



Computational Science and Engineering  
(International Master's Program)

Technische Universität München

Master's Thesis

**Schwarz Waveform Relaxation for a Single  
Column Climate Model**

Valentina Schüller







# Computational Science and Engineering (International Master's Program)

Technische Universität München

Master's Thesis

## Schwarz Waveform Relaxation for a Single Column Climate Model

Author: Valentina Schüller  
Examiner: Prof. Dr. Barbara Wohlmuth  
First Supervisor: Prof. Dr. Eric Blayo  
Second Supervisor: Dr. Florian Lemarié  
Submission Date: March 29, 2023





I hereby declare that this thesis is entirely the result of my own work except where otherwise indicated. I have only used the resources given in the list of references.

March 29, 2023

Valentina Schüller



---

## Acknowledgments

First and foremost, I would like to thank Eric Blayo and Florian Lemarié for their trust and close supervision. These six months were a period of immense learning and you encouraged me to ask all the questions I could come up with. Thank you to Barbara Wohlmuth for agreeing to examine this thesis, and for valuable advice along the way.

Thank you to everyone who has shared knowledge, opened doors, and helped make this thesis a reality, particularly: Sophie Valcke, Uwe Fladrich, Arnaud Caubel, Olivier Marti, Jan Streffing, Martin Schreiber, Hans-Joachim Bungartz, and Benjamin Rodenberg. I am deeply grateful to everyone in the EC-Earth SCM group and at the ECMWF, for patiently answering my questions and for developing the AOSCM in the first place. In my implementation, I relied heavily on excellent, open-source, software and documentation:

I want to especially acknowledge the use of xarray, ProPlot, and Jinja.

Thank you to all members of the Laboratoire Jean Kuntzmann for making me feel so at home during my time in Grenoble. And finally, my deepest thanks go out to everybody else who is part of the support network I am surrounded by: It would have been impossible without you.





---

## Abstract

Typical algorithms used to couple atmosphere and ocean models are computationally efficient but mathematically inconsistent, thus introducing a numerical error. Schwarz waveform relaxation is an iterative coupling method to restore consistency at the interface, allowing to investigate the magnitude and physical implications of this error. The large computational cost of Schwarz waveform relaxation prevents in-depth numerical studies with coupled general circulation models. On the other hand, the theoretical analysis of highly idealized problems is insufficient to draw conclusions about general circulation models. In this master's thesis, a coupled atmosphere-ocean model of intermediate complexity, the one-dimensional EC-Earth AOSCM, is modified to support different coupling schemes used in operational climate models, as well as Schwarz waveform relaxation. This makes it a tool to bridge the gap between previous theoretical and numerical studies. Multi-day simulations in this new setup illustrate how physical parameterizations in the atmosphere react to changes in interface boundary conditions at the sea surface. Small variations in sea surface temperature can yield potentially large differences between standard coupling algorithms and the more accurate result of the Schwarz method. The EC-Earth AOSCM shares a significant amount of code and architecture with the EC-Earth 3 climate model. This similarity allows us to address algorithmic aspects and implementation challenges for iterative coupling schemes in climate modeling.



# Contents

<b>Acknowledgements</b>	<b>vii</b>
<b>Abstract</b>	<b>ix</b>
<b>1. Introduction</b>	<b>1</b>
<b>2. Coupling Schemes in Atmosphere-Ocean Coupling</b>	<b>3</b>
2.1. The Generalized Atmosphere-Ocean Coupling Problem . . . . .	3
2.2. Standard Approaches for Coupling in Time . . . . .	5
2.2.1. The Parallel Algorithm . . . . .	6
2.2.2. The Sequential Algorithms . . . . .	7
2.3. Schwarz Waveform Relaxation . . . . .	9
2.3.1. Continuous Schwarz Waveform Relaxation . . . . .	9
2.3.2. Schwarz Waveform Relaxation in a Time-Discretized Setting . . . . .	12
<b>3. Overview of the EC-Earth AOSCM</b>	<b>13</b>
3.1. OpenIFS . . . . .	14
3.2. NEMO . . . . .	18
3.3. Interface Boundary Conditions . . . . .	19
3.4. Coupling with OASIS . . . . .	24
<b>4. Implementation</b>	<b>25</b>
4.1. Sequential Coupling Algorithms . . . . .	26
4.2. Schwarz Waveform Relaxation . . . . .	26
4.3. SWR Convergence Criteria . . . . .	29
4.3.1. Variables Included in Error Estimates . . . . .	29
4.3.2. Amplitude-Based Convergence Estimate . . . . .	30
4.3.3. Local Convergence Estimate . . . . .	30
4.3.4. Comparison with the Final Iterate . . . . .	31
<b>5. Numerical Experiments</b>	<b>33</b>
5.1. Control Experiment: Qualitative Comparison of Coupling Schemes . . . . .	33
5.1.1. Prognostic Variables . . . . .	34
5.1.2. Behavior of Atmospheric Thermodynamics . . . . .	36
5.1.3. Comparison of Convergence Criteria . . . . .	41

5.2. Case Study: Numerical Weather Prediction Setup . . . . .	43
5.2.1. Model Setup and Initial Conditions . . . . .	44
5.2.2. Results with the NWP Setup . . . . .	45
<b>6. Discussion</b>	<b>51</b>
6.1. Discussion of Numerical Results . . . . .	51
6.1.1. Impact of Coupling Schemes on Numerical Results . . . . .	51
6.1.2. Schwarz Waveform Relaxation Convergence . . . . .	52
6.1.3. Limitations of the Numerical Experiments . . . . .	53
6.2. Discussion of the Implementation . . . . .	54
<b>7. Conclusion and Outlook</b>	<b>57</b>
<b>Appendix</b>	<b>61</b>
<b>A. Appendix</b>	<b>61</b>
A.1. Reynolds Averaging . . . . .	61
A.2. Classification of AOSCM Interface Conditions . . . . .	62
A.3. Fundamentals of the OASIS Coupler . . . . .	63
A.3.1. Coupling Field Types in OASIS . . . . .	63
A.3.2. The LAG Parameter . . . . .	64
A.4. Sequential Coupling with OASIS . . . . .	65
A.5. OASIS Configuration Examples for SWR . . . . .	65
A.6. Multiplicative SWR with the AOSCM . . . . .	66
<b>Bibliography</b>	<b>69</b>

# 1. Introduction

To understand and predict how greenhouse gas emissions affect our planet and life on it, researchers use numerical climate models. These consist of different components which model various parts of our Earth system: oceanic and atmospheric dynamics, sea ice, vegetation, land surface, etc. The components are coupled in time and space to produce a comprehensive model of the climate on our planet.

Models for a system of this scale and complexity are inherently imperfect and it is not realistically possible to eradicate all errors which are part of climate models. Interesting questions which arise instead are thus where we introduce errors, how large they are, and how sensitive they are to changes in the model setup or initial conditions. In a coupled (Earth system) model, sources of error may be the model components, but also how we couple them. For instance, non-matching grids must be mapped to each other, introducing an interpolation error.<sup>1</sup> Even when the model components use matching grids, coupling introduces an error due to the fact that information between components is only exchanged at discrete points in time and not continuously. Which information is exchanged and at what points in time is referred to as the *coupling scheme* used by the model.

Most resources in Earth system modeling are spent on improving the model components either by increasing the physical and numerical accuracy or by reducing the computational cost. The latter allows researchers to use a higher resolution in time and space with the same computational resources. However, as noted by Gross et al. (2018, p. 3505): “As the error associated with each component decreases, the errors introduced by the coupling will eventually dominate.” For this reason, there have been increased efforts to study the impact of the coupling scheme on the numerical solution in Earth system models.

As a result of this research, it has been discussed that coupling schemes customary in climate modeling and numerical weather prediction are computationally efficient but mathematically inconsistent (Gross et al., 2018; Lemarié et al., 2015). Schwarz waveform relaxation (SWR) has been proposed as an iterative coupling approach to recover mathematical consistency, albeit with a high computational cost. It allows researchers to study how high the error introduced by conventional coupling methods is and whether the correct interface data is exchanged. Furthermore, model results produced with SWR can serve as reference solutions for model analysis (Gross et al., 2018). Previous studies using general circulation models showed that the difference between classical approaches and the converged SWR solution can be significant (Connors & Ganis, 2011; Lemarié et al., 2014), particularly around sunrise and sunset (Marti et al., 2021). In idealized atmosphere-ocean

---

<sup>1</sup>see for example Keyes et al. (2013, section 3.2.1) or Gatzhammer (2014, section 2.4) for a discussion of different data mapping techniques in multiphysics coupling.

models, Schwarz waveform relaxation serves as a framework to mathematically study the error propagation in time and find mathematically optimal interface conditions (Clement et al., 2022; Connors & Miloua, 2011; Lemarié et al., 2013b; They et al., 2022).

In this master's thesis, we use Schwarz waveform relaxation to study coupling errors in time for a single column climate model (SCM). These models are one-dimensional in space, simulating a single vertical column of, e.g., the atmosphere and the ocean. SCMs require significantly fewer computational resources than general circulation models and Earth system models. Since a single column model does not have a horizontal discretization, the grids match at the interface and this potential source of coupling error is removed. Compared to idealized climate models, SCMs include realistic physical parameterizations for subgrid-scale phenomena and are based on more complex governing equations by including, e.g., the conservation of internal energy, moisture, and salinity. In the context of studying physical processes in a model hierarchy (Maher et al., 2019) and developing Earth system models (Hourdin et al., 2017), SCMs can bridge a gap between idealized and full complexity models.

The model studied in this master's thesis is the EC-Earth coupled atmosphere-ocean single column model (AOSCM) (Hartung et al., 2018). We have modified the EC-Earth AOSCM to be a tool for comparing different atmosphere-ocean coupling schemes in a climate model of intermediate complexity. We conduct different numerical experiments in this new setup and show the sensitivity of multi-day forecasts to changes in the coupling algorithm. This master's thesis also provides a detailed formulation of the coupling problem solved by the model and describes how to implement Schwarz waveform relaxation in coupled models with a similar structure, particularly those making use of the OASIS3-MCT coupler. Our solution approach treats the atmosphere and ocean models as black boxes and is minimally invasive. However, it does not allow for higher order Schwarz waveform relaxation as presented, e.g., in Rütth et al. (2018).

The rest of this master's thesis is structured as follows: Chapter 2 gives an overview of coupling schemes commonly used in atmosphere-ocean coupling as well as the theoretical background for Schwarz waveform relaxation methods. Chapter 3 presents the EC-Earth AOSCM and the coupling problem it solves. The following Chapter 4 explains how we approached to implement different coupling configurations in the EC-Earth AOSCM. We continue with different numerical experiments at the PAPA station in the Northern Pacific Ocean: We test the impact of different coupling algorithms on multi-day simulations in Chapter 5. In Chapter 6, we discuss the numerical results and our implementation approach. The master's thesis closes in Chapter 7 with a conclusion and outlook.

## 2. Coupling Schemes in Atmosphere-Ocean Coupling

Before we move to the EC-Earth single column model, this chapter gives a general overview of potential algorithms to couple atmosphere and ocean models in time, so-called coupling schemes. It starts off with a symbolic definition of the atmosphere-ocean coupling problem in Section 2.1. Commonly used coupling schemes for this problem are presented in Section 2.2, while Schwarz waveform relaxation as an alternative approach to coupling is the topic of the concluding Section 2.3.

### 2.1. The Generalized Atmosphere-Ocean Coupling Problem

To discuss coupling approaches without limiting ourselves to a single model, we formulate the atmosphere-ocean coupling problem in a symbolic form, similar to Lemarié et al. (2014), Lemarié et al. (2015), and Marti et al. (2021). Particularly, we can view atmosphere-ocean coupling as a case of domain decomposition: The two models act on adjacent, non-overlapping domains  $\Omega_{\text{atm}}$  and  $\Omega_{\text{oce}}$  with a common interface  $\Gamma$ , the sea surface, cf. the illustration in Figure 2.1. Together, they solve a system of time-dependent, nonlinear partial differential equations to describe the state of Earth's atmosphere and ocean. Taking this holistic view will allow us to introduce Schwarz methods naturally later on in this chapter. In the interior of both subdomains, a partial differential operator  $\mathcal{L}$  acts on state variables  $\mathbf{U}$ . In both systems, this state vector includes the horizontal velocities and the temperature of the fluid. We will see in Chapter 3 that the atmospheric state vector  $\mathbf{U}^a$  additionally contains moisture and the oceanic state vector  $\mathbf{U}^o$  is completed by including salinity.

In addition to the system of partial differential equations, the atmospheric and oceanic states are subject to boundary and initial conditions:

$$\begin{cases} \mathcal{L}_{\text{atm}} \mathbf{U}^a = f_{\text{atm}} & \text{in } \Omega_{\text{atm}} \times [0, T], \\ \mathcal{B}_{\text{atm}} \mathbf{U}^a = g_{\text{atm}} & \text{in } \partial\Omega_{\text{atm}}^{\text{ext}} \times [0, T], \\ \mathbf{U}^a|_{t=0} = \mathbf{U}_0^a & \text{in } \Omega_{\text{atm}}, \end{cases} \quad (2.1)$$

$$\begin{cases} \mathcal{L}_{\text{oce}} \mathbf{U}^o = f_{\text{oce}} & \text{in } \Omega_{\text{oce}} \times [0, T], \\ \mathcal{B}_{\text{oce}} \mathbf{U}^o = g_{\text{oce}} & \text{in } \partial\Omega_{\text{oce}}^{\text{ext}} \times [0, T], \\ \mathbf{U}^o|_{t=0} = \mathbf{U}_0^o & \text{in } \Omega_{\text{oce}}. \end{cases} \quad (2.2)$$

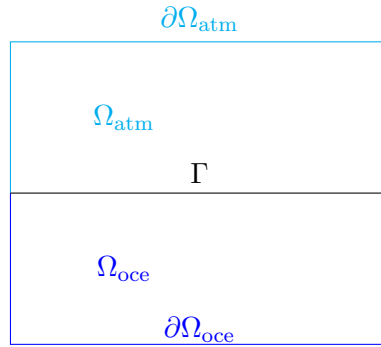


Figure 2.1.: Illustration of the domain decomposition problem in atmosphere-ocean interaction.

In these equations,  $f$  symbolizes forcing terms,  $g$  refers to boundary conditions, and  $U_0$  are the initial conditions, all of which we assume to be provided externally. We use the subscripts *atm*, *oce* and superscripts *a*, *o* to distinguish between the atmosphere and ocean when necessary.  $\mathcal{B}$  is a differential operator acting on the boundaries of the subdomains,  $\partial\Omega_{\text{atm}}$  and  $\partial\Omega_{\text{oce}}$ . Depending on the choice of the boundary operator, one can distinguish common types of boundary conditions:

- $\mathcal{B} = Id$  represents Dirichlet conditions,
- $\mathcal{B} = \frac{\partial}{\partial \mathbf{n}}$ , with  $\mathbf{n}$  the normal vector of  $\partial\Omega$ , represents Neumann boundary conditions,
- $\mathcal{B} = Id + r \frac{\partial}{\partial \mathbf{n}}$  represents Robin boundary conditions with a free parameter  $r$ .

Up to this point we have not discussed what happens at the interface  $\Gamma$ . This is the central question in coupling the two subproblems and the rest of this chapter will be concerned with this topic.

From the abstract perspective of the model problem, one introduces another differential operator  $\mathcal{C}$  which acts on  $U$  at the interface and thus denotes the *interface conditions*, also referred to as transmission conditions depending on the literature (Blayo et al., 2017; Gander, 2015; Gander & Halpern, 2007). Note that, in general,  $\mathcal{C}_{\text{atm}} \neq \mathcal{C}_{\text{oce}}$ . In many multiphysics applications,  $\mathcal{C}$  takes forms similar to the common types of boundary conditions presented above. The choice of the interface operators is used to classify the coupling problem, with a particularly popular choice being Dirichlet-Neumann algorithms in fluid-structure interaction problems (e.g., Gatzhammer, 2014).

