectrum. *Ocean Modelling*, 100, 49–56
- Breivik, Ø., P. A. E. M. Janssen, and J.-R. Bidlot, 2014: Approximate Stokes drift profiles in deep water. *Journal of Physical Oceanography*, 44 (9), 2433–2445
- Couvelard, X., Lemarié, F., Samson, G., Redelsperger, J.-L., Arduin, F., Benshila, R., and Madec, G.: Development of a two-way-coupled ocean–wave model: assessment on a global NEMO(v3.6)–WW3(v6.02) coupled configuration, *GMD*, 13, 3067–3090, 2020
- Kantha, L. H. and C. A. Clayson, 1994: An improved mixed layer model for geophysical applications. *Journal of Geophysical Research*, 99 (C12), 25 235–25 266
- Harcourt, R. R. (2013). A second-moment closure model of Langmuir turbulence. *Journal of Physical Oceanography*, 43(4), 673-697
- Jullien, S., Masson, S., Oerder, V., Samson, G., Colas, F., & Renault, L. (2020). Impact of ocean–atmosphere current feedback on ocean mesoscale activity: Regional variations and sensitivity to model resolution. *Journal of Climate*, 33(7), 2585-2602
- Renault, L., Masson, S., Arsouze, T., Madec, G., & McWilliams, J. C. (2020). Recipes for how to force oceanic model dynamics. *Journal of Advances in Modeling Earth Systems*, 12(2), e2019MS001715
- Lemarié, F., Samson, G., Redelsperger, J. L., Giordani, H., Brivoal, T., & Madec, G. (2021). A simplified atmospheric boundary layer model for an improved representation of air–sea interactions in eddying oceanic models. *GMD*, 14(1), 543-572
- Giordani, H., Bourdallé-Badie, R., & Madec, G. (2020). An Eddy-Diffusivity Mass-Flux Parameterization for Modeling Oceanic Convection. *Journal of Advances in Modeling Earth Systems*, 12(9), e2020MS002078
- Simmons, H.L., Jayne, S.R., St Laurent, L.C., Weaver, A.J. (2004). Tidally driven mixing in a numerical model of the ocean general circulation. *Ocean Modell.* 6, 245-263.
- Koch-Larrouy, A., Madec, G., Bouruet-Aubertot, P., Gerkema, T., Bessières, L., Molcard, R. (2007). On the transformation of Pacific Water into Indonesian Throughflow Water by internal tidal mixing. *Geophys. Res. Lett.* 34, L04604.
- de Lavergne, C., Vic, C., Madec, G., Roquet, F., Waterhouse, A.F., Whalen, C.B., Cuypers, Y., Bouruet-Aubertot, P., Ferron, B., Hibiya, T. (2020). A parameterization of local and remote tidal mixing. *Journal of Advances in Modeling Earth Systems*, 12(9), e2020MS002065.
- Jochum, M., Briegleb, B.P., Danabasoglu, G., Large, W.G., Norton, N.J., Jayne, S.R., Alford, M.H., Bryan, F.O. (2013). The impact of near-inertial waves on climate. *Journal of Climate*, 26, 2833-2844.
- Alford, M.H. (2020a). Revisiting Near-Inertial Wind Work: Slab Models, Relative Stress, and Mixed Layer Deepening. *Journal of Physical Oceanography*, 50, 3141-3156.
- Alford, M.H. (2020b). Global Calculations of Local and Remote Near-Inertial-Wave Dissipation. *Journal of Physical Oceanography*, 50, 3157-3164.
- Merryfield, W.J., Holloway, G., Gargett, A.E. (1999). A Global Ocean Model with Double-Diffusive Mixing. *Journal of Physical Oceanography*, 29, 1124-1142.



## 10 Sea ice

**Lead authors:** Martin Vancoppenolle<sup>1</sup>; Ed Blockley<sup>2</sup> and Clément Rousset<sup>1</sup>

<sup>1</sup> LOCEAN-CNRS, Paris, France

<sup>2</sup> Met Office, Exeter, UK

**Contributing authors:** Aksenov, Yevgeny (NOC, Southampton, UK); Feltham, Danny (CPOM, Reading, UK); Fichet, Thierry (UCLouvain, Be); Garric, Gilles (MO, Toulouse, Fr); Guemas, Virginie (CNRM, Toulouse, Fr); Holland, Paul (BAS, Cambridge, UK); Iovino, Dorotea (CMCC, Bologna, It); Keeley, Sarah (ECMWF, Reading, UK); Madec, Gurvan (CNRS, Parsi, Fr); Massonnet, François (UCLouvain, Be); Rampal, Pierre (CNRS, Grenoble, Fr); Ridley, Jeff (Met Office, Exeter, UK); Schroeder, David (CPOM, Reading, UK); Tietsche, Steffen (ECMWF, Bonn, Ge).

### Executive summary

SI<sup>3</sup>, the NEMO sea ice component, was put together over the last few years by unifying capabilities from several models previously used in NEMO. As a result, SI<sup>3</sup> is strongly modular and has a wide range of options for the representation of sea ice physics. However, it is not modular enough yet to accommodate all user applications and tackle current research questions. Therefore, a central aspect of the upcoming NEMO sea ice strategy is to further improve modularity. First the coupling to other components of the Earth System (snow, ocean waves, continental ice, biological and chemical tracers) will be enabled or facilitated. Another interface, currently in the core of the sea ice code, will be made explicit — the one between vertical physics (thermodynamics, halo-dynamics, optics) and horizontal drift and deformation processes. Fully splitting the 1D vertical processes from 2D horizontal dynamics will allow testing of the drastically different options under development. As well as the focus on modularity and interfacing of SI<sup>3</sup> there is also work expected on each of the horizontal and vertical sea ice physics components. Another priority area is documentation of the SI<sup>3</sup> model, which is important for building a strong user and developer community. Sea ice work is generally limited by resources, in particular by the lack of trained developers. For successful continuation of sea ice development, this issue should also be investigated.

### 10.1 General context, strategy and recent achievements for sea ice in NEMO

The representation of sea ice matters for many of the scientific and operational NEMO applications. From the signature of the consortium, NEMO assumes sea ice as an integral part of ocean dynamics and therefore needs not only an interface to sea ice but also a full sea ice model component. Until 2017, NEMO used to be interfaced with four different yet similar models (CICE, GELATO, LIM2 and LIM3) and has since transitioned to a unique sea ice model named SI<sup>3</sup> (Sea Ice modelling Integrated Initiative), based on the concept of integrating the best functionalities of its predecessors. SI<sup>3</sup> is part of the NEMO code, is compliant to many of the NEMO functionalities (AGRIF, BDY, ...), and is interfaced to atmospheric and marine biogeochemistry models.

SI<sup>3</sup> — as its predecessors and virtually all models in the community — follows the classical paradigm for sea ice modelling from the AIDJEX consortium (Arctic Ice Dynamics Joint Experiment; Coon et al., 1974). This paradigm has slowly evolved but not drastically changed since its inception, and relies upon:

- Splitting 2D horizontal (ice drift) and 1D vertical (growth and melt) processes;
- A continuum approach for 2D horizontal conservation of momentum; including a non-linear internal interaction term (referred to as *rheology*);
- Several parameterizations for 1D vertical physics (surface energy balance, growth and melt, snow, optics, halo-dynamics, subgrid-scale thickness distribution, ...);
- Finite-difference methods on *Eulerian* structured grids.

Based on discussions in the NEMO sea ice community between 2012 and 2017 and since the adoption of SI<sup>3</sup> in 2017, the strategy for sea ice in NEMO has followed these broad principles:

- A unique sea ice model in NEMO is the best response to duplication and resource waste.
- High modularity best addresses application-dependent needs and scientific uncertainties.
- Using the NEMO framework and principles at all possible levels (from consortium agreement, to coding rules, documentation, and output framework, ...) is highly efficient.
- In particular, following the NEMO numerical and computational choices is most natural. This is why SI<sup>3</sup> uses finite-difference, *Eulerian*, structured, C-grid numerical methods and the FORTRAN programming language.

With such principles in mind, several achievements have been completed since 2017.

- A modular code framework for SI<sup>3</sup> was developed.
- New physical options were implemented, many of which from the CICE code but not only. These include an alternative method for air-ice coupling; two representations of melt ponds; four sea ice rheological choices; and two landfast sea ice parameterizations.
- An evaluation package was developed, in python language.
- Sea ice-specific demonstration test cases were developed.
- Progress towards code documentation was accomplished.
- A long-term strategy was elaborated during SIWG workshops; and in particular through an international community workshop held in Laugarvatn, Iceland in 2019.
- Most NEMO sea ice research groups transitioned their systems to SI<sup>3</sup> and support has been provided to operational groups implementing SI<sup>3</sup>.

Many of the aforementioned developments stemmed from new collaborations between SIWG members, which contributed to developing team spirit. Three international projects must be acknowledged for contributing to structure and coordinate our activities: IMMERSE (H2020), IS-ENES3 (H2020) and SI<sup>3</sup> (CMEMS).

In parallel, a new actor has emerged in the sea ice modelling community, around the “Scale-Aware-Sea-Ice-Project” (SASIP), collecting major funds for 6 years (2021-2027) from the Schmidt Futures foundation. Among a large number of activities, SASIP is developing a new model based on neXtSIM using finite elements and discontinuous-Galerkin methods, and so therefore will be on a different grid, which they will couple to NEMO.

## 10.2 Issues and developments

Challenges in sea ice modelling resemble their ocean counterparts only in part. As ocean modellers, sea ice model developers face specific demands from sub-communities of users, in terms of target processes, output diagnostics and resolution. However, the physical understanding of sea ice is far less advanced than that of the ocean. Governing equations are under debate (we have only a few *rules*). Furthermore, increasing sea ice model resolution does not currently lead to better representation of the sea ice dynamics, as the continuum assumption arguably becomes invalid near kilometric resolutions, for which model grid-cells may no longer contain a representative sample of subgrid features.

The strengths and limitations of continuum sea ice models such as SI3, along with prospects for their evolution, are detailed in Blockley et al. (2020) and their suitability for operational forecasting applications is discussed in Hunke et al. (2020). These two publications summarise outcomes of a community workshop organised by the NEMO Sea Ice Working Group (SIWG) in Laugarvatn, Iceland, in 2019. In these papers, it is argued that:

- The continuum approach is useful for many years to come for climate applications;
- The continuum approach is questionable for high-resolution operational applications but no feasible better alternative currently exists;
- Discrete-element models (DEMs) are a possible alternative but are not ready and, for current HPC architectures at least, are prohibitively expensive.

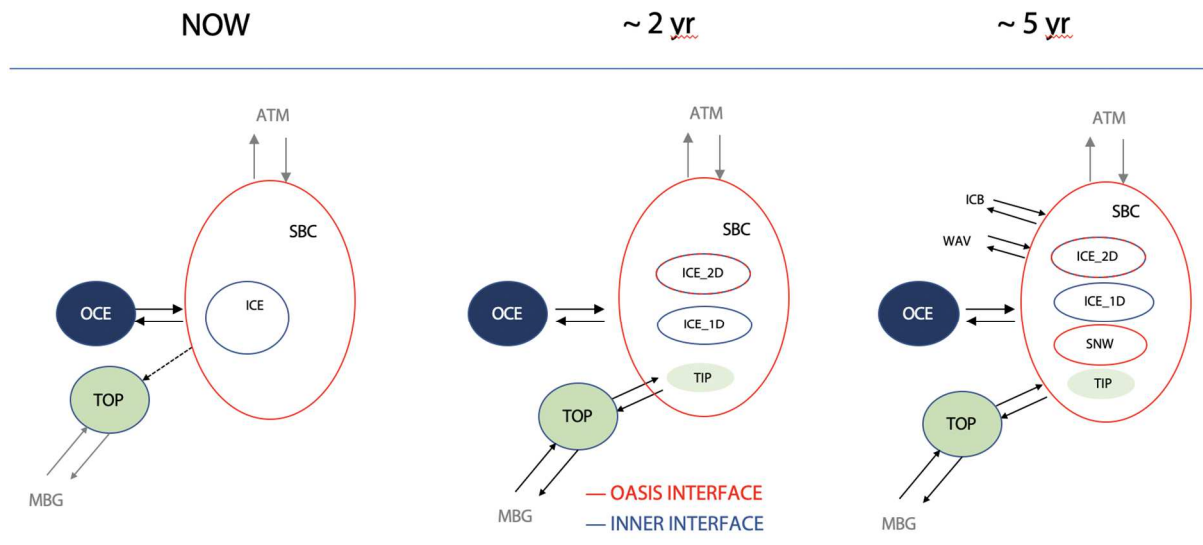
*In this context, we propose here a strategy for the evolution of SI3, mostly based on*

- The maintenance, update and development of the continuum approach;
- The study and improvement of the representation of sea ice dynamics, at  $\leq 1$ -km resolution
- The study and exploration of possible input from DEMs
- Increasing modularity of, and interfaces to, the current code to allow users to use SI3 for a wider range of applications

### 10.2.1 Modularity

Currently, sea ice in NEMO has sophisticated interfaces with the atmosphere and ocean, as well as a rudimentary interface with marine biogeochemistry. However, this does not match needs for many applications and so we face increasing requests from the scientific community to interface sea ice with an increasingly long list of ESM components.

An overarching question regarding more interfaces is to which extent the use of OASIS will be pushed forward and promoted as a standard option.



ICB = ICeBerg  
 MBG = Marine BioGeochemistry  
 TOP = Tracer Ocean Paradigm / TIP = Tracer in Ice Paradigm  
 SBC = Surface Boundary Condition

**Figure.** Schematic representation of the sea ice-related interfaces within NEMO, as it stands now (left), as planned within the next two years (centre), and our target over a 5-year horizon (right).

*Full split of horizontal and vertical physics (CNRS, 2021)*

Splitting the interface of SBC and sea ice between vertical and horizontal processes in NEMO will be useful for several reasons.

First, it is hard to decide upon physical, numerical or computational arguments whether the current approach (finite differences on a C-grid) or alternatives (finite-element with Discontinuous Galerkin framework) is best. Introducing an explicit interface for horizontal dynamics would allow for testing different options.

Second, some groups might see as positive the possibility to use other sea ice vertical physical packages (such as IcePack), which would also be facilitated with such an increased level of modularity.

Third, such horizontal-vertical modularity could also bring more flexibility for testing and debugging activities, of particular relevance for R&D and operational activities involving data assimilation.

Resolution: A split sea ice interface to SBC into vertical and horizontal processes is being implemented under the SASIP (CNRS). All technical points are not sorted yet. Most prominent issues are the shape of exchanged sea ice arrays and the point of view adopted for splitting - for example splitting can be based on what is computationally easier (i.e., 2D vs 1D processes in the model) or based on the physical processes (i.e., dynamics vs thermodynamics). In that respect, the place of ridging and rafting processes can be questioned because, although a 1D process in the model, it is strongly tied with the dynamics (drift and deformation). In either case we need to be clear about the terminology used.

### *Tracer in Ice Paradigm (TIP, CNRS, 2022-2024)*

The role of sea ice as an iron conveyor is also increasingly recognized (Lannuzel et al., Person et al.). Explicit representations of water isotopes or ice algae are also useful to tackle contemporary research questions.

The interface of sea ice with tracers is rudimentary in NEMO. There has been recent development work. For instance, Hayashida and Steiner (CCCMa, Victoria) have implemented an ice algae model in NEMO 3.6; whereas Person et al. (2020) have implemented a vertically-constant, category-dependent iron tracer in sea ice.

Resolution: An equivalent of TOP for sea ice (TIP for *Tracer in Ice Paradigm*) will be developed, first for vertically-constant passive tracers (CNRS, 2022-2023), based on Person et al. (2020), currently available in a separate branch of NEMO 3.6. TIP will support water isotopes, for which there are ongoing projects in the ESM community. Meanwhile, a specific discussion group with stakeholders (TOP-WG, SIWG, Victoria, CNRS) in order to foresee progress on other biogeochemical tracers. Possible contributions from ClimArctic are to be expected as well.

### *Interactions between sea ice and continental ice forms (UCLouvain, 2022-2027)*

Interactions between sea ice and continental cryosphere (large icebergs, shelves) have been put forward, in particular near the Antarctic Coast where they contribute to landfast sea ice development.

Antarctic land-fast ice does not currently emerge from model physics and is important for modelling coastal polynyas and dense water formation. An ad-hoc solution has been proposed by van Achter et al. (2022), combining tensile strength and large icebergs as part of the sea ice mask. However, as this approach requires observations of large icebergs to be available, it cannot be implemented globally. More generally, given the importance of iceberg-sea ice interactions, it is desirable for the sea ice and land-ice components to see each other in a more systematic and physical manner, which goes from basic conservation laws (area) to diagnostics. However, there are lots of open questions on the topic as to how iceberg and sea ice models should interact.

Resolution: A funded Belgian project (LICEPOD, UCL, 2022-2026) will explore means to improve the representation of iceberg-sea ice interactions, hopefully leading to improvements in Antarctic landfast ice in global configurations (UCL/CNRS) and contribute to improve the interface between sea ice and icebergs.

### *Wave-ice interface*

That sea ice attenuates waves, fractures upon wave action and moves in response to wave action is well established (Stopa et al. 2018; Boutin et al 2020). Interaction between waves and sea ice therefore needs to be included for NEMO to be fully coupled with waves.

Currently, however, there is no wave-ice interface within NEMO. Ad-hoc code has been developed in the past without much coordination. For instance, a wave contribution in the momentum equation has been implemented (and possibly other things) have been implemented in NEMO 3.6 at IFREMER.

NOC has implemented collisional rheology and floe-size (Rynders and Aksenov) accounting for wave effects and wave radiation pressure term in the momentum equation.

Some groups have argued for full floe size distribution (FSD), which is possible but implies extra complexity in the code and is expensive. However, none of the above have been sustainable developments and have not been included back into the NEMO trunk.

Resolution: Launch a specific discussion with stakeholders (Reading, NOC, Waves WG, SIWG, Brest, SASIP, Mercator, Met Office, and ECMWF). Two objectives: Identify current status of the wave-ice interface and define a strategy for wave-ice interactions that are agreeable to both waves and sea ice stakeholders. This task may benefit from a targeted workshop. There are 2 groups of stakeholders, with different foci and interests. The *wave* research community members are interested in the effect of sea ice on waves (attenuation, reflection, ...). The *sea ice* research community is interested in the effects of waves on sea ice (floe breaking, wave divergence, ...).

The ClimArctic project (led by IFREMER and with CNRS as a partner) will warrant progress on getting the wave-ice interface works back into NEMO, with identified NEMO developers as contributors.

Under the (UK) NERC DEFIANT project, which has a focus on Southern Ocean sea ice, an FSD will be brought back into SI3. The FSD model is based on Roach et al. (2018) but with the inclusion of a brittle fracture parameterisation (Bateson et al, 2022) that facilitates a realistic match against some observations of FSD in the central Arctic.

#### *Snow on sea ice as a separate medium*

Accounting for air-snow chemical exchanges over sea ice motivates improved representations of snow over sea ice in the CriSES project. More generally snow is also important for sea ice heat and mass balance, as well as for freeboard-based satellite retrievals of ice thickness. There are various advanced continental snow models (CROCUS, SnowTherm, ...). Model infrastructure is not ready to receive them, as snow is hard-coded into the sea ice model. If we want to benefit from such models, one needs to make the interface between sea ice and snow explicit, probably using OASIS.

Resolution: That work would require revising the snow-sea ice interface to better separate the two media. There are interested stakeholders (CriSES H2020 project, 2022-2026), which CNRM is part of. The CriSES group would take the lead on this task if they manage to recruit a suitable candidate.

#### 11.2.2 Horizontal dynamics

Horizontal ice dynamics are central in the representation of sea ice and present the largest challenges to the modelling community. In terms of horizontal dynamics, the continuum approach will be maintained and further developed. Issues we face are wide and many. For instance, how to represent km-scale drift and deformation processes? Is there a better sea ice rheology to accomplish that? How do the different advection schemes perform?

Resolution: Evaluating horizontal ice dynamics at high resolution using idealised and realistic configurations is underway through IMMERSE and will go on, both within the NEMO SIWG and well beyond inside the sea ice modelling community.

There is also work towards improving the different rheologies of the SI<sup>3</sup> code. The IMMERSE sea ice activities have highlighted numerical issues in EAP rheology, as well as uncertain behaviour in our VP

rheology implementation. Met Office has implemented the Rothrock (1975) formulation for ice strength and is currently exploring consequences. The CNRS-Grenoble group is seeking a numerically stable implementation of the BBM rheology (Olason et al., 2021) on a C-grid.

*Other issues:* Options for horizontal advection, like incremental remapping (Lipscomb and Hunke, 2004), or alternatives, should be examined. We should also consider progress in thickness redistribution (Roberts et al., 2019), alternative yield curves and landfast ice parameterization updates (JF Lemieux et al).

Finally, routines exist for the Tsamados et al. (2014) atmosphere-ice and ocean-ice drag and should be ported into SI3. This should be accomplished as part of the (UK) NERC CANARI project

### 10.2.3. Vertical physics

There are pending issues in the representation of vertical sea ice physics, reviewed hereafter.

#### *Thermodynamic inconsistencies and high liquid content ice types*

Several important problems related to supercooling and high-liquid content ice types are identified. Sub-freezing temperatures reach up to a few degrees below freezing in the worst cases. Ice towers form near ice-shelf boundaries, more so at high resolution and when under-ice-shelf cavities are open. These stem from a series of limitations: 1) ice nucleation in the ocean interior is not considered, 2) the minimum liquid fraction in sea ice is rather low, hence high-liquid fraction ice types (such as frazil and platelet ice) do not emerge, and newly formed ice piles instead of freely flowing in the ocean.

*Resolution:* A series of thermodynamic inconsistencies must be resolved to progress as follows:

(i) Ice thermodynamics should be formulated based on liquid fraction, following principles of the *mushy-layer* theory (Worster, 1992). This is partly done already, yet not systematically. Formulations for liquid fraction and *liquidus* (e.g. brine) salinity should be harmonised (Vancoppenolle et al., 2019). All properties should be written as weighted means of brine and pure ice contributions, and the same formulations for specific heat, thermal conductivity, enthalpy and permeability should be used throughout the code. Heat equation and ice growth and melt calculations should be written in terms of ice enthalpy, which needs to rewrite the numerical scheme for the diffusion of heat. This work has already been partly achieved in LIM1D.

(ii) Vertical variations in ice salinity should be accounted for, following developments over the last decade (Griewank and Notz, 2013; Rees-Jones and Worster, 2015). Such work has also been done in LIM1D (Thomas et al., 2020, Wongpan et al., 2021). Once done, the new ice liquid fraction can be made a parameter which allows for emergence of platelet ice (Wongpan et al., 2021).

(iii) To get rid of near-ice-shelf ice towers would require to consider the cycle of frazil ice: nucleation, frazil ice ascent and horizontal transport. There is no easily available solution, so we would need to develop one, which would take resources. There is expertise in the SIWG at U. Reading (Heorton et al., 2017, Mackie et al., 2020).

Part of these developments could occur in the framework of TRACCS, a French-ANR project under examination.

### *Light transmitted through sea ice*

Light transmission through sea ice in SI<sup>3</sup> follows a rather simple formulation at present. Surface albedo is formulated as a function of the sea ice and snow state as in many ESMs. Snow is assumed to be opaque, whereas light transmission through bare ice decays exponentially. Under sea ice, the same approach is used as in ice-free waters. Near-infrared is entirely absorbed in the near-surface ocean over the top 50 cm. Visible light is equally split into RGB bands.

There are several identified issues in this representation (Lebrun et al., in revision). First, accounting for snow depth and temperature is key to reproduce variations in light transmitted under sea ice. Second, ice-free water assumptions on spectral distribution of under ice light are invalid for ice-covered waters, and lead to underestimation of under-ice irradiance.

Lebrun et al. make propositions to resolve these two issues. First, surface transmission and attenuation coefficients can be tuned to observed transmittance. Second, spectral fractions must be adjusted under sea ice. Stroeve et al. (2021) adds that photosynthetically available radiation calculation from shortwave also needs to be slightly adjusted.

The tuned surface transmission and attenuation coefficients of Lebrun et al. are already implemented in SI3, however, they lead to spurious underestimation of surface melting.

#### Resolution:

- (i) Make sure the albedo scheme is satisfactory for all partners and evaluate options for future evolution (CNRS/Reading).
- (ii) Updating Lebrun et al. (2019) scheme for light attenuation with latest developments. Test multi-layer snow scheme as a solution for excess surface melt
- (iii) Revise infrared absorption and light fractionation under sea ice (ocean).
- (iv) Start discussions regarding further developments. If microstructure (gas and brines) was better resolved, sea ice optics could be less heavily parameterised and the current broadband scheme could be revised. Two-band and delta-Eddington schemes have been introduced in CICE, but what the advantages are of these approaches is not fully clear. Therefore, we need a preliminary evaluation before we move on.

### *Other issues*

Some melt pond processes are missing (under-ice ponds, refreezing). Resolution: Developments are ongoing in Reading (under-ice melt ponds, melt pond refreezing) and could come back into NEMO.

Snow formulation in SI3 is very simple and a constant snow density is used. Resolution: Easy progress is to implement the vertical density distribution by Lecomte et al. (UCLouvain).

We have no toy option for sea ice and this might be useful for idealised studies. Resolution: Implement a Semtner 0-layer model.



#### 10.2.4 Configurations, test beds data assimilation

- There are not many test cases for sea ice, yet they are useful and could be utilised more in the future. This is one of the recommendations of IMMERSE activities.
- There is a strong need for simplified ocean physics to help high-resolution testing of sea ice. Such a setup has been developed in the University of Reading and should come back into the trunk
- Ocean parameterizations under sea ice are heavily important to air-ice-sea exchanges and polar ocean state in ESMs, however they are rarely investigated.
- Data assimilation has been used with NEMO sea ice for several decades but is now becoming more popular with non-operational users. SI3 currently has access to the standard assimilation tools in NEMO, maintained by the DAWG, including the observation operator in 'OBS' and incremental analysis update (IAU) code in 'ASM'. The core data assimilation codes however (such as NEMOVAR) are developed and maintained separately outside of NEMO. DAWG needs to be connected to relevant SIWG members to ensure there are appropriate links between the ICE and OBS/ASM codes.
- We also have a few external tools (evaluation, etc..., e.g. Lin et al., 2021). How should they be shared and maintained?

#### 10.3 Documentation and dissemination

There is a strong need for documentation to facilitate the uptake of SI3 users. The current documentation is only an advanced draft and so this will need to be progressed as a top priority. This is also true of scientific papers describing the capabilities of SI3 code. In the end, both science and user guide aspects need to be covered. Sustainability aspects of the documentation should also be considered.

#### 10.4 High-Performance Computing

Most contemporary HPC issues internal to the sea ice code collapse down to rheology. Reducing iterations and global communications is the classical HPC concern for sea ice modellers. Such reductions are usually achieved by using methods for fast convergence (as in the adaptive EVP for instance) and by increasing the number of halos. Load balance is important but can now be handled by coupling sea ice and ocean through OASIS. In the prospect of transitioning towards GPU-based supercomputers while keeping CPU performance, array management and data transfer between arrays must be considered carefully and some modularity is probably needed here. There are ongoing talks with Andrew Porter and Chris Dearden from STFC in the UK regarding these issues. The SIWG only has few HPC experts and will essentially follow any practical recommendation from the HPC working group.