The form of  $\mathcal{C}$  depends on the underlying problem. For domain decomposition applications (where  $\mathcal{L}$  is the same in all subdomains), so-called natural or physical transmission conditions can be determined using the variational formulation of the problem, see for example Blayo et al. (2017). These impose constraints on  $\mathcal{C}$  which are required for






Figure 2.2.: Illustration of the multirate setting in the atmosphere-ocean coupling problem. In general, the atmosphere and ocean model exchange information whenever a coupling window is over, i.e., at the dashed lines in the figure.

well-posedness of the coupling problem.<sup>1</sup> For complex multiphysics problems such as atmosphere-ocean coupling, it is not possible to derive natural transmission conditions in this sense. Here, a common assumption is instead that the air-sea fluxes are continuous across the sea surface (Lemarié et al., 2015) and that they can be computed using a (nonlinear) function  $\mathbf{F}_{\text{oa}}$  which takes into account the atmospheric and oceanic states. This leads to the following generalized coupling problem solved in a general circulation model:

$$\begin{cases} \mathcal{L}_{\text{atm}} \mathbf{U}^a = f_{\text{atm}} & \text{in } \Omega_{\text{atm}} \times [0, T], \\ \mathcal{B}_{\text{atm}} \mathbf{U}^a = g_{\text{atm}} & \text{in } \partial\Omega_{\text{atm}}^{\text{ext}} \times [0, T], \\ \mathcal{L}_{\text{oce}} \mathbf{U}^o = f_{\text{oce}} & \text{in } \Omega_{\text{oce}} \times [0, T], \\ \mathcal{B}_{\text{oce}} \mathbf{U}^o = g_{\text{oce}} & \text{in } \partial\Omega_{\text{oce}}^{\text{ext}} \times [0, T], \\ \mathcal{C}_{\text{atm}} \mathbf{U}^a = \mathcal{C}_{\text{oce}} \mathbf{U}^o = \mathbf{F}_{\text{oa}}(\mathbf{U}^a, \mathbf{U}^o) & \text{on } \Gamma \times [0, T]. \end{cases} \quad (2.3)$$

We will discuss the forms of both  $\mathcal{C}$  and  $\mathbf{F}_{\text{oa}}$  in more detail in Section 3.3.

## 2.2. Standard Approaches for Coupling in Time

In order to discuss coupling schemes, we need to introduce a temporal discretization of the coupling problem: Instead of a continuous exchange of information, which would be the case in the real Earth system, the two numerical models for the ocean and atmosphere communicate at  $N$  points in time  $t_n$ , where  $t_0 = 0$  and  $t_N = T$ . We call the intervals  $[t_n, t_{n+1}]$  *coupling windows* and without loss of generality, we assume that all coupling windows have the same size during one simulation,

$$\Delta t = t_{n+1} - t_n = \frac{T}{N} \quad \forall n = 0, 1, \dots, N-1.$$

The model components can use time step sizes smaller than the coupling window size to integrate the subproblems in Equations (2.1) and (2.2). Similarly, we do not require the

<sup>1</sup>If the variational form imposes, e.g., flux conservation, Dirichlet interface conditions will not be sufficient.

models to use the same time step sizes, as long as the coupling time steps  $t_n$  are a subset of the model time steps. Allowing for this flexibility makes atmosphere-ocean coupling a *multirate* problem (Keyes et al., 2013). Figure 2.2 illustrates this property for one exemplary coupling window. This flexibility reflects the fact that the ocean is in general a slower system than the atmosphere and might not need the same temporal resolution to capture its dynamics sufficiently. Furthermore, it allows model time steps to be different depending on criteria such as accuracy, stability, and computational cost.

### 2.2.1. The Parallel Algorithm

At each coupling time step  $t_n$ , the air-sea fluxes  $\mathbf{F}_{\text{oa}}$  can be computed to obtain the interface conditions for the next coupling window. This choice introduces a coupling lag: The interface conditions are computed at  $t_n$  but used by both models until  $t_{n+1}$ . All coupling schemes discussed in this section are lagged, albeit to a different extent. In the context of multiphysics coupling, these coupling algorithms are also referred to as being *loose* (Keyes et al., 2013).

A simple, parallel coupling algorithm could take the following form: At a coupling time step  $t_n$ , both models send the average of their state vector over the previous coupling window  $[t_{n-1}, t_n]$ , which we denote by  $\langle \mathbf{U}^a \rangle_{t_{n-1}}^{t_n}$  and  $\langle \mathbf{U}^o \rangle_{t_{n-1}}^{t_n}$ . This is used to compute the air-sea fluxes for the next coupling window,

$$\mathcal{C}_{\text{atm}} \mathbf{U}^a(t) = \mathcal{C}_{\text{oce}} \mathbf{U}^o(t) = \mathbf{F}_{\text{oa}} \left( \langle \mathbf{U}^a \rangle_{t_{n-1}}^{t_n}, \langle \mathbf{U}^o \rangle_{t_{n-1}}^{t_n} \right) \quad \text{for } t \in [t_n, t_{n+1}]. \quad (2.4)$$

This algorithm is parallel because both models only depend on information from the previous coupling window. Therefore,  $\mathbf{U}^a$  and  $\mathbf{U}^o$  can be evolved to the next coupling time  $t_{n+1}$  in parallel. The disadvantage of this approach is that both models have to wait at  $t_n$  for the computation of  $\mathbf{F}_{\text{oa}}$  before continuing with the next coupling window.

However, the air-sea fluxes are usually computed by the atmospheric component in every time step instead of once per coupling window due to the faster nature of atmospheric dynamics (Marti et al., 2021). Thus, the two models do not see the same interface conditions in a given time step. Instead, the (relaxed) goal is to force both models by the same *mean* fluxes over a given coupling window to ensure their conservation (Lemarié et al., 2014). For the air-sea fluxes seen by the atmosphere, we thus obtain

$$\mathcal{C}_{\text{atm}} \mathbf{U}^a(t) = \mathbf{F}_{\text{oa}} \left( \mathbf{U}^a(t), \langle \mathbf{U}^o \rangle_{t_{n-1}}^{t_n} \right) \quad \text{for } t \in [t_n, t_{n+1}]. \quad (2.5)$$

At time  $t_n$ , the ocean model sends the average of its state vector (at the interface  $\Gamma$ ) over the previous coupling window  $[t_{n-1}, t_n]$  to the atmosphere model. The atmosphere model, on the other hand, sends the mean of the air-sea fluxes  $\langle \mathcal{C}_{\text{atm}} \mathbf{U}^a \rangle_{t_{n-1}}^{t_n}$  and we obtain:

$$\mathcal{C}_{\text{oce}} \mathbf{U}^o(t) = \langle \mathcal{C}_{\text{atm}} \mathbf{U}^a \rangle_{t_{n-1}}^{t_n} = \left\langle \mathbf{F}_{\text{oa}} \left( \mathbf{U}^a, \langle \mathbf{U}^o \rangle_{t_{n-2}}^{t_{n-1}} \right) \right\rangle_{t_{n-1}}^{t_n} \quad \text{for } t \in [t_n, t_{n+1}]. \quad (2.6)$$




Figure 2.3.: The parallel coupling scheme for atmosphere-ocean coupling. The diagonal arrows indicate communication of data between the model components. The braces illustrate that the averages  $\langle \cdot \rangle$  are based on the full coupling window. The interval  $[t_{n-2}, t_{n-1}]$  is included to visualize the coupling lag.

Because of the asymmetry introduced by computing the fluxes in the atmospheric component, the atmosphere and ocean see a lag of  $\Delta t$  and  $2\Delta t$  in the interface conditions, respectively. This coupling scheme has been referred to as the (asynchronous) parallel algorithm for atmosphere-ocean coupling (Lemarié et al., 2014; Marti et al., 2021) and is visualized in Figure 2.3.

### 2.2.2. The Sequential Algorithms

In contrast to the parallel algorithm, it is also possible to switch to a sequential coupling of atmosphere and ocean. Two variants exist, the sequential atmosphere-first and the sequential ocean-first algorithms. With sequential atmosphere-first coupling, the ocean waits with its computation of  $[t_n, t_{n+1}]$  until the atmosphere has computed  $\mathbf{U}^a(t_{n+1})$ . The ocean model then uses the new information provided by the atmospheric component:

$$\mathcal{C}_{\text{atm}} \mathbf{U}^a(t) = \mathbf{F}_{\text{oa}} \left( \mathbf{U}^a(t), \langle \mathbf{U}^o \rangle_{t_{n-1}}^{t_n} \right) \quad (2.7)$$

$$\mathcal{C}_{\text{oce}} \mathbf{U}^o(t) = \langle \mathcal{C}_{\text{atm}} \mathbf{U}^a \rangle_{t_n}^{t_{n+1}} = \left\langle \mathbf{F}_{\text{oa}} \left( \mathbf{U}^a, \langle \mathbf{U}^o \rangle_{t_{n-1}}^{t_n} \right) \right\rangle_{t_n}^{t_{n+1}}, \quad (2.8)$$

with  $t \in [t_n, t_{n+1}]$ . Note that, in this configuration, the coupling lag is reduced to  $\Delta t$  for both models.

In the sequential ocean-first case, the atmospheric component waits for the ocean model to compute the coupling window before evolving  $\mathbf{U}^a$  with the new interface data. Thus,

## 2. Coupling Schemes in Atmosphere-Ocean Coupling

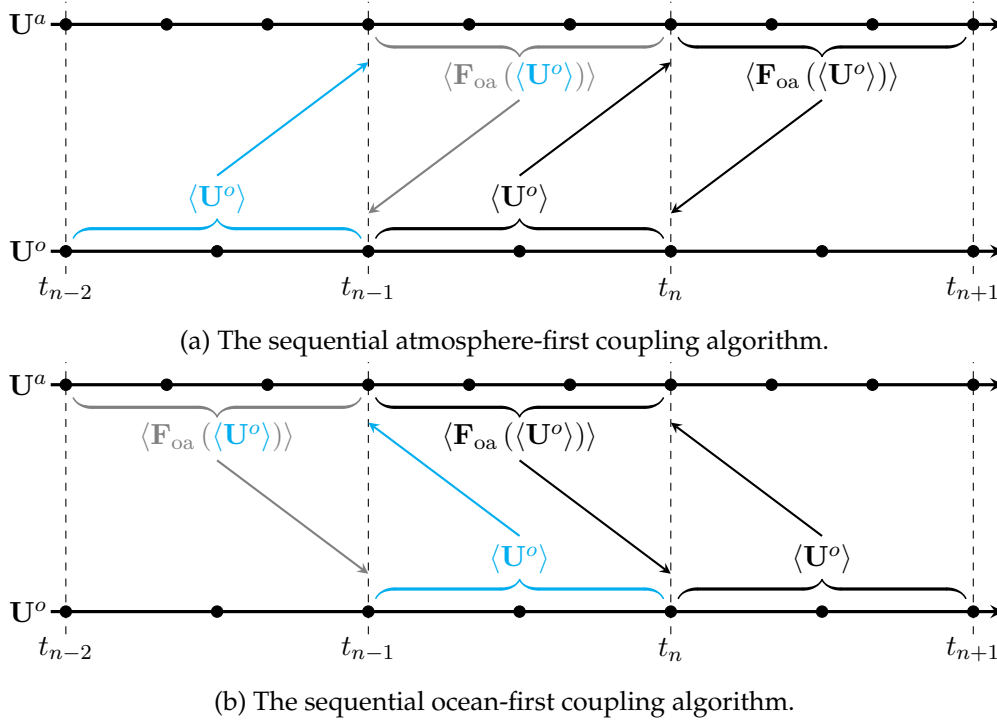


Figure 2.4: The sequential algorithms for atmosphere-ocean coupling. The interface data and communication arrows in black provide the boundary conditions for  $[t_n, t_{n+1}]$ . The respective data for the previous coupling window is shown in gray and cyan. In contrast to Figure 2.3,  $[t_n, t_{n+1}]$  does not depend on data from  $[t_{n-2}, t_{n-1}]$ , thus the coupling lag is reduced.

the fluxes take the following form for  $t \in [t_n, t_{n+1}]$ :

$$\mathcal{C}_{oce} \mathbf{U}^o(t) = \langle \mathcal{C}_{atm} \mathbf{U}^a \rangle_{t_{n-1}}^{t_n} = \left\langle \mathbf{F}_{oa} \left( \mathbf{U}^a, \langle \mathbf{U}^o \rangle_{t_{n-1}}^{t_n} \right) \right\rangle_{t_{n-1}}^{t_n} \quad (2.9)$$

$$\mathcal{C}_{atm} \mathbf{U}^a(t) = \mathbf{F}_{oa} \left( \mathbf{U}^a(t), \langle \mathbf{U}^o \rangle_{t_n}^{t_{n+1}} \right). \quad (2.10)$$

Again, this reduces the maximum coupling lag to  $\Delta t$ . Both sequential algorithms are illustrated in Figure 2.4.

In Marti et al. (2021), both sequential coupling schemes led to improved numerical results for a 3D coupled general circulation model compared to the parallel algorithm, with the sequential atmosphere-first version outperforming the ocean-first algorithm. The benefit of improved accuracy comes at the expense of a higher computational cost: Assuming that the computational effort for both models is similar, switching from parallel to sequential coupling increases the time to solution roughly by a factor of two.

According to the research by Marti et al. (2021), the parallel algorithm is used more frequently than the sequential atmosphere-first coupling scheme. As they state, "no model uses a sequential ocean-first algorithm" (Marti et al., 2021, p. 2960). Nevertheless, comparing all three variants can be of help to gain additional insight into the interaction of coupled fast and slow systems.

## 2.3. Schwarz Waveform Relaxation

### 2.3.1. Continuous Schwarz Waveform Relaxation

As has been stated in the literature, all three algorithms presented in the previous section are mathematically inconsistent (Gross et al., 2018; Lemarié et al., 2015; Lemarié et al., 2014; Marti et al., 2021). To justify this statement, we take on the perspective of Schwarz methods, which are based on the seminal work of Schwarz (1870). Schwarz methods encompass various flavors of iterative methods to approach domain decomposition problems. A comprehensive overview of their history and developments can be found in Gander (2008).

The central theory behind Schwarz methods states the following: Suppose we have a boundary value problem on a domain  $\Omega$  which we want to split up into problems on two subdomains  $\Omega_1$  and  $\Omega_2$ , where  $\Omega_1 \cup \Omega_2 = \Omega$ . These domains can overlap, as was the case in the original problem studied by Schwarz, i.e.,  $\Omega_1 \cap \Omega_2 \neq \emptyset$ , but they do not have to. We now solve the subproblems iteratively on each subdomain and use their solution to update the values on the interface between the subdomains, until the iterations converge. If the boundary value problems on  $\Omega$  and the subdomains are well-posed and appropriate interface conditions are chosen, the iterative method converges, and it converges to the solution of the original problem. While the ideas are very similar for different problems, the convergence proofs strongly depend on the problem at hand (Gander, 2008, Section 2.1), the subdomains, and the choice of interface conditions. The method can also be extended to more than two subdomains.

The iterative algorithm can take two different forms: Either both subdomains work with interface data from the previous iteration or they take turns and one waits for the other. The first method is referred to as the additive or parallel Schwarz method and was introduced in Lions (1988), whereas the other one is the original method proposed in Schwarz (1870) and is called the multiplicative or sequential Schwarz method. The multiplicative Schwarz method usually converges in fewer iterations than the additive version (Gander, 2008).

While Schwarz methods were originally designed for boundary value problems, they are straightforward to extend to time-dependent partial differential equations. The idea here is to solve the initial value problem on  $\Omega \times [0, T]$  by iterating over the two subproblems  $\Omega_1 \times [0, T]$  and  $\Omega_2 \times [0, T]$ . In the same way that it is possible to split  $\Omega$  into more than one subdomain, it is possible to split the time domain into smaller intervals, e.g.,  $[0, T_1]$  and  $[T_1, T]$ , so-called *Schwarz windows*. This is commonly solved by iterating until the

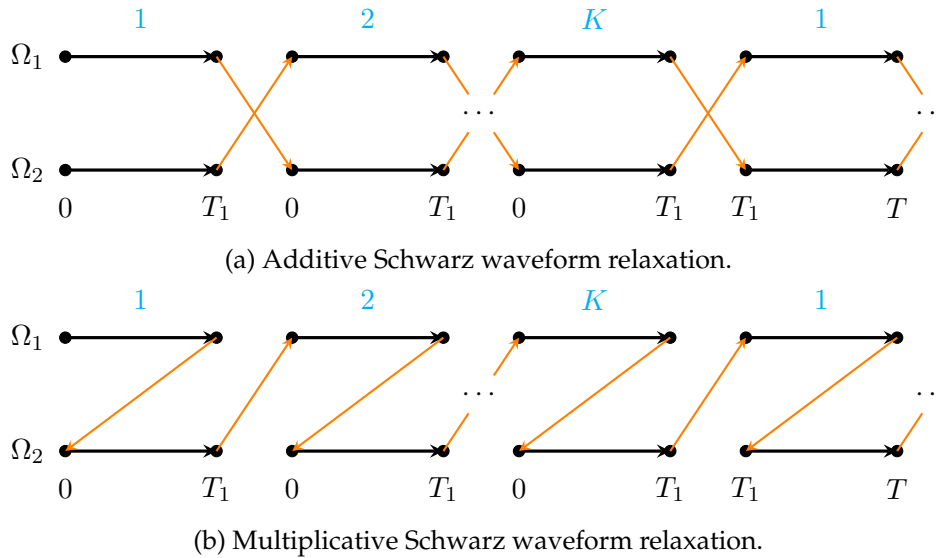


Figure 2.5.: Schwarz waveform relaxation for two subdomains depicted by the data dependency over time and across iterations, for two Schwarz windows  $[0, T_1]$  and  $[T_1, T]$ , with the iteration count in blue. Black arrows depict integration from one point in time to the next. Orange arrows represent communication of information between two iterations or from one Schwarz window to the next.

algorithm converges on  $\Omega_i \times [0, T_1]$  before moving on to the next time interval (Rüth et al., 2021). With a larger Schwarz window, more iterations are required for convergence (Gander, 2015). Schwarz methods for time-dependent partial differential equations are referred to as *Schwarz waveform relaxation* (SWR).<sup>2</sup>

In their original form, the additive and multiplicative Schwarz methods use Dirichlet conditions at the interface between subdomains. As noted by Gander and Halpern (2007), this leads to slow convergence of the iterations. By studying the underlying problem, one can come up with better choices for the interface conditions. Finding interface operators  $\mathcal{C}$  which drastically reduce the number of iterations is part of the branch of optimal or optimized Schwarz methods. Examples of how to approach such a task for Schwarz waveform relaxation particularly are Gander et al. (1999) for the wave equation, Gander and Halpern (2007) for advection–reaction–diffusion problems, Lemarié et al. (2013a, 2013b) for diffusion problems, and more recently Blayo et al. (2016) and Blayo et al. (2017) for the Stokes and shallow-water equations, respectively. In all of these studies, however, the problems are idealized cases of real world applications. Finding not only good but mathematically optimal interface conditions is unrealistic for complex, nonlinear problems such

<sup>2</sup>*Waveform relaxation* is a method developed in the context of circuit simulation and these roots are still reflected in the name. See, e.g., Gander (2015) for a description of the original method and how it is related to Schwarz waveform relaxation.

as atmosphere-ocean coupling.

Figure 2.5 illustrates the data dependency and flow of information for additive and multiplicative SWR. The model components exchange information between subsequent iterations but also from one Schwarz window to the next. In general, the Schwarz window can span multiple coupling periods  $\Delta t$ . If this is the case, the models communicate during an iteration at each coupling time step  $t_n$ , *in addition* to the communication between iterations. When to stop iterating could be determined by setting a fixed amount of iterations  $K$  or by choosing a criterion to determine convergence dynamically. We discuss this aspect in more detail in Section 4.3. By comparing Figures 2.3 to 2.5, one can see similar communication patterns in the parallel algorithm and additive SWR, and in the sequential algorithms and multiplicative SWR, respectively.

These observations reveal the following conclusions, previously stated, e.g., in Gross et al. (2018), Lemarié et al. (2014), and Marti et al. (2021): We can view the atmosphere-ocean problem as an example of domain decomposition with no overlap of the subdomains, cf. Figure 2.1. By taking the perspective of Schwarz methods, the parallel algorithm can be seen as the first iteration of an additive SWR method. Similarly, the sequential atmosphere-first and ocean-first coupling schemes are equivalent to the first iteration of a multiplicative SWR algorithm. Since even optimal Schwarz methods take as many iterations as there are subdomains to converge, the solutions produced by all three coupling schemes are mathematically inconsistent: The subproblems  $\Omega_i \times [t_n, t_{n+1}]$  are only solved approximately before moving on to the next Schwarz window.

Solving Equation (2.3) approximately is not a problem in general—in fact, it is the norm. Numerical methods per se result in approximate solutions and Earth system models are numerical models. The question is, as mentioned in the introduction, how large the error due to the approximation is. Schwarz waveform relaxation serves as a tool to answer this question: By repeating the time integration from  $t_n$  to  $t_{n+1}$  until the air-sea fluxes converge, one can generate a reference solution and compare it to the results with standard coupling schemes. The converged solution uses fully re-synchronized fluxes during  $t \in [t_n, t_{n+1}]$ :

$$\mathcal{C}_{\text{atm}} \mathbf{U}^a(t) = \mathbf{F}_{\text{oa}} \left( \mathbf{U}^a(t), \langle \mathbf{U}^o \rangle_{t_n}^{t_{n+1}} \right) \quad (2.11)$$

$$\mathcal{C}_{\text{oce}} \mathbf{U}^o(t) = \langle \mathcal{C}_{\text{atm}} \mathbf{U}^a \rangle_{t_n}^{t_{n+1}} = \left\langle \mathbf{F}_{\text{oa}} \left( \mathbf{U}^a, \langle \mathbf{U}^o \rangle_{t_n}^{t_{n+1}} \right) \right\rangle_{t_n}^{t_{n+1}}. \quad (2.12)$$

If the convergence speed of the iterations is very slow, or if the method does not converge, the SWR perspective implies that the interface conditions could be incompatible and should be worked on, an issue discussed by Lemarié et al. (2014). In the case of non-convergence, the underlying coupling problem might not “obey regularity” or could even be ill-posed (Gross et al., 2018, p. 3523). Such a result is not unthinkable considering the fact that Earth system model components are usually developed independently of each other. Particularly the implementation of subgrid-scale physics is mostly based on semi-empirical laws and the interaction of the physical parameterizations with dynamics

is non-trivial. This is already the case in a single model component and gets amplified in a coupled situation, potentially impairing the regularity of solutions. See also Gross et al. (2018, Section 5) for a more in-depth discussion of this issue.

### 2.3.2. Schwarz Waveform Relaxation in a Time-Discretized Setting

The mathematical theory behind Schwarz waveform relaxation is based on the model components exchanging functions of time, not discrete values. That is, over one Schwarz window, the subdomains exchange time-continuous interface conditions. The SWR algorithm we have defined based on the parallel and sequential coupling schemes notably uses constant fields in the interface conditions: the average of  $\mathbf{U}^o$  or  $C_{\text{atm}}\mathbf{U}^a$  over the time intervals  $[t_n, t_{n+1}]$ . One can view these single values as approximations of the "true" interface conditions using a constant function.

Again, this is a source of error in the coupling algorithm. One can demonstrate that using single value coupling limits the achievable order of convergence in time of the numerical solution. Furthermore, this can affect the stability and energy conservation of the coupled model, particularly for multirate applications. See, e.g., R uth et al. (2018) and references therein for numerical studies demonstrating these issues. It is possible to mitigate these problems by exchanging higher order interpolants, a technique referred to as higher order waveform iterations or relaxation (R uth et al., 2018).

The topic of convergence order is its own, separate issue in weather and climate simulations: The ocean model studied in this thesis uses a first-order time integration method as part of its time integration scheme (Madec et al., 2017). Thus the overall achievable convergence order of the model is not just limited by the coupling scheme, but also by the individual models. But even if this were not the case, the physical parameterizations in Earth system models make convergence studies in the classical sense difficult. One reason for this is that the parameterization schemes are based on semi-empirical models optimized for certain time ranges, reference solutions might thus "violate physical assumptions" (Gross et al., 2018, p. 3512). Gross et al. (2018, Section 2e) discuss these issues and propose ways of going forward.

For the purposes of this master's thesis, it is thus clear that implementing quadratic or cubic interpolations of the interface conditions is not necessary from the point of view of numerical convergence order. Nevertheless, it might have benefits regarding the stability of solutions for larger time step sizes, especially seeing the multirate nature of atmosphere-ocean coupling. Implementing higher order waveform relaxation is not a trivial task as indicated by the discussion in R uth et al. (2021). In the context of atmosphere-ocean coupling, some additional care has to be taken to respect the conservation of fluxes and energy by the interpolations. Additionally, one would have to make sure that coupling fields do not move out of physical bounds due to interpolation. We will discuss some of these aspects in more detail in Section 6.2.



### 3. Overview of the EC-Earth AOSCM

This chapter presents the main components of the EC-Earth coupled atmosphere-ocean single column model, first described in Hartung et al. (2018). We omit everything related to the sea ice component LIM since we focus on pure atmosphere-ocean coupling in this thesis. Furthermore, we disregard terms which do not play a significant role near the model interface, i.e., the sea surface. We begin with the atmosphere and ocean models in Section 3.1 and Section 3.2, respectively: OpenIFS and NEMO.<sup>1</sup> We specify the systems of equations in detail with Hartung et al. (2018) as a starting point, with adjustments and expansions based on ECMWF (2014), Lauritzen et al. (2022), and Olbers et al. (2012). Each section ends with a short overview of the model discretizations. In Section 3.3, we specify the interface boundary conditions to complete the definition of the coupling problem. A short overview of the coupling setup with OASIS3-MCT is given in Section 3.4.

Both OpenIFS and NEMO are based on a modified version of frequently used conservation laws in geophysical fluid dynamics which are called *primitive equations*. These encompass at least the conservation of mass (i.e., the continuity equation), momentum (i.e., the Navier-Stokes equations), and internal energy. The term *primitive* here refers to the fact that these equations contain the minimal set of assumptions such that they “could practically be integrated numerically” (Vallis, 2017, p. 104). As given by Vallis (2017, Section 2.2.4), three approximations are thus implied for both models: the hydrostatic approximation, the shallow-fluid approximation, and the traditional approximation. The latter two have the effect that in some terms in the equations of motion, the vertical extent or vertical velocities of the ocean/atmosphere are neglected. The hydrostatic approximation states that the vertical momentum equation is reduced to a balance between vertical pressure gradient and buoyancy force, leading to the hydrostatic equation:

$$\frac{\partial p}{\partial z} = -\rho g. \quad (3.1)$$

This also implies that (vertical) convective processes are removed from the Navier-Stokes equations and need to be parameterized instead.

In contrast to Hartung et al. (2018), we explicitly state how physical parameterizations affect the equations solved by OpenIFS and NEMO. Like this, the necessary boundary conditions at the sea surface become visible.

---

<sup>1</sup>We use the AOSCM with the same model components as described in Hartung et al. (2018): cycle 40R1 of the OpenIFS single column model and version 3.6 of NEMO.

### 3.1. OpenIFS

The primitive equations for OpenIFS in the single column setup are given in Equations (3.2) to (3.5). Equations (3.2) and (3.3) are the momentum equations for the horizontal wind components, consisting of the zonal and latitudinal wind velocities  $u$  and  $v$ , respectively. Equation (3.4) is the thermodynamic equation for the temperature  $T$ , i.e., it represents the conservation of internal energy. Equation (3.5) is the conservation equation for moisture  $q$ .

$$\frac{\partial u}{\partial t} = -\omega \frac{\partial u}{\partial p} + F_u - g \frac{\partial \mathcal{J}_u}{\partial p} + f(v - v_g) \quad (3.2)$$

$$\frac{\partial v}{\partial t} = -\omega \frac{\partial v}{\partial p} + F_v - g \frac{\partial \mathcal{J}_v}{\partial p} - f(u - u_g) \quad (3.3)$$

$$\frac{\partial T}{\partial t} = -\omega \frac{\partial T}{\partial p} + F_T - g \frac{\partial \mathcal{J}_T}{\partial p} + \frac{RT\omega}{c_p p} + \frac{g}{c_p} \frac{\partial}{\partial p} (\mathcal{F}_{\text{SW}} + \mathcal{F}_{\text{LW}}) + P_T \quad (3.4)$$

$$\frac{\partial q}{\partial t} = -\omega \frac{\partial q}{\partial p} + F_q - g \frac{\partial \mathcal{J}_q}{\partial p} + P_q \quad (3.5)$$

We use the symbol  $\phi$  to refer to a conserved quantity, which can be either one of the four atmospheric state variables  $u$ ,  $v$ ,  $T$  or  $q$ . We abbreviate partial differentiation  $\frac{\partial}{\partial \gamma} \phi$  with respect to a coordinate  $\gamma$  with the shorthand  $\partial_\gamma \phi$ . In the single column model, we only encounter derivatives with respect to time,  $\partial_t$ , or two variants of vertical partial derivatives: We can express the vertical levels of the atmosphere with respect to pressure  $p$  or altitude  $z$ , thus leading to partial derivatives  $\partial_p$  and  $\partial_z$ .<sup>2</sup> The first two terms on the right hand side in all four equations,  $-\omega \partial_p \phi$  and  $F_\phi$ , represent the vertical and horizontal advection, respectively. Here,  $-\omega$  is the vertical velocity in pressure coordinates. In the single column model, both advection terms are supplied as so-called forcing, i.e., their values are read in as boundary conditions from a file. The fourth term in the momentum equations combines the Coriolis effect with the balance of geostrophic wind and pressure gradient force,  $f$  being the Coriolis parameter<sup>3</sup>: Since the geostrophic horizontal wind  $u_g$ ,  $v_g$  is in balance with the pressure gradient force, it does not lead to a momentum tendency in the Coriolis term. By removing the geostrophic wind contribution from the Coriolis effect, one can omit the pressure gradient force in the momentum equation (e.g., They et al., 2022).

The fourth term in the thermodynamic equation represents the change of internal energy due to work on the volume.<sup>4</sup> Therein,  $R$  is the moist air gas constant,  $c_p$  is the heat capacity of moist air at constant pressure, and  $\omega$  is the vertical velocity in pressure coordinates. This

<sup>2</sup>Contrary to ECMWF (2014) and Hartung et al. (2018), we formulate the equations in vertical pressure instead of  $\eta$ -coordinates. The  $\eta$ -coordinate used in OpenIFS was first introduced in Simmons and Burridge (1981) to circumvent numerical issues around steep orography (e.g., mountains). This is not needed for pure atmosphere-ocean coupling. Directly using pressure coordinates makes the equations shorter and reduces the amount of newly introduced variables.

<sup>3</sup> $f = 2\Omega \sin(\theta)$ , where  $\Omega$  is the angular velocity and  $\theta$  denotes the latitude.

<sup>4</sup>For more background on why it takes this form, cf. Vallis (2017, Sections 1.6.1 and 2.6.2).

term has been derived by making use of the equation of state which relates density  $\rho$  and pressure  $p$ :

$$p = \rho RT. \quad (3.6)$$

In the OpenIFS single column model, the air pressure  $p$  is read in from initial conditions and kept constant for the whole simulation. This leaves a closed system of five prognostic variables ( $u, v, T, q, \rho$ ), along with five equations (3.2 to 3.6). As opposed to the 3D version of the model, no continuity equation is solved.

All other terms are related to physical parameterizations in OpenIFS.<sup>5</sup> As mentioned by Hartung et al. (2018, p. 4120), these are: "radiation, turbulence, cloud and convection parameterisation schemes as well as the non-orographic gravity wave drag, orographic gravity wave drag and surface drag." From this list, we have omitted the gravity wave drag schemes: The orographic gravity wave drag scheme is only relevant over land (ECMWF, 2014, p. IV.11). Additionally, non-orographic gravity waves are important mainly in the "middle atmosphere, comprising the stratosphere and the mesosphere" (ECMWF, 2014, Part IV, p. 67). We can thus disregard them in the atmospheric boundary layer and especially near the sea surface. The surface drag parameterization is part of the turbulence scheme and we will not discuss it in more detail here. Thus, the radiation, turbulence, cloud, and convection parameterization schemes remain. Although not mentioned in Hartung et al. (2018), the default setup includes three additional parameterizations at the ocean surface. We will now present each parameterization briefly and describe how it appears in Equations (3.2) to (3.5).

**Radiation** The radiation scheme used in OpenIFS is described in detail in ECMWF (2014, Part IV, Chapter 2). For our purposes, it suffices to acknowledge the following: In the atmosphere, radiation leads to heating or cooling of the atmospheric layers, which can be represented as an additional term in the thermodynamic equation,

$$\left(\frac{\partial T}{\partial t}\right)_{\text{rad}} = \frac{g}{c_p} \frac{\partial \mathcal{F}}{\partial p} = \frac{g}{c_p} \frac{\partial \mathcal{F}_{\text{SW}}}{\partial p} + \frac{g}{c_p} \frac{\partial \mathcal{F}_{\text{LW}}}{\partial p}, \quad (3.7)$$

where  $\mathcal{F} = \mathcal{F}_{\text{SW}} + \mathcal{F}_{\text{LW}}$  is the net radiative flux made up of a net short-wave and net long-wave radiative flux,  $c_p$  is the specific heat at constant pressure of moist air, and  $g$  is the gravitational acceleration.