#### 10.5 Networking and community

Current sea ice team involves permanent and non-permanent personnel mostly from consortium institutions, but not only. Most of these are also sea ice working group members. Here is a recent list of contributors to sea ice code development and/or evaluation:

- CMCC (It) — Dorotea Iovino
- CNRS/LOCEAN (Fr) — Gurvan Madec, Clément Rousset, Martin Vancoppenolle
- CNRS/IGE (FR) — Pierre Rampal, Laurent Brodeau
- ECMWF (Eur) — Sarah Keeley, Steffen Tietsche
- Mercator Océan (Fr) — Gilles Garric
- Met Office (UK) — Ed Blockley, Emma Fielder, Ann Keen
- NOC (UK) — Yevgeny Aksenov, Stefanie Rynders
- UCLouvain (Be) — Thierry Fichefet, François Massonnet, Xia Lin
- University of Reading (UK) — Danny Feltham, David Schroeder, Rebecca Frew, Adam Bateson

This list is dynamic. In the coming years, personnel associated with external projects, in particular ClimArctic (French ANR), CriSes (H2020), and SASIP (Schmidt Futures), to name a few, will contribute to sea ice developments in NEMO. The SIWG should remain open to contributions from all groups working with NEMO and interested in sea ice. We should also seek to enhance collaborations with Canadian and American research groups working on CICE, with whom we could share scientific ideas and to reduce duplication where relevant.

Funding assumptions are detailed in this chapter, they are probably not exhaustive. We will also need to coordinate with other working groups in particular:

- Land ice
- Waves
- HPC
- Data assimilation

## 10.6 Summary and perspectives

The key point of our strategy is to foster modularity of the NEMO sea ice code, not only by better defining interfaces between sea ice and other components of the Earth System, but also among different categories of sea ice processes.

In two years we target a more objective splitting between horizontal and vertical physics, and a tracer-in-ice module.

In five years, we foresee improved interfaces with snow, waves, bergs. There are lots of possible other developments. Most of these depend upon the availability of project funding, as well as on the contribution of trained scientists and developers, more of which can hopefully use the SI3 code in the near future.

Increased modularity will enable a more objective evaluation of the possible choices on horizontal ice dynamics, based on physical, numerical and HPC criteria. It is expected that this will help us make further progress for the next round of our strategy.

In the long run (10 yr), discrete element approaches for ice modelling could take more importance, bringing up issues such as the coupling between continuum approaches for ocean and atmosphere or even in the sea ice and discrete elements, and we could seek to harmonise the treatment of continental ice and sea ice, which would help to treat the pack ice / landfast ice / ice shelf transition.

## References

- Blockley, Ed, Martin Vancoppenolle, Elizabeth Hunke, Cecilia Bitz, Daniel Feltham, Jean-François Lemieux, Martin Losch, et al. 2020. « The future of sea ice modelling: where do we go from here? » *Bulletin of the American Meteorological Society*. <https://doi.org/10.1175/BAMS-D-20-0073.1>.
- Boutin, Guillaume, Camille Lique, Fabrice Ardhuin, Clément Rousset, Claude Talandier, Mickael Accensi, et Fanny Girard-Ardhuin. 2020. « Towards a Coupled Model to Investigate Wave–Sea Ice Interactions in the Arctic Marginal Ice Zone ». *The Cryosphere* 14 (2): 709-35. <https://doi.org/10.5194/tc-14-709-2020>.
- Briegleb, P., et Bonnie Light. 2007. « A Delta-Eddington Multiple Scattering Parameterization for Solar Radiation in the Sea Ice Component of the Community Climate System Model ». <https://doi.org/10.5065/D6B27S71>.
- Heorton, H., Radia, N., and Feltham, D.L., 2017. « A Model of Sea Ice Formation in Leads and Polynyas ». *Journal of Physical Oceanography* 47 (7). <https://doi.org/10.1175/JPO-D-16-0224.1>.
- Hunke, Elizabeth, Richard Allard, Philippe Blain, Ed Blockley, Daniel Feltham, Thierry Fichefet, Gilles Garric, et al. 2020. « Should Sea-Ice Modeling Tools Designed for Climate Research Be Used for Short-Term Forecasting? » *Current Climate Change Reports* 6 (4): 121-36. <https://doi.org/10.1007/s40641-020-00162-y>.
- Lebrun, Marion. 2019. « De l'interaction entre banquise, lumière et phytoplancton arctique ». Thèse de Doctorat, Sorbonne Université.
- Lecomte, O., T. Fichefet, M. Vancoppenolle, et M. Nicolaus. 2011. « A new snow thermodynamic scheme for large-scale sea-ice models ». *Annals of Glaciology* 52: 337-46. <https://doi.org/10.3189/172756411795931453>.
- Lemieux, Jean-François, Frédéric Dupont, Philippe Blain, François Roy, Gregory C. Smith, et Gregory M. Flato. 2016. « Improving the Simulation of Landfast Ice by Combining Tensile Strength and a Parameterization for Grounded Ridges ». *Journal of Geophysical Research: Oceans* 121 (10): 7354-68. <https://doi.org/10.1002/2016JC012006>.
- Lin, Xia, François Massonnet, Thierry Fichefet, et Martin Vancoppenolle (2021). « SITool (v1.0) – a New Evaluation Tool for Large-Scale Sea Ice Simulations: Application to CMIP6 OMIP ». *Geoscientific Model Development* 14 (10): 6331-54. <https://doi.org/10.5194/gmd-14-6331-2021>.
- Lipscomb, W. H., et E. C. Hunke. 2004. « Modeling sea ice transport using incremental remapping ». *Monthly Weather Review* 132: 1341-54.
- Mackie, Shona, Patricia J. Langhorne, Harold D. B. S. Heorton, Inga J. Smith, Daniel L. Feltham, et David Schroeder. 2020. « Sea Ice Formation in a Coupled Climate Model Including Grease Ice ». *Journal of Advances in Modeling Earth Systems* 12 (8): e2020MS002103. <https://doi.org/10.1029/2020MS002103>.
- Olason, Einar, Guillaume Boutin, Anton Korosov, Pierre Rampal, Timothy Williams, Madlen Kimmritz, Véronique Dansereau, et Abdoulaye Samaké. 2021. « A New Brittle Rheology and Numerical

Framework for Large-Scale Sea-Ice Models ». Preprint. Earth and Space Science Open Archive. Earth and Space Science Open Archive. World. 24 septembre 2021.  
<https://doi.org/10.1002/essoar.10507977.2>.

Person, R., M. Vancoppenolle, et O. Aumont. 2020. « Iron Incorporation From Seawater Into Antarctic Sea Ice: A Model Study ». *Global Biogeochemical Cycles* 34 (11): e2020GB006665.  
<https://doi.org/10.1029/2020GB006665>.

Roberts, Andrew F., Elizabeth C. Hunke, Samy M. Kamal, William H. Lipscomb, Christopher Horvat, et Wieslaw Maslowski. 2019. « A Variational Method for Sea Ice Ridging in Earth System Models ». *Journal of Advances in Modeling Earth Systems* 11 (3): 771-805.  
<https://doi.org/10.1029/2018MS001395>.

Rothrock, D. 1975. « The energetics of the plastic deformation pack ice by ridging ». *Journal of Geophysical Research* 80: 4514-19.

Squire, V. A. 2007. « Of ocean waves and sea-ice revisited ». *Cold Regions Science and Technology* 49: 110-33. <https://doi.org/10.1016/j.coldregions.2007.04.007>.

Stopa, J, P Sutherland, et F Arduin. 2018. « Strong and highly variable push of ocean waves on Southern Ocean sea ice ». *Proceedings of the National Academy of Sciences* 115 (23): 5861-65.

Stroeve, Julienne, Martin Vancoppenolle, Gaele Veysiere, Marion Lebrun, Giulia Castellani, Marcel Babin, Michael Karcher, Jack Landy, Glen E. Liston, et Jeremy Wilkinson. 2021. « A Multi-Sensor and Modeling Approach for Mapping Light Under Sea Ice During the Ice-Growth Season ». *Frontiers in Marine Science* 7. <https://doi.org/10.3389/fmars.2020.592337>.

Thomas, M., M. Vancoppenolle, J. L. France, W. T. Sturges, D. C. E. Bakker, J. Kaiser, et R. von Glasow. 2020. « Tracer Measurements in Growing Sea Ice Support Convective Gravity Drainage Parameterizations ». *Journal of Geophysical Research: Oceans* 125 (2): e2019JC015791.  
<https://doi.org/10.1029/2019JC015791>.

Tsamados, M, D. Feltham, D. Schroeder, et D. Flocco. 2014. « Impact of Variable Atmospheric and Oceanic Form Drag on Simulations of Arctic Sea Ice ». in: *Journal of Physical Oceanography* 44 (5).  
<https://doi.org/10.1175/JPO-D-13-0215.1>

Van Achter, Guillian, Thierry Fichet, Hugues Gousse, Charles Pelletier, Jean Sterlin, Pierre-Vincent Huot, Jean-François Lemieux, Alexander D. Fraser, Konstanze Haubner, et Richard Porter-Smith. 2022. « Modelling Landfast Sea Ice and Its Influence on Ocean–Ice Interactions in the Area of the Totten Glacier, East Antarctica ». *Ocean Modelling* 169 (janvier): 101920.  
<https://doi.org/10.1016/j.ocemod.2021.101920>.

Vancoppenolle, Martin, Gurvan Madec, Max Thomas, et Trevor J. McDougall. 2019. « Thermodynamics of Sea Ice Phase Composition Revisited ». *Journal of Geophysical Research: Oceans* 0 (ja). <https://doi.org/10.1029/2018JC014611>.

Wongpan, P., M. Vancoppenolle, P. J. Langhorne, I. J. Smith, G. Madec, A. J. Gough, A. R. Mahoney, et T. G. Haskell. 2021. « Sub-Ice Platelet Layer Physics: Insights From a Mushy-Layer Sea Ice Model ». *Journal of Geophysical Research: Oceans* 126 (6): e2019JC015918.  
<https://doi.org/10.1029/2019JC015918>.

Worster, M. G. 1992. « Interactive Dynamics of Convection and Solidification ». In *Interactive dynamics of convection and solidification*, edited by S. H. Davis et al, 113-38. Kluwer.

## 11 Land ice / Ocean interactions

**Lead Author:** Pierre Mathiot<sup>1</sup>

**Co-authors (alphabetical order):** J. Deshayes<sup>3</sup>, P. Holland<sup>2</sup>, K. Hutchinson (*NEMO System Team member*)<sup>3</sup>, A. Jenkins<sup>4</sup>, N. Jourdain<sup>1</sup>, C. Rousset (*NEMO System Team member*)<sup>3</sup>, R.S. Smith<sup>5</sup>

**Other participants of the working group:** T. Fichefet<sup>6</sup>, J. Harles<sup>7</sup>, J. Marson<sup>8</sup>, Jeff Ridley<sup>9</sup>

<sup>1</sup> Institut de géophysique de l'environnement, France

<sup>2</sup> British Antarctic Survey, United Kingdom

<sup>3</sup> LOCEAN, France

<sup>4</sup> Northumbria University Newcastle, United Kingdom

<sup>5</sup> National Centre for Atmospheric Science, United Kingdom

<sup>6</sup> Université Catholique de Louvain-La-Neuve, Belgium

<sup>7</sup> National Oceanographic Center, Southampton, United Kingdom

<sup>8</sup> University of Manitoba, Canada

<sup>9</sup> Met Office, Exeter, United Kingdom

**Priority categories definition:**

**Category 1:** this category includes the themes/topics we think are critical.

**Category 2:** Will be very welcome within 5 years.

**Category 3:** Relevant but for later.

### 11.1 Executive summary

Land ice / ocean interactions include **ice-shelves**, **glacier termini** and **icebergs** melting, as well as **surface** and **sub-glacial runoff** from the ice sheet. Those interactions have diverse influences on the ocean, climate and sea level rise (Schloesser et al. 2019; Bronselaer et al. 2018; Oppenheimer et al., 2019). Due to the missing representation of the interactions between ice sheets, oceans and the atmosphere in many of the current Earth System Models (ESMs), the influence of glacial fresh water sources remains highly uncertain. In order to include an ice sheet component in NEMO-based ESMs, icebergs and explicit circulation under the ice shelves were implemented within NEMO several years ago (Marsh et al., 2015; Mathiot et al., 2017). Since then, these aspects have become more commonly used in processes and climate studies (Merino et al. 2016; Jourdain et al., 2017; Storkey et al., 2018; Haussman et al., 2020; Huot et al., 2021). More recently, NEMO has been coupled to various ice sheet models (Smith et al., 2021; Pelletier et al., 2021 and Favier et al., 2019). Despite these efforts, modelling these processes is at the early stages and remains challenging.

The 10 year roadmap of this working group is dedicated to supporting the development of ESMs that are able to model Antarctic and Greenland ice sheet / ocean interactions. The objective is to release robust and easy to use NEMO based ESMs able to simulate realistically the oceanic state and variability on the Greenland and Antarctica continental shelves, as well as the polar ice sheets states and variability over the recent and future centuries. In such ESMs, there are 2 sources of issues: missing physics and external biases. External biases are outside the scope of this document as they come from a lack of tuning, realism of the forcings and other components of ESMs. We therefore focus only on how to improve physics of the land ice ocean interactions within NEMO. In terms of major priorities, our short-term highlights focus on exploring and testing ice shelf cavity and sub-ice boundary layer parameterizations, schemes for the migration of a calving front and conservation for ice sheet coupling, improving the HPC performance of the iceberg module and the design of test cases to

accompany the associated new developments. Despite the top priority of such topics, fundings have been identified only for exploring new ice shelf cavity parametrisations and enhancing the ocean/ice sheet coupling method.

## 11.2 Timeline

### Estimated NEMO Land Ice / Ocean interaction strategy timeline

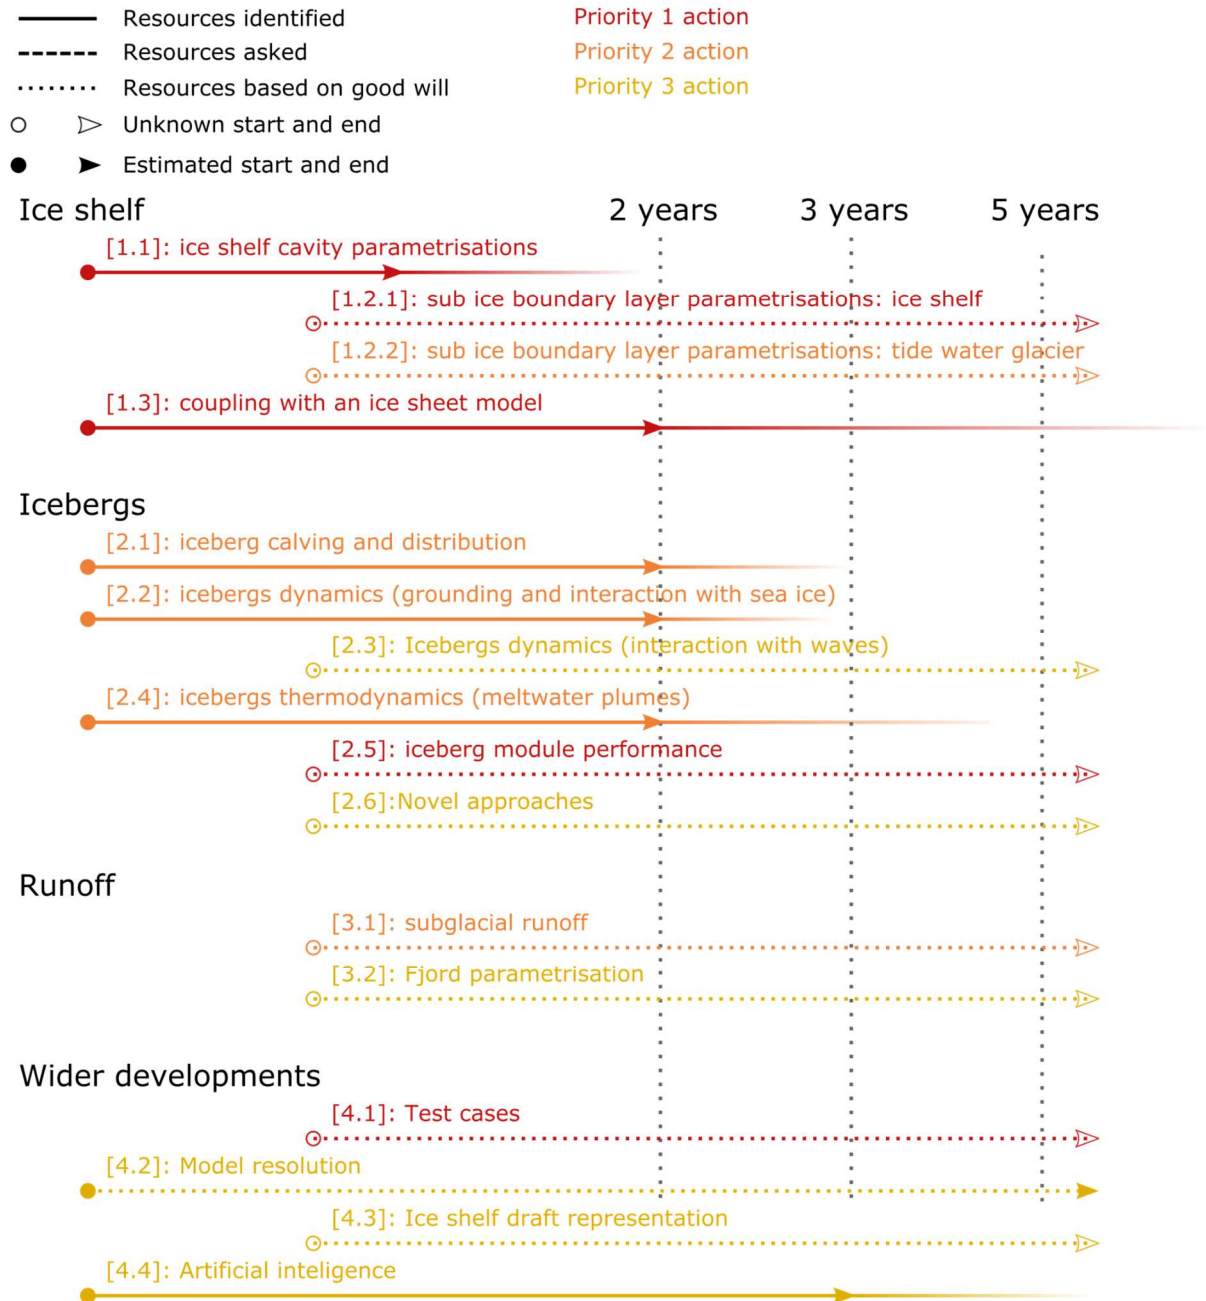


Figure 1: Timeline of the various actions described in this chapter. Numbers in brackets are referring to the section number in the current chapter where a detailed description of the task is available. Arrows indicate the time frame to reach each milestone, where a faded line after the arrow shows a short-term goal followed by longer term continuous work on the topic. Task numbers mentioned in brackets on the figure are also reported in the text of each section ([X.Y.Z])

In this chapter, the main focus is on short to medium term actions (2-5 year timescale). It is worth noting that some of these actions are “never ending” tasks. The research and engineering work

described in this chapter are split into different streams. The estimated timing and priority associated with each topic for the next 5 years is described in the timeline in Fig. 1.

The short-term actions are strongly dependent on what funding applications are successful. It is worth noting that a significant portion of the work identified in this chapter does not yet have dedicated funding assured. Furthermore, the “champions” of many of the major deliverables are not in the NST and so without dedicated human resources, our milestones are vulnerable to external factors.

### 11.3 Context

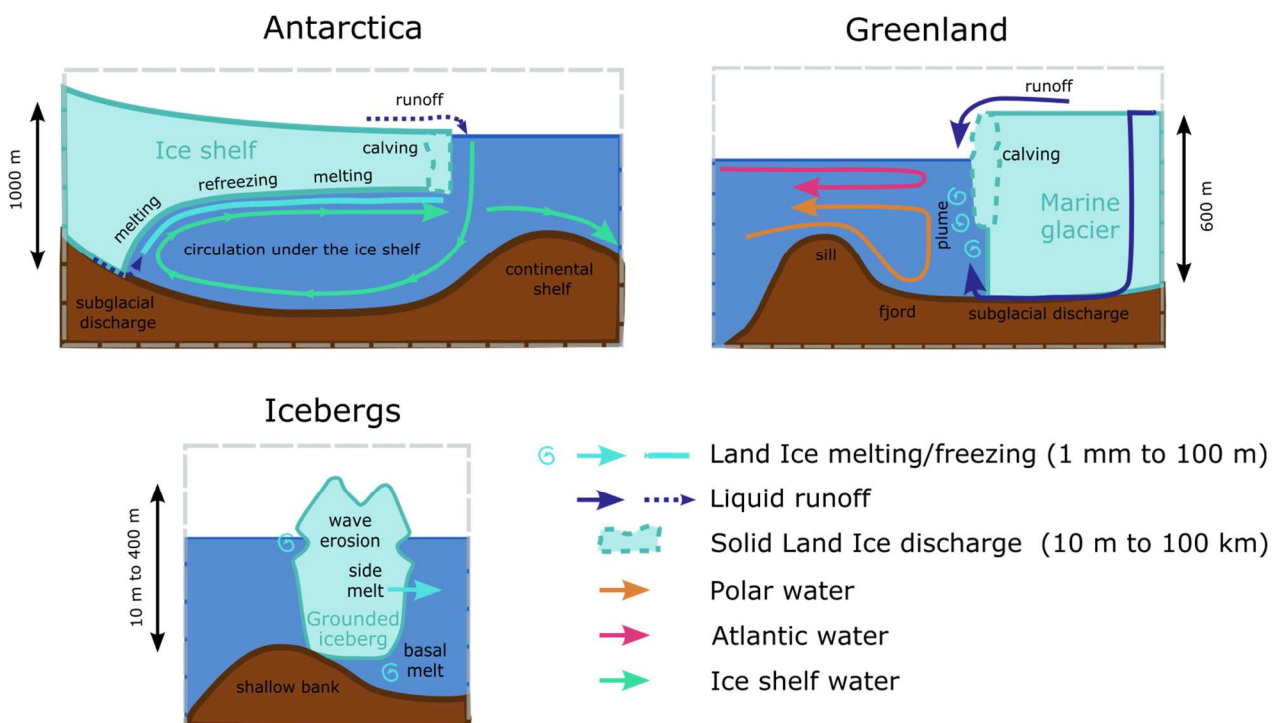
*It is important to note that most NEMO applications that include the polar regions must carefully consider the fresh water inputs from land ice melting. Development of NEMO must therefore involve significant work on the representation of the interaction between land ice and ocean.*

The land ice includes ice sheets, icebergs, and ice-shelves. **Land ice** builds up through the accumulation of snowfall over Greenland and Antarctica. It influences the ocean through the melting **ice-shelves** or **glacier termini** at the edge of the continents, via calving **icebergs** that slowly drift at the ocean surface, and seasonally at the **surface** and **subglacial runoff** induced by ice sheet surface melting (Fig. 2).

Icebergs in NEMO are handled through the ICB module and surface runoff is handled through the runoff module. Both are included within the Surface Boundary Condition (SBC) code. Basal ice-shelf melt is handled through the ISF module. ICB and ISF modules are relatively new developments within NEMO (Marsh et al., 2015 and Mathiot et al., 2017).

Although land ice and sea ice bear some physical similarities, the modelling components to handle them are usually drastically different due to the differing scales of the problem and the distinctly separate processes at play.

This chapter is designed for applications ranging from Earth System Models (Ocean / Atmosphere / Ice sheet coupled together) to ocean global and regional configurations forced by atmospheric fields (with static ice or with a coupled ice sheet model). The horizontal resolution considered in this document spans from 50 km to 1 km.





*Figure 2: Schematic description of the Land Ice / Ocean interactions. Antarctica and Greenland boxes show the typical kinds of interactions at play, but note that these are not exclusive to the denoted regions (marine glaciers are present in Antarctica and ice shelves are present in Greenland).*

## 11.4 Ice shelves

Ice sheet mass loss accounts for around a third of the present rate of global mean sea level rise, and this contribution is expected to increase and eventually dominate global mean sea level change in the coming decades and centuries (Oppenheimer et al., 2019). Additionally, ice sheet mass change is also persistently the most uncertain term in the future global mean sea level budgets (Oppenheimer et al., 2019). It is thus important to understand the physical detail of how the ocean and ice sheets interact. Accurate representation of these interactions is therefore necessary in order to either adequately capture the physics of the sub-ice boundary layer, represent circulation and transport within the ice shelf cavity if permitted by the model resolution, or adequately parametrize these interactions if the whole (or part of) the ice shelf cavities are missing.

The first implementation of the ice shelf module was done about 6 years ago and published in Mathiot et al. (2017). Since then, it has been widely used by various institutes as a basic parametrization of closed ice shelf cavities (as in Storkey et al. 2018; Merino et al. 2016; Boucher et al. 2020), for interactive cavities (as in Haussman et al., 2020; Huot et al., 2021) or for its coupling interface (as in Smith et al., 2021; Pelletier et al., 2021; Favier et al., 2019). Furthermore, in 2019 the ice shelf module was re-written to allow for a mix of parametrized ice shelf cavities and explicit ice shelf cavities, a more stable ice sheet coupling interface (Smith et al., 2021) and finally to facilitate the easy inclusion of a new parametrization for the ice shelf cavity or sub-ice boundary layers.

### 11.4.1 Ice shelf cavity parametrizations: **Category 1**

*Resources: IGE as part of ESM2025 project*

**[1.1]** Most global climate models, such as the ones used in CMIP6 do not resolve ice-shelf cavities. The entirety of ocean / ice shelf interactions needs to therefore be parametrized. Several parameterizations of varying complexity have been developed in the last 20 years to derive melt rates from far-field ocean properties. However, assumptions in the various formulations differ, giving rise to a large variety of melt rates (Favier et al. 2019). In the latest version of NEMO, only (Beckmann and Goosse 2003) is available and its performance is known to be poor (Favier et al. 2019 and Burgard et al., 2022, submitted). The evaluation of the various ice shelf parametrizations in a realistic test bed (Burgard et al., 2022, submitted) will help to define the few promising parametrizations to be adapted, implemented and tested within NEMO. Such parametrizations could also benefit ‘cavity resolving’ configurations so as to include the contribution of part of the cavities poorly represented like the area close to the grounding line.

**Key paper:** Burgard, C.; N. Jourdain; R. Reese; A. Jenkins; P. Mathiot: An assessment of basal melt parameterisations for Antarctic ice shelves, 2022, Cryosphere Discussion.

### 11.4.2 sub-ice boundary layer parametrizations:

*Resources:*

- *NUN, no specific funding identified*
- *BAS, funded secure to work on the scientific problem (but not with NEMO)*
- *NOC, phd proposal and NERC CLASS project*

**[1.2.1]** As detailed in Asay-Davis et al. (2017), the current ice shelf melt formulations are still immature. They suffer from many deficiencies and lack of knowledge. Preliminary results from ISOMIP+ (Asay-Davis et al. 2016) raised the issue that current treatments of sub-ice-shelf

thermodynamics do not converge with vertical resolution, either within a given model or between models. Different choices about how T, S, heat, and freshwater fluxes are treated in the sub-ice-shelf boundary layer led to disagreements between models and a lack of convergence within a given model with increasing vertical resolution. Recommendations on how to achieve such convergence with increasing resolution, changing vertical coordinates or sampling method will be welcome to achieve more robust science. Furthermore, the representation of the buoyant plume dynamics along the ice interface and the associated entrainment of ambient water should not be overlooked. Whilst this can be especially difficult to simulate in z coordinate models, there is evidence that this process may be improved by implementing sigma coordinates or other more flexible vertical coordinates.

In the latest flux formulation, ice shelf melt is proportional to the top friction velocity so melt rate is very sensitive to the drag, surface roughness and tides (Gwyther et al. 2015; Hausmann et al. 2020; Jourdain et al. 2017). Knowledge on these key parameters needs to be improved. Regarding the core of the current melt formulations, the latest idealized studies show that the top boundary layer consists of an inner, friction-dominated boundary layer and an outer geostrophic flow, with buoyancy playing a dominant role in both (Jenkins 2016). This double-layer structure is not accounted for in the current sub-ice-shelf / ocean parametrization. A turbulent closure scheme would however be needed before such scheme could be applied to realistic problems.

These questions are open research questions and some works are underway. **Category 1**

**Key paper:** Asay-Davis, X.S., Jourdain, N.C. & Nakayama, Y. Developments in Simulating and Parameterizing Interactions Between the Southern Ocean and the Antarctic Ice Sheet. *Curr Clim Change Rep* 3, 316–329 (2017). <https://doi.org/10.1007/s40641-017-0071-0>

**[1.2.2]** Similar approach to the one used to represent ice shelf / ocean interactions cannot be used for vertical ice faces. The typical horizontal length scale of the buoyant plumes along a vertical ice face (~10 m) is far smaller than the model horizontal resolution considered in this chapter (~ 1km to ~50 km). Thus, present day regional and global configurations do not have sufficient resolution to capture the ice sheet / ocean interactions at play on glacier termini (ice shelf front or marine glacier). These interactions and the ocean circulation they trigger need to be parametrized. (Cowton et al. 2015; Jenkins 2011; Slater et al. 2016; Rignot et al. 2016). However, it is not clear yet what is the most adequate choice of parametrization for NEMO. So, before the implementation any of these solutions, we recommend a detailed analysis of performance of what is available, what is needed and at what NEMO resolution these parametrizations are relevant. Having such parametrization for Greenlandic glacier termini will also benefit the ice shelf / ocean interaction communities by its potential application to the vertical ice shelf front. **Category 2**

#### 11.4.3 Coupling with an ice sheet model: **Category 1**

*Resources:*

- UKESM as part of ESM2025
- IGE as part of ESM2025 and EIS projects

**[1.3]** Despite the limitations mentioned above, models with ice shelf/ocean interaction have advanced to the point where they are being used not only in hindcasts or sensitivity studies with static ice shelves, but also in Earth System Models with evolving geometry based on the response of an ice sheet model to ice shelf melt and surface mass balance (Smith et al. 2021). These tools will be of great help in order to estimate future contributions of Antarctica to sea level rise. In NEMO, an asynchronous ice sheet / ocean coupling method is implemented and has been successfully used in various configurations (Favier et al. 2019; Smith et al. 2021; Pelletier et al. 2021).

Two points need a careful evaluation. First, migration of the calving front is allowed by this method but needs to be tested in a realistic configuration to verify its stability and the coherency with the iceberg module. Additionally, by construction, the procedure used to move the ice shelf draft and grounding line is significantly non-conservative. This issue could be critical for climate applications. An option is available to correct the model state in order to remove any trend created by the coupling method. This scheme has only been tested in idealized test cases and needs to be evaluated in realistic applications.

To mitigate the conservation limitation and to allow for a high frequency coupling, synchronous ice sheet / ocean coupling methods are available (Jordan et al. 2018; Goldberg et al. 2012). However, we (chapter's authors) suggest to not engage any work on the synchronous coupling method until a detailed evaluation of the available method has been made and ongoing work on this at BAS shows encouraging results.

**Key paper:** Smith R.S., Mathiot P., Siahaan A., Lee V., Cornford S.L., Gregory J.M., Payne A.J., Jenkins A., Holland P.R., Ridley J.K., Jones C.G., Coupling the U.K. Earth System Model to dynamic models of the Greenland and Antarctic ice sheets, *J. Adv. Modeling Earth Systems*, accepted, (2021) <https://doi.org/10.1029/2021MS002520>

**Possible overlaps/dependencies:** wave working group, sea ice working group, vertical mixing, tides, vertical coordinates (kernel working group), machine learning.

**Expected code changes:** localized to ice shelf module (ice shelf cavity parametrization sub-module, sub-ice boundary layer parametrization sub-module and ice shelf coupling interface) and domaincfg tools (all changes related to vertical coordinates).

## 11.5 Icebergs

Ice sheet acceleration has increased the flux of icebergs over the last 30 years, which will accelerate further in future. Icebergs have been included as lagrangian particles in several ocean models (review in Asay-Davis et al. 2017) and have been shown to significantly impact the intrusions of CDW towards ice shelves (Bett et al. 2020) and the Southern Ocean in general (Schloesser et al. 2019). However, the influence of these changes cannot be assessed because current iceberg models are based on overly simple physics, with little consideration of links between icebergs and ice-shelf thickness, bathymetry or sea-ice stress.

The lagrangian iceberg module in NEMO had been implemented in 2015 (Marsh et al., 2015). Since then, because of its wide usage in various institutes, its stability has been improved and the icebergs thermo and dynamics has improved (Merino et al., 2016). However, as mentioned above, the lagrangian iceberg module in NEMO is not ready yet to tackle the future key questions. In the following, we described the key improvement needed in the lagrangian icebergs model to make it able to represent interaction with sea ice and bathymetry, and thus its evolution in a changing climate.

### 11.5.1 Calving and distribution: **Category 2**

*Resources:*

- *BAS as part of Ocean/Ice project*
- *UCL as part of LICEPOD project*

**[2.1]** In NEMO, calving rate is prescribed as a forcing. Improving physics of the calving itself is outside the scope of this document. This being said, the fresh water distribution from iceberg melting is very dependent of their size distribution (Stern et al., 2016). Furthermore, most of the volume is located in the largest icebergs (Tournadre et al. 2016). It is therefore critical that large icebergs are represented.

Calving of such large icebergs is rare and quasi-random. It is not adapted to the current iceberg generation scheme. Once such icebergs are generated, a fragmentation scheme (England et al., 2020; Bouhier et al., 2018) is needed to avoid an excessive life time and unrealistic melt pattern (Bouhier et al. 2018)

#### 11.5.2 Dynamics:

*Resources: BAS as part of Ocean/Ice project*

**[2.2]** The dynamic interactions between icebergs and surrounding sea ice is essential to reproduce the observed trajectory pattern. Lichey and Hellmer (2001) suggest a formulation of the sea-ice force that includes free drift of iceberg in low concentrated areas and the locking of icebergs in the sea ice pack. The locking formulation was recently included in the FESOM ice-ocean model (Rackow et al. 2017) and already tested in NEMO (Marson et al. 2018). **Category 2**

On shallow banks, fields of isolated grounded icebergs are critical when endeavoring to represent both landfast ice (supported by Olason (2016) with isolated islands) and polynyas at the lee side (Massom et al. 1998; 2001; Nihashi and Ohshima 2015). Experiments with crude representation of grounded icebergs show large improvements in the representation of landfast ice, pack ice and polynya representation (Huot et al. 2021; Bett et al. 2020). A landfast ice scheme is already available within NEMO (Lemieux et al. 2016) with sea-ice keels as anchor points. Modification of this scheme will be needed to add iceberg keels as extra source of anchor points. Thus, the realism of iceberg triggered landfast ice is highly dependent of the realism of the grounded icebergs fields. Therefore, work is needed to improve the realism of the iceberg generation (size distribution, thickness, calving sites ...) and the grounding scheme (see Vaňková and Holland, 2017). **Category 2**

**[2.3]** Wave stress is known to have an impact on the iceberg drift (only tested on the Grand Banks area). There is little published investigation of the response of icebergs to wave motion or the consequences for the wave field. Despite this lack of literature on the subject, the Canadian Iceberg Forecasting model does include such stress in its iceberg dynamical equation (Kubat, 2005). Furthermore, wave observations have shown that icebergs damp the wave field. It has also been clearly demonstrated that including a parametrization for icebergs in the wave model reduces all the large biases in the Southern Ocean (Ardhuin et al., 2011). **Category 3** (*no resources*)

#### 11.5.3 Thermodynamics: **Category 2**

*Resources: University of Manitoba, Canadian projects "Arctic ice, freshwater marine coupling and climate change" and a NSERC Discovery Grant.*

**[2.4]** Recent works show that the plume generated along the sidewall of an icebergs has different regime depending on the background velocity relative to the plume velocity: attached (meltwater is channelled directly to the surface and 'shield' the icebergs) or detached (meltwater is mixed over a broader layer). Each regime drives different melt rates and leads to different impacts on ocean stratification and upwelling of nutrients (FitzMaurice et al., 2017). In NEMO, meltwater is injected only at the surface and the distinction between the two regimes is currently not available. Furthermore, the various canonical iceberg melt formulations (wave erosion, lateral and basal melt) from (Gladstone et al., 2001) need a thorough analysis.

#### 11.5.4 Performance: **Category 1**

*Resources: NEMO ST, no specific funding identified.*

[2.5] During the development of the latest version of NEMO (4.2) the issue of performance of the iceberg modules was identified. The cost of the Lagrangian icebergs model depends mostly on the maximum number of icebergs within sub-domains (controlled by the calving rate, melting rate and advection within the sub-domain). So, for similar geographical size of the model sub-domains, the coarser the model resolution is, the larger the relative cost of the icebergs model compared to the total cost of a simulation. First analysis of the iceberg performance carried out by Met Office shows that the iceberg code is entirely serial and not performant on CPU as well as on GPU. It probably requires a re-write of the linked-list logic. It is, however, hard to tell beforehand what the speed up benefit such changes may produce. Furthermore, by nature, there is a large load balance issue among the 'iceberg domains' and the 'iceberg free' domains.

#### 11.5.5 Novel approaches:

*Resources: no funding*

#### [2.6]

**Simplified model:** In order to mitigate the issue related to the iceberg model cost, some could think about a simplified iceberg model. So before scoping for such a model, we will explore an alternative solution based on what is currently available. Stern et al. (2016) showed that generating large category icebergs only gives very similar results to a model with a mix of various sizes. This result is mostly explained by the fact that most of the iceberg mass is concentrated in the largest icebergs. As the cost is proportional to the number of iceberg categories, with the existing code, there are two solutions to decrease the cost on the current icb model:

- Track multiple icebergs as one unit (scaling factor in the namelist). Increasing the scaling factor will reduce the number of icebergs but we don't have any information regarding its impact on the spatial icb distribution. This warrants testing.
- Simulate only large icebergs and no medium or small icebergs as done in Stern et al. (2016). If the results are reproducible, it could be a viable alternative to decrease the cost. This also warrants testing. **Category 2**

**Eulerian model:** An iceberg Eulerian model has been developed for the MPI-ESM model as part of a PhD project (Erokhina et al., 2020). It has been developed for paleo-applications. The main goal was to be able to simulate a Heinrich event where very large numbers of icebergs are released in a short amount of time for a reasonable cost. No comparison in terms of cost or results with respect to a Lagrangian iceberg model has been made. The implementation and tuning cost of such a model in NEMO may be large in terms of FTE and elapsed time. Therefore, we suggest that first an inspection is conducted to verify if the performance (speed and result) of the refactored code and the limited number of categories are together able to make the Lagrangian iceberg model cheap enough for wider application. **Category 3**

**AI model:** Modelling iceberg drift is still challenging after almost 40 years of development. Current operational drift models are based on the momentum equation and use wind and ocean currents to calculate the drag forces (Kubat et al., 2005). They are still prone to large forecasting errors when compared to observations. Novel approaches based on AI are in development in order to improve iceberg drift forecasting (Yulmetov and Freeman, 2019). In

collaboration with the ML working group, we may evaluate the benefit of such techniques for the NEMO community. **Category 3**

**Key paper:** Asay-Davis, X.S., Jourdain, N.C. & Nakayama, Y. Developments in Simulating and Parameterizing Interactions Between the Southern Ocean and the Antarctic Ice Sheet. *Curr Clim Change Rep* 3, 316–329 (2017). <https://doi.org/10.1007/s40641-017-0071-0>

**Possible overlaps/dependencies:** Wave working group, Sea Ice working group, vertical mixing.

**Expected code changes:** iceberg distribution, calving, dynamics and thermodynamics work are limited to the corresponding sub-modules. The work to improve performance could lead to a change in the kernel of the iceberg module. Furthermore, interactions with other working groups could lead to minor changes in their respective modules.

## 11.6 Runoff

*Resources: UKESM, funding not identified*

**[3.1]** It is known that the emergence of fresh subglacial runoff at glacier or ice sheet grounding lines generates buoyant turbulent plumes that enhance heat transfer across the ice–ocean boundary and the submarine melt rate for the portion of the glacier face (Jenkins, 2011) or ice shelf (Wei et al. 2020) in direct contact with the plume. It is also a key process in the fjord ocean dynamics (Gladish et al. 2015) and in the transport of nutrients to the surface (Hopwood et al. 2018). Furthermore, estimates of the subglacial runoff for Greenland (still very uncertain for Antarctica) are now available using a regional atmospheric model (IMOTHEP project). However, most global models (NEMO included) neglect the input of subglacial runoff because of the lack of data or because model capability to inject fresh water in depth are missing. **Category 1**

**[3.2]** It is worth noting, mostly for Greenland, that because of the resolution of the targeted configurations, most of the Fjords where the Greenlandic Marine Glaciers sit cannot be explicitly represented. The modeled circulation and freshwater inputs (glacier melt, icebergs melts, ice mélange) within a 2D (x-z) fjords therefore needs to be evaluated. If it appears that such a simple representation is not fit for purpose, such fjords will need to be parametrized. Such parametrization has been developed at Oxford based on MIT-GCM and adapted for HadGEM3 climate model via an external toolbox. The core of the parametrization is based on a three-layer box model with subglacial runoff as input. **Category 3**

**Key paper:** Gladish, C. V., Holland, D. M., Rosing-Asvid, A., Behrens, J. W., & Boje, J. (2015). Oceanic Boundary Conditions for Jakobshavn Glacier. Part I: Variability and Renewal of Ilulissat Icefjord Waters, 2001–14, *Journal of Physical Oceanography*, 45(1), 3-32. Retrieved Sep 19, 2021, from <https://journals.ametsoc.org/view/journals/phoc/45/1/jpo-d-14-0044.1.xml>

**Possible overlaps/dependencies:** Ice Shelf section of Land Ice strategy, vertical mixing.

**Expected code changes:** Inclusion of subglacial runoff could require a re-write of the runoff module to have it more generic. Parametrization of lateral melt of the marine glacier termini will fit in the ice shelf module as a new sub-module. Finally, it is difficult to say what the impact of a fjord parametrization into the NEMO code will be as a discussion on this topic has not yet happened.

## 11.7 Wider model developments:

### 11.7.1 Test case

**[4.1]** Since a couple of years, NEMO has included more and more test cases for evaluation, development and debugging purposes. The land ice / ocean interaction is not well represented in these test cases (only few capabilities are tested). To assist the development of the land ice /ocean interaction in the future we strongly encourage developers to develop and join a test case with any development made on this topic (icebergs, ice-shelf, coupling, tide water glacier ...). **Category 1** (*Resources: all developers*)

### 11.7.2 Model resolution

**[4.2]** Horizontal resolution is key in multiple points mentioned above (Fjord, ice shelf cavities). Furthermore, realistic representation of Antarctic shelf properties needs representation of the small-scale heat and mass exchange across the Antarctic continental shelf (Nakayama et al. 2014). Finally, kilometre-scale variations in melt are a key component of the complex ice–ocean interactions taking place beneath the ice shelf and so require high horizontal resolution models (Dutrieux et al. 2013). Two way nesting (AGRIF) could thus be an important tool to decrease bias related to representation of small-scale processes within ESMs, evaluate processes at play in these regions and help design the parametrization of such processes. Such applications will likely use the latest AGRIF development and test its robustness (multi-zooms within ESMs). However, AGRIF is not compatible with the lagrangian iceberg module which could limit its usage for Arctic and Antarctic applications. Finally, getting AGRIF refined areas accepted and functional as part of the ocean configuration of a full ESM will require code maturity, changes to infrastructure and time convincing a lot of different subject-experts that it's robust and scientifically worthwhile.

### 11.7.3 Ice shelf draft representation

**[4.3]** In the future, NEMO will include a new representation of the interaction between bathymetry and ocean using a penalization method. As mentioned in the conclusion of Debreu et al. (2020), this method could lead to some improvement and simplification in the representation of ice shelf cavities and coupling with an ice sheet model. However, this is a long-term feature. It is still at the test case stage and will probably not be mature enough within the next 5 years (ie not beyond the period covered by the strategy). Our position is to wait until it is mature enough so that any benefits from the penalization method will be welcome in the representation of subgrid scale bathymetry features, ice shelves and ice sheet coupling. **Category 3** (*Resources: NEMO System Team*)

### 11.7.4 Artificial intelligence

**[4.4]** Furthermore, based on the first results of machine learning based parametrizations (Rosier et al., 2022) that we are aware of, Deep Learning based parametrization of ocean / land ice interactions will likely be developed and evaluated against a more conventional parametrization. Therefore, potentially (depending of the comparison outcome) there will be a need for an interface to send data in/out between NEMO and a Deep Learning environment such as SmartSim (Partee et al., n.d.) to use such parametrization within the NEMO framework. If relevant, development of such an interface is out of the scope of the Land Ice / Ocean interaction chapter and should be addressed to the HPC working group or Machine learning working group. **Category 3** (*Resources: NEMO ST, French project AIAI submitted*)

**Possible overlaps/dependencies:** Kernel/AGRIF/HPC strategy

## 11.8 People

This working group gathers much expertise on the various domains needed to carry out the work presented in this chapter of the NEMO development strategy. Here is a description of the expertise needed and which institutes from the working group have it:

- Expertise in ice shelf cavity parametrization: IGE and Northumbria University
- Expertise in sub-ice boundary layer parametrization: BAS and Northumbria University
- Expertise in ice shelf/ocean coupling: IGE, Reading University, BAS and Northumbria University
- Expertise in climate modelling: LOCEAN and Reading University
- Expertise on vertical coordinates in ocean model: NOC
- Expertise on iceberg modelling: University of Manitoba, BAS, UCL, IGE
- Expertise on sea-ice modelling: UCL, LOCEAN
- Expertise on NEMO development: LOCEAN, NOC and IGE

As a group, we informally have been working together for about 5 years, based on good will. Since more recently, with the move toward Earth System Models and the inclusion of an ice sheet component in such models, this community of NEMO users and developers has worked together supported by various projects (EIS French project, *LICEPOD* Belgium project, H2020 projects PROTECT, TiPACCs, ESM2025). Furthermore, most of the groups involved in this working group are also collaborating together via FRISP (Forum for Research Into Ice Shelf Processes). Despite the ongoing collaborations between the various teams involved, it is worth noting as a risk for the realization of the described strategy on Land Ice / Ocean interaction that in the project mentioned above, a non-negligible part of the work does not have any identified funding and for the others, it is not explicitly mentioned in the deliverables that new development must feedback into the NEMO distribution.

This working group also recognizes that international collaboration is needed to make a breakthrough in modelling land ice / ocean interactions. Understanding the physics of the ice shelf/ocean interaction and their parametrization in ocean models is a ‘grand challenge’ that would benefit from the large international community. In order to foster such needed international collaborations, people from this working group have in the past contributed to MISOMIP (Asay-Davis et al., 2016) and ISMIP6 (Jourdain et al. 2020). Now, they contribute to build MISOMIP2 and participate to ISMIP6 extension and plan to work actively on shaping future IS- and C-MIP7 exercises.

## 11.9 Bibliography

- Ardhuin F., J. Tournadre, P. Queffeulou, F. Girard-Ardhuin, F. Collard (2011) Observation and parameterization of small icebergs: Drifting breakwaters in the southern ocean, *Ocean Modelling*, Volume 39, <https://doi.org/10.1016/j.ocemod.2011.03.004>.
- Asay-Davis, Xylar S., Stephen L. Cornford, Gaël Durand, Benjamin K. Galton-Fenzi, Rupert M. Gladstone, G. Hilmar Gudmundsson, Tore Hattermann, et al. 2016. “Experimental Design for Three Interrelated Marine Ice Sheet and Ocean Model Intercomparison Projects: MISOMIP v. 3 (MISOMIP +), ISOMIP v. 2 (ISOMIP +) and MISOMIP v. 1 (MISOMIP1).” *Geoscientific Model Development* 9 (7): 2471–97. <https://doi.org/10.5194/gmd-9-2471-2016>.
- Asay-Davis, Xylar S., Nicolas C. Jourdain, and Yoshihiro Nakayama. 2017. “Developments in Simulating and Parameterizing Interactions Between the Southern Ocean and the Antarctic Ice Sheet.” *Current Climate Change Reports* 3 (4): 316–29. <https://doi.org/10.1007/s40641-017-0071-0>.
- Beckmann, A, and H Goosse. 2003. “A Parameterization of Ice Shelf–Ocean Interaction for Climate Models.” *Ocean Modelling* 5 (2): 157–70. [https://doi.org/10.1016/S1463-5003\(02\)00019-7](https://doi.org/10.1016/S1463-5003(02)00019-7).



- Bett, David T., Paul R. Holland, Alberto C. Naveira Garabato, Adrian Jenkins, Pierre Dutrieux, Satoshi Kimura, and Andrew Fleming. 2020. "The Impact of the Amundsen Sea Freshwater Balance on Ocean Melting of the West Antarctic Ice Sheet." *Journal of Geophysical Research: Oceans* 125 (9): 1–18. <https://doi.org/10.1029/2020JC016305>.
- Boucher, Olivier, Jérôme Servonnat, Anna Lea Albright, Olivier Aumont, Yves Balkanski, Vladislav Bastrikov, Slimane Bekki, et al. 2020. "Presentation and Evaluation of the IPSL-CM6A-LR Climate Model." *Journal of Advances in Modeling Earth Systems* 12 (7). <https://doi.org/10.1029/2019MS002010>.
- Bouhier, Nicolas, Jean Tournadre, Frédérique Rémy, and Rozenn Gourves-Cousin. 2018. "Melting and Fragmentation Laws from the Evolution of Two Large Southern Ocean Icebergs Estimated from Satellite Data." *The Cryosphere* 12 (7): 2267–85. <https://doi.org/10.5194/tc-12-2267-2018>.
- Bronselaer, Ben, Michael Winton, Stephen M. Griffies, William J. Hurlin, Keith B. Rodgers, Olga V. Sergienko, Ronald J. Stouffer, and Joellen L. Russell. 2018. "Change in Future Climate Due to Antarctic Meltwater." *Nature* 564 (7734): 53–58. <https://doi.org/10.1038/s41586-018-0712-z>.
- Cowton, Tom, Donald Slater, Andrew Sole, Dan Goldberg, and Peter Nienow. 2015. "Modeling the Impact of Glacial Runoff on Fjord Circulation and Submarine Melt Rate Using a New Subgrid-Scale Parameterization for Glacial Plumes." *Journal of Geophysical Research: Oceans* 120 (2). <https://doi.org/10.1002/2014JC010324>.
- Debreu, L., N. K.R. Kevlahan, and P. Marchesiello. 2020. "Brinkman Volume Penalization for Bathymetry in Three-Dimensional Ocean Models." *Ocean Modelling*. <https://doi.org/10.1016/j.ocemod.2019.101530>.
- Dutrieux, P., D. G. Vaughan, H. F.J. Corr, A. Jenkins, P. R. Holland, I. Joughin, and A. H. Fleming. 2013. "Pine Island Glacier Ice Shelf Melt Distributed at Kilometre Scales." *Cryosphere* 7 (5). <https://doi.org/10.5194/tc-7-1543-2013>.
- England, Mark R, Till J W Wagner, and Ian Eisenman. 2020. "Modeling the Breakup of Tabular Icebergs." *Science Advances* 6 (51): eabd1273. <https://doi.org/10.1126/sciadv.abd1273>.
- Erokhina, Olga (2020) A new Eulerian iceberg module for climate studies: Formulation and application to the investigation of the sensitivity of the AMOC to iceberg calving.
- Favier, L., N.C. Jourdain, A. Jenkins, N. Merino, G. Durand, O. Gagliardini, F. Gillet-Chaulet, and P. Mathiot. 2019. "Assessment of Sub-Shelf Melting Parameterisations Using the Ocean-Ice-Sheet Coupled Model NEMO(v3.6)-Elmer/Ice(v8.3)." *Geoscientific Model Development* 12 (6). <https://doi.org/10.5194/gmd-12-2255-2019>.
- FitzMaurice, A., C. Cenedese, and F. Straneo. 2017. "Nonlinear Response of Iceberg Side Melting to Ocean Currents." *Geophysical Research Letters* 44 (11). <https://doi.org/10.1002/2017GL073585>.
- Gladish, Carl V., David M. Holland, Aqqalu Rosing-Asvid, Jane W. Behrens, and Jesper Boje. 2015. "Oceanic Boundary Conditions for Jakobshavn Glacier. Part I: Variability and Renewal of Ilulissat Icefjord Waters, 2001–14." *Journal of Physical Oceanography* 45 (1): 3–32. <https://doi.org/10.1175/JPO-D-14-0044.1>.
- Gladstone, Rupert M., Grant R. Bigg, and Keith W. Nicholls. 2001. "Iceberg Trajectory Modeling and Meltwater Injection in the Southern Ocean." *Journal of Geophysical Research: Oceans* 106 (C9):

- 19903–15. <https://doi.org/10.1029/2000JC000347>.
- Goldberg, D. N., C. M. Little, O. V. Sergienko, A. Gnanadesikan, R. Hallberg, and M. Oppenheimer. 2012. "Investigation of Land Ice-Ocean Interaction with a Fully Coupled Ice-Ocean Model: 1. Model Description and Behavior." *Journal of Geophysical Research: Earth Surface* 117 (2): 1–16. <https://doi.org/10.1029/2011JF002246>.
- Gwyther, David E., Benjamin K. Galton-Fenzi, Michael S. Dinniman, Jason L. Roberts, and John R. Hunter. 2015. "The Effect of Basal Friction on Melting and Freezing in Ice Shelf–Ocean Models." *Ocean Modelling* 95 (November). <https://doi.org/10.1016/j.ocemod.2015.09.004>.
- Hausmann, U., J.-B. Sallée, N. C. Jourdain, P. Mathiot, C. Rousset, G. Madec, J. Deshayes, and T. Hattermann. 2020. "The Role of Tides in Ocean-Ice Shelf Interactions in the Southwestern Weddell Sea." *Journal of Geophysical Research: Oceans* 125 (6). <https://doi.org/10.1029/2019JC015847>.
- Hopwood, M. J., D. Carroll, T. J. Browning, L. Meire, J. Mortensen, S. Krisch, and E. P. Achterberg. 2018. "Non-Linear Response of Summertime Marine Productivity to Increased Meltwater Discharge around Greenland." *Nature Communications* 9 (1): 3256. <https://doi.org/10.1038/s41467-018-05488-8>.
- Huot, Pierre-Vincent, Thierry Fichefet, Nicolas C. Jourdain, Pierre Mathiot, Clément Rousset, Christoph Kittel, and Xavier Fettweis. 2021. "Influence of Ocean Tides and Ice Shelves on Ocean–Ice Interactions and Dense Shelf Water Formation in the D’Urville Sea, Antarctica." *Ocean Modelling* 162 (June): 101794. <https://doi.org/10.1016/j.ocemod.2021.101794>.
- Hsiung C.C. and A.F. Aboul-Azm, (1982) Iceberg drift affected by wave action, *Ocean Engineering*, Volume 9, Issue 5, 1982, Pages 433-439, ISSN 0029-8018, [https://doi.org/10.1016/0029-8018\(82\)90035-X](https://doi.org/10.1016/0029-8018(82)90035-X).
- Jenkins, Adrian. 2011. "Convection-Driven Melting near the Grounding Lines of Ice Shelves and Tidewater Glaciers." *Journal of Physical Oceanography* 41 (12): 2279–94. <https://doi.org/10.1175/JPO-D-11-03.1>.
- . 2016. "A Simple Model of the Ice Shelf–Ocean Boundary Layer and Current." *Journal of Physical Oceanography* 46 (6): 1785–1803. <https://doi.org/10.1175/JPO-D-15-0194.1>.
- Jordan, James R., Paul R. Holland, Dan Goldberg, Kate Snow, Robert Arthern, Jean-Michel Campin, Patrick Heimbach, and Adrian Jenkins. 2018. "Ocean-Forced Ice-Shelf Thinning in a Synchronously Coupled Ice-Ocean Model." *Journal of Geophysical Research: Oceans* 123 (2): 864–82. <https://doi.org/10.1002/2017JC013251>.
- Jourdain, Nicolas C., Xylar Asay-Davis, Tore Hattermann, Fiammetta Straneo, Hélène Seroussi, Christopher M. Little, and Sophie Nowicki. 2020. "A Protocol for Calculating Basal Melt Rates in the ISMIP6 Antarctic Ice Sheet Projections." *The Cryosphere* 14 (9): 3111–34. <https://doi.org/10.5194/tc-14-3111-2020>.
- Jourdain, Nicolas C., Pierre Mathiot, Nacho Merino, Gaël Durand, Julien Le Sommer, Paul Spence, Pierre Dutrioux, and Gurvan Madec. 2017. "Ocean Circulation and Sea-Ice Thinning Induced by Melting Ice Shelves in the Amundsen Sea." *Journal of Geophysical Research: Oceans* 122 (3): 2550–73. <https://doi.org/10.1002/2016JC012509>.
- Kubat, I. and M. Sayed, S. B. Savage and T. Carrieres, (2005) An Operational Model of Iceberg Drift,