**Vertical Turbulent Transport** An in-depth explanation of vertical turbulent transport in OpenIFS is given in ECMWF (2014, Chapter IV.3). We will introduce here how turbulent transport is generally modeled for the boundary layers of both the atmosphere and the ocean, which makes it easy to reuse ideas and notation in Section 3.2. A central concept

<sup>5</sup>We decided to omit the nudging (relaxation) term in the primitive equations. In typical simulations with the EC-Earth AOSCM, nudging is only utilized above the atmospheric boundary layer, cf. Hartung et al. (2018, Table 1), and is therefore not relevant near the surface.

in this context is so-called Reynolds averaging: By separating a conserved quantity  $\phi$  into a mean state  $\bar{\phi}$  and deviations  $\phi' = \phi - \bar{\phi}$  (the Reynolds decomposition), one can study how the fluctuations  $\phi'$  affect the conservation equation for  $\phi$ . This leads to an additional term  $\overline{w'\phi'}$ , which represents vertical turbulent transport.<sup>6</sup> Since the fluctuations act on subgrid scales, we have to approximate this term. The turbulent closure hypothesis states that these turbulent fluxes can be expressed in terms of large scale features, i.e., using information available on the resolved scales. This results in the following common way of expressing vertical turbulence:

$$\left(\frac{\partial\phi}{\partial t}\right)_{\text{turb}} = \frac{1}{\rho} \frac{\partial}{\partial z} \mathcal{J}_\phi. \quad (3.8)$$

Here,  $\mathcal{J}_\phi$  is the vertical turbulent flux of  $\phi$ , positive downwards (we choose the sign convention analogously to ECMWF, 2014).<sup>7</sup>  $\rho$  is the fluid density. Both NEMO and OpenIFS use the Eddy-Diffusivity Mass-Flux (EDMF) framework to define the vertical turbulent flux,

$$\mathcal{J}_\phi = \left( \rho K_\phi \frac{\partial\phi}{\partial z} - M(\phi_u - \phi) \right). \quad (3.9)$$

The EDMF framework, first described by Siebesma et al. (2007), consists of two terms: The first-order turbulence closure  $\rho K_\phi \partial_z \phi$  where  $K_\phi$  is commonly called *eddy viscosity* if it relates to the momentum equations and *eddy diffusivity* otherwise (Lemarié et al., 2014).<sup>8</sup> The second term is the advective mass flux  $M(\phi_u - \phi)$ , “used for convective transport in the cumulus cloud layer” (Siebesma et al., 2007, p. 1231). We will go into more details on the definition of  $\mathcal{J}_\phi$  in Section 3.3, since it plays an important part in the interface conditions. For the primitive equations, we stick with the shorter form of Equation (3.8), reformulated from height into pressure coordinates.<sup>9</sup>

In the OpenIFS documentation,  $\mathcal{J}_\phi$  is in general defined for the generalized liquid water static energy  $s_l$  and the specific total water  $q_t$ , as opposed to  $T$  and  $q$ . Near the surface, it is defined for the dry static energy  $s$  and  $q$ . In this work, we consider the atmosphere near the surface and use  $\mathcal{J}_T$  instead of  $\mathcal{J}_s$ . Since  $s = gz + c_p T$ , the two can be related, and this transformation must be done inside OpenIFS since the thermodynamical equation is given in terms of  $T$ , not  $s$ .

---

<sup>6</sup>A detailed derivation of this term is given in Appendix A.1.

<sup>7</sup>In contrast to the IFS documentation, we denote the flux by  $\mathcal{J}$  and not  $J$  following, e.g., Lauritzen et al. (2022) and Olbers et al. (2012).

<sup>8</sup>In the IFS documentation,  $K_\phi$  is referred to as the exchange coefficient (ECMWF, 2014, Chapter IV.3).

<sup>9</sup>To move from height to pressure coordinates, we use the chain rule and the hydrostatic balance, yielding:

$$\frac{\partial\phi}{\partial z} = \frac{\partial\phi}{\partial p} \frac{\partial p}{\partial z} = -\rho g \frac{\partial\phi}{\partial p}.$$

**Convection** The formation of clouds due to moist convection, i.e., atmospheric vertical motion where “phase changes of water play an appreciable role” (American Meteorological Society, 2022, Moist Convection), is parameterized in OpenIFS using a bulk mass flux scheme, cf. ECMWF (2014, Chapter IV.6) and Fitch (2022). This affects the mass flux term in the vertical turbulent flux  $\mathcal{J}_\phi$ , see Equation (3.9) and does not need further treatment in the primitive equations.

**Clouds** The cloud scheme in OpenIFS influences the thermodynamic and moisture equations. Equations (7.12) and (7.13) in ECMWF (2014, Part IV) specify this microphysics parameterization in more detail. As can be seen therein, the terms related to the cloud scheme are purely local, i.e., they do not contain additional vertical derivatives. This means that no (interface) boundary conditions are necessary for the cloud scheme. An explicit formulation of the cloud scheme would significantly increase the complexity of the conservation equations for temperature and humidity, without adding information related to coupling. For this reason, we use  $P_q$  and  $P_T$  in the OpenIFS equation set as a placeholder for the impact of clouds.

**Parameterizations for Near Surface Ocean Effects** OpenIFS contains three additional parameterizations near the sea surface which affect the temperature and humidity boundary conditions (ECMWF, 2014, Section IV.8.9). They capture the cool skin and warm layer effect, as well as salinity effects on the saturation specific humidity  $q_{\text{sat}}$  at the sea surface. For our numerical experiments, we decided to turn off the warm layer effect, corresponding to setting the parameter `LEOCWA` to `FALSE`. Since the vertical resolution of NEMO is high enough to resolve the typical size of the warm layer (Fairall et al., 1996), the warm layer parameterization is not necessary in a coupled atmosphere-ocean simulation.<sup>10</sup> The other two parameterizations correspond to multiplying the sea surface temperature and humidity seen by the atmosphere in the vertical turbulent flux parameterization by scalar factors. We do not need to adjust the primitive equations to incorporate these effects.

**Model Discretization** We use the OpenIFS-SCM with 60 vertical levels where the resolution decreases with higher altitudes (lower pressure). To compute the dynamic tendencies in the primitive equations, we use the two-time-level semi-Lagrangian scheme in all experiments (a Eulerian scheme is also available, cf. Hartung et al., 2018). In general, the time step size is shared by all parts of the model, except for the vertical turbulent fluxes, which are computed twice per time step (ECMWF, 2014, Chapter IV.3). The time step size for the computationally intensive radiation scheme can be set independently.

---

<sup>10</sup>In fact, the warm layer effect parameterization is turned off for coupled simulations in later versions of the IFS (ECMWF, 2021, Part IV, p. 167).

## 3.2. NEMO

The primitive equations used in the one-dimensional setup of NEMO 3.6 are given in Hartung et al. (2018), Madec et al. (2017), and Reffray et al. (2015). We present a modified version of the equations given therein to have a more consistent formulation for both ocean and atmosphere.

In addition to the already mentioned assumptions for the primitive equations, given at the beginning of this chapter, NEMO makes use of the Boussinesq hypothesis (i.e., density variations are neglected except in their contribution to the buoyancy force) and the assumption of incompressibility. These simplify the continuity equation to  $\nabla \cdot \mathbf{u} = 0 = \partial_x u + \partial_y v + \partial_z w$ , where  $\mathbf{u} = (u, v, w)^T$  is the velocity vector. In the single column model, vanishing horizontal derivatives are assumed, and thus

$$\frac{\partial}{\partial z} w = 0 \quad (3.10)$$

remains as the continuity equation. With the boundary conditions of vanishing vertical velocity at the sea surface and sea floor, i.e.,  $w|_{z=0} = w|_{z=-H} = 0$ , we thus obtain that the vertical velocities in the NEMO single column model are zero everywhere.

The single column version of NEMO solves the 1D Navier-Stokes equations along with a nonlinear equation of state (polyEOS80-bsq),  $\rho = \rho(\theta, S, p)$ , which couples the conservation equations for the two tracers  $\theta$  and  $S$  to the momentum equations.

$$\frac{\partial u}{\partial t} = -\frac{1}{\rho} \frac{\partial}{\partial z} \mathcal{J}_u \quad +fv \quad (3.11)$$

$$\frac{\partial v}{\partial t} = -\frac{1}{\rho} \frac{\partial}{\partial z} \mathcal{J}_v \quad -fu \quad (3.12)$$

$$\frac{\partial \theta}{\partial t} = -\frac{1}{\rho} \frac{\partial}{\partial z} \mathcal{J}_\theta + \frac{1}{\rho_{\text{ref}} c_p} \frac{\partial I(Q_{\text{sr}}, z)}{\partial z} \quad (3.13)$$

$$\frac{\partial S}{\partial t} = -\frac{1}{\rho} \frac{\partial}{\partial z} \mathcal{J}_S \quad (3.14)$$

The four oceanic prognostic variables are the horizontal velocity components  $u$  and  $v$ , the potential temperature  $\theta$ , and the practical salinity  $S$ .<sup>11</sup> As in Equations (3.2) to (3.5), we use  $\partial_z \mathcal{J}_\phi$  to refer to the vertical turbulent transport of a conserved quantity  $\phi$  which can be one of the four state variables. While  $\rho$  denotes density computed via the equation of state, the constant ocean reference density  $\rho_{\text{ref}} = 1035 \text{ kg m}^{-3}$  is used in parts of the equation system.  $I(Q_{\text{sr}}, z)$  is the penetrative part of the solar surface heat flux, signifying solar heating of the ocean near the sea surface due to the net short wave radiation from the atmosphere,  $Q_{\text{sr}}$ .

---

<sup>11</sup>In contrast to the absolute temperature  $T$ ,  $\theta$  is the temperature a fluid parcel would have if it were raised adiabatically to the sea surface (American Meteorological Society, 2022).

As before,  $f$  is the Coriolis parameter.<sup>12</sup>

Both Reffray et al. (2015) and Hartung et al. (2018) include a term  $\mathcal{E} - \mathcal{P}$  in the salt conservation equation, where  $\mathcal{E}$  and  $\mathcal{P}$  are the evaporation and precipitation fluxes, respectively. However, these fluxes naturally only act at the sea surface as part of the salinity boundary condition, cf. Equation (3.22). They do not affect the whole vertical column, as would be implied by including them in the conservation equation, therefore we discard them here.

**Physical Parameterizations** Contrary to Hartung et al. (2018), we omit the terms for physical parameterization  $P_\phi$ . Three parameterizations are mentioned therein, all of which we have moved into other terms:

- The mentioned turbulent kinetic energy-dependent eddy coefficient/“1.5 turbulent closure” (Hartung et al., 2018, p. 4121) precisely refers to the parameterization used to define the vertical turbulent flux  $\mathcal{J}_\phi$ , i.e., the turbulent closure scheme of NEMO.
- Similarly, the parameterization of Langmuir circulations is part of the turbulent closure scheme (Couvelard et al., 2020).
- The effect of chlorophyll on heating is considered in the computation of the penetrative solar surface heat flux  $I(Q_{sr}, z)$  (Madec et al., 2017, Section 5.4.2).

**Model Discretization** We use NEMO in the configuration with 75 fixed vertical levels, where the resolution is substantially higher near the ocean surface than towards the sea floor (Madec et al., 2017, Section 4.3). The primitive equations are discretized using second-order finite differences on a staggered grid, which is why, in the 1D configuration, NEMO still uses  $3 \times 3$  grid cells horizontally. Due to the vanishing horizontal derivatives in 1D, these nine cells have identical values. Different time integration schemes are used for the various terms in the primitive equations and we refer to the NEMO documentation for more information (Madec et al., 2017).

### 3.3. Interface Boundary Conditions

We now bridge the gap from the generalized atmosphere-ocean coupling problem in Section 2.1 to the equations solved by the EC-Earth AOSCM. Equations (3.2) to (3.5) and Equations (3.11) to (3.14) represent  $\mathcal{L}_{\text{atm}} \mathbf{U}^a = f_{\text{atm}}$  and  $\mathcal{L}_{\text{oce}} \mathbf{U}^o = f_{\text{oce}}$ , respectively. We do not discuss the boundary conditions at the sea floor and top of atmosphere since these do not play a significant role in atmosphere-ocean coupling. Viewing the problem as an example of domain decomposition (the perspective of Chapter 2) brought about that we require a continuity of air-sea fluxes across the model interface  $\Gamma$ . If we just consider the equations

<sup>12</sup>In comparison to Hartung et al. (2018), we have left out the nudging terms in the primitive equations as this feature is not implemented in the current version of the model (Deppenmeier et al., 2020).

seen in this chapter, the interface conditions are necessary to properly define all vertical derivatives appearing in the primitive equations at the sea surface: Simply stated, for a term of the form  $\partial_z \bullet = -\rho g \partial_p \bullet$ , the associated interface boundary condition provides a value  $\bullet|_\Gamma$ . In the EC-Earth AOSCM such terms appear for all vertical turbulent fluxes  $\mathcal{J}_\phi$ , the radiative fluxes  $\mathcal{F}_{\text{SW}}$  and  $\mathcal{F}_{\text{LW}}$ , and the solar surface heat flux  $I(Q_{\text{sr}}, z)$ . These should be formulated in a way that ensures flux continuity across the sea surface  $\Gamma$ .

Although the primitive equations in the atmosphere are given with respect to pressure, we switch back to the  $z$ -coordinate at the sea surface for consistency between both models. For this section, we adopt the following notation: The  $z$ -coordinate is defined as positive *upwards*, with the sea surface being located at  $z = 0$ . The lowest grid point in the atmospheric grid is located at  $z \approx 10\text{m}$ . We will refer to this point as  $z_1$ . NEMO's computational grid begins directly at the surface, but the values of the state variables are actually computed in the center of each grid cell (Madec et al., 2017, Chapter 4). Since the first ocean grid cell has a length of 1m, the highest grid point in the ocean grid is located at  $z \approx -0.5\text{m}$ , which we will refer to as  $z_{-1}$ . We distinguish atmospheric and oceanic quantities by the superscripts  $a$  and  $o$  where necessary, as in Chapter 2.

**General Approach for  $\mathcal{J}_\phi$**  As an interface condition, it should generally hold that

$$\mathcal{J}_\phi^a|_{z=0} = \mathcal{J}_\phi^o|_{z=0}. \quad (3.15)$$

For the turbulent fluxes  $\mathcal{J}_\phi$ , we need to take into account five oceanic and atmospheric state variables  $\phi$ : the horizontal velocities  $u$  and  $v$ , temperature  $T$ , moisture  $q$ , and salinity  $S$ . The fluxes for the first four are computed inside OpenIFS, using the EDMF scheme introduced in Equation (3.9). At the surface, the mass flux term  $M(\phi_u - \phi)$  is assumed to be zero (ECMWF, 2014, Section IV.3.1). Thus,

$$\mathcal{J}_\phi = \rho K_\phi \frac{\partial \phi}{\partial z}. \quad (3.16)$$

The common approximation for  $\mathcal{J}_\phi|_{z=0}$  is based on a nonlinear bulk formula which, in its simplest form, depends on the jump between both media

$$\mathcal{J}_\phi|_{z=0} = \rho K_\phi \frac{\partial \phi}{\partial z} \Big|_{z=0} = \rho C_\phi \|\mathbf{u}^a(z_1) - \mathbf{u}^o(z_{-1})\| (\phi^a(z_1) - \phi^o(z_{-1})), \quad (3.17)$$

with  $\mathbf{u} = (u, v)^T$  and  $C_\phi$  the *transfer coefficient*, a scalar factor which nonlinearly depends on, e.g.,  $\mathbf{u}$  and  $\phi$ .<sup>13</sup> This type of equation stems from boundary layer theory for fluid mechanics, central concepts used here are the law of the wall in general (Schlichting & Gersten, 2017) and Monin-Obukhov similarity theory (Monin & Obukhov, 1954) for the

---

<sup>13</sup>In other works, e.g., Lemarié et al. (2014),  $C_\phi$  is called exchange coefficient, which could lead to confusion since this is also a name for  $K_\phi$  (ECMWF, 2014, Section IV.3).



atmosphere in particular.<sup>14</sup> The nonlinear function  $\mathbf{F}_{\text{oa}}(\mathbf{U}^a, \mathbf{U}^o)$  we introduced in Chapter 2 is an abstraction of Equation (3.17).

The bulk formulas solved in OpenIFS differ from Equation (3.17). For instance, the surface currents in the ocean are not actually used in the computation of the turbulent fluxes, i.e., it is assumed that  $\mathbf{u}_o(z_{-1}) = 0$ .<sup>15</sup> The norm of the surface winds  $\|\mathbf{u}_a(z_1)\|$  is replaced with the *wind speed*

$$|U(z_1)| = \sqrt{\|\mathbf{u}^a(z_1)\|_2^2 + w_*^2}, \quad (3.18)$$

where  $w_*$  is the *free convection velocity scale* which depends on atmospheric temperature and moisture near the surface (ECMWF, 2014, Part IV, p.36).