- Lemieux, Jean-François, Frédéric Dupont, Philippe Blain, François Roy, Gregory C. Smith, and Gregory M. Flato. 2016. "Improving the Simulation of Landfast Ice by Combining Tensile Strength and a Parameterization for Grounded Ridges." *Journal of Geophysical Research: Oceans* 121 (10): 7354–68. <https://doi.org/10.1002/2016JC012006>.
- Lichey, Christoph, and Hartmut H. Hellmer. 2001. "Modeling Giant-Iceberg Drift under the Influence of Sea Ice in the Weddell Sea, Antarctica." *Journal of Glaciology* 47 (158): 452–60. <https://doi.org/10.3189/172756501781832133>.
- Marsh, R., V. O. Ivchenko, N. Skliris, S. Alderson, G. R. Bigg, G. Madec, A. T. Blaker, et al. 2015. "NEMO–ICB (v1.0): Interactive Icebergs in the NEMO Ocean Model Globally Configured at Eddy-Permitting Resolution." *Geoscientific Model Development* 8 (5): 1547–62. <https://doi.org/10.5194/gmd-8-1547-2015>.
- Marson, Juliana M., Paul G. Myers, Xianmin Hu, and Julien Le Sommer. 2018. "Using Vertically Integrated Ocean Fields to Characterize Greenland Icebergs' Distribution and Lifetime." *Geophysical Research Letters* 45 (9): 4208–17. <https://doi.org/10.1029/2018GL077676>.
- Massom, R. A., P.T. Harris, Kelvin J. Michael, and M.J. Potter. 1998. "The Distribution and Formative Processes of Latent-Heat Polynyas in East Antarctica." *Annals of Glaciology* 27 (January): 420–26. <https://doi.org/10.3189/1998AoG27-1-420-426>.
- Massom, R. A., K. L. Hill, V. I. Lytle, A. P. Worby, M.J. Paget, and I. Allison. 2001. "Effects of Regional Fast-Ice and Iceberg Distributions on the Behaviour of the Mertz Glacier Polynya, East Antarctica." *Annals of Glaciology* 33 (September). <https://doi.org/10.3189/172756401781818518>.
- Mathiot, Pierre, Adrian Jenkins, Christopher Harris, and Gervan Madec. 2017. "Explicit Representation and Parametrised Impacts of under Ice Shelf Seas in the &Lt;&Gt;&Z&Lt;/I&Gt;&Lt;Sup&Gt;\*&Lt;/Sup&Gt; Coordinate Ocean Model NEMO 3.6." *Geoscientific Model Development* 10 (7): 2849–74. <https://doi.org/10.5194/gmd-10-2849-2017>.
- Merino, Nacho, Julien Le Sommer, Gael Durand, Nicolas C. Jourdain, Gervan Madec, Pierre Mathiot, and Jean Tournadre. 2016. "Antarctic Icebergs Melt over the Southern Ocean : Climatology and Impact on Sea Ice." *Ocean Modelling*. <https://doi.org/10.1016/j.ocemod.2016.05.001>.
- Nakayama, Y., R. Timmermann, M. Schröder, and H. H. Hellmer. 2014. "On the Difficulty of Modeling Circumpolar Deep Water Intrusions onto the Amundsen Sea Continental Shelf." *Ocean Modelling* 84: 26–34. <https://doi.org/10.1016/j.ocemod.2014.09.007>.
- Nihashi, Sohey, and Kay I. Ohshima. 2015. "Circumpolar Mapping of Antarctic Coastal Polynyas and Landfast Sea Ice: Relationship and Variability." *Journal of Climate* 28 (9): 3650–70. <https://doi.org/10.1175/JCLI-D-14-00369.1>.
- Olason, Einar. 2016. "A Dynamical Model of Kara Sea Land-fast Ice." *Journal of Geophysical Research: Oceans* 121 (5): 3141–58. <https://doi.org/10.1002/2016JC011638>.
- Oppenheimer, M., et al. 2019. *Chapter 4: Sea Level Rise and Implications for Low Lying Islands, Coasts and Communities*. Edited by UK Cambridge University Press, Cambridge. IPCC speci. <https://www.ipcc.ch/srocc/chapter/chapter-4-sea-level-rise-and-implications-for-low-lying->

islands-coasts-and-communities/.

- Partee, Sam, Matthew Ellis, Alessandro Rigazzi, Scott Bachma, Marques Gustavo, Shao Andrew, and Benjamin Robbins. n.d. "Using Machine Learning at Scale in HPC Simulations with SmartSim: An Application to Ocean Climate Modeling." <https://arxiv.org/pdf/2104.09355v1.pdf>.
- Pelletier, C, T Fichefet, H Goosse, K Haubner, S Helsen, P.-V. Huot, C Kittel, et al. 2021. "PARASO, a Circum-Antarctic Fully-Coupled Ice-Sheet - Ocean - Sea-Ice - Atmosphere - Land Model Involving f.ETISH1.7, NEMO3.6, LIM3.6, COSMO5.0 and CLM4.5." *Geoscientific Model Development Discussions* 2021: 1–59. <https://doi.org/10.5194/gmd-2021-315>.
- Rackow, Thomas, Christine Wesche, Ralph Timmermann, Hartmut H. Hellmer, Stephan Juricke, and Thomas Jung. 2017. "A Simulation of Small to Giant Antarctic Iceberg Evolution: Differential Impact on Climatology Estimates." *Journal of Geophysical Research: Oceans* 122 (4): 3170–90. <https://doi.org/10.1002/2016JC012513>.
- Rignot, E., Y. Xu, D. Menemenlis, J. Mouginot, B. Scheuchl, X. Li, M. Morlighem, et al. 2016. "Modeling of Ocean-Induced Ice Melt Rates of Five West Greenland Glaciers over the Past Two Decades." *Geophysical Research Letters* 43 (12): 6374–82. <https://doi.org/10.1002/2016GL068784>.
- Rosier, S. H. R., Bull, C. Y. S., and Gudmundsson, G. H.: Predicting ocean-induced ice-shelf melt rates using a machine learning image segmentation approach, *The Cryosphere Discuss.* [preprint], <https://doi.org/10.5194/tc-2021-396>, in review, 2022.
- Schloesser, Fabian, Tobias Friedrich, Axel Timmermann, Robert M. DeConto, and David Pollard. 2019. "Antarctic Iceberg Impacts on Future Southern Hemisphere Climate." *Nature Climate Change* 9 (9): 672–77. <https://doi.org/10.1038/s41558-019-0546-1>.
- Slater, Donald A., Dan N. Goldberg, Peter W. Nienow, and Tom R. Cowton. 2016. "Scalings for Submarine Melting at Tidewater Glaciers from Buoyant Plume Theory." *Journal of Physical Oceanography* 46 (6): 1839–55. <https://doi.org/10.1175/JPO-D-15-0132.1>.
- Smith, Robin S., Pierre Mathiot, Antony Siahann, Victoria Lee, Stephen L. Cornford, Jonathan M. Gregory, Antony J. Payne, et al. 2021a. "Coupling the U.K. Earth System Model to Dynamic Models of the Greenland and Antarctic Ice Sheets." *Journal of Advances in Modeling Earth Systems*, September. <https://doi.org/10.1029/2021MS002520>.
- . 2021b. "Coupling the U.K. Earth System Model to Dynamic Models of the Greenland and Antarctic Ice Sheets." *Journal of Advances in Modeling Earth Systems*, September. <https://doi.org/10.1029/2021MS002520>.
- Stern, A. A., A. Adcroft, and O. Sergienko. 2016. "The Effects of Antarctic Iceberg Calving-Size Distribution in a Global Climate Model." *Journal of Geophysical Research: Oceans* 121 (8): 5773–88. <https://doi.org/10.1002/2016JC011835>.
- Storkey, David, Adam T. Blaker, Pierre Mathiot, Alex Megann, Yevgeny Aksenov, Edward W. Blockley, Daley Calvert, et al. 2018. "UK Global Ocean GO6 and GO7: A Traceable Hierarchy of Model Resolutions." *Geoscientific Model Development* 11 (8): 3187–3213. <https://doi.org/10.5194/gmd-11-3187-2018>.
- Tournadre, J., N. Bouhier, F. Girard-Ardhuin, and F. Rémy. 2016. "Antarctic Icebergs Distributions 1992–2014." *Journal of Geophysical Research: Oceans* 121 (1): 327–49. <https://doi.org/10.1002/2015JC011178>.

- Vaňková, Irena, and David M. Holland. 2017. "A Model of Icebergs and Sea Ice in a Joint Continuum Framework." *Journal of Geophysical Research: Oceans* 122 (11): 9110–25. <https://doi.org/10.1002/2017JC013012>.
- Wei, Wei, Donald D. Blankenship, Jamin S. Greenbaum, Noel Gourmelen, Christine F. Dow, Thomas G. Richter, Chad A. Greene, et al. 2020. "Getz Ice Shelf Melt Enhanced by Freshwater Discharge from beneath the West Antarctic Ice Sheet." *The Cryosphere* 14 (4). <https://doi.org/10.5194/tc-14-1399-2020>.
- Yulmetov R. and Freeman R. (2019) Machine learning for tactical iceberg drift forecasting, Proceedings of the 25th International Conference on Port and Ocean Engineering under Arctic Conditions, June 9-13, 2019, Delft, The Netherlands

## 12 Tides

### 12.1 Executive Summary

Tides contribute to the meridional heat transport (by providing some of the mechanical mixing necessary to the return flow) and are important for the “shelf-enabling” driver aiming at kilometeric resolution near the coast. The Tidal Working Group (WG) was formed in mid-2020 with members from the UK, France and Canada and started regular monthly meetings in November 2020 with the goal of discussing current issues, good practices and suggesting potential improvements in NEMO with a particular focus on representing the tides and their energy conversion to baroclinic modes. We have targeted so far two implementations that are feasible within two years: Internal wave drag parameterization (that improves the general tides by parameterizing some of the unrepresented barotropic-baroclinic conversion) and proposing an alternative implementation to the Self-Attraction and Loading (SAL) potential, which is simpler than the existing one but which can allow for a feedback of the OGCM own sea surface height. Other recommendations are made and some discussions of the impact of the numerics on tides are discussed. Given the recent formation of the WG, we have not drawn any more firm roadmap.

### 12.2 Introduction

Munk et Wunsch (1998) stress the importance of the mechanical work done by tides --i.e., mixing mostly along mid-ocean ridges-- to support the meridional heat transport. NEMO already includes a module for adding tidal mixing without explicit representation of the tides (Laverne et al., 2016) which aims precisely at representing this effect. Another approach is to improve the explicit representation of tides with the benefit of resolving the associated processes (tidal rectification and other interactions with the rest of the dynamics), the only drawback being a potential reduction of the allowable timestep. These two approaches are debatable and will likely continue to be.

However, with increased resolution (starting somewhat arbitrarily from 1/10th deg), global OCGMs are successfully able to represent explicitly barotropic tides, some form of barotropic to baroclinic energy conversion and their interactions with the general circulation. Their introduction helps to represent an important part of the energy spectrum and is more consistent with high temporal frequency forcings which themselves include atmospheric tides. Moreover, the “shelf-enabling” NEMO driver where most centers aim at kilometeric effective resolution cannot be achieved without accurate representation of the tides. We will therefore concentrate in the following sections on the second approach, i.e., the explicit representation of tides.

In a numerical ocean model, surface tides are the barotropic response to the astronomical forcing [and open boundary forcing if the domain is limited], and dissipated by bottom friction or converted to baroclinic motions along steep topography and can interact with the general circulation. Tides were first introduced in NEMO 2.3 and then improved substantially in 3.6 with a better time-splitting scheme, the addition of astronomical, self-attraction-and-loading potentials and better boundary conditions available in the module BDY.

The Tidal Working Group (WG) was formed in mid-2020 with members from the UK, France and Canada and started regular monthly meetings in November 2020 with the goal of discussing current issues, good practices and suggesting potential improvements in NEMO (see <https://forge.nemo-ocean.eu/wgs/tides/home> for members’ list and meeting summaries). This WG has an ambiguous position as it is a user of other NEMO WG’s innovations. Some actions are simply to apply and test

them (and recommend the better ones). In fact, tidal applications are quite interesting benchmarks as they are in general very demanding on the stability of the overall model.

Below is a summary of the recommendations that could be relevant to the NEMO Development Strategy. Section 3 describes code implementations that are feasible in a two year time frame, while Section 4 enumerates other potential or longer-term developments and applications.

## 12.3 Two-year implementation plan

### 12.3.1 Internal wave drag parameterization

Arbic et al. (2010) note that even a 1/12th degree resolution tide-enabled OGCM require some additional internal dissipation of the tides for obtaining the best accuracy, which cannot be provided by bottom friction alone, unless the drag coefficient is increased to unrealistic values. In order to represent the dissipation of internal waves and release of the energy into mean potential energy, several groups tested different parametrizations. The most likely candidate is that of Jaynes and St. Laurent (2001) which was adopted in NEMO by Kodaira et al. (2016) and Wang et al. (2021). This is especially useful in barotropic simulations (as originally intended). We however cannot rule out that it can improve the accuracy of the tides in 3D baroclinic simulations as OCGMs can only represent a few modes (at best) of internal waves and therefore some dissipation is required for the higher modes or the conversion of energy to other un-represented physics such as non-hydrostatic processes. There are still some unknowns on how to implement the internal drag in NEMO as some filtering is required (24 or 25h) that we hope to clarify in the coming year. We also note that the parametrized dissipation of energy should also translate into an additional mixing of tracers (i.e., a source of turbulent kinetic energy). Jérôme Chanut started the implementation work with feedback from other group members.

### 12.3.2 Self-Attraction and Loading (SAL)

The SAL term can contribute to approximately 5-10% of the tidal water level. Therefore, its accurate representation is paramount to a precise tidal simulation. It is implemented as a potential used in the surface pressure gradient (SPG). As presented by many authors (e.g. Ray, 1998; Stepanov and Hughes, 2004), it requires the calculation of a convoluted and therefore rather expensive expression. Some approximations exist. The simplest one is to assume that the SAL term is a linear response of the surface elevation,  $\beta\eta$  [where  $\eta$  is sea surface height and  $\beta$  is a scalar coefficient; e.g., Arbic et al., 2010]. Another and more precise one is to take advantage of the very accurate tidal atlases (i.e., tide-assimilative solutions), from which the SAL term can be diagnosed as a decomposition by spatially varying tidal constituents. In the latter case, some missing contributions would be that of the mean circulation ---affecting mainly itself, and the absence of some minor tidal constituents from the SAL atlas. A simple approach would be to allow for a feedback of the OGCM own sea surface height (tidally-filtered if the tidal constituent SAL method is already used; or in full as in Kodaira et al. (2016) and Wang et al., 2021, if not) into the SPG via a user-defined spatially varying feedback coefficient  $\beta$ , as for instance derived by Stepanov and Hughes (2004). Note though that all

these options should be tested for their accuracy before a more precise work plan can be developed. Chris Wilson and Jeff Polton will likely be working on this with feedback from other group members.

## 12.4 Other Recommendations

### 12.4.1 Numerics

Numerical representation of the different contributions and interactions of tides (external ,i.e., depth-averaged, or internal , i.e., 3D, modes) with the rest of the dynamics is a field overlapping with other WGs but we think important to list of few items worth pursuing in the future:

- Vertical coordinate development will continue to be paramount: transition between z-levels to sigma-like levels closer to shore with the idea of better representing the bottom dissipative boundary layer and minimizing the blocking of cross-shelf transport by staircase topography) and related numerics (pressure gradient errors, high-order and monotonic advection schemes)
- Testing the new RK3 time-stepping scheme in the context of tidal application will be critical
- Internal/external interactions during time-splitting: after the external mode sub-cycling, we only correct the 3D velocity as to ensure the volume conservation between the 3D internal and external modes. However, some terms may have likely diverged, such as bottom stress. Moreover, Demange et al. (2019) show that the present external mode sub-cycling is not fully consistent with the barotropic mode which reduces the stability of the time-splitting methodology. Their correction requires the use of an eigenvalue decomposition done on the fly, therefore adding a dependency to an external library while building the executable.
- Vertical turbulence is assumed to be locally generated and dissipated in NEMO. As the horizontal resolution increases, this assumption will no longer be valid. Some models take advection into account, although likely none guarantee conservation, since the turbulence variables (TKE and possibly other moments) are discretized at the W-level.
- Moreover, as mentioned in Section 2.1, the dissipation of internal waves should appear as a source in the turbulent kinetic energy equation. Comparison with in-situ microstructure data is a key aspect of this tuning.
- Non-hydrostatic wave generation and wet-and-drying capabilities in very high-resolution configurations will be more and more important.
- A fair fraction of tidal and coastal/nearshore modelling is conducted with unstructured grid models. While numerical modes remain a serious issue, improvement of the latter, flexibility in representing complex coastlines and variable resolution make them serious contenders for future coastal and more generally ocean applications. From a NEMO perspective, it is critical to improve the multiple nesting capabilities in order to keep an edge with more flexible boundary treatment/vertical coordinate between the parent and child grids and more flexible computing distributions... or envision some convergence with the unstructured grid methodology.
- Harmonic analysis online in NEMO should be available even for long period tides. This implies coding restarts output to be managed in diaharm.F90



## 12.4.2 Definition

Historically the motivation driving this working group has come from the desire to better represent subgridscale processes in global models. However, as NEMO increasingly becomes used for shallow water sea level applications, such as deterministic forecasting of sea level over a range of possible timescales, explicit calculation of tidal processes are increasingly critical.

The term “tide” is loosely defined with some contributions from the atmosphere, or long-period tides that are strongly aliased with the mesoscale dynamics or the seasonality of the ocean. Therefore, moving forward, some applications will require a more precise definition of “tides”.

## 12.4.3 Applications

In consideration of the NEMO end users and downstream applications of NEMO tides a couple of items arise that could add value to tidal outputs:

- Data assimilation (or corrections methods) OCGMs were not at their inception meant to resolve tides although with improved resolution they are now getting close to represent tides reasonable well in the deep ocean but with some difficulty in shelf areas and around Antarctica (where tidal resonance and ice-shelf interactions are an issue). To remedy the situation, some corrections are required and possible using an external source such as a tidal atlas (e.g., FES or OSU). Two member groups have experimented with their own approaches, one consisting with spectrally nudging the tides in the momentum equations (Wang et al., 2021) and the second being a correction method adding tides as source terms in a coarse OGCM. Both approaches offered promising results.
- Varying bottom roughness (why having one value for the whole ocean when the morphology of the sea bottom is known to vary?)
- Tidal harmonic outputs can be improved by postprocessing of an ensemble run. E.g. Byrne et al (2021).
- Some improvements are possible in astronomical tides (or “equilibrium tide” as in Ray, 1998). One is to explore the use of the full astronomical potential (via only ephemeris; see Ray and Cartwright, 2007, as an example or a recent application in Logemann et al. 2018) instead of the more typical decomposition by tidal components.
- The SAL full computation is very expensive, as already mentioned above. NOC is nonetheless investigating methods to speed up an on-the-fly Green’s function.
- Energetics: Normal mode decomposition is important to follow the energy flowing between the different dynamical components of the ocean. However, it was noted that an offline diagnostic was difficult. One suggestion is to either output the 3D fields at high frequency (including the pressure term which is not an option at this time) or doing the diagnostic online which implies relying on an additional library for the eigenvector decomposition (Nurgoho et al. 2017, chap5 thesis).

## 12.5 References

- Arbic, B. K., Wallcraft, A. J., & Metzger, E. J. (2010). Concurrent simulation of the eddying general circulation and tides in a global ocean model. *Ocean Modelling*, 32(3-4), 175-187. <https://doi.org/10.1016/j.ocemod.2010.01.007>
- Byrne, D., Polton, J., & Bell, C. (2021). Creation of a global tide analysis dataset: Application of NEMO and an offline objective analysis scheme. *Journal of Operational Oceanography*, 1-14. doi:10.1080/1755876X.2021.2000249
- De Lavergne, C., Madec, G., Le Sommer, J., Nurser, A. G., & Garabato, A. C. N. (2016). The impact of a variable mixing efficiency on the abyssal overturning. *Journal of Physical Oceanography*, 46(2), 663-681. <https://doi.org/10.1175/JPO-D-14-0259.1>
- Demange, J., Debreu, L., Marchesiello, P., Lemarié, F., Blayo, E., & Eldred, C. (2019). Stability analysis of split-explicit free surface ocean models: implication of the depth-independent barotropic mode approximation. *Journal of Computational Physics*, 398, 108875. <https://doi.org/10.1016/j.jcp.2019.108875>
- Jayne, S. R., & St. Laurent, L. C. (2001). Parameterizing tidal dissipation over rough topography. *Geophysical Research Letters*, 28(5), 811-814.
- Kodaira, T., Thompson, K. R., & Bernier, N. B. (2016). Prediction of M2 tidal surface currents by a global baroclinic ocean model and evaluation using observed drifter trajectories. *Journal of Geophysical Research: Oceans*, 121(8), 6159-6183.
- Logemann, K., Linardakis, L., Korn, P., & Schrum, C. (2021). Global tide simulations with ICON-O: testing the model performance on highly irregular meshes. *Ocean Dynamics*, 71(1), 43-57.
- Munk, W. (1997). Once again: once again—tidal friction. *Progress in Oceanography*, 40(1-4), 7-35. <https://doi.org/10.1007/s10236-020-01428-7>
- Munk, W., & Wunsch, C. (1998). Abyssal recipes II: Energetics of tidal and wind mixing. *Deep Sea Research Part I: Oceanographic Research Papers*, 45(12), 1977-2010.
- Nugroho, 2017, PhD, Tides in an OGCM in the Indonesian seas, Université Paul Sabatier
- Ray, R. D. (1998). Ocean self-attraction and loading in numerical tidal models. *Marine Geodesy*, 21(3), 181-192. <https://doi.org/10.1080/01490419809388134>
- Ray, R. D., & Cartwright, D. E. (2007). Times of peak astronomical tides. *Geophysical Journal International*, 168(3), 999-1004. <https://doi.org/10.1111/j.1365-246X.2006.03293.x>
- Stepanov, V. N., & Hughes, C. W. (2004). Parameterization of ocean self-attraction and loading in numerical models of the ocean circulation. *Journal of Geophysical Research: Oceans*, 109(C3). <https://doi.org/10.1029/2003JC002034>
- Wang, P., Bernier, N. B., Thompson, K. R., & Kodaira, T. (2021). Evaluation of a global total water level model in the presence of radiational S2 tide. *Ocean Modelling*, 168, 101893. <https://doi.org/10.1016/j.ocemod.2021.101893>

## 13 Marine Biogeochemistry and TOP Interface

Authors: T. Lovato, O. Aumont, G. Lessin

### 13.1 Executive Summary

Marine biogeochemistry within the NEMO framework is addressed through a built-in biogeochemical component (PISCES, Aumont et al., 2015) and the TOP (Tracers in the Ocean Paradigm) module that provides a seamless, hardwired coupling interface with non-legacy marine ecosystem models such as MEDUSA (Yool et al., 2013), BFM (Lovato et al., 2020), ERSEM (Skakala et al., 2020) and BAHMBI (Palazov et al., 2021).

In this chapter, the foreseen evolution of NEMO components dealing directly and indirectly with marine biogeochemistry are outlined by considering the need to sustain the orthogonality between physical processes and oceanic tracers' dynamics (14.2), to extend the TOP workflow consolidated in previous years (14.3), and to foster the code readiness to handle future evolutions in marine biogeochemical models (14.4). These three main themes will be coordinated by the TOP working group representative and discussed/developed along with external experts from the European marine ecosystem modelling community.

As the NEMO modelling system is a multifarious space, additional emerging issues relating marine biogeochemistry and the different components of the framework are detailed within specific chapters. Here, only highlights of key cross chapter synergies will be provided to trace their relation with the main development themes of marine biogeochemistry and TOP interface.

### 13.2 Orthogonality between physical and biogeochemical components

Physical processes represent an essential driver in shaping the spatio-temporal distribution of living and non-living oceanic properties and the improvement toward both more accurate and up-to-date representation will benefit the simulation of marine biogeochemistry.

The following processes are foreseen to enable for major orthogonality between physical and biogeochemical components:

- a. Enhance **particle dynamics in the water column** by allowing the selection of numerical schemes for vertical sinking with an increasing degree of accuracy (e.g. by using a technical design similar to the physical advection). A certain degree of flexibility is highly desirable within the modelling system to enable the balance between computational costs and accuracy (see e.g., Aumont et al., 2015). Moreover, this physical process directly applies to a variety of marine 'particles', like planktonic organisms and particulate organic/inorganic matter.
- b. Improve **optical properties in the water column** by considering the potential contribution of remote sensing data in providing new insights on the role of coloured dissolved substances and particles. This will involve the revision of current schemes to ingest more complex definitions of the light spectrum and it will provide a more articulated representation of coastal zone dynamics (e.g. for CMEMS end-users).

- c. Complement **seawater temperature and salinity definition** obtained from the two main formulations of the equation of state in NEMO: EOS80 provides the potential temperature and practical salinity, and TEOS10 the conservative temperature and absolute salinity. As a wide number of biogeochemical parametrizations derive from experimental evidence, it would be useful to extend the biogeochemistry interface with at least the use of in-situ temperature fields. Further enhancements suggested as 'best practices' in Orr and Epitalon (2015) and recent literature should also be considered.

#### Synergy with **HPC Chapter**

A close collaboration with the HPC working group will be necessary to achieve a more effective and less computationally expensive solution to speed up marine quantities transport, namely by improving the numerical performance of advection and diffusion schemes inherited from the physical core.

### 13.3 Extend and consolidate TOP workflow

The TOP interface was soundly revised in the previous five years of development, such that the workflow modularity was largely consolidated and a number of handlers were created to advance in the integration with non-legacy biogeochemical models.

The following issues should be tackled to maintain the TOP workflow and further expand it:

- a. **Interface technical developments** will be carefully evaluated to ensure a contained maintenance for the coupling of built-in and non-legacy biogeochemical models over the long term. However, new elements are still needed to further increase the interface modularity, such as the user-defined handling of restarts and outputs (namely in MY\_TRC sub-module) and the possibility to use also three-dimensional forcing, e.g., to reproduce the release of tracer quantities within the model domain beyond the system boundaries.
- b. **TOP workflow resilience** will benefit from the setup of a dedicated test case to verify the consistency of all data handlers and processes inherited from the NEMO core. This test case will likely be a new idealised configuration to evaluate the correct simulation of passive tracers' dynamics due to physical schemes and prescribed surface, coastal, and lateral boundary conditions. In addition, this simplified configuration will provide a useful example of the generalised coupling interface to new users.