**Horizontal velocities** For  $\mathcal{J}_u$  and  $\mathcal{J}_v$ , the computations are combined into a vertical turbulent momentum flux vector  $\mathcal{J}_M = \rho^a \boldsymbol{\tau} = (\mathcal{J}_u, \mathcal{J}_v)^T$ . The vector  $\boldsymbol{\tau}$  is often called the *wind stress*. The momentum flux at the surface is computed as follows:

$$\mathcal{J}_M|_{z=0} = \rho^a C_M |U(z_1)|^2 \frac{\mathbf{u}^a(z_1)}{\|\mathbf{u}^a(z_1)\|}, \quad (3.19)$$

i.e.,  $\mathcal{J}_M$  and  $\mathbf{u}^a(z_1)$  are assumed to be collinear.<sup>16</sup>

**Sensible and latent heat flux** For the temperature and moisture boundary conditions, the sensible and latent heat flux are computed. The sensible heat flux  $\mathcal{J}_{T,a}$  at the surface is defined as

$$\mathcal{J}_T^a|_{z=0} = \rho^a c_p^a C_H |U(z_1)| (T^a(z_1) - T_{\text{skin}}), \quad (3.20)$$

where  $T_{\text{skin}}$  is the *skin temperature* of the ocean.<sup>17</sup> For our purposes, we can consider this to be equal to the sea surface temperature sent by NEMO, corresponding to the absolute temperature on the highest model level:  $T^o(z_{-1})$ .<sup>18</sup> The conversion from potential temperature  $\theta$  to absolute temperature  $T$  happens inside NEMO before sending the values to OpenIFS.

The latent heat flux  $\mathcal{J}_q$  at the sea surface is computed as

$$\mathcal{J}_q|_{z=0} = \rho^a C_Q |U(z_1)| (q(z_1) - q_{\text{sat}}(T_{\text{skin}})). \quad (3.21)$$

<sup>14</sup>We note that To find  $K_\phi \partial_z \phi$ , an implicit equation has to be solved iteratively. For some of the quantities, values from the previous time step are used. For example,  $C_\phi$  depends on the Obukhov length  $\mathcal{L}$ , which is computed based on a friction velocity  $u_*$  defined in terms of  $\mathcal{J}_\phi$ .

<sup>15</sup>Neglecting the ocean currents has implications on the stability and accuracy of the coupling scheme, cf. Connors and Ganis (2011) and Renault et al. (2019).

<sup>16</sup>This is not explicitly stated but can be seen in ECMWF (2014, Part IV, Equations 3.9, 3.10).

<sup>17</sup>For the sake of completeness, we mention that the parameterization for vertical turbulent fluxes uses the specific heat capacity of *dry* air, i.e.,  $c_p = c_{p,\text{dry}}$  (ECMWF, 2014, Section IV.3.1). In the radiation scheme and the dynamical core,  $c_p = c_{p,\text{moist}}$ .

<sup>18</sup>To be precise, OpenIFS slightly modifies the value received by NEMO in the cool skin parameterization (ECMWF, 2014, Section IV.8.9).

The saturation humidity at the surface  $q_{\text{sat}}$  is again defined in terms of skin temperature. The computation of  $q_{\text{sat}}$  is affected by the salinity effect parameterization mentioned at the end of Section 3.1.

**Salinity flux** The boundary condition for the salinity flux is computed in NEMO as:

$$\mathcal{J}_S|_{z=0} = (\mathcal{E} - \mathcal{P}) \cdot S(z_{-1}), \quad (3.22)$$

where  $\mathcal{P}$  and  $\mathcal{E}$  are the precipitation and evaporation fluxes computed in OpenIFS. This boundary condition incorporates the salinity change due to freshwater influx and outflux at the sea surface into a *virtual* salt flux across the interface.<sup>19</sup> This ensures salt conservation while keeping vertical velocities equal to zero at the sea surface. More information on this boundary condition can be found in (Huang, 1993), where it is referred to as the mixed boundary condition for the salinity balance.

**Solar surface heat flux** For the term  $\partial_z I(Q_{\text{sr}}, z)$  in NEMO's thermodynamic equation, eq. (3.13), OpenIFS has to provide the net shortwave radiation at the sea surface. The radiative fluxes in OpenIFS are computed at "half-levels", where the lowest one corresponds to surface pressure, i.e.,  $z = 0\text{m}$ . Thus,

$$Q_{\text{sr}} = (1 - \alpha) \mathcal{F}_{\text{SW}}|_{z=0}, \quad (3.23)$$

where  $\alpha$  denotes the albedo of the ocean surface.

**Radiation boundary condition** The boundary conditions used for the radiation scheme in the EC-Earth AOSCM are, to our knowledge, not documented. The ocean is not a source of shortwave radiation but impacts it depending on the surface albedo (which determines how much atmospheric shortwave radiation is reflected). In ice-free conditions, NEMO does not provide albedo values to OpenIFS. On the other hand, OpenIFS provides the *net* shortwave radiation to NEMO. Thus, we can safely assume that OpenIFS computes its own boundary condition for  $\mathcal{F}_{\text{SW}}$  based on an internal value for the ocean surface albedo.

As a large body with thermal energy, the ocean acts as a source of longwave radiation  $\mathcal{F}_{\text{LW}}$  for the atmosphere. While this is not clearly stated in the IFS documentation or the AOSCM code, we assume that the boundary condition for  $\mathcal{F}_{\text{LW}}|_{z=0}$  considers the ocean as a black body  $\sigma T_{\text{skin}}^4$ , where  $\sigma$  denotes the Stefan-Boltzmann constant.

**Synthesis** Although the interface conditions between ocean and atmosphere could be treated by a third component, most of the boundary conditions needed by both models are computed in OpenIFS, i.e., the atmosphere model. The momentum flux  $\mathcal{J}_M|_{z=0} =$

---

<sup>19</sup>In comparison to Hartung et al. (2018), we have omitted  $S_t$  in the salinity boundary condition: The rate of change of the sea ice thickness budget  $S_t$  is zero in ice-free simulations.

$(\mathcal{J}_u|_{z=0}, \mathcal{J}_v|_{z=0})^T$  is computed by taking into account only the surface wind, not the ocean currents. The resulting flux is used as a boundary condition in both models:

$$\mathcal{J}_u|_{z=0} = \rho^o K_u^o \frac{\partial u^o}{\partial z} \Big|_{z=0} = \rho^a K_u^a \frac{\partial u^a}{\partial z} \Big|_{z=0} = \rho^a C_M |U(z_1)|^2 \frac{u^a(z_1)}{\|\mathbf{u}^a(z_1)\|} \quad (3.24)$$

$$\mathcal{J}_v|_{z=0} = \rho^o K_v^o \frac{\partial v^o}{\partial z} \Big|_{z=0} = \rho^a K_v^a \frac{\partial v^a}{\partial z} \Big|_{z=0} = \rho^a C_M |U(z_1)|^2 \frac{v^a(z_1)}{\|\mathbf{u}^a(z_1)\|}. \quad (3.25)$$

For the temperature boundary condition in the atmosphere, the sensible heat flux  $\mathcal{J}_T^a|_{z=0}$  is computed in OpenIFS. For moisture  $q$ , the latent heat flux  $\mathcal{J}_q|_{z=0}$  is used. We cannot use the same boundary condition for ocean temperature as in the atmosphere: The turbulent flux of temperature in the ocean is affected by the sensible and latent heat fluxes, but also by the net longwave radiation. These are summarized into a non-solar heat flux  $Q_{\text{ns}}$  which is accumulated in OpenIFS and sent to NEMO:

$$Q_{\text{ns}} = \mathcal{J}_T^a|_{z=0} + \mathcal{J}_q|_{z=0} + \mathcal{F}_{\text{LW}}|_{z=0} - \sigma T_{\text{skin}}^4. \quad (3.26)$$

Thus we arrive at the following interface boundary conditions for the vertical turbulent temperature and moisture fluxes:

$$\mathcal{J}_T^a|_{z=0} = \rho^a c_p^a K_T^a \frac{\partial T^a}{\partial z} \Big|_{z=0} = \rho^a c_p^a C_H |U(z_1)| (T^a(z_1) - T_{\text{skin}}) \quad (3.27)$$

$$\mathcal{J}_T^o|_{z=0} = \rho^o c_p^o K_T^o \frac{\partial T^o}{\partial z} \Big|_{z=0} = Q_{\text{ns}} \quad (3.28)$$

$$\mathcal{J}_q|_{z=0} = \rho^a K_q^a \frac{\partial q}{\partial z} \Big|_{z=0} = \rho^a C_Q |U(z_1)| (q(z_1) - q_{\text{sat}}(T_{\text{skin}})). \quad (3.29)$$

For radiative fluxes, OpenIFS uses an internally defined boundary condition based on  $T_{\text{skin}}$  and the ocean surface albedo. Evaporation  $\mathcal{E}$ , precipitation  $\mathcal{P}$ , and net short wave radiation  $Q_{\text{sr}}$  are sent from OpenIFS to NEMO to complete the boundary conditions for the vertical turbulent transport of salinity and penetrative solar radiation (Equations 3.22 and 3.23).<sup>20</sup> Figure 3.1 summarizes the variables which are exchanged between the models in order to compute the interface boundary conditions  $C_{\text{atm}} \mathbf{U}^a$  and  $C_{\text{oce}} \mathbf{U}^o$ .<sup>21</sup>

<sup>20</sup>These interface conditions cannot simply be classified in terms of Dirichlet, Neumann, or Robin boundary conditions. We do not go into detail here but have included a short discussion in Appendix A.2.

<sup>21</sup>In Hartung et al. (2018), it was stated that an additional variable, the temperature sensitivity of non-solar heat fluxes, is sent from OpenIFS to NEMO. In fact, this variable is only relevant for areas with ice, i.e., for the sea ice model LIM which we do not consider here.

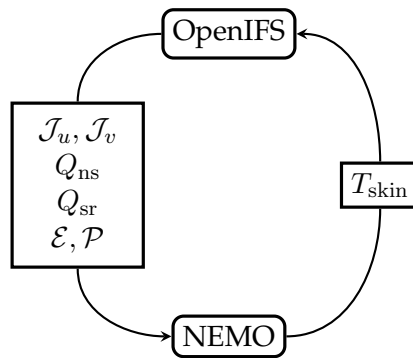


Figure 3.1.: Overview of exchanged data between OpenIFS and NEMO during a coupled run of the EC-Earth AOSCM.

### 3.4. Coupling with OASIS

OpenIFS and NEMO are coupled with the OASIS3-MCT coupling software, which we will refer to as OASIS. Detailed information on this tool can be found in the corresponding papers (Craig et al., 2017; Valcke, 2013) and the user guide (Valcke et al., 2015). We give a short overview of some fundamental concepts in OASIS in Appendix A.3. The OASIS coupler takes care of communicating data between model components and can, in that process, apply various transformations to the coupling fields. This includes taking time averages and regridding data using interpolation or various other techniques to map non-matching grids onto each other. OASIS is designed to take care of these tasks while ensuring high computational performance and minimizing the error introduced by transformations. In addition to communication, it supports I/O of variables defined by each model component.

From a technical standpoint, the model components call the OASIS library to instantiate and finalize the coupler, declare coupling fields, and trigger communication. These calls are part of the model code and changes here (e.g., adding new coupling fields) necessitate recompilation of the EC-Earth AOSCM. During a coupled run, OASIS decides whether data has to be sent at a given model time step, based on user input provided in a configuration file, the *namcouple*. Significant aspects of the coupling setup can be changed at runtime using this file, including the coupling window size and transformation options. Overall, essential parts of the coupling logic are transparent to the model components.

The algorithm used by default in the EC-Earth AOSCM is the parallel algorithm described in Section 2.2. As NEMO uses a  $3 \times 3$  grid, whereas OpenIFS uses a  $1 \times 1$  grid, OASIS also takes care of remapping the interface values onto each other. From NEMO to OpenIFS, only the value in the center of the grid is sent to OpenIFS. Vice versa, the single value supplied by OpenIFS is broadcast to all nine grid points of the NEMO grid.

## 4. Implementation

The EC-Earth AOSCM uses the parallel algorithm described in Section 2.2 as a coupling scheme. In the context of this master’s thesis, we want to test the impact of more advanced coupling algorithms on the numerical solution produced by the model. Namely, the goal is to support both the sequential atmosphere-first and ocean-first algorithms, as well as Schwarz waveform relaxation (SWR). Regarding SWR, the goal is to support the iterative version of the three non-iterative coupling schemes, i.e., exchange the average of the interface values over one coupling window. Based on this, more advanced versions of Schwarz waveform relaxation can be implemented, as we will discuss in Section 6.2.

The AOSCM is based on the Earth system model EC-Earth 3, thus a lot of the implementation challenges of Schwarz waveform relaxation already appear in the single column model. The model components of the EC-Earth AOSCM are large Fortran codes, continuously and independently developed by different research groups. Our implementation is for the most part independent of the model components, OpenIFS and NEMO. We treat them as black boxes as much as possible to be minimally invasive, cf. the discussion in Gatzhammer (2014, Chapter 3). Instead, we base our SWR implementation on concepts of the OASIS coupler and on the typical workflow when using climate models coupled with OASIS: Runtime configuration scripts determine settings for the compiled model binaries and the simulation output is stored at a predetermined destination with a known file structure. We hope that this approach allows developers of other climate models to reuse ideas of our implementation.

Our solution allows model users to choose coupling schemes at runtime, not compile time, similar to other simulation settings like simulation length, time step size, vertical resolution, etc. Especially considering the size and complexity of the source code, this is a far more user-friendly approach than requiring an in-code change and recompilation.

In this chapter, we give an overview of the implementation aspects relevant to support different coupling schemes for the EC-Earth AOSCM. We restrict ourselves to a broad overview of the necessary changes, while technical notes and examples are given in the appendix. The chapter starts with a short description how sequential algorithms can be supported for a model coupled using the OASIS3-MCT coupler. We continue with a description of our SWR implementation in Section 4.2. In Section 4.3 we tackle the question of suitable convergence criteria for the Schwarz iterations.

### 4.1. Sequential Coupling Algorithms

Whether an AOSCM experiment uses the parallel algorithm or one of the sequential ones described in Section 2.2.2 can be controlled using a configuration file for OASIS, the *namcouple*. More specifically, one has to change the `LAG` parameters provided therein. This was previously mentioned and utilized in Marti et al. (2021) and Streffing et al. (2022). In Appendix A.4, we explain how to set parameters based on the desired coupling scheme.

Users of the AOSCM do not adapt the *namcouple* directly, but use an XML file to set simulation parameters. There is no reason to deviate from the workflow for this feature, thus we have simply implemented a new XML parameter for OASIS where users can set the coupling scheme to an integer value. The default value is 0 and represents the parallel algorithm. 1 and 2 denote the sequential atmosphere-first and ocean-first coupling scheme, respectively. In case a user chooses a different value, the program aborts with an error explaining which parameter in the XML to adapt. This provides additional functionality to users while keeping the added complexity minimal, similar to adding a toggle for a physical parameterization scheme.

### 4.2. Schwarz Waveform Relaxation

Preceding the work in this thesis, there have been two implementations of Schwarz waveform relaxation in coupled Earth system models, both of which use OASIS as coupling software. Marti et al. (2021) adapted the IPSL-CM6 model, a three-dimensional coupled general circulation model. To support Schwarz waveform relaxation, the time loops of the ocean and atmosphere model were extended. The resulting algorithm uses a constant number of iterations for each Schwarz window. The Schwarz window size can be adjusted to span a single coupling window, multiple coupling windows, or even the full simulation time. As gets clear from the published code, this implementation is not model-agnostic and requires significant adjustments in each component.

The EC-Earth AOSCM uses the same ocean and sea-ice component as IPSL-CM6 (NEMO and LIM/SI3) but a different atmosphere model. Reusing the approach of Marti et al. (2021) would thus require a significant amount of implementation work which would have to be repeated each time a component of the EC-Earth AOSCM is updated.

A second implementation approach was developed as part of the COCOA project, although it has not been published at this point. As described in the project report (Valcke, 2021), Schwarz waveform relaxation was tested for the single column version of the CNRM-CM6 model. In this approach, OASIS itself is utilized to support repeated evaluations of the same time interval. For this reason, no significant adjustments of the model code are required. Instead, the runtime settings of OASIS are adapted using external scripts which take control of executing Schwarz iterations. Since OASIS is used by “more than 35 different climate modeling groups around the world” (Valcke, 2013, p. 374), basing the implementation on OASIS concepts might lead to easier reuse of different coupling

approaches than coming up with an implementation based on OpenIFS and NEMO.