#### Synergy with **Tools, V&V Chapter**

The foreseen development of a dedicated TOP test case overlaps with the main activities of the V&V working group and it represents a useful interaction to increase the reliability of the code and support its long-term robustness.

### 13.4 Readiness for future biogeochemical complexity

The overall structure of the TOP interface is founded on the support of the marine pelagic component (arrays for state variables, time integration, etc.) and only a few elements are available to handle additional dynamical components. Nowadays biogeochemical models are increasingly addressing ecological processes occurring in other marine compartments (see e.g., Vancoppenolle & Tedesco, 2017; Lessin et al., 2018) and the following actions should be taken in the medium term:

- a. **Infrastructure for marine sea-ice and benthic components** has to be designed in a more generalised, compatible framework as the existing one for the pelagic compartment. These elements should be integrated within the TOP interface to enable a coherent structure of the coupling framework, by designing dedicated sub-modules to provide access for shared memory arrays, initial and boundary conditions, and data saving. In addition, relevant physical processes should be inherited from the general NEMO framework (e.g. from SI3) and passed to the sub-module(s).
- b. **Interfaces at the boundaries with the pelagic compartment** need to be included in the development of the new dynamical components. This would translate into the identification of suitable parameterizations and schemes to resolve the exchanges of biogeochemical quantities (e.g. inorganic nutrients, organic matter) at the seaice- and benthic-pelagic interfaces.

#### Synergy with **Sea-ice Chapter**

The proposed development of a specific TOP interface to handle biogeochemical quantities within the marine sea-ice would benefit from the interaction with the SEA-ICE working group not only to design the interface, but also to identify key physical processes interacting with the sympagic ecosystem.

### 13.5 References

- Aumont, O., Éthé, C., Tagliabue, A., Bopp, L., & Gehlen, M. (2015). PISCES-v2: an ocean biogeochemical model for carbon and ecosystem studies. *Geoscientific Model Development*, 8(8), 2465-2513.
- Lessin, G., Artioli, Y., Almroth-Rosell, E., Blackford, J. C., Dale, A. W., Glud, R. N., ... & Yakushev, E. (2018). Modelling marine sediment biogeochemistry: Current knowledge gaps, challenges, and some methodological advice for advancement. *Frontiers in Marine Science*, 5, 19.
- Lovato T., Vichi M., Butenschön M. (2020). Coupling BFM with Ocean models: the NEMO model V3.6 (Nucleus for the European Modelling of the Ocean). BFM Report series N. 2, Release 1.1, June 2020, Bologna, Italy, <http://bfm-community.eu>, pp. 31
- Orr, J. C., & Epitalon, J. M. (2015). Improved routines to model the ocean carbonate system: mocsy 2.0. *Geoscientific Model Development*, 8(3), 485-499.
- Palazov, A., Ciliberti, S. A., Lecci, R., Gregoire, M., Staneva, J., Peneva, E., ... & Agostini, P. (2021, April). The BS-MFC service and system evolutions within Copernicus Marine Service. In EGU General Assembly Conference Abstracts (pp. EGU21-5435).

Skakala, J., Bruggeman, J., Brewin, R. J., Ford, D. A., & Ciavatta, S. (2020). Improved representation of underwater light field and its impact on ecosystem dynamics: A study in the North Sea. *Journal of Geophysical Research: Oceans*, 125(7), e2020JC016122.

Vancoppenolle, M., and Tedesco, L. (2017). *Numerical Models of Sea Ice Biogeochemistry*. Hoboken, NJ: John Wiley & Sons, 492–515. doi: 10.1002/9781118778371.ch20

Yool, A., Popova, E. E., & Anderson, T. R. (2013). MEDUSA-2.0: an intermediate complexity biogeochemical model of the marine carbon cycle for climate change and ocean acidification studies. *Geoscientific Model Development*, 6(5), 1767-1811.

## 14 Two-way nesting capability (AGRIF)

WGLs: Sébastien Masson, Jérôme Chanut, Adam Blaker

### 14.1 Executive Summary

*The majority of development planned for AGRIF during this period concerns addressing known weaknesses in the present implementation (external mode coupling, nest with a complex geometry, conservation, and preprocessing) and improving accessibility and documentation with a view to strengthening the user base. The three key developments will be the use of AGRIF for online coarsening, increasing flexibility in the definition of refined grids by moving away from purely rectangular refined regions, and improvements in computational performance. Each of these developments will deliver computational cost benefits, either through runtime performance (time to solution) or through smaller computational resource requirements (carbon footprint).*

### 14.2 Team involved

The core team leading the AGRIF development will be:

- Sébastien Masson: leading the technical development aspects
- Adam Blaker: leading development of documentation and user base, meeting coordination
- Jérôme Chanut: technical development, plus support and assistance as former AGRIF development lead
- Laurent Debreu: AGRIF library development, support.
- Rachid Benchila: technical development.
- Gaston Irrman: Performance.
- Others: TBA

The team will meet quarterly to discuss activities, and will review progress and reassess priorities as required. Most coordination should be possible virtually. When and where goals align with active user groups they will be engaged to ensure best use of resources.

### 14.3 Summary from previous strategic plan

Key aims were:

- Development of AGRIF to efficiently and consistently tackle use of high resolution meshes
- Consolidating AGRIF integration within NEMO, unlocking potential for shelf seas and climate processes
- Realising the potential of AGRIF to improve representation of key processes (e.g. dense overflows)

Key issues identified were:

- Inability to overlap nested grids (previously had to have a single parent level and could not exchange with sibling grids)
- Inability to assign separate MPP resources for each refinement level (resulted in difficulties with load balancing). Also lacks capability to handle mixed MPI-OpenMP parallelism.
- Imperfect 1:1 nesting due to issues with: baroclinic vs barotropic timestep coupling, capacity to handle higher order numerical schemes, global domain boundaries (cyclic, north fold)

- Conservation issues - hard to implement corrective methods with leapfrog time stepping scheme
- Usability/accessibility problems arising from separate pre-processing tools
- Inability to make vertical refinements or vertical coordinate changes within nests
- Potential to ‘properly’ degrade/coarsen grids there in principle, but not exploited

Issues resolved during previous development period:

- Merging of DOMAINcfg and AGRIF pre-processing tools
- Ability to change vertical coordinate and/or increase vertical resolution within nest domains
- Capability to use higher order schemes (clarified ghosts points definition)
- Ability to run agrif zoom on small (e.g. 3x3 points) mpp sub-domains
- Capability to handle global domain boundaries implemented
- Capability to run sibling grids in parallel (80% complete)
- Much improved user guide

#### 14.4 AGRIF is not as used as it should be!

The main concern regarding AGRIF is that this feature is used only by a few users although most users say they would be interested to use it and although AGRIF is presented as the solution to go at higher resolution at a reasonable cost. We need to facilitate and promote AGRIF use.

##### 14.4.1 Facilitate AGRIF usage:

Preprocessing tools have been strongly improved in the past years with the development of DOMAINcfg. It is now much easier to set up a configuration with nested zooms. This tool also ensures the consistency of the parent/child grids near the dynamical interface, which is a key point for the success of grid nesting procedure.

Some work is still needed to ensure the robustness and the sustainability of the DOMAINcfg tools. A first task should focus on use by DOMAINcfg of the current NEMO modules instead of old ones dating back to version 3.6. This cleaning step is necessary to maintain DOMAINcfg in accordance with any new developments of the nesting scheme as for example the connection between s-z grids or the possibility to benefit from the Brinkman’s penalization (*Debreu et al. 2020*, probably used in NEMO in the coming years) to define more “continuous” topography connection between the different grids.

##### 14.4.2 Promote AGRIF usage :

NEMO-AGRIF user guide has been added (<https://sites.nemo-ocean.io/user-guide/zooms.html>). It will be expanded to include more step-by-step examples of “how to set up realistic use cases” hoping it will make AGRIF more appealing and users more confident to use it. The user guide must also be kept up-to-date and enhanced for example by describing how to define a hierarchy of nested grids or use the vertical refinement.

NEMO-AGRIF still lacks proper documentation of the nesting procedure itself. This will help to broaden the knowledge base around AGRIF to a larger developer community and ensure we are not dependent solely on the expertise of one or two individuals. Subtleties on the external mode coupling in a Leapfrog context are, for example, not explained elsewhere which is not facilitating the transition to RK3.



## 14.5 Develop AGRIF New functionalities

### 14.5.1 Use of AGRIF as an online coarsening tool

Coarsening physical variables to run a BGC model greatly speeds up the modelling of tens of tracer variables. The existing coarsening tool is no longer maintained in versions 4 and later for multiple reasons and a new solution has to be found.

Technically speaking, the coarsening process is based on fluxes and divergence conservation principles that are identical to the ones used in the nesting restriction step. It seems therefore natural to envision AGRIF as an online coarsening tool that will in addition directly benefit from higher order, more selective schemes, already shipped with the AGRIF library (see *Debreu et al., 2012*). The recent adaptation of the code to deal with global cyclic child grids should make this adaptation even more straightforward.

The physical quality and the computation cost of this solution should be compared with a solution based on the coupling of the ocean dynamics with the BGC model through OASIS.

### 14.5.2 Nests with a complex nest geometry

AGRIF Grid refinement with complex geometries, for instance to follow complex coastlines (see *Holt et al, 2017*) or dynamically active regions (*Sein et al, 2016*), provides a functionality usually reserved to unstructured grid strategies.

The current solution is based on the use of multiple “sister” grids that are placed side by side and can run in parallel. Today, this method has two main drawbacks. First, adjacent sister grids cannot directly exchange data and must communicate through their common parent grid. Second, the deployment of a (very) large number of sister grids running in parallel must be facilitated to make this solution really usable: automatic defining of the nests mosaic, analysis of outputs spread over numerous nests, computational load balance...

Another solution that should be explored is the possibility to use the existing BDY modules to set unstructured boundaries for AGRIF nests. This solution would also benefit from the various open boundary schemes shipped with BDY modules (Flather, radiation methods, ...). This would represent a convergence of some parts of AGRIF and BDY modules which is good from a code maintenance perspective but could require a significant amount of work (to be precisely estimated).

### 14.5.3 Other functionalities to be added

Today, AGRIF is incompatible with the use of icebergs, ice cavities, wetting and drying. This is clearly a limitation for some applications (e.g. icebergs are needed for global climate simulations). These issues could be addressed according to the needs expressed by users and the input of new developers.

## 14.6 AGRIF physical robustness and performances

The quality of AGRIF results (conservation properties, truly robust multi-grid mesh coupling, ...) can be improved and must be adapted to the evolutions of NEMO dynamical core (e.g. RK3). These gains may however have significant computational costs that must be evaluated to quantify the relative benefits/costs of the different nesting strategies.

### 14.6.1 External mode coupling

Dealing with the external mode is certainly the most difficult part to obtain a truly robust multi-grid mesh coupling. In the case of a split-explicit treatment of free surface as in NEMO, the question arises whether data exchange between grids should occur at barotropic or baroclinic level. The second option (i.e. baroclinic) has been chosen in NEMO mainly because it does not require a deep reengineering of the model flowchart. It does however lead to possibly

growing errors at the grid interface, because of diverging barotropic mode solutions within their sub-integrations.

One solution would be to use radiative or Flather type (potentially from BDY modules), in place of clamped, open boundary schemes (*Penven et al. 2006; Herzfeld and Rizwi, 2019*) which can help to minimise the mismatch, but at the expense of departure from exact volume conservation. The second solution could be to perform the coupling at the barotropic time-step. Recent progress in understanding leading mode splitting errors (*Demange et al 2019*) has led to the design of barotropic time stepping schemes with ad-hoc built-in dissipation. This makes external mode exchanges even more simple. The other advantage of sub-step exchange, is that it makes possible exact coupling with adjacent grids having the same spatio-temporal refinement ratio. Hence, it opens up possibilities to truly multigrid nesting. This gain may however be compensated by the additional overhead of frequent grid exchanges which must be carefully quantified.

#### 14.6.2 Conservation properties

Conservation issues in NEMO, as in many other models, mainly come from the time dimension.

Volume is perfectly conserved in the present implementation, taking advantage of the forward nature of the barotropic model between two consecutive baroclinic steps. This is nevertheless not the case for tracers, for which the model advection and diffusion schemes compute their own set of fluxes at the grids interface from exchanged tracer values (this still guarantees monotonicity if required). One can still retrieve exact conservation thanks to “refluxing” methods (*Debreu et al, 2012*), but with time refinement, the Leapfrog time stepping greatly complexifies the exchange of time-integrated fluxes over the correct time interval (*Herrnstein, 2005*). Ensuring perfect conservation should be re-considered in the upcoming, two-time level, RK3 framework.

#### 14.6.3 Performance

AGRIF is proposed as a solution for future configurations allowing to reach, at lower cost, very high resolution (km-scale) in places where it is needed. The viability of such a solution requires good HPC performances of AGRIF that must therefore be investigated and optimised.

The 2-way nesting strategy implies interpolations/extrapolations between child and parent grids, which intrinsically generates communications that could slow down simulations and limit the model scalability. We recently started to quantify the impact of AGRIF on NEMO HPC performances and to identify the HPC bottlenecks. Different solutions have to be explored in the coming years: improve interpolation methods, test other MPI communication schemes (e.g. neighbourhood collective), try to reduce/avoid the update of 3d fields over the whole overlapping region or even exclude nested areas from the parent domain.

These optimizations must be considered in conjunction with the work on the numerical properties of AGRIF. In short, what is the HPC cost of the perfect conservation properties of AGRIF? Can we find better numerical schemes that cost less? Could we consider, at least for some applications, to downgrade the numerical properties of AGRIF if this allows us to significantly improve its HPC performances?

### 14.7 Priorities and timescales

#### 14.7.1 2 years:

- Coarsening of physical variables (for BGC modelling)
- Adapt nesting methodology to RK3 (1st working draft is already there), and review conservation properties.
- Improve mpp optimization (G. Irrmann PhD) and load balancing among child grids.

- Clean DOMAINcfg tool and share some of its modules with NEMO core (+ eventually work on the connection between s-z grids).
- Maturation of the user guide

#### 14.7.2 5 years:

- Implement sub-step (barotropic) exchanges
- Neighbouring grids communications or convergence of BDY/AGRIF functionalities
- Publish in-depth description of NEMO-AGRIF implementation (developer guide)
- Outreach activities (e.g. training sessions, user/developer workshop, perhaps linked to Drakkar)
- Further improve AGRIF HPC aspects (e.g. compatibility with an efficient GPU usage?)

#### 14.7.3 10 years:

- Compatibility with icebergs, ice cavities, wetting and drying
- Real ability of using zooms with complex geometries (configuration setup and HPC)
- Make AGRIF work with ALE ?

## 14.8 References

Debreu, L., N. Kevlahan, and P. Marchesiello (2020): Brinkman volume penalization for bathymetry in three-dimensional ocean models. *Ocean Modelling*, 145, 1-13.

Debreu, L., Marchesiello, P., Penven, P. and G. Cambon (2012): Two-way nesting in split-explicit ocean models: Algorithms, implementation and validation. *Ocean Modelling*, 49-50, 1-21.

Demange, J., L. Debreu, P. Marchesiello, F. Lemarie, E. Blayo and C. Eldred (2019): Stability analysis of split-explicit free surface ocean models: implication of the depth-independent barotropic mode approximation. *Journal of Computational Physics*, 398, 108875.

Herrnstein, A. R. (2005): A Parallel Ocean Model With Adaptive Mesh Refinement Capability for Global Ocean Prediction. Phd thesis, LLNL, 258p.

Herzfeld, I. and F. Rizwi (2019): A two-way nesting framework for ocean models. *Environmental Modelling and Software* 117, 200-213.

Holt, J., P. Hyder, M., Ashworth, J., Harle, H. T., Hewitt, H., Liu, A. L., New, S., Pickles, A., Porter, E., Popova, J. I., Allen, J., Siddorn, and R. Wood (2017): Prospects for improving the representation of coastal and shelf seas in global ocean models, *Geosci. Model Dev.*, 10, 499–523.

Penven, P., L. Debreu, P. Marchesiello and J. C. McWilliams (2006): Evaluation and application of the ROMS 1-way embedding procedure to the central california upwelling system. *Ocean Modelling*, 12, 157-187.

Sein, D. V., S. Danilov, A. Biastoch, J. V. Durgadoo, D. Sidorenko, S. Harig and Q. Wang (2016): Designing variable ocean model resolution based on the observed ocean variability. *Journal of Advances in Modeling Earth systems*. 8, 904-916.

## 15 Adaptations for High Performance Computing

WGL: Italo Epicoco

### 15.1 Executive Summary

Continuous adaptation of the NEMO model is required to retain excellent computational performance in the rapidly changing landscape of HPC technologies and architectures. A key challenge is to find solutions that guarantee performance portability on heterogeneous architectures whilst allowing the code to be efficiently maintained and developed even by people who are not experts in HPCs. In the next 5 years we propose to continue to focus on three main directions: (i) the improvement and completion of the tiling-based implementation, overcoming the current performance limitations; (ii) the transition to DSL based solutions enabling support of parallelisation for GPU-based architectures; (iii) lowering the time-to-solution by means of techniques based on the reduction of numerical precision while ensuring an adequate level of accuracy of the results also for long-term simulations. Finally, it is emphasised that the performance optimization activities are transversal and the computational efficiency must be considered from the beginning each time new features are introduced in the model. This requires that the HPC WG works in close collaboration with all other working groups.

### 15.2 Background and context

The rapid development of HPC architectures, the continuous growth of computing capabilities (which today has reached up to exaflops), including the evolution of processing units (CPU, GPU, FPGA, etc.), memory, storage, and networking technologies require a continuous adaptation of the NEMO code to different and changing computing architectures. Moreover, the continuous development of new features in the model, new numerical schemes and the development of new physical processes poses a big challenge to preserve the computational performance together with code maintainability, readability and portability. This chapter aims at identifying the main HPC aspects to be addressed in the next five years and outlines our vision for the next 10 years.

In outlining the HPC priorities we must take into account the current HPC context, the technology trend and the current bottlenecks identified in the NEMO code. The current technology trend demonstrates a consolidation and ever greater use of GPUs into high performance architectures; 8 of the top 10 most powerful parallel HPC platforms exploit GPUs (see top500 list). Moreover, even traditional CPUs are evolving towards a many-core approach making support of the shared-memory paradigm by numerical models extremely important. The shared-memory approach and GPU programming slightly change the computational approach, shifting the focus on per-thread parallelism and on instruction level parallelism. The current official version of NEMO is based on a pure message passing paradigm; it does not support multi-thread parallelism for either CPU or GPU. This represents one of the major weaknesses of NEMO, which not only prevents the model from fully exploiting GPU-based architectures and many-core CPU processors, but also prevents it from being efficiently used in coupled models where other components support hybrid parallelization (distributed- and shared-memory parallelization).

Moreover, in the last decade more specialised and specific processing units have started to appear in the context of parallel architectures, such as FPGA, ASICs components, TPUs, etc. These new technologies may generate a plethora of programming environments and paradigms to follow in order to exploit their capabilities. The big challenge here is to develop a codebase able to support heterogeneous architectures which preserves its maintainability and readability. In this regard, the

promise of Domain Specific Languages (DSL) must be carefully considered. On the other hand, the specialisation of the processing units leads also to changes in the numerical precision used for the representation of the floating point values. Several studies proved that lowering the numerical representation of floating point values does not lead to a significant loss of accuracy in the results. At ECMWF a version of the IFS atmospheric model [1] at single precision was developed proving that for annual integrations and medium-range ensemble forecasts no noticeable reduction in accuracy, and an average gain in computational efficiency by approximately 40% can be achieved. Paxton et al. [2] reached similar conclusions in a study of the round-off error produced by the use of single precision and even half precision numbers in a shallow water model. The adoption of a mixed numerical precision was also evaluated with some test cases for NEMO by the Computational Earth Sciences group at the BSC; they proved that with a careful selection of the variables to be lowered to single precision, the model can still produce results with an accuracy comparable to that obtained with the double precision version but with a significant improvement of the computational efficiency.

Finally, it is worth mentioning here the role of the cloud computing technologies which are gaining ever more interest for the HPC applications. The cloud paradigm was born mainly to provide the user with a powerful computational platform at low cost, mostly suited for handling big data and applying lazy-coupled parallel computation. In the last few years, the interest of the cloud providers has moved to the HPC world, offering solutions suitable for running tightly-coupled HPC applications. In this regard, the key word is containerization. A study was conducted, within the ESIWACE2 project by CMCC, to port NEMO on a Sarus container (developed by CSCS) and to evaluate the computational performance compared with the native execution of the model. The results [3] were really encouraging, demonstrating that the computational performance of NEMO, executed through a Sarus container in a HPC architecture, is comparable with the execution directly on the machine without the container.

In the last three years, many HPC optimizations were done in NEMO according to the NEMO Development Strategy defined for the 2018-2022 time period, mainly supported by the IMMERSE, IS-ENES3 and ESIWACE2 project's fundings. Among the main optimizations we mention here the use of a wider halo region; the re-organisation of the operations into 2D tiles; rationalisation of I/O operations extending the use of XIOS to reading restart files; investigation of mixed precision; offloading of the diagnostics computation to GPU through a CUDA based approach; experimentation with PSyclone DSL to run NEMO entirely on a multi-node GPU-based parallel architecture. Moreover, some specific optimizations tailored for some NEMO modules included: a new support for macro task parallelization disentangling the execution of the biogeochemistry model from the ocean dynamics; a new representation of vertical layers by means of Quasi-eulerian Coordinates (QCO) which replaces Vertical Varying Layer (VVL); and an improved algorithm to balance the workload balancing when computing nested grids in AGRIF.

The recent computational performance analysis of NEMO v4.2 reported a good improvement with respect to NEMO v4.0. Some examples relate to: eNEATL configuration which is a regional configuration at  $1/36^\circ$  of the North-East Atlantic able to scale up to 14,000 cores reaching up to 15 SYPD (Simulated Years Per Day) with a performance improvement of 50% w.r.t. NEMO v4.0 mainly due to the use of QCO instead of VVL (QCO greatly reduces the number of 3D fields used to represent the model grid); a global configuration, GO8-ORCA025 developed at MetOffice, at  $1/4^\circ$  runs at 4 SYPD with QCO, tiling and extended halo activated; STFC achieved 2 SYPD with 100 GPUs on JUWEL Booster for a global ORCA12 configuration.

### 15.2.1 Main Issues

Although a lot of effort has been put into it in the past three years to improve the computational performance of NEMO, it still suffers from some relevant limitations.

The first evident limitation is the lack of shared memory parallelism which limits an efficient exploitation of multi- and many-core architectures and it limits an efficient use of NEMO coupled with other climate components, like IFS atmospheric model, which are parallelized with shared memory and distributed memory parallel approach.

Moreover, the computational performance of NEMO is bounded by the memory bandwidth. The tiling approach, the reduction in the number of 3D arrays, the experimentation of the loop fusion technique, the use of mixed precision lowering the numerical precision as much as possible are all optimizations aimed at reducing the memory access and hence enhancing the arithmetic intensity. The tiling optimization resulted in an average improvement of more than 20% (reaching 50% for some routines), but the use of the extra halo, mandatory for the tiling to be applied, introduced a computational penalty which hid the benefit of the tiling.

Moreover, NEMO cannot be efficiently executed on parallel architectures that are primarily GPU based. Really promising investigations were conducted to execute NEMO on GPUs but today the official version of NEMO does not support GPU.

Finally, the performance portability over different architectures and the code maintainability and readability represent some of the crucial factors which limit the adoption of disruptive solutions. The separation of concern between the HPC optimization tasks and the climate process modelling would help having a code readable by climate scientists, efficient on high-end architectures and easily portable on different machines.

### 15.2.2 Priorities for the 2023-2027

Considering the current limitations of the NEMO code we can put as first priority the support of per-thread parallelization and support for mixed parallelization based on shared- and distributed-memory paradigm as well as porting on GPU-based parallel architectures.

Moreover, even if a lot of effort was devoted to optimizations in the last three years, many of the code changes, such as tiling and extended halo management, need to be consolidated in order to lead to better performance results.

The exploitation of the DSL approach is considered of paramount importance in order to support heterogeneous architectures (including GPUs, but also specialised processing units at long term perspective) ensuring performance portability and a separation of concerns between HPC and physics and ocean dynamics representation.

Finally, the support for mixed precision computation (double, single and half precision) is considered the natural evolution to better exploit the new generation of processing units.

These envisioned actions are better detailed in the next sections

## 15.3 Focus areas

### 15.3.1 Tiling-based approach

The rationale for tiling on CPU based HPCs is fairly easy to explain. There is a memory hierarchy in HPCs that typically has 3 relevant levels of memory. The rate at which data can be passed between the caches and main memory is termed memory bandwidth. In order to fully exploit the computational capacity of the processor, it is necessary to do as many calculations as possible using cached data

without recourse to the main memory. Tiling allows us to divide the calculation into chunks of work that can remain cache-resident for as long as possible. The technique leaves the tile size and shape as tunable parameters, which can be tuned appropriately for cache sizes on any platform.

For a typical deployment of a model configuration on current HPCs, only a small number of (the local) 3D arrays can be stored in L3. Within NEMO, the calculations for 3D fields are done within triply-nested (3D) DO loops over the whole local domain. The number of calculations within each triply-nested loop is very limited. This leads to a lot of data being transferred per calculation. If the processing is re-organized so that the DO loops perform calculations for smaller subsets of the local domain, data transfers can be greatly reduced.

The last NEMO release (v.4.2) allows the number of halo points in an MPI domain to be specified (the main choices being 1 or 2). When a 2-point halo is chosen the number of exchanges between halos is greatly reduced and a 2D horizontal tiling of the domain can be used for most 3D calculations. In practice on SIMD processors the first (ji) index cannot be tiled without degrading the model performance. The computational cost of some subroutines in some configurations (typically those with relatively low vertical resolution) is reduced by 30-50%. Somewhat larger and more consistent improvements in performance can be achieved for some subroutines by tiling the calculations also in the vertical direction. We intend to implement this for the most costly routines where we can.

An additional advantage of the tiling is that it allows OpenMP threads to perform the calculations for the tiles in parallel. During the initial serial implementation, it became clear that this was not easy to achieve due to overlap between tiles and “non-overlapping” DO loops were introduced in some subroutines to ensure correctness of the calculations. A new code refactorization, that allows the tiles to be calculated independently (and hence be suitable for OpenMP) and is not difficult for code developers to understand, has been formulated and tested but needs to be agreed and consolidated.

### 15.3.2 Moving towards DSL

One approach to being able to develop a model that can satisfy Portability, Performance and Productivity of the code is that of Domain-Specific Languages. Traditionally, this means that a domain scientist writes their model in a language specifically designed to be expressive and powerful for their particular field. The main advantages of this approach are in its "separation of concerns": domain scientists can concentrate on the scientific aspects of the code while computational scientists can optimise the code that is generated or target entirely new hardware by working on aspects of the compiler.

What is required is a way of evolving an existing code base such that it can take advantage of DSL technology without having to be rewritten from scratch. This approach has been explored in the ISENES2, ESiWACE2 and ExCALIBUR Marine Systems Projects which have worked and are working to extend the 'PSyclone' code-generation and transformation system so that it is able to work with existing, unmodified NEMO source code. In a sense, this treats the NEMO code with its associated coding standards as a DSL and thus no (or only very minor) re-writing of the model is required: the use of PSyclone is transparent to NEMO developers since there is no need to switch from Fortran.