The approach is significantly simpler to implement and less dependent on the underlying model components. It works as follows: The first Schwarz iteration is equivalent to a regular coupled AOSCM run, with the same available settings (including switching between the parallel or one of the sequential algorithms). During the simulation, OASIS saves the values of the interface variables at every coupling time step to an output file.

At  $t_0$  of the second Schwarz iteration, the model components read in the coupling fields previously exchanged at  $t_1$ , using the output file saved by OASIS.<sup>1</sup> Generally in the second iteration at a coupling time step  $t_n$ , the models read in the values sent at  $t_{n+1}$  during the first iteration. Recalling the parallel algorithm (Figure 2.3), this reduces the coupling lag: The values exchanged at a coupling time step  $t_n$  in the parallel algorithm represent the average of the respective coupling variable over the interval  $[t_{n-1}, t_n]$ . By shifting the coupling data in time, we reduce the time shift and re-synchronize the air-sea fluxes seen by the model, as described in Section 2.3. Interface data computed during the second iteration is not directly sent from one component to the other. Instead, it is once again written out to a file, for use during the next iteration. This way, model components still run in parallel and are controlled by OASIS, but they do not directly exchange data during any iteration except for the first one. Information flows between the iterations, as seen in Figure 2.5.

In the SWR implementation by Marti et al. (2021), the model executable contains the full logic for doing Schwarz iterations, since this functionality is fundamentally part of the model code. By contrast, the approach by Valcke (2021), which we adopt here, keeps the Schwarz algorithm separate from the model code. Doing an SWR experiment thus implies running the model executable *multiple times* with different configuration settings. Each time the compiled model is executed, a single Schwarz iteration is done. In between the iterations, an outer layer written in Python is responsible for controlling the Schwarz algorithm. This includes the following tasks:

- Select the correct OASIS configuration file (see Appendix A.5 for an example illustrating the differences in configuration files between the first and later Schwarz iterations).
- Shift the time stamps of the OASIS output files to read them in correctly in the next iteration.
- Remap the coupling variables from the  $1 \times 1$  OpenIFS grid to the  $3 \times 3$  NEMO grid and vice versa. This is a task that OASIS can take care of directly when communicating data between components. However, OASIS does not support regridding or remapping when coupling fields are only read in or written out (Valcke et al., 2015, pp. 31–32).

<sup>1</sup>We use the same notation for time steps as in Chapter 2:  $t_n = n \cdot \Delta t$  is the  $n$ -th coupling time step, and  $\Delta t$  is the length of a coupling period.

## 4. Implementation

---

- Rename model output directories to avoid name clashes in subsequent iterations.
- Decide whether to continue iterating based on some criterion (see Section 4.3).

The functionality of the outer layer can also be extended to include other features which are not part of the basic Schwarz algorithm described in Section 2.3. In the current implementation, we use the Python wrapper to also reduce the amount of model output and store general information about the experiment. In addition, one could use it to test acceleration methods for SWR, examples being dynamic Aitken or Quasi-Newton approaches (Gatzhammer, 2014; R uth et al., 2021).

In our implementation, OpenIFS and NEMO only exchange interface data averaged over one coupling period, instead of higher-order interpolants of the coupling variables. In Section 6.2, we discuss to what extent a more accurate approximation of the true flux could be exchanged with the implementation idea we use in this thesis. The implemented solution is equivalent to Schwarz waveform relaxation with piecewise constant interface data averaged over each coupling period  $\Delta t$ . The Schwarz window size is equal to the simulation time. If the parallel algorithm is chosen in the first iteration, the iterations are equivalent to an additive Schwarz method. If a sequential coupling scheme is used in the first iteration, the subsequent iterations are equivalent to those of a multiplicative Schwarz method, albeit one which is inefficiently implemented. See Appendix A.6 for an explanation of this latter observation. We have checked for example simulations that the converged SWR solutions do not differ between the multiplicative and additive SWR method.

Because the simulation time determines the Schwarz window size, more iterations are necessary for convergence when doing longer simulations. If the algorithm converges, it yields the same solution as an SWR implementation where the Schwarz window size is chosen equal to the coupling period: The output at time  $t$  depends only on previous model time steps, prescribed forcing, and the interface conditions. Thus, if the coupling time step sizes are equal,  $[t_0, t_1]$  produces the exact same iterations, no matter the size of the Schwarz window.<sup>2</sup> The converged result of  $[t_0, t_1]$  propagates to later time steps, which is equivalent to first iterating over  $[t_0, t_1]$  until convergence, then continuing with  $[t_1, t_2]$ , etc.

In summary, we have implemented basic Schwarz waveform relaxation as described in Section 2.3 with minimal overhead and while treating OpenIFS and NEMO as black boxes, except for manually remapping the two grids between iterations. The functionality is available at runtime, easily extensible and, if necessary, reversible.

---

<sup>2</sup>This is also a way to verify our implementation: We have confirmed that the model output during the first coupling window is exactly equal in every Schwarz iteration, independent of the simulation length, i.e., the Schwarz window size. This is not the case for later coupling windows until the iterations have converged.



### 4.3. SWR Convergence Criteria

For the single column model, the computational effort is low enough to do a fixed “large” number of iterations which ensures that the Schwarz method will have converged by the end. Nevertheless, we would like to use a runtime convergence estimator to save computational resources and as a case study for other coupled climate models. We connect convergence criteria previously used in the context of atmosphere-ocean coupling (Lemarié et al., 2014; Marti et al., 2021; Valcke, 2021) with those used in other multiphysics coupling libraries (e.g., Mehl et al., 2016; R uth et al., 2021)

#### 4.3.1. Variables Included in Error Estimates

To measure convergence, we only use variables exchanged at the coupling interface, similar to other implementations of SWR (R uth et al., 2021). We take into account all interface data already exposed to OASIS, i.e., the variables depicted in Figure 3.1, averaged over each coupling window.

Marti et al. (2021) only take into account the sea surface temperature, since that is the only coupling variable sent by the ocean to the atmosphere. They motivate this by stating that, as soon as the SST converges, the atmosphere “by construction computes the same fluxes” (Marti et al., 2021, p. 2972) in each iteration. This argument holds especially when the changes of the coupling variable across iterations are very small, e.g., at a magnitude near machine precision. However, Marti et al. (2021) used convergence tolerances based on whether differences are physically negligible or not. It is unclear a priori whether a physically negligible change in SST might cause (or be caused by) a non-negligible change of, e.g., wind stress or atmospheric heat fluxes. In this case, it is a safer choice to include more variables in computing a convergence criterion.

As soon as the ocean model sends more coupling fields than just sea surface temperature (e.g., surface currents or sea-ice-related quantities), it is not sufficient to restrict convergence analysis to an SST-based criterion. In our implementation, we use all variables which are part of the OASIS *namcouple*. This covers cases with more sophisticated coupling setups as well and is, in principle, model-agnostic.

Valcke (2021) used other output variables besides the coupling fields to estimate convergence, e.g., temperature or velocity further away from the air-sea interface. This requires looking at model output data (instead of OASIS output data) or reading the values directly from inside the model. The latter choice would mean code changes inside the AOSCM components, which we want to avoid. The former has another restriction in the case of the EC-Earth AOSCM: While both NEMO and OpenIFS use double precision internally and during coupling, the model output is restricted to single precision. This can lead to an error estimate which is different than one based on the internal model state. Finally, a computation based only on coupling fields can be carried over more easily to a 3D model. There, reading in full model output files is significantly more costly than the comparably small 2D coupling fields.

### 4.3.2. Amplitude-Based Convergence Estimate

The first convergence criterion we implement is similar to the one discussed by Valcke (2021): For each coupling variable  $c$ , we compute the maximum absolute difference between two subsequent iterations, i.e., the discrete maximum norm. We consider  $c^k \in \mathbb{R}^N$  to be the vector of coupling values of iteration  $k$ . It has dimension  $N$ , equal to the number of coupling windows in the simulation. The maximum absolute difference is thus given by:

$$\|c^{k+1} - c^k\|_\infty = \max_{n \in \{0, \dots, N-1\}} (|c_n^{k+1} - c_n^k|) \quad (4.1)$$

We compare this difference to the spread of the coupling variable over the Schwarz window. Since our numerical experiments are on the order of days and the Schwarz window is equal to the simulation length, this is an estimate of the diurnal cycle of the coupling variable. We compute the amplitude  $\mathcal{A}$  as

$$\mathcal{A}(c^{k+1}) = \max_{n \in \{0, \dots, N-1\}} c_n^{k+1} - \min_{n \in \{0, \dots, N-1\}} c_n^{k+1}. \quad (4.2)$$

Note that we estimate the spread using iteration  $k + 1$  instead of  $k$ . This is a somewhat arbitrary choice; we assume that iteration  $k + 1$  is closer to the converged Schwarz solution and thus gives a more realistic value for the amplitude of  $c$ .

We consider an iteration  $k$  to have converged if

$$\|c^{k+1} - c^k\|_\infty \leq \text{TOL} \cdot \mathcal{A}(c^{k+1}), \quad (4.3)$$

is fulfilled for all coupling variables  $c$  (cf. Figure 3.1). Here,  $\text{TOL} > 0$  is a fixed tolerance which we choose as  $\text{TOL} = 10^{-3}$  to signify a physically negligible Schwarz correction.

We can use previously published AOSCM results (Hartung et al., 2018, Figure 7) to give expected value ranges for the convergence threshold: For sea surface temperature, we see diurnal cycles on the order of 0.1 K to 1 K, giving a convergence threshold of  $10^{-4}$  K to  $10^{-3}$  K. For surface shortwave radiation, we see daily amplitudes on the order of  $10^2 \text{ Wm}^{-2}$  to  $10^3 \text{ Wm}^{-2}$ . Thus, the convergence tolerance for shortwave radiation will be between 0.1  $\text{Wm}^{-2}$  and 1  $\text{Wm}^{-2}$ .

### 4.3.3. Local Convergence Estimate

As stated by Valcke (2021, p. 15), “variables having a marked amplitude of the diurnal cycle will appear as converging more easily” with the above method than those which change very little over the course of a simulation. In addition to this, we have assumed in Equation (4.3) that the full simulation data (or at least one full day of simulation data) is available when computing the convergence threshold. This is fine in our case because the Schwarz window is equal to the length of the full simulation, and because the numerical experiments we conduct span multiple days.

However, to achieve faster convergence of the iterations, shorter Schwarz windows are preferable. In other coupling libraries such as preCICE, the Schwarz window is typically equal to the coupling window (Rüth et al., 2021). In this case,  $\mathcal{A}(c^k)$  is much smaller than the amplitude of the diurnal cycle (typical sizes for coupling windows are between 1h and 3h). Alternatively, one could prescribe typical amplitudes for each variable, which comes with its own set of difficulties.

For these reasons, we propose a second criterion which is not based on the amplitude of the coupling variable but its magnitude in each coupling window. For each coupling variable  $c$  and every coupling time step  $t_n$  with  $n = 0, \dots, N - 1$  it must hold that

$$\left| c_n^{k+1} - c_n^k \right| \leq \text{TOL} \cdot \left| c_n^{k+1} \right|. \quad (4.4)$$

Again,  $\text{TOL} > 0$  is a tolerance which we choose equal to  $10^{-3}$  in the following numerical experiments.

Returning to numerical examples from Hartung et al. (2018, Figure 7), a typical SST value is around 280 K, thus the difference between iterations would have to be below roughly 0.3 K.<sup>3</sup> For shortwave solar radiation, the criterion is dynamically adjusting: Near the peak of solar radiation, the upper bound is similar as in the amplitude-based criterion (between  $0.1 \text{ Wm}^{-2}$  and  $1 \text{ Wm}^{-2}$ ). At night, solar surface radiation is close to 0, thus the convergence bound will be significantly stricter.

Equation (4.4) is simple to use in a setup where the Schwarz window size is equal to the coupling window size: Instead of checking Equation (4.4) for the full vector  $c^k$ , one only estimates the error at the current time step  $t_n$  before moving on to the next coupling window. In that sense, it only depends on information which is available locally in time, whereas Equation (4.3) is based on global-in-time information.

In total, we would expect the two convergence criteria to give comparable results, while small differences in the resulting amount of iterations might stem from different magnitudes and diurnal cycles of a coupling variable.

#### 4.3.4. Comparison with the Final Iterate

Both Marti et al. (2021) and Lemarié et al. (2014) do not compare subsequent iterations. Instead, they compute the error of an iterate  $c^k$  with respect to a final iteration  $c^M$ , with  $M > k$  and the assumption that  $M$  is large enough to be the converged Schwarz solution. We can adapt Equation (4.3) and Equation (4.4) to follow a similar approach, resulting in

$$\|c^{k+1} - c^M\|_\infty \leq \text{TOL} \cdot \mathcal{A}(c^M), \quad (4.5)$$

$$\left| c_n^{k+1} - c_n^M \right| \leq \text{TOL} \cdot \left| c_n^M \right|. \quad (4.6)$$

<sup>3</sup>The sea surface temperature is converted to Kelvin before sending it to OpenIFS although NEMO uses °C otherwise.

#### 4. Implementation

---

As before, we check these criteria for each coupling variable exchanged between OpenIFS and NEMO, taking into account every coupling time step  $t_0, \dots, t_{N-1}$ . We use the comparison with the final iterate to test whether Equation (4.3) and Equation (4.4) are good runtime estimators for SWR convergence in Section 5.1.3.

## 5. Numerical Experiments

In this chapter, we investigate the impact of different coupling algorithms on the numerical solution produced by the EC-Earth AOSCM. As a test case, we place the single column model at the PAPA station in the Northeastern Pacific Ocean (50.033°N, 145.205°W). This location was also studied in Hartung et al. (2018) and is a standard example in the context of climate and weather model development.

We use reanalysis data for the initial conditions and atmospheric forcing. For the ocean, we obtained initial data from the GLORYS12V1 data set by the Copernicus Marine and Environment Monitoring Service (CMEMS), which contains daily averages in NetCDF format.<sup>1</sup> We extracted data at the grid cell closest to the PAPA station (50°N, 145°W) and linearly interpolated it from 50 to 75 vertical levels. Using this data set kept the amount of preprocessing minimal while ensuring realistic initial conditions. For OpenIFS, we use an input file based on the ERA-Interim data set (Dee et al., 2011), with six-hourly data from 00:00 on July 1 to 18:00 on July 30, 2014 (UTC). At the time of writing, no other input data for the atmosphere was obtainable, something we will get back to in Chapter 6. Because of the available initial data, we restrict ourselves to multi-day simulations in July 2014 at the PAPA station. The local time zone at the station during this time of the year is Pacific Daylight Time (PDT), seven hours behind UTC. In all plots of this chapter, we show output data after applying a time shift of -7h.

We study two numerical experiments in this chapter: The control experiment in Section 5.1 is a single four-day simulation for which we examine the mechanisms at play when switching the coupling algorithm. In Section 5.1.3, we compare the performance of different SWR convergence criteria in the control experiment. As a second experiment, we compare 1,248 two-day simulations with the AOSCM to get statistical insight into the impact of coupling schemes on a typical weather forecast. The setup and results of this case study are detailed in Section 5.2.

### 5.1. Control Experiment: Qualitative Comparison of Coupling Schemes

For the control experiment, we use the same time step sizes as in two previous studies with the AOSCM, Deppenmeier et al. (2020) and Hartung et al. (2018):  $\Delta t_{\text{atm}} = \Delta t_{\text{oce}} = 15$  min for both OpenIFS and NEMO. Due to the considerations in Lemarié et al. (2014), we choose

---

<sup>1</sup><https://doi.org/10.48670/moi-00021>

a coupling period larger than the model time step size:  $\Delta t_{\text{cpl}} = 1$  h. We use the forcing and initial data described above to do a four-day simulation at the PAPA station, from July 1, 2014 at 00:00 until July 5, 2014 at 00:00 (UTC), resulting in a local simulation time from June 30 at 17:00 until July 4 at 17:00.

With this setup, we compare the parallel algorithm and the two sequential algorithms with the converged SWR solution. In this first experiment, we do twenty Schwarz iterations to test the convergence criteria discussed in Section 4.3. The output data of iterations 13 until 20 is exactly equal up to the floating-point precision of the output files.<sup>2</sup> Thus, we can safely consider iteration 20 to be the converged SWR result. As described in Chapter 2, the non-iterative coupling algorithms are the first iteration of an additive/multiplicative SWR method. We therefore consider them approximations to the converged SWR result, which takes on the role of a reference simulation.

We spend the first parts of this section to present an in-depth comparison of model results for this experiment. In Section 5.1.3, we use the Schwarz iterations produced during the experiment to compare the convergence criteria presented in Section 4.3.