PSyclone is developed by STFC's Hartree Centre, the UK Met Office and the Australian Bureau of Meteorology. It forms a key part of the build system for the UK Met Office's new LFRic atmosphere model, due to go operational within the next two years. As such, the UK Met Office is committed to its ongoing development and support.

PSyclone has been extended and developed such that it is now able to process a complete NEMO configuration (based on the GO8 configuration from the Met Office) and transform it such that it can

be run on GPU-based machines. When this is done, the ORCA1, ocean-only version of the configuration runs on a single NVIDIA V100 GPU at 3.5 times the speed of a Skylake socket. The PSyclone-processed, ORCA12 configuration has been run on 92 A100 GPUs on the JUWELS Booster machine with performance equivalent to 270 Intel Skylake sockets.

The tiled code will be largely “transparent” to PSyclone since the tiling implementation introduces little bespoke code at the subroutine level. As such, PSyclone transformations should only need to reflect corresponding updates to the NEMO coding convention. In PSyclone the tiling is implemented at the loop level, such that nested DO loops over the MPI sub-domain are each enclosed by a loop over tiles and these loops then fused (where possible). In NEMO the manual tiling is instead implemented at the timestep level, such that a loop over tiles will enclose whole sections of code. Despite this, PSyclone and the NEMO tiling do not necessarily conflict in purpose. Parallelisation of the NEMO tiling could still be handled by PSyclone’s OpenMP transformations. More generally, existing components of the tiling framework could be “offloaded” to PSyclone to simplify the code base.

Finally, it is worth mentioning that PSyclone opens the way for a wide range of hardware- and/or configuration-specific optimisations to be developed since these optimisations take the form of separate Python transformation scripts instead of having to be hardcoded in Fortran, as an example we mention the loop-fusion transformation. This is not only quicker to do (one script can be applied across the whole code base) but can be done independently of the ongoing scientific development of NEMO. The Python transformation scripts, used by PSyclone to transform the code, could be also used as an automatic tool to check if the NEMO code complies with the code conventions.

### 15.3.3 Mixed precision

During the last decades, the available computational performance has been steadily increasing (Moore's law), while the increase in CPU memory speed has been lagging. As a result, many computational codes, such as NEMO, have seen that CPU speed is no longer the main limiting factor for their performance, becoming memory-bound applications.

On the other hand, most modern processors implement vector operations which allow the number of floating-point operations per cycle to be doubled by halving the size of the operands.

As a result, mixed-precision approaches emerge as a powerful solution to improve application efficiency by improving the speed at which variables are read from memory and increasing the degree of parallelism in a single core.

In the last few years, several teams have applied precision reduction techniques to improve the performance of their codes, from the routine level to the whole application, including Earth Sciences codes. One clear example is the IFS model, whose code was migrated to mixed-precision so that most of the fields are represented in single precision and finally put in production in 2021.

Nevertheless, entirely moving a computational model to mixed-precision can be an arduous task. Unavoidably, a decrease in the precision used to represent the operands will lead to different results from the operations in which they are involved, with a high probability of generating numerical errors and instabilities, especially in computational models of a chaotic nature that perform a considerable number of operations at different scales and in which small perturbations can be propagated and amplified, leading to different process representations. As a result, one of the biggest challenges that this kind of work poses is identifying which variables can be safely demoted to lower precision, especially if the software is intended to provide results comparable to the higher precision counterpart.



These questions were very briefly exposed in the previous development strategy document, written at the moment when the NEMO community was starting to pay attention to this problem. Since then, the Barcelona Supercomputing Center (BSC) has been working on a methodology to move FORTRAN codes to mixed-precision in an automatic way, intending to simplify the most challenging steps of transferring a complete model to mixed-precision. This involves identifying the variables that can safely be moved to a lower precision and performing the necessary changes in the code to avoid affecting the interoperability with other variables or operands represented at the same or another precision level.

The methodology has been implemented by a set of tools that do all the necessary analysis to classify the different fields into two groups (those whose precision can be reduced and those which can't) to then create an actual implementation in the target precision.

This analysis uses a precision emulator and is based on a set of tests (variable, threshold pairs) to decide if the original precision is required or not. These variables are model diagnostics and internal fields whose value is obtained from actual model runs. As a result, the methodology is dependent on the model configuration used in the analysis.

Consequently, the BSC and the NEMO consortium agree that the accessibility to this set of tools is a key asset to promote the use of mixed-precision in NEMO.

In addition to the BSC tools, other approaches can be evaluated such as the PROMISE [4] tool which returns a subset of variables that can be transformed into single precision, taking into account a required accuracy on the computed result. It is based on the delta debugging search algorithm and Discrete Stochastic Arithmetic (DSA). CADNA (Control of Accuracy and Debugging for Numerical Applications) [5] allows the use of stochastic numerical types in numerical models. In practice, classic floating-point variables are replaced by the corresponding stochastic variables, which are composed of three floating-point values and an integer to store the accuracy. During the execution CADNA is able to evaluate the round-off error and to detect numerical instabilities, providing useful information about the numerical precision that can be used for the variables.

#### 15.4 Impact on code base

Many of the HPC and performance optimizations impact the whole NEMO code. Here, the real challenge is to do the code transformations while preserving the code maintainability. The tiling approach requires revisiting the whole code to change the DO LOOPS ranges. Fortunately, most of these transformations can be applied using precompiler macros which are already part of the latest version of the NEMO code.

Making use of a DSL typically would require changing the programming language, but since PSyclone is tailored around the code structure of NEMO, the use of PSyclone does not heavily impact the NEMO codebase. However, some actions should be taken into account such as: full integration of PSyclone processing into the build system and SETTE suite; addressing those parts of the NEMO code base that do not work well on GPU (e.g. the Iceberg component, statement functions).

Finally, the support for the mixed-precision following the BSC approach would require adding a minimum set of changes to NEMO like a core set of function interfaces for different precision to facilitate the operation of the automatic tool. A considerable part of those modifications was developed in collaboration with the ECMWF and are currently already available. Likewise, the number of code changes needed to run a new configuration in lower precision can be smaller if there is already a reference configuration prepared to run in that precision. For this reason providing mixed-precision

support for a configuration of the ORCA1 family will be beneficial in the long run, and the plan is to do it once the automatic tool is ready to be used by the consortium members.

A few HPC optimizations would impact only on specific NEMO modules such as: optimization on AGRIF; optimization on SI3 to make the code suitable for per-thread parallelization; optimization on iceberg module.

### 15.5 Cross-cutting issues

Taking into account the rapid evolution of the NEMO code, the analysis and monitoring of the computational performance of NEMO during the model evolution represents one of the most time-consuming parts within the HPC activities. Also, the debugging process often requires a big effort in terms of person-hours. In this regard, the development of an automatic tool which allows to perform scalability tests and performance analysis can help. Exploiting the approach used for the SETTE test, largely used to perform a first stage of code validation, an automatic tool able to execute scalability tests can be used to gather and collect the evolution of the computational performance along with the evolution of the NEMO versions highlighting the bottlenecks.

Moreover, the exploitation of optimised numerical libraries should be considered and evaluated; layering the computation on optimised libraries will also increase the performance portability on different architectures.

Finally, it is worth mentioning that the data handling and I/O operations constitute one of the major factors which impacts on the computational performance. This aspect has not been explicitly addressed in this chapter because the I/O management is performed by the XIOS library developed jointly at IPSL and CEA. XIOS manages output of diagnostics and other data produced by climate component codes and offers temporal (average, minimum, maximum, etc.) and spatial post-processing operations. In this regard, the NEMO HPC-WG will also be in charge of monitoring the XIOS performance, maintaining a tight collaboration with the XIOS team, reporting any issue regarding I/O and proposing the development of new features for XIOS in order to better support NEMO. The data handling should also take into account the interfaces towards Machine Learning modules and approaches that will be considered in the near future (see next Chapter)

### 15.6 References

- [1] Váňa, F., Düben, P., Lang, S., Palmer, T., Leutbecher, M., Salmond, D., & Carver, G. (2017). Single Precision in Weather Forecasting Models: An Evaluation with the IFS, *Monthly Weather Review*, 145(2), 495-502. Retrieved Sep 1, 2022
- [2] Paxton, E. A., Chantry, M., Klöwer, M., Saffin, L., & Palmer, T. (2022). Climate Modeling in Low Precision: Effects of Both Deterministic and Stochastic Rounding, *Journal of Climate*, 35(4), 1215-1229. Retrieved Sep 1, 2022
- [3] <https://github.com/eth-cscs/ContainerHackathon/blob/master/NEMO/report/report.pdf>
- [4] Stef Graillat, Fabienne Jézéquel, Romain Picot, François Févotte, Bruno Lathuilière, Auto-tuning for floating-point precision with Discrete Stochastic Arithmetic, *Journal of Computational Science*, Volume 36, 2019, ISSN 1877-7503, <https://doi.org/10.1016/j.jocs.2019.07.004>.
- [5] P. Eberhart, J. Brajard, P. Fortin, and F. Jézéquel. High performance numerical validation using stochastic arithmetic. *Reliable Computing*, 21:35-52, 2015

## 16 Machine Learning

WGL: Julien Le Sommer and Andrea Storto

### 16.1 Executive Summary.

The combination of machine learning with scientific computing is an active area of research which could eventually improve geoscientific models and their integration into broader numerical systems, such as climate models and operational forecasting systems.

These emerging approaches could help quantify and reduce systematic model errors, improve the representation of unresolved processes, reduce the numerical cost of model simulations, improve the representation of uncertainty propagation and allow to better leverage observations in model design.

But, while the potential of ML for improving ocean/sea-ice models is large, the *Technological Readiness Level* of these applications is still relatively low. Yet, given the possible implications on some of the NEMO consortium requirements, we think that the NEMO consortium should be proactive on these topics.

We therefore propose to set up a new NEMO working group (WG) focusing on “*machine learning and model uncertainties*” in order to coordinate and accelerate development in this area. The ambition of this working group is to gather a group of interested parties for fostering innovation and exchanges of expertise on these interdisciplinary topics.

In parallel, we propose to leverage resources from several recently funded research projects for moving forward on an explicit list of actions for the next 5 years. These actions will overall aim at :

- (A1) introducing several ML-based components into NEMO,
- (A2) improving the representation of model uncertainties in NEMO,
- (A3) investigating the potential of differentiable emulation.

The roadmap, which is described in more detail below, will be implemented in two successive phases with several key actions to be completed before 2025. Their successful implementation will require strong coordination with other NEMO working groups and with the NEMO ST. Community and capacity building activities will also be pivotal to our success in this area.

Additionally, we propose that the *working group on machine learning and model uncertainties* acts as a link with the data assimilation community which leverages NEMO in operational systems. We propose to monitor new developments in this field and to identify key needs that would help users to run NEMO within data assimilation frameworks.

### 16.2 Context and purpose of this chapter

While advanced statistical methods and statistical learning have long been used in geosciences and remote sensing for solving inverse problems (Larry et al. 2016), we have witnessed over the past 5 years a very fast increase in the number of applications of *machine learning* (ML, see appendix), and

specifically deep learning (DL, Goodfellow et al. 2016), to the field of fluid mechanics (Brunton et al. 2020) and computational fluid dynamics (Kochkov et al. 2021). This acceleration reflects a more general trend with a growing number of applications of ML to physical sciences (Carleo et al. 2019) and scientific computing (Thurey et al. 2021).

Several new usages of machine learning relevant to the design and the usage of geoscientific models have emerged over this period, with published proof of concepts of applications for calibrating model parameters (Couvreur et al. 2021), for designing subgrid closures (Bolton and Zanna, 2019), for downscaling model data (Stengel et al. 2020), for accelerating the execution of specific code components (Chantry et al. 2021), for guiding the design of numerical schemes (Zhuang et al. 2021, Magiera et al. 2020), for learning underlying equations of motions (Champion et al. 2019), or for building representation of model errors (Bonavita and Laloux, 2020).

How (and how fast) ML will eventually affect the landscape of numerical tools used for studying and predicting oceanic flows and sea-ice dynamics is still unclear at this stage. Indeed, many of the works cited above are still exploratory proofs of concepts which do not yet exhibit the technological readiness for being implemented and maintained in production codes like NEMO. But the field is moving fast with many on-going research projects across the world. It is therefore reasonable to anticipate that the technological readiness of these applications will increase rapidly, and that new areas of applications could emerge over the period covered by this strategy (2023-2027).

We anticipate that, by 2027, physics-based models, such as NEMO, will still be widely used and that their structure will not be deeply affected by ML. Still, by then, ML will probably be more often used for analyzing their output and for calibrating their parameters. It is also likely that offloading some specific compute-intensive code components to GPUs through ML-based emulation will be a mature and viable option by then. We anticipate that ML will at that stage provide realistic opportunities for improving prediction systems involving data assimilation, and practical options to better exploit hybrid computer architectures. We also anticipate that ML will provide a framework for more systematically leveraging observations in the design of geoscientific models and prediction systems (Schneider et al. 2017).

All aspects, components and use-cases of physics-based models could eventually be affected by ML but the most radical innovations will arguably concern the estimation (and the correction) of model errors, and the (probabilistic) representation of model uncertainties. At present in NEMO, model uncertainties are accounted for with a combination of stochastic physics schemes and ensemble simulations. Both are essential tools for the integration of NEMO in prediction systems, but also for understanding oceanic variability within the climate systems. As described in the next section, we anticipate that ML will eventually deeply affect our ability to characterize model errors and to represent uncertainties and their propagation. We therefore propose to *jointly address machine learning and model uncertainties*, with the vision that two topics should naturally converge at some stage in the future.

In this context, the ambition of this chapter is to identify practical actions aiming at: (i) fostering the exploration of ML applications for the design and usage of the NEMO code and (ii) moving towards a more robust probabilistic representation of model uncertainties in NEMO simulations. Our general ambition is to prepare NEMO development in these areas beyond 2027 with the vision that the two topics will naturally merge in the future.

This chapter focuses in priority on applications of ML that require specific developments in the NEMO code. Indeed, implementing a subgrid closure or a numerical scheme designed with ML but expressed as a closed form equation (as the ones obtained with equation discovery approaches, see e.g. Zanna and Bolton 2020), would not require major changes to the NEMO code. Similarly, using ML-based approaches for calibrating model parameters would not a priori require any change of the NEMO code itself. On the contrary, implementing and maintaining parameterizations expressed as NN would require a dedicated interface and may have large implications on the NEMO codebase.

### 16.3 Target areas of applications of ML relevant to the development of NEMO

Based on the recently published literature, we identify four areas of high potential application of ML to the development of ocean/sea-ice models in the coming years which are relevant to NEMO.

#### 16.3.1 Better account for the impact of unresolved processes on resolved scales

A first area of application of ML to ocean-sea-ice models is the design of subgrid parameterizations, and more generally the representation of unresolved scales and processes with ML, which has attracted quite some attention over recent years in the geoscientific and climate modeling community. Published works relevant to ocean model development have mostly focused on the representation of ocean macro-turbulence (see Zanna and Bolton 2021 for a recent review of mesoscales eddy closures with ML) but applications can be thought of for many different processes and scales. Given recent advances in the computational fluid mechanics (CFD) and ocean modeling communities, the design of subgrid closures for ocean macro-turbulence appears as a reasonably low hanging fruit for applying ML to ocean-sea-ice model design. It should be noted however that, at the time of writing, the most advanced *interactive* ocean simulations with ML are still based on idealized flow configurations (as for instance Bolton and Zanna 2019, Guillaumin and Zanna 2021), while realistic ocean simulations have only been used for *non-interactive* inference so far (Partee et al. 2021), but this limitation should most likely soon be overcome. Current challenges are associated with how to account for the different flow regimes encountered at different locations across ocean basins, how to optimally define the filtering operator used to formulate the ML problem, how to bring prior physical or mathematical knowledge in the learning process (Frezat et al. 2021) and how to combine deterministic and stochastic components of eddy closures. Besides the representation of ocean macro-turbulence, ML could also probably be used for improving the representation of vertical physics in the OSBL, of fine scale processes at the air-sea interface, and of unresolved processes at the ice-sheet/ocean interface. All these examples would a priori use information drawn from finer resolution models (possibly down to LES simulations). Depending on the specific problem, the *technology readiness level* of applications of ML to represent unresolved scales range from intermediate to high.

#### 16.3.2 Accelerate the execution of specific code components with deep emulators

Fast emulation is another important area of application of machine learning relevant to ocean/sea-ice models development that has emerged over recent years. *Emulators* (aka *surrogate* models) are statistical models that learn to mimic the behavior of pre-existing numerical codes at reduced numerical cost. Emulators are used quite extensively for sensitivity analysis or for calibrating model parameters (see for instance Salter and Williamson, 2016, Williamson et al. 2017). In this context, emulators are generally aiming at reproducing some summary statistics of model trajectories (as for

instance spatially and temporally averaged temperature bias). More recently, thanks to the versatility of deep neural networks as general purpose approximators, machines have been trained to emulate not only summary statistics of model trajectories but the entire model state and its evolution with time along model trajectories (Nonnenmacher & Greenberg, 2021; Kasim et al. 2022). Such *deep emulators* have successfully been used for reducing the cost of existing parameterization in atmospheric models and porting specific code components to GPUs (Chantry et al. 2021). Moreover, because automatic differentiation is readily available in ML libraries, deep emulation can also be used as a strategy for approximating the linear tangent and adjoint operators of model components or entire models (Hatfield et al. 2021)<sup>7</sup>. Deep emulation could therefore eventually open the possibility to formulate inverse problems for adjusting specific model parameters or for guiding the development of new code components with observations (Schneider et al. 2017), while allowing more versatility to better exploit future computing architectures. In this sense, deep emulation may offer an alternative route to a full rewriting of our models in differentiable frameworks (see appendix), an undertaking that would be required for designing code components through end-to-end learning (Frezat et al. 2022). The *technology readiness level* of deep emulation approaches ranges from low (approximation of the adjoint operator for entire HPC codes) to intermediate (porting specific code components to GPU).

### 16.3.3 Quantify and reduce parametric and systematic model errors

A third area of application where ML could be leveraged for NEMO development is related to the quantification and the reduction of parametric and systematic model errors. Indeed, as with any physics-based computer model, uncertainties in the formulation of the NEMO model, or in its parameters, result in errors in practical simulations (see for instance Allen et al. 2003 for a discussion on model errors and uncertainties). Several approaches have been developed over past decades for quantifying and reducing model errors, in particular in the data assimilation community. Over recent years, we have also witnessed the emergence of many new approaches leveraging ML in this context. ML is for instance now often used for calibrating physics-based model parameters (for recent examples with NEMO, see for instance Falls et al., 2021; Williamson et al. 2021). In practice, these approaches involve training cheap emulators which predict some summary statistics that are used for formulating inverse problems for optimizing model parameters (see Salter et al. 2016; Couvreur et al. 2021; Clearly et al. 2021). Another emerging area is the quantification of systematic model errors with ML, as proposed for instance by Bonavita and Laloyaux (2020). Approaches are being developed for training representations of model errors with ML models in the context of DA systems.

There would certainly be some benefit in using such pre-trained ML-models at the run-time in model simulations as online bias correction procedures. More generally, one could think of using more systematically ancillary datasets such as analysis increments, high-resolution simulations, reanalyses or observations, for inferring and correcting systematic model errors at the run-time from NEMO model simulations. Several of the applications listed above could a priori be leveraged from NEMO model output without any change to the NEMO code-base but the ML-based correction of model

---

<sup>7</sup> It should be noted that we are here referring to an adjoint operator describing not only the sensitivity of model solutions to the model state but also to model parameters (as opposed to what was previously available with NEMO-TAM).

errors would require calling ML-models from the NEMO code. The *technology readiness level* of the above approaches ranges from low (for bias correction techniques) to intermediate (for model parameter estimation).

#### 16.3.4 Improve the representation of model uncertainties and their propagation

As mentioned in the previous section, the explicit representation of uncertainties in NEMO simulations is another area of application of ML. Model uncertainties are presently represented with stochastic physics schemes and ensemble simulations (Brankart et al. 2015). These methodologies provide a practical way to describe the evolution of the probability distribution of plausible oceanic states given uncertainties in the forcing, in the model formulation or in the model parameters. This probabilistic framework is essential to the integration of NEMO into prediction systems with data assimilation, but also for uncertainty quantification in the context of climate projection and predictability studies (Bessières et al. 2017). The ability to generate optimal ensembles, at given numerical cost, in these contexts is becoming more and more necessary. For the time being, available stochastic physics schemes in the NEMO community are based on analytical formulations that account for the sub-grid ocean variability, approximated by perturbing model tendencies or parameters. The methods for designing and optimizing subgrid parameterizations with ML described above could also be generalized for controlling the parameters of stochastic parameterizations with ML. One may for instance leverage lightweight emulators calibrated on ensemble simulations for designing data-driven stochastic parameterizations. On a more long term level, one may also think about leveraging ML for optimizing and orchestrating the production of ensemble simulation and therefore turn deterministic geoscientific models into probabilistic modelling frameworks. Over the past few years, several methods based on ML have indeed been proposed for post-processing and calibrating ensemble forecasts (Grönquist et al. 2021; see also Schulz and Lerch 2021 for a review). In this context, ML is used for learning the parameters of the distribution of possible states (including the ensemble spread and bias). But similar approaches could also be used for ensemble simulations in more general contexts, as for instance for guiding the orchestration of ensemble simulations or resampling optimally probability distributions at fixed computational cost. We stress that such methods should be considered as internal to NEMO as they would require communication at the run-time across the members of the ensembles. Overall, the *technology readiness level* of the ideas mentioned above is pretty low (except for the approaches for parameter calibration), but the potential is certainly large, and the field is fast moving. We advocate that preparing NEMO to possibly leverage upcoming future development for better representing model errors and uncertainties could be done with minimal efforts.

### 16.4 Overall ambition and proposed roadmap

#### 16.4.1 Overall ambition :

Our ambition with respect to machine learning and model uncertainties is to :

- A. prepare the ground for robust deployments of ML-based components in full-scale production systems using NEMO by 2027,
- B. allow for more robust (probabilistic) representation of model uncertainties in NEMO simulations, leveraging ML when appropriate, and

- C. nurture and engage a community of NEMO users and candidate developers for preparing future evolution in these areas.

Advances in the above areas will require work on the practical implementation of ML within NEMO while fostering new interdisciplinary collaborations to fuel NEMO development in this fast moving field in the future.

Our short term priority [2 years] will therefore be to define how to implement, distribute and maintain pre-trained ML-based components for *interactive inference* into NEMO.

We also would like to simplify the testing of new ideas involving ML with NEMO in research projects and reduce the time for transition of these new ideas into the NEMO codebase.

#### 16.4.2 Proposed approach:

We propose to set up a new NEMO working group on “*machine learning and model uncertainties*”. The working group will be jointly led by Julien Le Sommer (CNRS) and Andrea Storto (CNR). The working group will gather a relatively large number of interested parties<sup>8</sup>, for sharing recent scientific results on topics ranging from uncertainty quantification, differentiable emulation, ML-based parameterizations and the explicit representation of uncertainty propagation in ocean models.

Subgroups will be invited to contribute to more topical meetings on dedicated NEMO-related questions. The group's practical contributions to NEMO will mostly be based on developments undertaken in research projects (see list below) which will be articulated with the relevant NEMO WG and the NEMO ST as needed.

#### 16.4.3 Main expected milestones

We list below the main actions with direct implication onto NEMO development that have been identified.

##### (A1) ***Introducing several ML-based components into NEMO***

- **A1.1** [2 years] Define how ML-based components should be included and delivered with the NEMO codebase (interface, distribution/versioning);
- **A1.2** [2 years] Deliver a proof-of-concept practical implementation of a ML-based subgrid parameterization in NEMO ;
- **A1.3** [5 years] NEMO reference codebase to comprise several ML-based components usable in full scale production simulations ;

##### (A2) ***Improving the representation of model uncertainties in NEMO***

- **A2.1** [2 years] Define a rationale and a roadmap for improving the explicit representation of model uncertainties within NEMO (with ensemble, stochastic parameterizations) ;

---

<sup>8</sup> A group of 30 participants has been identified and the group has been officially launched in Sept. 2022.



- **A2.2** [4 years] Deliver an up-to-date state-of-the-art stochastic physics package for NEMO

**(A3) Investigating the potential of differentiable emulation**

- **A3.1** [3 years] Deliver a proof-of-concept demonstration of the differentiable emulation of (one of) NEMO (components);
- **A3.2** [5 years] Inform whether deep emulation is a viable option for accelerating NEMO simulations on various architectures, and for approximating a tangent-linear model for DA

As described in more detail below, this list of actions is based on activities planned in several (funded) research projects.

## 16.5 Analysis of the need and options for introducing ML into NEMO

### 16.5.1 Types of tasks involved

In practice, ML models should be trained (*learning* step) before being used (*inference* step). Most applications of ML to geoscientific models so far are based on supervised learning tasks formulated from datasets of higher resolution model output. There are yet very few applications of inference at the runtime in full complexity realistic models (to our knowledge, none with ocean models beside Partee et al. 2021).

In terms of tasks at the runtime in a NEMO application, three very different tasks can be encountered:

- *Interactive inference* : requires calling a pre-trained ML model and modifying the NEMO model state based on the output of the ML model (ML model can be run on GPU or CPU with good performance)
- *Online learning* : involves optimizing parameters of a ML model with a training objective based on the NEMO model state as a simulation (or an ensemble of simulation) is progressing (ML model should be trained on GPU).
- *Interactive learning* : involves optimizing the parameters of a ML model with a training objective which evaluation requires to call specific components of the NEMO code or to run one (or several) NEMO simulations with input parameters determined by the training process (ML model should be trained on GPU).

### 16.5.2 The question of the interface

A key question is to define how ML models can be encoded and maintained into NEMO for *interactive inference*. There are several options, each coming with pros/cons :

- (a) Implement NN in FORTRAN (as proposed by Curcic et al. 2019 with the Neural-Fortran library<sup>9</sup>). Note that this approach is not adapted to training tasks, and yields code which may be difficult to maintain and change over time. Leveraging GPUs may also not be straightforward with this approach.

---

<sup>9</sup> See <https://github.com/modern-fortran/neural-fortran>

- (b) Call Python ML-libraries from FORTRAN leveraging Python C-bindings (as proposed by Noah Brenowitz<sup>10</sup>). This approach should in principle be adapted to both interactive inference and online learning, but may be computationally suboptimal, scarcely portable, and requires significant code for the wrapper functions.
- (c) Leverage the existing C/C++ bindings of specific ML libraries (as PyTorch or TensorFlow) from FORTRAN. This approach was proposed by Ott et al. (2020) with the FORTRAN-Keras Bridge<sup>11</sup>, but several other frameworks now develop this idea as for instance the Inero library<sup>12</sup> (ECMWF) and the FORTRAN-ML Bridge<sup>13</sup> (ICCS). This approach is adapted to interactive inference (with pre-trained ML-models), and, in principle, to on-line learning. A caveat is that it is specific to existing ML-frameworks, which may be a hazardous choice in such a fast moving technological landscape<sup>14</sup>.
- (d) Leverage a more generic (ML-framework agnostic) interface between FORTRAN and high-abstraction-level languages (as for instance Python or Julia). At the time of writing, the SmartSim library<sup>15</sup> (see also appendix below) seems a promising option that should allow to perform efficient interactive learning, online learning and interactive learning at scale, with a variety of ML-libraries and differentiable programming languages (inc. Jax).

We think that we should opt for option d, because :

- ML is a fast moving field and we should not be attached too tightly to a specific library as the technological landscape may evolve very fast;
- this would allow to optimize the orchestration of the inference on multiple processors more easily, therefore probably better exploit the available resources on hybrid supercomputers;
- this would allow us to investigate online learning strategies (short term), and prepare interactive learning (longer term), therefore opening the possibility to investigate the design of deep emulators.

## 16.6 Status of stochastic physics and probabilistic modelling tools into NEMO

### 16.6.1 Probabilistic modelling with stochastic physics and ensembles

A large variety of applications rely on a probabilistic representation of the ocean state. Such representations aim at modelling explicitly uncertainties and their propagation. These uncertainties can be due to unresolved processes, to model resolution, to model parameters or to the model underlying assumptions (see e.g. Palmer, 2012; Berner et al., 2017). In practice, this is usually achieved with a combination of stochastic physics schemes and ensemble simulations (Bessieres et

---

<sup>10</sup> See [https://github.com/nbren12/call\\_py\\_fort](https://github.com/nbren12/call_py_fort) and the associated blog post ar <https://www.noahbrenowitz.com/post/calling-fortran-from-python/>

<sup>11</sup> See <https://github.com/Cambridge-ICCS/fortran-ml-bridge>

<sup>12</sup> See <https://github.com/ecmwf-projects/infero>

<sup>13</sup> See <https://github.com/Cambridge-ICCS/fortran-ml-bridge>

<sup>14</sup> Although the recent creation of the PyTorch foundation (see <https://pytorch.org/foundation>) under the Linux foundation provides a bit more long term visibility to PyTorch.

<sup>15</sup> See <https://github.com/CrayLabs/SmartSim> and the appendix below.

al. 2017). Ensembles can be generated with perturbations of the model initial states, or perturbations of the model parameters according to some explicit assumptions as to the uncertainty of these parameters. In addition, ensemble simulations can leverage procedures for accounting for the uncertain nature of the representation of physical processes in physics-based models. These procedures, known as *stochastic physics schemes*, use stochastic processes for describing the probability distribution of possible system states given the assumption as to the uncertainty in the formulation of the models. The most commonly used perturbation schemes are the stochastically perturbed parametrization tendencies (SPPT) where model tendencies are perturbed collinearly to their unperturbed tendencies; stochastically perturbed parameters (SPP) where model parameters are modulated stochastically in space and/or time; stochastic kinetic energy backscatter (SKEB) schemes that aim to stochastically re-inject at larger scales the kinetic energy dissipated at the grid scale (Storto et al. 2021; Berner et al. 2017). More process-based perturbations have also been proposed for oceanic applications (such as for instance the uncertainty of small-scale density fluctuations on large-scale horizontal density gradients, see Brankart 2013 and Stanley et al. 2022). Interestingly some ML based eddy closures are also formulated as stochastic physics schemes (eg. Guillaumin and Zanna 2021). Some practical applications may also leverage stochastic physics schemes without ensemble simulations.

#### 16.6.2 Status of probabilistic modelling tools into NEMO

Up to now, the evolution of probabilistic tools used for describing model uncertainties in NEMO has not been coordinated at the NEMO developers committee level, although some application use-cases of these tools fall into the scope of previous NEMO WGs on Data assimilation and Data interface. In the absence of a pro-active planning process and clear development roadmap in this area, only a small fraction of the developments undertaken in the NEMO ecosystem have actually been incorporated into NEMO official releases.

In terms of stochastic physics schemes, the NEMO model, since version 3.6, embeds a stochastic physics scheme (STO) based on the work of Brankart (2013) and Brankart et al. (2015). However, several other packages have been developed over the past years for a variety of target applications (e.g. Andrejczuk et al., 2016; Juricke et al., 2017; Perezhogin et al., 2019; Storto and Andriopoulos, 2021, Leroux et al. 2022). None of the later developments has yet been included into NEMO releases, although some of them are of crucial importance for the forecasting community.

Tools for orchestrating the production of ensemble simulations exist in several groups in the NEMO users community but the NEMO codebase is not officially equipped with such functionality. Indeed in most practical use-cases, ensembles are generated through external orchestration protocols, generally embedded in broader data assimilation workflows. But some application use-cases require communication and IOs across ensemble members, which is notably simplified if the ensemble production is steered from a single executable. Bessieres et al (2017) have proposed a FORTRAN implementation of such code, which allows to produce ensemble simulations and compute ensemble statistics from a single NEMO executable. This approach leverages an extra layer of MPI communications. This code is currently used in several groups in the NEMO users community but is not officially supported in NEMO releases.

### 16.6.3 Analysis of the needs

Based on the above elements and on the consortium requirements, we identify the need to plan the evolution of probabilistic modelling capabilities in NEMO, in order (i) to foster their broader adoption by NEMO users and (ii) to prepare the use of ML in this context. We stress that this need is partly, but not exclusively, driven by the needs of the operational applications.

In terms of stochastic physics schemes, there is a short term need to survey and rationalize the packages available in the community and to analyse what part of them is generic enough to be ported sustainably into future NEMO releases. On the ensemble side, we identify the need to clarify whether the computation of ensemble statistics should eventually be orchestrated from NEMO itself, and if so, how the communications should be orchestrated. We note the SmartSim solution to interface ML frameworks to NEMO could probably be leveraged to this purpose.

Because these questions have not been officially covered into the NEMO development process up until now, the main short term milestone of the proposed working group will be to establish a rationale and a roadmap for NEMO development in this area (A2.1). We hope that, as the activity of the group develops, more practical actions will be identified and proposed for funding.

## 16.7 Implementation plan and identified issues

### 16.7.1 Large scale impacts on the NEMO codebase

The contributions to the NEMO code-base envisioned in this chapter for the coming years will be mostly local and should not lead to large scale code changes. Still, the implementation of an ML interface could require changes in the code API in the future. On a longer time scale, the emulation of specific code components would indeed benefit from further improving the modularity and the orthogonality of individual NEMO code components when possible. In that sense, the vision underlying this chapter is to *consider NEMO as a system*, that can be used as a component of wider systems. This implies in particular that NEMO develops and maintains clear and stable protocols (APIs) for inputting and outputting information with other systems. The integration of ML into NEMO would also benefit from considering each sub-component of NEMO as independent systems themselves whenever possible.

### 16.7.2 Resources and funding assumptions

The strategy described in this chapter will leverage resources provided by several (funded) research projects. This includes in particular :

- M2LINES [2021-2025] : project funded by VESRI, which will develop ML-based submesoscale closures for NEMO.
- EDITO-Model Lab [2023-2025] : project funded by Horizon Europe which will explore deep emulation strategies of NEMO components in idealized set-up.
- MEDIATION [2022-2026] : project funded by the (french agency) ANR, which will explore and benchmark options for coupling ocean models with ML frameworks.
- TRACCS [2023-2030] : project funded by the (french agency) ANR, which will implement the NEMO ML interface and explore approaches for deep emulation of NEMO components.
- ERGO2 [2022-2025] : project funded by the Copernicus Climate Change Service (call C3S2\_601), which will focus on the update of the stochastic physics schemes, data-driven

systematic model error corrections, and update of the tangent-linear and adjoint version of NEMO.

- ACCIBERG [2023-2026] project funded by Horizon Europe which will develop a stochastic physics scheme for the sea-ice model component of NEMO.

We provide below an *unconsolidated* list of other on-going or planned activities involving embedding ML into physical models that we are aware of and that may have implications for NEMO:

- Representation of fine scale air-sea interactions (F. Lemarié, INRIA, Grenoble; as part of the INRIA-AirSea/ATOS collaboration)
- Coupling interface between ice sheet models and coarse resolution ocean model (N. Jourdain, IGE, Grenoble; as part of the ANR AIAI project)
- Surface wave model emulation (Oxford, ECMWF)
- Improving lateral diffusion in NEMO with ML (Tallinn, TallTech)
- Several initiatives are also ongoing at CMCC for integrating ML with NEMO

This list is still preliminary and will be consolidated as the activities of the NEMO working group on Machine Learning and model uncertainties develop.

### 16.7.3 Community and capacity building

Besides the practical milestones listed in this chapter, we have also identified some actions that would foster NEMO community and capacity building within the next 5 years. Progress on these topics will be reported back to the NEMO developers committee.

- Establish new external collaborations on questions related to the orchestration of online and interactive learning (eg with INRIA/DataMove)
- Establish new external collaborations on the methodologies for deep emulation (eg with ECMWF, Hereon)
- Establish long term collaborations with super-computer vendors (eg. HPE/Cray, ATOS/Bull, NVidia, ...) on questions related to ML libraries/performance (to be coordinated with the NEMO HPC WG)
- Organize dedicated training on ML/DL and ML/HPC to NEMO developers and other interested parties in the NEMO ecosystem.

### 16.7.4 Cross-cutting issues

We provide below a list of cross-cutting questions that will require specific coordination with other NEMO working groups and developers in the NEMO broader ecosystem.

- Investigating whether deep emulation is an interesting option for maintaining a linear tangent model for NEMO will require interactions with the ocean variational data assimilation community (NEMOVAR, OceanVar, MOI DA expert team).
- The proof-of-concept practical implementation of a ML-based subgrid parameterization in NEMO will a priori focus on ocean macro-turbulence. This will require coordination with the NEMO Eddy closure working group.

- Defining how to maintain and deliver code components encoded as NN will require interactions with the NEMO Tools, Verification & Validation WG and NEMO Community Engagement WG.
- Investigating whether deep emulation is a viable strategy for porting NEMO to various architectures will require dedicated liaison with the NEMO HPC WG, in particular for identifying the proof of concept demonstration.
- The use of SmartSim is somewhat redundant with some of the functionalities provided by XIOS2. The definition of an ML interface for NEMO should therefore be discussed with XIOS developers in order to define a coherent roadmap.
- The planned activities for improving the coupling interface between ice-sheet models and coarse resolution ocean models with ML will require specific liaison with the NEMO Land Ice / Ocean WG to be established.

## 16.8 Appendices :

### **Machine Learning**

Machine learning (ML) is a vast field, which entails a large range of methods and algorithms for building numerical codes that learn how to accomplish their tasks. The behavior of the resulting numerical codes is therefore not prescribed a priori, but rather depends on parameters that should be learned in order for the numerical code to meet a prescribed objective. ML algorithms can be leveraged for different sorts of tasks as for instance : clustering, dimensionality reduction, classification or regression problems. ML entails a vast zoo of methods which differ in speed and accuracy. A key dimension of the modern ML landscape is that ML algorithms are encoded in ML libraries as for instance PyTorch, TensorFlow, Scikit-learn, which allows to easily combine and reuse pre-existing building blocks for solving new problems.

### **Neural networks**

A neural network (NN) is a specific type of machine consisting of a series of mathematical operators (called “layers”) which parameters can be trained in order to meet a prescribed training objective. Each layer combines an affine transformation (defined by its weights and biases as parameters) and a nonlinear operator (called activation). This structure implies that all the layers are piecewise differentiable. One can therefore compute analytically the derivative of the training objective with respect to the network parameters (weights and biases). Neural networks are generally encoded in dedicated software packages (eg. PyTorch, TensorFlow). The automatic differentiation and the gradient descent algorithms available in these packages allow optimizing the NN parameters on any training objective.

### **Differentiable programming**

Differentiable programming (DP) is an emerging paradigm at the interface between scientific computing, computational physics and machine learning in which software modules can be trained with gradient-based optimization (Lavin et al. 2021; Thuerey et al. 2021). DP allows the design of programs that can be tuned to achieve a given objective, provided the derivative of the program can

be computed. It can be seen as a generalization of Deep Learning (Goodfellow et al. 2016), where programs can be composed of NN building blocks and procedural programming in arbitrary algorithmic structures using control flow. This paradigm allows to augment procedural codes with trainable components or to constrain trainable components with physical or mathematical priors, similarly to Physics Informed Deep Learning techniques (Raissi et al. 2019). These approaches allow the design of simulators that can be trained *end-to-end* and generally lead to algorithms that can be trained with less data and generalize better to unseen conditions. In practice, DP requires system-wide automatic differentiation, which is readily available in DL frameworks or some high-level languages (as Julia or Jax). Oceananigans (Ramadhan et al. 2020) and Veros (Häfner et al. 2021) are two examples of modern ocean models allowing to leverage DL, while achieving good performance on GPU.

### The SmartSim library

This open-source library developed by HPE/Cray-lab is one of the several available software options for interfacing pre-existing scientific codes with ML libraries (e.g. TensorFlow or PyTorch). SmartSim (<https://github.com/CrayLabs/SmartSim>) is specifically aiming at providing a lightweight, non-intrusive and efficient interface for C, C++ or FORTRAN simulators using MPI. SmartSim relies on an in-memory data structure store (Redis) which allows diskless IOs. SmartSim has been used with the MOM6 ocean model for online inference (Partee et al. 2022). This library could also be used for orchestrating the production of (interactive) ensemble simulations or for outsourcing computations on GPUs (e.g. diagnostics). Several other open source libraries for orchestrating complex HPC workflows and ensemble simulations are also available, including for instance Melissa developed by INRIA (<https://gitlab.inria.fr/melissa>).

## 16.9 References

- Allen, M., Kettleborough, J., & Stainforth, D. (2003). Model error in weather and climate forecasting. In Seminar on Predictability of weather and climate. <https://www.ecmwf.int/node/7686>
- Andrejczuk, M., Cooper, F. C., Juricke, S., Palmer, T. N., Weisheimer, A., & Zanna, L. (2016). Oceanic stochastic parameterizations in a seasonal forecast system. *Monthly Weather Review*, 144(5), 1867–1875. <https://doi.org/10.1175/MWR-D-15-0245.1>
- Berner, J., et al. (2017). Stochastic parameterization: Toward a new view of weather and climate models. *Bulletin of the American Meteorological Society*, 98(3), 565–588. <https://doi.org/10.1175/BAMS-D-15-00268.1>
- Bessières, L., Leroux, S., Brankart, J.-M., Molines, J.-M., Moine, M.-P., Bouttier, P.-A., Penduff, T., Terray, L., Barnier, B., & Sérazin, G. (2017). Development of a probabilistic ocean modelling system based on NEMO 3.5: Application at eddying resolution. *Geoscientific Model Development*, 10(3), 1091–1106. <https://doi.org/10.5194/gmd-10-1091-2017>

Bolton, T., & Zanna, L. (2019). Applications of Deep Learning to Ocean Data Inference and Subgrid Parameterization. In *Journal of Advances in Modeling Earth Systems* (Vol. 11, Issue 1, pp. 376–399). American Geophysical Union (AGU). <https://doi.org/10.1029/2018ms001472>

Bonavita, M., & Laloyaux, P. (2020). Machine learning for model error inference and correction. *Journal of Advances in Modeling Earth Systems*, 12(12). <https://doi.org/10.1029/2020MS002232>

Brankart, J.-M. (2013). Impact of uncertainties in the horizontal density gradient upon low resolution global ocean modelling. *Ocean Modelling*, 66, 64–76. <https://doi.org/10.1016/j.ocemod.2013.02.004>

Brankart, J.-M., Candille, G., Garnier, F., Calone, C., Melet, A., Bouttier, P.-A., Brasseur, P., & Verron, J. (2015). A generic approach to explicit simulation of uncertainty in the NEMO ocean model. *Geoscientific Model Development*, 8(5), 1285–1297. <https://doi.org/10.5194/gmd-8-1285-2015>

Brunton, S. L., Noack, B. R., & Koumoutsakos, P. (2020). Machine Learning for Fluid Mechanics. In *Annual Review of Fluid Mechanics* (Vol. 52, Issue 1, pp. 477–508). Annual Reviews. <https://doi.org/10.1146/annurev-fluid-010719-060214>

Carleo, G., Cirac, I., Cranmer, K., Daudet, L., Schuld, M., Tishby, N., Vogt-Maranto, L., & Zdeborová, L. (2019). Machine learning and the physical sciences. *Reviews of Modern Physics*, 91(4), 045002. <https://doi.org/10.1103/RevModPhys.91.045002>

Champion, K., Lusch, B., Kutz, J. N., & Brunton, S. L. (2019). Data-driven discovery of coordinates and governing equations. In *Proceedings of the National Academy of Sciences* (Vol. 116, Issue 45, pp. 22445–22451). Proceedings of the National Academy of Sciences. <https://doi.org/10.1073/pnas.1906995116>

Chantry, M., Hatfield, S., Dueben, P., Polichtchouk, I., & Palmer, T. (2021). Machine learning emulation of gravity wave drag in numerical weather forecasting. *Journal of Advances in Modeling Earth Systems*, 13(7). <https://doi.org/10.1029/2021MS002477>

Cleary, E., Garbuno-Inigo, A., Lan, S., Schneider, T., & Stuart, A. M. (2021). Calibrate, emulate, sample. *Journal of Computational Physics*, 424, 109716. <https://doi.org/10.1016/j.jcp.2020.109716>

Couvreux, F. et al. (2021). Process-Based Climate Model Development Harnessing Machine Learning: I. A Calibration Tool for Parameterization Improvement. In *Journal of Advances in Modeling Earth Systems* (Vol. 13, Issue 3). American Geophysical Union (AGU). <https://doi.org/10.1029/2020ms002217>

Curcic, M. (2019). A parallel Fortran framework for neural networks and deep learning. *ACM SIGPLAN Fortran Forum*, 38(1), 4–21. <https://doi.org/10.1145/3323057.3323059>

Falls, M., Bernardello, R., Castrillo, M., Acosta, M., Llorca, J., & Galí, M. (2021). Use of genetic algorithms for ocean model parameter optimisation [Preprint]. *Biogeosciences*. <https://doi.org/10.5194/gmd-2021-222>



Frezat, H., Balarac, G., Le Sommer, J., Fablet, R., & Lguensat, R. (2021). Physical invariance in neural networks for subgrid-scale scalar flux modeling. *Physical Review Fluids*, 6(2), 024607.

<https://doi.org/10.1103/PhysRevFluids.6.024607>

Frezat, H., Le Sommer, J., Fablet, R., Balarac, G., & Lguensat, R. (2022). A posteriori learning for quasi-geostrophic turbulence parametrization. *Journal of Advances in Modeling Earth Systems*.

<https://doi.org/10.1029/2022MS003124>

Goodfellow, I., Bengio, Y., & Courville, A. (2016). *Deep learning*. The MIT Press.

Guillaumin, A. P., & Zanna, L. (2021). Stochastic-deep learning parameterization of ocean momentum forcing. *Journal of Advances in Modeling Earth Systems*, 13(9).

<https://doi.org/10.1029/2021MS002534>

Grönquist, P., Yao, C., Ben-Nun, T., Dryden, N., Dueben, P., Li, S., & Hoefler, T. (2021). Deep learning for post-processing ensemble weather forecasts. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 379(2194), 20200092.

<https://doi.org/10.1098/rsta.2020.0092>

Häfner, D., Nuterman, R., & Jochum, M. (2021). Fast, cheap, and turbulent—Global ocean modeling with gpu acceleration in python. *Journal of Advances in Modeling Earth Systems*, 13(12).

<https://doi.org/10.1029/2021MS002717>

Hatfield, S., Chantry, M., Dueben, P., Lopez, P., Geer, A., & Palmer, T. (2021). Building Tangent-Linear and Adjoint Models for Data Assimilation With Neural Networks. In *Journal of Advances in Modeling Earth Systems* (Vol. 13, Issue 9). American Geophysical Union (AGU).

<https://doi.org/10.1029/2021ms002521>

Juricke, S., Palmer, T. N., & Zanna, L. (2017). Stochastic subgrid-scale ocean mixing: Impacts on low-frequency variability. *Journal of Climate*, 30(13), 4997–5019. <https://doi.org/10.1175/JCLI-D-16-0539.1>

Kasim, M. F., Watson-Parris, D., Deaconu, L., Oliver, S., Hatfield, P., Froula, D. H., Gregori, G., Jarvis, M., Khatiwala, S., Korenaga, J., Topp-Mugglestone, J., Viezzer, E., & Vinko, S. M. (2022). Building high accuracy emulators for scientific simulations with deep neural architecture search. *Machine Learning: Science and Technology*, 3(1), 015013. <https://doi.org/10.1088/2632-2153/ac3ffa>

Kochkov, D., Smith, J. A., Alieva, A., Wang, Q., Brenner, M. P., & Hoyer, S. (2021). Machine learning–accelerated computational fluid dynamics. In *Proceedings of the National Academy of Sciences* (Vol. 118, Issue 21, p. e2101784118). *Proceedings of the National Academy of Sciences*.

<https://doi.org/10.1073/pnas.2101784118>

Lary, D. J., Alavi, A. H., Gandomi, A. H., & Walker, A. L. (2016). Machine learning in geosciences and remote sensing. In *Geoscience Frontiers* (Vol. 7, Issue 1, pp. 3–10). Elsevier BV.  
<https://doi.org/10.1016/j.gsf.2015.07.003>

Lavin, A., Zenil, H., Paige, B., Krakauer, D., Gottschlich, J., Mattson, T., Anandkumar, A., Choudry, S., Rocki, K., Baydin, A. G., Prunkl, C., Paige, B., Isayev, O., Peterson, E., McMahon, P. L., Macke, J., Cranmer, K., Zhang, J., Wainwright, H., ... Pfeffer, A. (2021). Simulation intelligence: Towards a new generation of scientific methods. ArXiv:2112.03235 [Cs]. <http://arxiv.org/abs/2112.03235>

Leroux, S., Brankart, J.-M., Albert, A., Brodeau, L., Molines, J.-M., Jamet, Q., Le Sommer, J., Penduff, T., & Brasseur, P. (2022). Ensemble quantification of short-term predictability of the ocean dynamics at kilometric-scale resolution: A Western Mediterranean test-case. <https://doi.org/10.5194/os-2022-11>

Magiera, J., Ray, D., Hesthaven, J. S., & Rohde, C. (2020). Constraint-aware neural networks for Riemann problems. In *Journal of Computational Physics* (Vol. 409, p. 109345). Elsevier BV.  
<https://doi.org/10.1016/j.jcp.2020.109345>

Nonnenmacher, M., & Greenberg, D. S. (2021). Deep Emulators for Differentiation, Forecasting, and Parametrization in Earth Science Simulators. In *Journal of Advances in Modeling Earth Systems* (Vol. 13, Issue 7). American Geophysical Union (AGU). <https://doi.org/10.1029/2021ms002554>

Ott, J., Pritchard, M., Best, N., Linstead, E., Curcic, M., & Baldi, P. (2020). A fortran-keras deep learning bridge for scientific computing. *Scientific Programming*, 2020, 1–13.  
<https://doi.org/10.1155/2020/8888811>

Palmer, T. N. (2012). Towards the probabilistic Earth-system simulator: A vision for the future of climate and weather prediction: Towards the Probabilistic Earth-System Simulator. *Quarterly Journal of the Royal Meteorological Society*, 138(665), 841–861. <https://doi.org/10.1002/qj.1923>

Partee, S., Ellis, M., Rigazzi, A., Shao, A. E., Bachman, S., Marques, G., & Robbins, B. (2022). Using Machine Learning at scale in numerical simulations with SmartSim: An application to ocean climate modeling. *Journal of Computational Science*, 62, 101707.  
<https://doi.org/10.1016/j.jocs.2022.101707>

Perezhogin, P. (2019) Deterministic and stochastic parameterizations of kinetic energy backscatter in the NEMO ocean model in double-gyre configuration. *IOP Conference Series. Earth and Environmental Science*, 386, 012025. <https://doi.org/10.1088/1755-1315/386/1/012025>

Ramadhan, A., Wagner, G., Hill, C., Campin, J.-M., Churavy, V., Besard, T., Souza, A., Edelman, A., Ferrari, R., & Marshall, J. (2020). Oceananigans.jl: Fast and friendly geophysical fluid dynamics on GPUs. *Journal of Open Source Software*, 5(53), 2018. <https://doi.org/10.21105/joss.02018>

Raissi, M., Perdikaris, P., & Karniadakis, G. E. (2019). Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *Journal of Computational Physics*, 378, 686–707.

<https://doi.org/10.1016/j.jcp.2018.10.045>

Salter, J. M., & Williamson, D. (2016). A comparison of statistical emulation methodologies for multi-wave calibration of environmental models. *Environmetrics*, 27(8), 507–523.

<https://doi.org/10.1002/env.2405>

Schneider, T., Lan, S., Stuart, A., & Teixeira, J. (2017). Earth system modeling 2. 0: A blueprint for models that learn from observations and targeted high-resolution simulations. *Geophysical Research Letters*, 44(24). <https://doi.org/10.1002/2017GL076101>

Schulz, B., & Lerch, S. (2021). Machine learning methods for postprocessing ensemble forecasts of wind gusts: A systematic comparison. *ArXiv:2106.09512 [Physics, Stat]*.

<https://doi.org/10.1175/MWR-D-21-0150.1>

Stanley, Z., Grooms, I., Kleiber, W., Bachman, S. D., Castruccio, F., & Adcroft, A. (2020).

Parameterizing the impact of unresolved temperature variability on the large-scale density field: Part 1. Theory. *Journal of Advances in Modeling Earth Systems*, 12(12).

<https://doi.org/10.1029/2020MS002185>

Stengel, K., Glaws, A., Hettinger, D., & King, R. N. (2020). Adversarial super-resolution of climatological wind and solar data. *Proceedings of the National Academy of Sciences*, 117(29), 16805–16815. <https://doi.org/10.1073/pnas.1918964117>

Storto, A., & Andriopoulos, P. (2021). A new stochastic ocean physics package and its application to hybrid-covariance data assimilation. *Quarterly Journal of the Royal Meteorological Society*, 147(736), 1691–1725. <https://doi.org/10.1002/qj.3990>

Thuerey, N., Holl, P., Mueller, M., Schnell, P., Trost, F., & Um, K. (2021). Physics-based deep learning.

<https://doi.org/10.48550/ARXIV.2109.05237>

Williamson, D. B., Blaker, A. T., & Sinha, B. (2017). Tuning without over-tuning: Parametric uncertainty quantification for the NEMO ocean model. *Geoscientific Model Development*, 10(4), 1789–1816. <https://doi.org/10.5194/gmd-10-1789-2017>

Zanna, L., & Bolton, T. (2020). Data-Driven Equation Discovery of Ocean Mesoscale Closures. In *Geophysical Research Letters* (Vol. 47, Issue 17). American Geophysical Union (AGU).

<https://doi.org/10.1029/2020gl088376>

Zanna, L., & Bolton, T. (2021). Deep learning of unresolved turbulent ocean processes in climate models. In G. Camps-Valls, D. Tuia, X. X. Zhu, & M. Reichstein (Eds.), *Deep Learning for the Earth Sciences* (1st ed., pp. 298–306). Wiley. <https://doi.org/10.1002/9781119646181.ch20>

Zhuang, J., Kochkov, D., Bar-Sinai, Y., Brenner, M. P., & Hoyer, S. (2021). Learned discretizations for passive scalar advection in a two-dimensional turbulent flow. *Physical Review Fluids*, 6(6), 064605.

<https://doi.org/10.1103/PhysRevFluids.6.064605>

## 17 Data Assimilation interfaces

WGL: Dan Lea (Met Office)

WG members: Ali Aydogdu (CMCC), Mike Bell (Met Office), Jean-Michel Brankart (CNRS), Marcin Chrust (ECMWF), Pierre De Mey- (LEGOS), David Ford (Met Office), Matt Martin (Met Office), Elisabeth Remy (Mercator-Ocean), Giovanni Ruggiero (Mercator-Ocean), Andrea Storto (CNR), Jennifer Waters (Met Office), Anthony Weaver (Cerfacs)

A Data Assimilation Interfaces WG has coordinated the development of the OBS and ASM sections of the NEMO code for many years.

The observation operator utility in OBS has been developed within NEMO over more than 10 years and maintained primarily by the Met Office. Unfortunately, Mercator and CMCC have not made use of this code (though ECMWF does) and the research community has made little use of it. [NOC recently developed a COaST assessment tool; which calculates model / observation differences; it does not use the observation operator code]. In order to avoid internal duplication of effort, the Met Office will probably transfer to an observation operator based on the JEDI system, so there may not be a group willing to maintain the observation operator tool in future.

The OBS code needs to be updated to accommodate the new equation of state (TEOS10) as well as EOS80.