### 5.1.1. Prognostic Variables

For both models, we show the prognostic variables over time at the model level closest to the air-sea interface. Figure 5.1a contains the prognostic variables of NEMO: the zonal current  $u$ , meridional current  $v$ , the temperature  $T$ , and the salinity  $S$ . The level closest to the surface is at a height of  $z_{-1} \approx -0.5$  m. In NEMO it is assumed that these values are constant up to the surface. In particular, sea surface temperature (SST) is defined as

$$\text{SST} = T^o(z_{-1}).$$

For the ocean variables, all four coupling schemes result in a similar overall behavior. This is most striking for the salinity  $S$ , where differences become visible only near the end of the simulation. Here, the parallel and ocean-first coupling scheme result in a slight phase shift compared to the atmosphere-first and converged SWR solution. Such a phase shift also appears for the other ocean variables.


The ocean surface currents  $u$  and  $v$  are relatively weak throughout the simulation, on the order of 0.1 m/s. For these variables, differences between the coupling schemes are visible throughout the full simulation, although their absolute value is fairly small (on the order of 0.01 m/s). The error in the parallel and ocean-first case is more prominent in amplitude compared to the atmosphere-first coupling scheme. Combined with the phase shift, this makes them appear to be a worse candidate than the atmosphere-first algorithm.

We can draw similar conclusions from the sea surface temperature evolution: Here, the difference between coupling schemes seems to increase over time, reaching a magnitude of 0.1 °C around days two and three of the simulation.

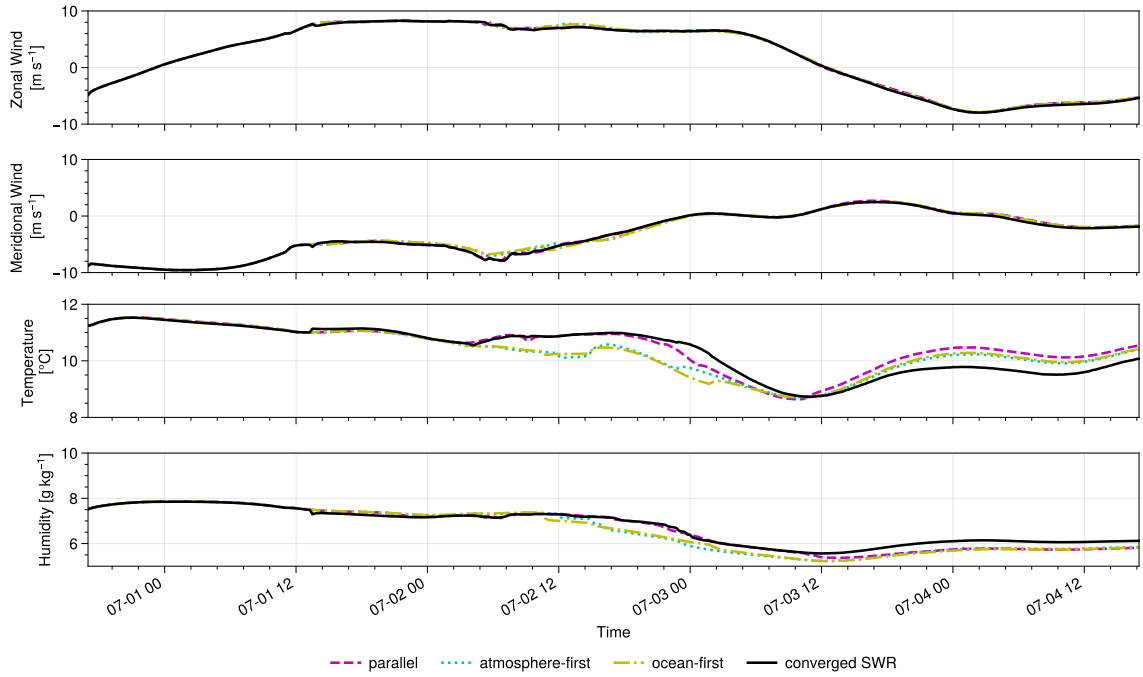
---

<sup>2</sup>Note that the model results are written out in single precision while the AOSCM internally uses double precision.

5.1. Control Experiment: Qualitative Comparison of Coupling Schemes



(a) Prognostic variables of NEMO at the highest model level ( $z_{-1} \approx -0.5$  m).



(b) Prognostic variables of OpenIFS at the lowest model level ( $z_1 \approx 11.9$  m).

Figure 5.1.: Prognostic variables of the AOSCM during the control experiment.

Judging only from the oceanic results, all coupling schemes give similar results while the atmosphere-first algorithm reduces the phase error and is thus the best of the non-iterative coupling algorithms. This plays an even bigger role when increasing the size of the coupling period, as was previously seen by Marti et al. (2021) and Voldoire et al. (2022). We have confirmed their results for a coupling period of  $\Delta t_{\text{cpl}} = 4$  h but do not include them here.

Figure 5.1b shows the prognostic variables of OpenIFS: the zonal and meridional wind speeds ( $u$  and  $v$  in the OpenIFS primitive equations), temperature  $T$ , and moisture  $q$ . In the OpenIFS-SCM, the lowest model level (level 60) does not correspond to a constant height but to a constant pressure level, equal to 1012.15 hPa in the control experiment. Over the course of the simulation, this corresponds to a height of  $z_1 \approx 11.9$  m (with a standard deviation below 3 cm).

The horizontal wind speeds are significantly larger in magnitude than the surface currents, reaching up to about 10 m/s. Slight differences are visible between the coupling algorithms, on the order of 0.1 to 1 m/s. Note that this is a similar relative difference as for the ocean surface currents (about 10%).

The output for  $T$  and  $q$  shows significant differences between the coupling schemes, starting on July 2. For temperature, the non-iterative coupling schemes produce solutions which are different from the converged SWR solution by a full degree Celsius in magnitude. In both variables, the two sequential algorithms perform worse than the parallel algorithm at daytime during July 2. Around 00:00 on July 3 for temperature, the parallel algorithm "detaches" from the SWR trajectory and behaves more like the sequential algorithms. At the end, the error of the parallel algorithm is even larger than that of the sequential atmosphere-first and ocean-first coupling schemes. A similar result can be seen for moisture, although the parallel algorithm stays close to the SWR result for longer, until July 3 around noon.

As seen in Figure 5.1a, the differences in SST are relatively small, while the resulting spread of atmospheric temperature is a lot more significant. The origins for this phenomenon lie in the parameterizations for atmospheric thermodynamics, which we proceed to investigate using additional output data of the OpenIFS single column model.

### 5.1.2. Behavior of Atmospheric Thermodynamics

We start off with the surface heat fluxes computed by OpenIFS. These play an important role in the coupling conditions, particularly for temperature and moisture, cf. Section 3.3. The surface sensible and latent heat fluxes, shown in panels a) and b) of Figure 5.2 show a similar pattern as temperature and moisture itself: Around July 2, the two sequential coupling schemes behave very differently than the converged SWR solution. While the parallel scheme performs better initially, it joins the other non-iterative solutions towards the end of the simulation. This happens first for the sensible heat flux (which is related to temperature), and then on July 4 for the latent heat flux (which is related to moisture). Notably, all four coupling schemes produce similar results around noon of July 3.



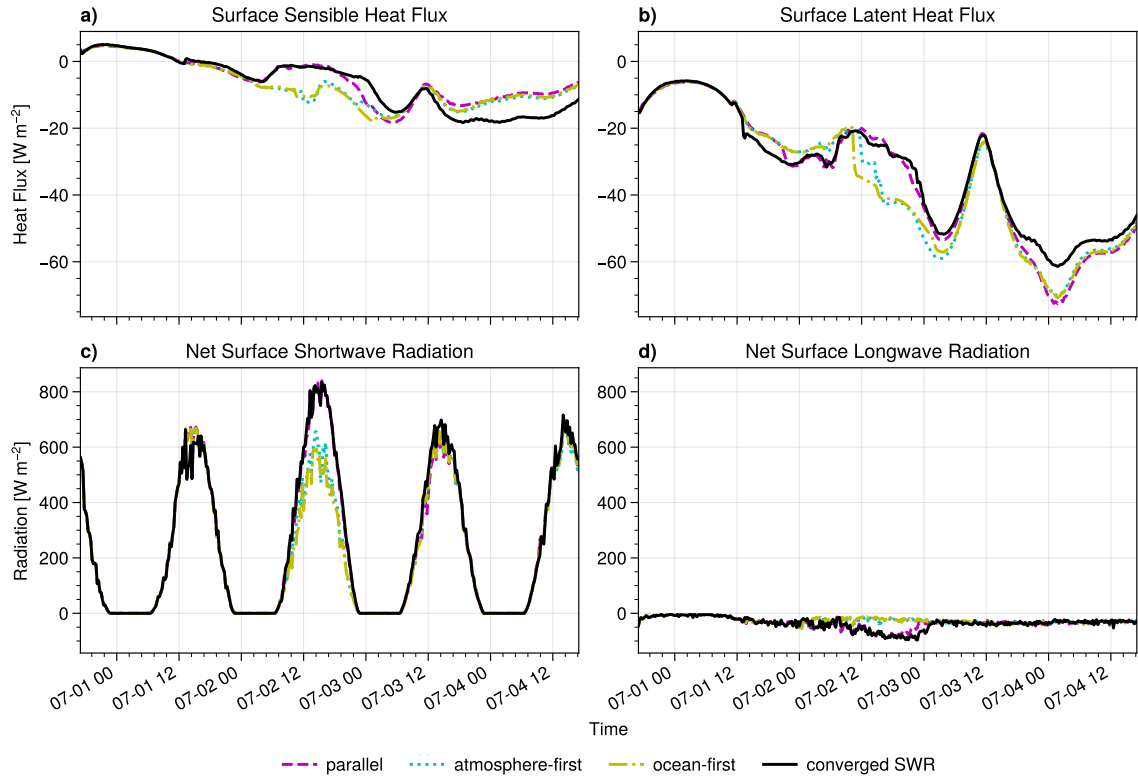


Figure 5.2.: Atmospheric surface heat fluxes in the control experiment for all four coupling schemes over time. The top row shows the sensible and latent heat fluxes,  $\mathcal{J}_T$  and  $\mathcal{J}_q$ , the bottom row shows net shortwave and longwave radiation, as defined in Section 3.3.

The surface radiation, cf. panels c) and d) of Figure 5.2, reveals more information on the differences on July 2: While all coupling schemes behave very similarly for most of the simulation, the sequential algorithms produce visibly different radiative fluxes around this time. For shortwave radiation, they give fluxes which are approximately  $200 \text{ W/m}^2$  lower than the parallel and converged SWR solution. The net longwave radiation, on the other hand, is up to  $100 \text{ W/m}^2$  higher for the sequential algorithms. The sea surface temperature does not show strong differences at this point of the simulation, see Figure 5.1a, thus this change is caused by the radiation scheme itself, not by ocean longwave radiation.

These results re-emphasize that the atmospheric thermodynamics react significantly to the comparatively small change in sea surface temperature, depending on the coupling algorithm. However, the question remains what triggers such a strong response. We found an answer to this by looking at the boundary layer type used inside OpenIFS to switch between different formulations for vertical turbulence.

As described by Fitch (2022), OpenIFS distinguishes four types of atmospheric boundary layers which determine the type of turbulence closure used by the model (i.e., how to approximate  $\overline{w'\phi'}$  for the prognostic variables  $\phi$ ). Possible states are dry stable (DS), dry convective (DC), a boundary layer topped by stratocumulus clouds (Sc), or one topped by shallow cumulus clouds (Cu). While a first-order purely diffusive turbulence closure is used for the dry stable boundary layer, the other three cases yield different variants of an EDMF scheme. In panel a) of Figure 5.3, we show the boundary layer type over time for all four coupling schemes in the control experiment.

While the simulation starts off with a dry stable atmosphere, all four coupling schemes switch to an unstable boundary layer around 12:00 on July 1. The boundary layer type remains unstable until the end of the simulation and all four coupling schemes end up with a stratocumulus-topped boundary layer around July 3. However, in the time between these two regimes, very different patterns emerge depending on the coupling algorithm. We first note that even in the converged SWR solution, the boundary layer type can change very often in short time intervals. For instance, the model switches between all three unstable boundary layer types around 12:30 on July 1. Since the boundary layer type is determined in each model time step using only instantaneous information, rapid changes of the chosen type can occur. There is no incentive for this parameterization to converge to a more “continuous” behavior, i.e., one which avoids individual time steps which use a different turbulence scheme than the model time steps before or after.

In general, we see a similar pattern emerge as in the previous plots: In the control experiment, the parallel scheme seems to be the most successful at obtaining a similar boundary layer type as Schwarz waveform relaxation. The atmosphere-first and ocean-first algorithms result in the same boundary layer type for most of the simulation. Particularly, they spend a longer time in the cumulus-topped regime and switch to the stratocumulus-topped boundary layer earlier than the other two algorithms. The parallel and SWR solution, on the other hand, result in a dry convective boundary layer for more time steps.

We can also study how often the boundary layer type of a given model time step changes between two Schwarz iterations, see panel b) of Figure 5.3. This is another indicator for the sensitivity of this parameterization to slight changes of interface conditions due to Schwarz corrections. Recall that we do 20 Schwarz iterations in total, while the numerical solution stops changing after iteration 13. In the same regime as seen above, between July 1 at 12:00 and July 3 at 00:00, the boundary layer type switches the most often, with a maximum of five changes. Comparing the two plots, we can conclude that the boundary layer type changes across Schwarz iterations are particularly low whenever all three non-iterative methods yield the same type. Changes of boundary layer type occur during the first six Schwarz iterations in this experiment, where the magnitude of corrections is also expected to be larger than later on.

Around July 2, the model atmosphere appears to be in a state where slight changes in interface conditions yield a different boundary layer type in the decision tree described by Fitch (2022). The discrete change of turbulence parameterization can then result in the significant differences in atmospheric heat fluxes, temperature, and moisture seen in the

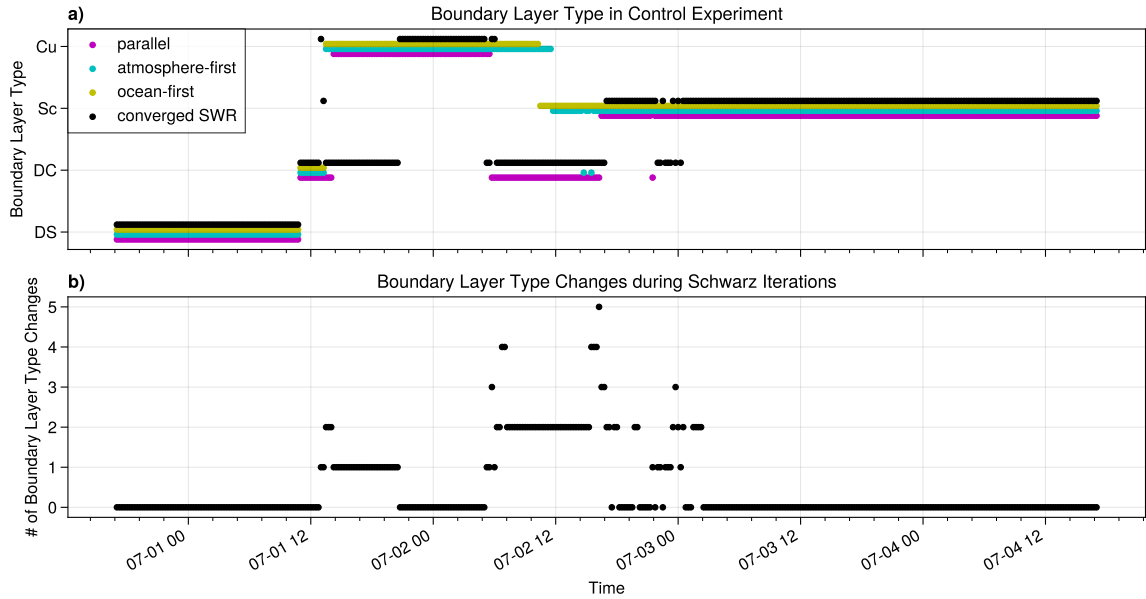


Figure 5.3.: Boundary layer types as determined by OpenIFS over time. Panel (a) shows the boundary layer type at each time step depending on the coupling scheme. Panel (b) shows the number of total boundary layer type changes when using SWR. Whenever the type changes from one iteration to the next, we increment the counter.

same time span.

However, this does not explain why all three non-iterative coupling schemes deviate from the converged SWR solution at the end of the simulation: Here, the radiation scheme, as well as the boundary layer type are insensitive to the coupling algorithm (Figure 5.1b and Figure 5.3). Nevertheless,  $T$ ,  $q$ ,  $\mathcal{J}_s$ , and  $\mathcal{J}_q$  give different results.