The incremental analysis update (IAU) code within ASM needs to be tidied up and adapted for RK3. This should be completed before the end of 2022. Again, it is not clear which groups use this code, but the Met Office will maintain the primary IAU option.

Is there a requirement for a tangent linear (or linearised perturbation) version of NEMO? The TAM code has not been updated since Version 3.6. but the NEMOVAR consortium have plans to develop 4DVAR capabilities using a more recent version of NEMO (probably 4.0.X) which will involve updating the TAM to be relevant to this later version. The new MLMU WG will explore options for differentiable emulation of NEMO. This may offer a viable (and more general) alternative to NEMOTAM.

The Mercator Ocean International Marine Data Assimilation (MDA) Expert Team has set up a working group to assess the feasibility of developing a shared MDA framework. The first approach to assist shared development considered by the working group is the potential of the JEDI and PDAF frameworks. The interface between NEMO and PDAF puts few constraints on the NEMO model, though to use NEMO within PDAF one would need to change the top-level NEMO subroutines. Were JEDI to be used as the framework, and the NEMO model itself run within JEDI, a JEDI-compatible interface between NEMO and the assimilation software would need to be developed and maintained. One would need to be able to restart the NEMO model from and output it as data in the form of the assimilation software's representation of the NEMO state vector.

A second approach to developing a shared MDA framework is to agree on how the main objects used by the data assimilation codes should be represented and to transition toward them. The geometry of the model grid and the model state vector are two of the main objects. The arrays constituting these objects within NEMO can be encapsulated (wrapped up) within Fortran90 derived types. Introducing the interface between these derived types and NEMO's internal representation of them into the NEMO code-base would define a standard for the representation of these objects

within data assimilation codes. This could be done within an assimilation directory (either the ASM directory or a new one). A SETTE test-case to verify this interface at each NEMO release would be highly desirable. This approach would support the use of NEMO for data assimilation with JEDI.

We have decided that, at least for now, the DAI WG will continue to coordinate the work on the OBS and ASM sections of NEMO and take the lead on defining interfaces between NEMO and data assimilation systems (such as NEMOVAR, OceanVAR and SAM). It should do this in close collaboration with the MLMU WG.

## 18 Tools, Verification and Validation

### 18.1 Executive Summary

Tools to build NEMO and check whether updates to the NEMO trunk change the results are well established. This group will develop such essential NEMO support tools in line with the priorities identified in the NEMO V&V roadmap<sup>16</sup>. The plans are:

- **Within 2 years:** To have the system team fully proficient with all aspects of a GitLab-based development environment. Including: the use of GitLab runners for automated tasks and continuous integration and regular updating of Wiki and web-based support material
- **Within 5 years:** To have Unit-testing capabilities in most code areas. Frequent, automated testing of the code base. Full support for exascale and heterogeneous computing environments (e.g. GPU co-processors) via Domain Specific Language pre-processing tools.
- **Within 10 years:** To have code testing carried out by AI-enabled agents

Working Group members: Andrew Coward (Chair and chapter lead), Mike Bell, Claire Levy, Simon Müller, Nicolas Martin, Daley Calvert, Sibylle Techene, Jerome Chanut, Julien le Sommer, Sebastien Masson

### 18.2 Introduction

Users require access to a number of tools and guidance material about them in order to make effective and efficient use of NEMO. With so many potential applications and use cases, it is important to determine what tools can be supported as part of a central development strategy. At the very least, the strategy must empower the community to support itself by providing frameworks for the discovery and retention of the supporting software and documentation for these tools.

The previous NDS contained a NEMO validation and user support chapter that recognised that new tools and platforms were needed to improve the user experience and ensure robust testing before releases. The Verification and Validation (V&V) WG, set up in 2018, developed a [NEMO V&V roadmap](#), which provides a comprehensive assessment of the status, opportunities and priorities for improvement of the tools and processes for verification and validation by NEMO developers and the NEMO System Team. It includes proposals for relatively short-term development within a fairly general framework. Expanding the role of the V&V WG to cover tools development is a natural progression because improvement of the verification methods now requires development of the tools and there would be significant overlap between the work and membership of the V&V and Tools WGs if they were separate groups.

**[Action: Expand remit of V&V WG to cover all actions arising from this chapter]**

---

<sup>16</sup> [forge.ipsl.jussieu.fr/nemo/attachment/wiki/WorkingGroups/Verification/nemo-validation-and-verification-roadmap\\_draft\\_1.0\\_Nov\\_2020.pdf](https://forge.ipsl.jussieu.fr/nemo/attachment/wiki/WorkingGroups/Verification/nemo-validation-and-verification-roadmap_draft_1.0_Nov_2020.pdf)

This chapter describes the current NEMO tools and V&V processes and their proposed evolution. Areas covered include:

- Code repositories and web information systems (section 4)
- Build tools and basic code verification tools (section 5)
- Additional configuration support tools (sections 6 & 7)
- Containerisation as a method for providing 'ready to run' models (section 8)

Ambitions for NEMO applications continue to grow to encompass more complex and larger scenarios whilst hardware platforms and operating environments also become more varied and subject to frequent changes. With this background, more thorough and routine testing of the developing code base is required. The rest of this chapter outlines plans, based on the [NEMO V&V roadmap](#), for providing such testing. Options discussed in sections 9-11 include:

- Improvements to the System Team's basic regression testing (SETTE)
- Use of Continuous Integration facilities within the new GitLab framework
- A phased introduction of unit testing

NEMO development has an established workflow including some relevant verification and validation tests. Still, the need to improve this V&V shared practices is recognized by all developers as a high priority, both for developers (to facilitate and accelerate collaborative developments, and reduce the time spent on bug fixing) and for the users (to make the future releases more reliable). The methodology improvements proposed in this chapter will improve NEMO and its development. However, they will also - even with more automatic processes - increase the time needed to finalise a development and include it in the NEMO reference. The choices made need to take into account their impact on the rate of progress by and the resources that need to be allocated to the NST.

### 18.3 Code repositories and web information systems

The most basic requirement is that users and developers have access to version-controlled code in a public facing repository. The enforced (but timely) move from the IPSL-hosted subversion server to a GitLab server hosted by Mercator Ocean has provided a modern, long-lasting solution to the basic need and supports additional services. This move complies with the recommendation of the V&V WG.

This repository remains the central location for the traditional release branches and main trunk. However, the Git-enabled ability to fork the entire repository into an independent (but still traceable) copy means that even the most speculative or fringe developments can maintain a managed pathway back into the central repository without impacting the central repository until deemed worthy.

Sites that have, in the past, maintained mirrored copies with local developments that rapidly slip out of phase with central changes should find better options for maintaining local developments whilst retaining the ability to return those developments to the central copy.

Beyond the basics, the main repository also needs to support bug reporting and feature requests with clear allocation of responsibilities and traceable resolutions. The GitLab issues functionality provides all the necessary infrastructure and it is unlikely any external solutions will be required for these aspects. Issue boards may also be used to pair issue tracking and project management for organising the annual workplan.

GitLab's Wiki capabilities will be used to host project-based information as a replacement for the previous Trac-based wiki pages. Furthermore, the purchase of a static IP that has been associated with

the NEMO GitLab service (<https://sites.nemo-ocean.io>), means that GitLab runners can be used to build and deploy webpages automatically as RST files as soon as they are pushed onto the repository. Importantly, these web pages are publicly available and are already being used to improve the user experience by, for example, providing a user-guide which is tied directly to each release (e.g.: <https://sites.nemo-ocean.io/user-guide/>).

External services will continue to be used for user forums (discourse) and conversational, topic-based, exchanges between developers (zulip).

#### 18.4 Makenemo and SETTE

The makenemo script will continue to be maintained as part of the system and will retain its underlying use of the FCM build system. The subset of FCM components actually used will be held as part of the NEMO repository to protect against external developments of FCM that may affect its compatibility. Architecture files in the form of templates for common compilers and versions for the major supercomputing platforms available to consortium members will also continue to be maintained.

The introduction of Domain Specific Language pre-processing is expected during the development period. Tools such as PSyclone have already been tested with NEMO and are likely to provide the best method of supporting NEMO on heterogeneous HPC platforms by, for example, refactoring code for use on GPUs. A method of adding PSyclone as a pre-processing step within makenemo has been developed and will form part of the default provision.

#### **[Action: incorporate PSyclone into the makenemo build system]**

SETTE remains an important component of the development process. A SETTE suite capable of performing necessary tests on supported reference configurations on at least one HPC platform available to each consortium member will be provided. Developments for NEMO will continue to be considered ineligible for inclusion if they fail any SETTE tests.

SETTE has evolved into a more versatile testing system that can be adapted for tasks outside the required testing regime. Some further, limited development of its capabilities may continue in coordination with the recommendations of the Verification and Validation Working Group.

Other improvements to the SETTE testing proposed by the V&V WG include:

- o Use of containers to make SETTE easier to transfer between systems.
- o Use of cylc (or a similar tool) to allow SETTE to test a wider range of options or to be run more quickly (using a large number of parallel tasks)
- o Re-organisation to make it easier to incorporate demonstration cases within SETTE testing (both by developers and as part of the SETTE tool).

#### **[Action: further evolution of SETTE to make it easier to incorporate additional test configurations]**

GitLab's Continuous Integration capabilities open up the possibility of more rigorous and routine testing. Options for regular deployment of tests to consortium member's HPC platforms via GitLab Runners will be explored. With minimal adaptation, the SETTE suite could be used in such a scenario. Such capabilities may also influence the merge protocols. For example, the successful result of such an automated testing/trusting tool may become a mandatory requirement before a reviewer can accept a new development.



**[Action: explore GitLab Runners for deployment of larger scale, regular testing]**

### 18.5 Critical tools

DOMAINcfg: This tool allows the creation of a domain configuration file (domain\_cfg.nc) containing the ocean domain information required to define an ocean configuration from scratch.

**[Action: Maintain compatibility with any changes in vertical coordinate capabilities]**

WEIGHTS: This directory contains software for generating and manipulating interpolation weights for use with the Interpolation On the Fly (IOF) option. **[No actions expected]**

REBUILD\_NEMO: This is a tool to rebuild NEMO output files from multiple processors (mesh\_mask, restart or XIOS output files) into one file.

**[Action: A refresh may be needed; especially if tiling changes impose a change in restart file organisation]**

### 18.6 External tools

NEMO benefits from sites and packages maintained by individuals and organisations outside of the NST. Important examples include:

CDFTOOLS is a diagnostic package written in fortran 90 for the analysis of NEMO model output, initialized in the frame of the DRAKKAR project (<https://www.drakkar-ocean.eu/>).  
[https://github.com/pmathiot/CDFTOOLS\\_4.0\\_ISF](https://github.com/pmathiot/CDFTOOLS_4.0_ISF)

SOSIE is an interpolation package for NetCDF files (full description available <https://brodeau.github.io/sosie/>). It is mainly used to pre-process files in order to start a NEMO simulation or to post-process output files.

**[Action: Maintain links to such externals in all documentation and user-guides]**

### 18.7 Containers

The almost total control of NEMO via its namelists provides opportunities for NEMO to be deployed in containerised environments. A totally containerised version is of little interest, other than for demonstration purposes, but a containerised NEMO with externally mounted experiment directories may appeal to non-traditional users. This has been proposed as a way of empowering local users in developing countries to configure and use limited area models without needing all the usual computing infrastructure. Some container systems, such as Singularity, even allow containerised executables to link to external libraries. This makes containers a serious option even within established communities since a single compiled case can be moved between platforms and still utilise libraries that are optimised for specific hardware.

If support for containerisation continues, future releases of NEMO may include containerised versions, precompiled for the most common processor architectures. Containers may also enable wider testing as part of a Continuous Integration system by offering simpler deployment to a range of diverse HPC platforms.

Containerisation also offers the prospect of “reproducible science” since it is possible to provide a complete containerised solution that is guaranteed to reproduce published results. These ideas are

explored further in a presentation by James Harle available [here](#) which even includes links to a GitHub-based repository with an associated workflow that can be used to build your own singularity containers with the minimum of fuss.

**[Action: Explore options for providing containerised versions with each official release]**

### 18.8 Unit testing

One recommendation from the [NEMO V&V roadmap](#) was to seek an approach to enable the incremental introduction of unit testing of the NEMO subroutines/modules: A unit testing framework could be developed to support tests at a module level. Simple test coding conventions such as a 'test\_' prefix for each test could enable a unit testing framework to be used within the SETTE system to parse the code, orchestrate the subroutine tests and extract and analyse the results (e.g. differences from the expected results). The unit testing framework would need to provide tools to generate inputs (using USR configuration codes and random or analytical fields). Code developers would need to define input data, parameters, expected results and pass/fail criteria. The feasibility of this approach could be studied by trialling it for one or two "representative" modules. One of the main issues to address is how to define and implement the input fields for the representative modules.

**[Action: Explore options for providing a unit testing framework within NEMO]**

### 18.9 Other issues/opportunities

Testing of code "in situ" has significant advantages and could be achieved by running the code using USR configurations. There are choices to be made on whether we need logicals to turn off other processes (as in MPAS), how we document the results (e.g. using more formal methods as in FESOM) and how we incorporate the tests in regression checks. We should experiment with alternative approaches so that we find one that works well (balancing the cost and the importance).

MOM6 have pioneered a number of useful additional tests of the code. For example: checking symmetry properties (solutions should not depend on swapping i and j); checking that there are no inconsistencies in the dimensions of quantities used by the code.

Could we take more advantage from the test/demonstration cases – perhaps in collaboration with the COMMODORE community?

Could we make better use of the "real-world" validation carried out by CMCC, the Met Office and MOi? This was viewed as outside the scope of the V&V roadmap but at the moment it seems to be a missed opportunity (Perhaps it happens "anyway" but could it happen more shortly after new releases; what would be the costs and benefits?) The MOi (Mercator Ocean International) METOF Expert Team could perhaps play a rôle in coordinating this.

## 19 NEMO Community engagement

WG members: Jeffrey Polton (NOC), Katherine Hutchinson (CNRS), Mike Bell (UK MetOffice), Enda O’Dea, Andrew Shao, Daley Calvert

*“Today’s users are tomorrow’s expert contributors”*

### 19.1 Executive Summary

Supporting users and enlarging the group of people using the NEMO ocean model for research and development as well as for operational frameworks is key for the next generation of ocean and climate services and applications. Being a community model, NEMO has already put in place some initiatives, devoted to providing comprehensive information about the mathematical model and numerical schemes as well as instructions to install and run a NEMO configuration. That said, there are many gaps in the available resources and work needs to be done to organize them efficiently and present the material in a practical and accessible way, especially for debutant users.

The scope of this chapter is to develop a roadmap of potential actions for the NEMO community to engage in, so as to further improve dissemination capacity and better connect with the multitude of both potential and existing users. At present, a major challenge for this working group is finding the personnel hours to dedicate to outreach, education and engagement. As such, we (the authors) have attempted to identify the major priorities for the coming years and focus our limited resources on these. Training of new users has been singled out as being a priority at present and we believe that uniting efforts on this front is where the greatest return on investment time can be made. In order to better connect users to the proposed online training resources, we recognize that the website needs to be revamped and the wikipedia page needs to be rehailed. To accompany the users once they are on their feet, there is the NEMO user support forum, Discourse. We plan on providing guidelines for developers to turn their test cases into useful training resources by simply adding some explanation and uploading the resource onto Zenodo where they can get a citable DOI and be recognised for their efforts. We hope to attract more members to this group and plan to add action items accordingly as our support grows.

For the 2023-2027 period we plan to focus on the following items:

- Elaborate training material for debutant users and developers
- Revamp online platforms: website and Wikipedia
- Promote user engagement via the user forum Discourse
- Incentivise proactive and preemptive generation of worked-examples

It is worth noting that the following items are the ongoing responsibility of each member of the NST:

- Regular update of the Reference Manuals with every major development
- Update NEMO users’ guide before each new release

### 19.2 Current Status and Gaps in NEMO Community Engagement

#### 19.2.1 Reference Manuals

As of today, NEMO documentation is maintained by the Consortium members through the System Team and offered to users through the NEMO webpage <https://www.nemo-ocean.eu/>. Here one can access the Reference Manuals, see their description and how to cite them. The manuals include NEMO ocean engine, sea ice, and biogeochemistry components. The documentation has specific DOIs and is available as pdf and as html.

The Reference Manual provides a comprehensive and scientifically-based description of the NEMO ocean model. Each chapter corresponds to a specific submodule of NEMO code and reports a description of the numerical schemes as implemented in the code. This document is managed by the NEMO System Team. In principle each PI should create a ticket on the specific development and, once the action is completed, draft the section/chapter which is then assessed by the internal reviewer before its publication. With this process, the contents of the reference manual are in theory kept updated and in-line with the code release with the developer taking responsibility to update the documentation and the reviewer checking that this is done before approving the merge. The NEMO officer also has a responsibility here to ensure that this is taking place for the main code changes undertaken by their team members. In practice, due to high activity in the developments and the fact that only a few chapters have dedicated Chapter Leaders, the guides are often not up-to-date and may not document all the major upgrades of the code. While there is work to be done on this front, this does not fall under the expertise of this Working Group. The developers of the code are the most appropriate people to update the associated chapters. We encourage the NEMO Systems Team to consider that the update of the associated documentation is a necessity for every major new development.

### 19.2.2 User Guides and Demonstrations

In addition to the project website (<https://www.nemo-ocean.eu>) which provides a description of the NEMO modelling framework and related engines, NEMO has dynamically evolving user documentation that is tied to the code base (<https://sites.nemo-ocean.io/user-guide/index.html>). This offers the following functionalities:

- Getting started
- Setup up a testing framework
- Setup your configuration
- Advanced use with embedded zooms, coupling, data assimilation and tracers
- Guidance on contribution

Specifically, the NEMO Users' Guide (latest version obtainable from <https://sites.nemo-ocean.io/user-guide/>) provides information about system requirements, how to extract and install XIOS and NEMO, how to create and compile a new configuration, run it and use CPP keys. The available documentation is a good starting point for users. It also provides a description of the available reference configurations and list of working ones, with links to input data to use for their executions. Problems, however, lie in the fact that the material assumes that the reader already has a basic training in ocean modelling. As a result the debutant user is left unable to proceed without more hands-on basic training from a local NEMO user or developer. In this process the local contact spends extensive time on introductory training. This is likely happening at many institutions in parallel. Energy could be optimized by making this online material more user friendly for true beginners and adding further resources and worked examples for the 'student' to get some practice before setting up their own study configuration. Having a section on the expected results from

reference configurations would be very useful for users to ground-truth their output and make sure that they are doing the right thing. Furthermore, information is missing in the users' guide regarding the main changes between versions and releases and how to migrate from one version to another (e.g. namelist parameter/bloc that changed, obsolete features, new file convention). This prohibits 'students' from understanding how to adapt the available instructions to their own version and means that the guidelines quickly go out of date. Of great value here would be material that is not NEMO version specific, a general guide of tips and tricks, plus some pointers to help users migrate to the most recent official release.

### 19.2.3 User Forums

The NEMO Community offers the Discourse platform (<https://nemo-ocean.discourse.group/>) for exchanging information among users and developers. This new platform is a major improvement on the previous forum and is being actively used to:

- address issues as faced in the working releases
- support users for the usage of the code, configurations and demonstrators
- answer user questions
- understand the users needs for potential new developments

The NEMO System Team should continue to support and evolve its use of the Discourse platform. We recognise that a hand-full of developers take on the majority of the burden to answer questions on Discourse and ideally the workload should be more evenly shared. One of the potential solutions for this is to make it easier to answer questions by having a library of configurations and test cases that users could be referred to through Discourse. The idea here would be that developers convert a selection of their configurations and test cases into worked examples by adding some descriptive information and instructions and upload these onto Zenodo where they get a citable DOI. The developers would then get recognition for their outreach and dissemination efforts as the DOI could be added to their CV, as well as optimizing their efforts on Discourse by referring users to the worked examples where relevant. Conversely the users would then have additional resources that may be relevant in answering some of their "frequently asked questions" (FAQs).

### 19.2.4 Training (for users and developers; on-the-job & courses; continuous development)

Supporting the scientific and research sector is key for the development of the NEMO code and community. Academia may greatly benefit from the advancements of NEMO from the numerical point of view. Similarly the NEMO code base could greatly benefit from Academia participating more closely in, or leading aspects of its development. Additionally the operational sector is actively engaged with their users and NEMO could benefit from sharing some of the tools they use. A document of NEMO development coding rules exists, and so should be updated and shared online. Furthermore, as NEMO is a growing community of developers, having some resources aimed at training new Systems Team members could be highly beneficial.

The current approach to training provision comes from four main directions: firstly, a centralized NEMO provision through documentation and NEMO user guides, which are discussed above and are very technical in nature. Secondly, locally managed provision; thirdly community driven support, and finally workshops. In this discussion locally managed training is delivered peer-to-peer whereas community driven support represents a one-way exchange of user defined information through web based technologies.

#### *19.2.4.1 Locally managed (peer-to-peer) provision:*

Typically in this model, training is gleaned from key individuals who serve as “fountains of knowledge” or mentors. In some instances these individuals might typically also be on the NEMO System Team. The experienced individual will work with colleagues giving them working examples (configured for their architecture), and introductory guidance on how to get started (this is the namelist, change these timestepping parameters / forcing files / etc). The “apprentice” will gradually learn as they go and receive tailored guidance according to their knowledge and skill requirements. This on-the-job apprenticeship approach is more sustainable in large modelling teams, where natural turn-over of staff can be locally enriching (bringing new HPC techniques or analysis techniques and export NEMO expertise elsewhere) rather than devastating (e.g. in the loss of a single key individual) to the group’s skill base. For small research groups in Academia, where there is no top down strategic support for NEMO, this situation presents a high risk barrier to engagement with NEMO. Additionally, as this locally provided training is most likely taking place at many of the Consortium institutions, there could be an optimization of efforts by pooling resources in a central library/database.

#### *19.2.4.2 Community driven support (web based content):*

Community Driven Support encompasses ad-hoc web-based content typically and addresses the following types of problems, in the form of user written notes: ‘How to set up realistic regional configurations’, including ‘How to compile and run NEMO on particular architectures’, Grass roots examples include:

- Worked example in 1/12 degree South Asia domain (<https://doi.org/10.5281/zenodo.6423211>)
- Worked example of the 500m Severn Estuary configuration (<https://doi.org/10.5281/zenodo.6469990>)
- The Coastal Ocean Assessment Toolbox (COAsT): <https://british-oceanographic-data-centre.github.io/COAsT/docs/>

With less focus on How-to and more emphasis on archival best practice, other grass roots examples of documented configurations include:

- Caribbean, (<http://doi.org/10.5281/zenodo.3228088>) set up with forcing data on JASMIN and provided with scripts designed to auto build and run clean configurations.

At a higher strategic and international level community resources are created to facilitate countries building their own national centers for operational oceanography. For example resources from the Expert Team on Operational Ocean Forecasting Systems (ETOOFS practical session, Jun 2021) include publicly accessible training material:

- Institutional website: <https://www.mercator-ocean.eu/en/oofs-guide/>
- Tutorials: <https://www.surf-platform.org/tutorial.php>.

Additionally, tutorials have been shared and organized through OceanTeacher Global Academy (<https://classroom.oceanteacher.org/>) with participants.

Long term projects that truly embrace the open source philosophy should make a significant investment in “how-to” materials. On this front NEMO is lacking and we feel that a greater level of effort should be paid to develop resources specifically on the following topics:

- Beginners introduction to NEMO ocean model: what is an ocean model, fundamental equations, description of the blue-white-green components, how does it work in the HPC environment. Discussion on the rationale for how and why.
- Using NEMO: improvement of existing resources to download and install, a more holistic explanation of the NEMO structure in terms of the different files and folders and how it all fits together. How to compile and run NEMO for setting up a realistic configuration
- Contributing to the development of NEMO: a clearly outlined protocol of development do's and don'ts, an introduction to the git-lab environment, the procedure on how to submit a candidate new development to the NEMO team for an "outside"/"unofficial" developer.
- Using NEMO to understand ocean processes: description of a process (e.g. a wind driven gyre or sub-ice shelf cavity circulation), hands-on practice tutorial on how to compile and run the associated test case, some guidelines on dealing with output, plotting, and guidelines of what expected key results should be.
- Information highlighting the main novelties in terms of functionality and updates offered by NEMO.

### 19.2.4.3 Workshops:

Some organizations deliver NEMO specific training through workshop formats but the true extent to which these workshops systematically help is hard to ascertain. The following are workshops known to the authors:

- An introduction to ocean modelling: Running NEMO in Docker, targeted for interested environmental scientists with command line access to their laptop. Run twice (Belize, Merida, Mexico). (Material: [10.5281/zenodo.6417227](https://zenodo.org/record/6417227))
- Coastal Ocean Assessment Toolkit: Python diagnostics package for high resolution regional NEMO model, targeted at new NEMO data users / potential NEMO diagnostics developers. Run several times in the UK. Hosted by the UK Joint Marine Modelling Programme. (Material with examples: [10.5281/zenodo.4041413](https://zenodo.org/record/4041413))
- Relocatable NEMO: how to build and configure regional NEMO model, targeted at scientists who want to do this, but have been put off, or haven't yet. This was planned to be hosted by the UK Joint Marine Modelling Programme. It had significant UK interest but was postponed indefinitely. (Example configurations with documentation: [10.5281/zenodo.3228087](https://zenodo.org/record/3228087), [10.5281/zenodo.6423211](https://zenodo.org/record/6423211))
- Expert Team for Operational Ocean Forecasting (ETOofs, [https://www.gooscean.org/index.php?option=com\\_oe&task=viewGroupRecord&groupID=198](https://www.gooscean.org/index.php?option=com_oe&task=viewGroupRecord&groupID=198)), joint IOC-WMO and GOOS: practical workshop for understanding the benefits and implementing operational ocean forecasting systems. Sessions have been organized by Mercator Ocean International with the support of CMCC and University of Bologna (<https://www.mercator-ocean.eu/en/news/etoofs-workshops/>) and consist on demonstrations on how to set up high resolution regional to coastal configurations based on NEMO. The practical session in particular focused on the implementation of the Southeastern Brazilian coastal model by using SURF (Structured and Unstructured grid Relocatable ocean platforms for Forecasting, Trotta et al. 2021 and <https://www.surf-platform.org/tutorial.php>). Targeted students and scientists who want to learn on how to set up an operational ocean forecasting system by using NEMO.

Historically these were designed and delivered within organizations, though the COVID19 pandemic broadened the potential reach of this type of content. As the workshops are often constructed around web based material, this material later becomes a point of reference for the delegates and the material can be updated and recycled with time. The major benefit of these workshops is that they provide a point in time where users can schedule the often delayed learning of new skills. A practical step here to improve use of time would be to work to ensure that these workshops could be advertised and offered more widely across the user community.

### 19.3 Priorities for 2023-2025 and [champions]

1. Better communication of new releases. [MB + KH]
2. Review the Getting Started material in the Users' Guide [KH + JP]
3. Collate all existing training material on NEMO and identify what we can use from these external resources. [KH + JP]



4. Explore using containers for a user-friendly NEMO for workable examples [in collaboration with Daley Calvert and JamesHarle]..
5. Identity the gaps in the available NEMO educational resources and identify local champions to work on multimedia material to fill these gaps. [KH + JP]
6. Encourage the dedication of System Team personnel hours to actively monitoring the pulse of the Discourse and explore ways to give credit to the team members that do so. [KH]
7. Encourage developers to convert a selection of their configurations and test cases into worked examples by adding a description and instructions and posting onto Zenodo. [JP]
8. Establish a web space to house the NEMO tutorials and workable examples, the “NEMO zoo” [KH and need IT support for this]..
9. Revamp the NEMO website adding links to new user resources [KH + Nicolas Martin]
10. Update the NEMO Wikipedia page to better connect with the general public [KH]

## 19.4 Timeline