We get additional insight by looking at the vertical profile of atmospheric temperature in the boundary layer, given in Figure 5.4. Therein, we show the temperature on the last 15 model levels (46-60): This corresponds to a height of around 2.3 km in the control experiment at a constant pressure level of 763 hPa. The interval fully contains the boundary layer, which does not exceed a height of 850 m.

We show the temperature profile only at specific points in time during the simulation. At 12:00 on July 2, panel a), the atmosphere-first and ocean-first profile are already visibly different from the SWR result, whereas the parallel algorithm's solution closely follows the converged profile. At this time, the former two have the cumulus-topped boundary layer type whereas the latter are classified as dry convective, cf. Figure 5.3.

Panel b) shows the profile 12 hours later. The two sequential methods have developed a vertical stratification which does not change significantly until the end of the simulation,

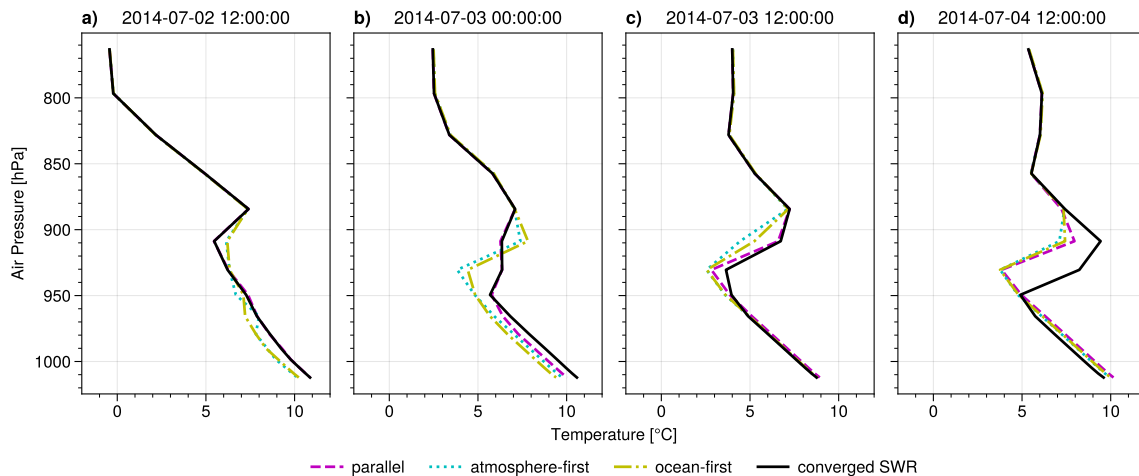


Figure 5.4.: Snapshot of the vertical air temperature profile at four different times during the control experiment. We show the last ten model levels which contain the boundary layer.

see panels c) and d). The solution with the parallel coupling scheme seems to depart from the SWR solution near the surface, while still following the shape of the curve further upwards. Another 12 hours later, at 12:00 on July 3, the surface temperatures are very close to each other, corresponding to the lower difference we also see in the other plots in this chapter. However, we see that the three non-iterative solutions start to group around 930 hPa, a process that is complete by the end of the simulation. In panel d), five hours before the final model time step, the new profiles are fully developed. We see here that the already significant difference in surface temperatures is linked to a large temperature difference of 4.5 °C at 930 hPa.

Overall, Figure 5.4 shows that the small changes in interface conditions due to the coupling method lead to a significant change in vertical stratification of the atmospheric boundary layer in the control experiment. One reason for this appears to be that OpenIFS distinguishes boundary layer types in each time step, which are sensitive to the surface temperature and can yield strong deviations in the structure of the planetary boundary layer. The surface variables capture a part of this signal. Nevertheless, it is necessary to study the whole boundary layer in order to discuss the potential physical implications of changing the coupling scheme and, therefore, the interface conditions seen by the model.

The OpenIFS configuration settings contain a toggle for the convective mass flux scheme<sup>3</sup> which is part of the EDMF framework, cf. Equation (3.9). In Figure 5.5, we show the air temperature profile for the control experiment if the mass flux term is turned off. Here, only the ocean-first coupling scheme visibly deviates from the converged SWR solution.

<sup>3</sup>LECUMF in the NAEPHY namelist

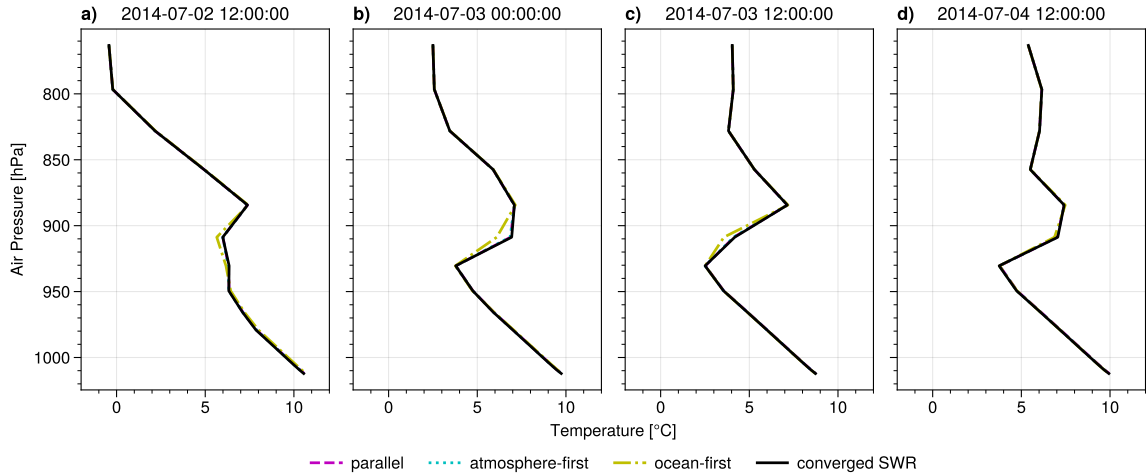


Figure 5.5.: Same as Figure 5.4 but we turn off the convective mass flux scheme during the experiment.

However, the differences are never large enough to trigger a different vertical structure of the boundary layer. This further indicates that the parameterization for vertical turbulence, specifically the sensitivity of the mass flux scheme to small SST changes can cause large differences between the non-iterative and SWR algorithms, even on time scales of a few days.

### 5.1.3. Comparison of Convergence Criteria

For the previous visualizations, we used the 20th Schwarz iteration as the reference solution. We have verified that at this iteration, the AOSCM output has converged up to floating-point precision. We use the control experiment to study the convergence criteria we introduced in Section 4.3. The aim here is to find out whether runtime estimates of the convergence error are close to those taking into account the reference solution. We also compare the two different convergence thresholds, one depending on the amplitude of a coupling variable over a full simulation, the other depending on its value at the current time step  $t_n$ . Note that for this assessment, we do not use the model output but instead the debug files written out by OASIS (which use double instead of single precision).

In the example simulation, all four convergence criteria—Equations (4.3) to (4.6)—give the same result: They are fulfilled for all variables starting from the eleventh iteration. We have selected two coupling variables to visualize the convergence behavior in Figure 5.6: the SST, which is sent from NEMO to OpenIFS, and the non-solar heat flux  $Q_{ns}$ , which is sent from OpenIFS to NEMO. Dots denote the two different ways of estimating the error in iteration  $k$  while lines represent the convergence thresholds.

How the error is estimated corresponds to the left hand side of the convergence criteria.

## 5. Numerical Experiments

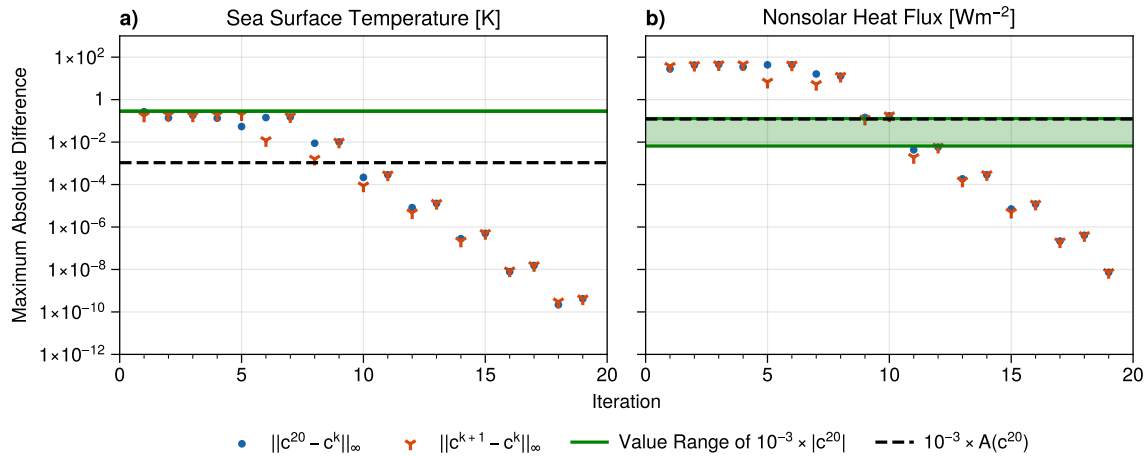


Figure 5.6.: Convergence behavior of Schwarz iterations for two coupling variables in the control experiment.

The circular dots show the maximum absolute difference of each Schwarz iteration  $k$  with respect to the reference solution of iteration 20. This corresponds to the computation used in Equation (4.5) and Equation (4.6). It assumes that the reference value  $c^{20}$  is already available and known to be converged, which is why these two criteria can not be evaluated at runtime. The Y-shaped data points correspond to the maximum value (across all time steps  $t_n$ ) of the difference between two subsequent iterations  $k$  and  $k+1$ . This is the largest value of the left hand side in Equation (4.3) and Equation (4.4). Since no reference solution is required, this error estimate can be computed at runtime.

Both estimates show results of similar magnitude and behavior for each coupling variable. From the data in Figure 5.6, we see that the convergence error declines stepwise, staying at the same magnitude for two consecutive iterations before dropping noticeably. Note also that the location of the downward step is shifted by one in the SST results compared to  $Q_{\text{ns}}$ . The strongest difference between the two error estimates arises when  $\|c^{k+1} - c^{20}\|_{\infty}$  and  $\|c^k - c^{20}\|_{\infty}$  have a similar magnitude. This is the case in, e.g., iterations  $k = 6, 8$  for SST or  $k = 5, 7, 9$  for  $Q_{\text{ns}}$ . At such a point, the runtime criterion could be met for a coupling variable while the comparison with the reference solution would yield a significantly higher error. However, since there is a phase shift between the error estimates depending on the coupling variable, this will not be a problem if all convergence thresholds are of suitable magnitude. After all, an iteration is only marked as converged if Equation (4.3) and Equation (4.4) are fulfilled for all coupling variables.

Figure 5.6 also shows the convergence thresholds for the different criteria. The black dashed line represents  $\text{TOL} \cdot A(c)$ , i.e., the amplitude convergence threshold. This threshold is used in Equation (4.5) and Equation (4.3). In the plots we have only included the amplitude of iteration 20 for legibility and since the amplitude over the full 4-day simula-

tion is not significantly affected by the coupling scheme (this is also indicated for the SST and solar heat flux in Figure 5.1b). The green shaded area represents the range of  $\text{TOL} \cdot |c|$ , which we again only display for  $c^{20}$  in all four cases. As can be seen in Figure 5.6, and as was already mentioned in Section 4.3, the two thresholds can give quite different results depending on the variable.

It is not clear a priori that small absolute values of a coupling variable coincide with a small difference between iterations. This could lead to very rapid or very slow convergence of the “local” convergence criteria (Equations 4.4 and 4.6). In our experiments, all convergence criteria behaved similarly and we did not see any indication for such an issue.

We can conclude that the two runtime criteria (Equations 4.3 and 4.4) give similar results as the ones using a reference solution for comparison. The weaknesses of both criteria are counteracted by the fact that we take all coupling variables into account. In the forthcoming analysis, we only accept a solution as converged if both Equation (4.3) and Equation (4.4) are met. Especially for implementations where a global amplitude of  $c$  is not available, paying more attention to the individual thresholds will be necessary. For future experiments, using a lower value of TOL for SST (or °C instead of Kelvin) would make sense. Since we rely on two criteria and check them for all coupling variables, having one relatively loose threshold is not a big problem as long as the threshold for other variables is sufficiently strict.

## 5.2. Case Study: Numerical Weather Prediction Setup

Up to now, we have only looked at a single, four-day simulation with the EC-Earth AOSCM. While the coupling algorithms did have a visible impact on the model results for the control experiment, it is not clear whether this is representative in general. To this end, we want to study the effect of Schwarz waveform relaxation on a typical application of coupled atmosphere-ocean models more robustly. These are, in general, numerical weather prediction (NWP) and climate simulations. In this thesis, we restrict ourselves to numerical weather prediction, for a variety of reasons.

First of all, our implementation of SWR uses a Schwarz window equal to the total simulation time (see Chapter 4). Thus, the amount of iterations necessary for convergence scales with the length of the simulation. This makes Schwarz iterations on longer time scales infeasible for now. Potential solution strategies are discussed in Chapter 6.

Secondly, model drift plays a bigger role for time scales longer than a couple of days. We would have to include nudging above the boundary layer, as previously done by Deppenmeier et al. (2020) and Hartung et al. (2018). However, atmospheric nudging “interferes with the performance” of the boundary layer parameterization (Hartung et al., 2018, p. 4130). Since we are precisely interested in the behavior of the boundary layer, we want to avoid using the relaxation terms in the single column model.

Finally, at the stage of writing this thesis, we only have atmospheric initial conditions and forcing data available for the month of July, 2014. For climate applications, we would

Variable name	Symbol	Value
Simulation time	$T$	48h
Coupling time step	$\Delta t_{\text{cpl}}$	1h
Ocean time step	$\Delta t_{\text{oce}}$	30min
Atmosphere time step (dynamics, physics without radiation)	$\Delta t_{\text{atm}}$	12min
Radiation time step	$\Delta t_{\text{rad}}$	1h

Table 5.1.: Time stepping setup for the NWP case study.

need forcing data for longer time spans to compare at least monthly or seasonal means produced by the model.

### 5.2.1. Model Setup and Initial Conditions

Representative for a typical weather forecast, we decide to do two-day simulations. Two days are still well within the atmospheric predictability limit but we expect differences due to the diurnal cycle and interplay of different parameterizations. Since NEMO is mainly forced by the atmosphere, two days of simulation also give the ocean some spin-up time to develop a diurnal SST cycle, cf. Figure 5.1a.

For the experiments in this section, we use the same time step sizes as used for the operational ensemble forecasts by the ECMWF. Of course, the ECMWF’s NWP model differs from the AOSCM significantly: for instance, it is a global, three-dimensional model with a higher vertical resolution which includes an additional wave model. However, the OpenIFS-SCM is based on their atmospheric model, while the ocean model used in ECMWF predictions is a version of NEMO. Suitable time step sizes vary for each atmosphere and ocean model due to the implemented numerical methods; (Open)IFS time step sizes are comparatively large, which is why it makes sense to take the host model as a reference instead of, e.g., the significantly smaller time step sizes chosen by Marti et al. (2021). An overview of the time step sizes for this case study is given in Table 5.1.<sup>4</sup>

The limiting factor for the amount of forecasts is the available input data for the atmospheric component, which ranges from July 1, 00:00 until including July 30, 18:00, in six-hourly intervals. To do two-day forecasts, the last possible initial condition is on July 28, 18:00. This leaves 112 available initial conditions to compare the four coupling schemes.

However, the atmosphere is a chaotic system, which is why NWP is commonly done with ensemble forecasts: The same simulation is repeated with small perturbations in the initial conditions, to allow for a statistical analysis of the resulting state of the atmosphere. How to perturb initial conditions in a physically and statistically motivated sense is not a trivial task. For our case study, we have decided on the following method: We use the initial data from the ERA input to run a first set of two-day simulations. There, we

<sup>4</sup>In operational forecasts, the time step of ice model is  $\Delta t_{\text{ice}} = 1\text{h}$ , but we consider an ice-free case here.






Figure 5.7.: Schema of how we create initial data for ensemble forecasts.

use the three non-iterative coupling schemes to produce three different final states of the atmosphere. These are the new initial conditions we use to initiate the ensemble forecast.

The first set of initial conditions can be generated using a simulation from July 1 to July 3, 00:00 (UTC). Thus, the final set of initial conditions consists of six-hourly initial conditions from July 3, 00:00 to July 28, 18:00. We therefore have  $4 \cdot 26 = 104$  initial dates with three initial conditions per date. We do not use the original input data as initial conditions: The two day simulations from the AOSCM give very different results from the original ERA data, which would prevent any meaningful statistical analysis. This is unsurprising as the ERA data was created with a different atmospheric model and because we force the AOSCM only every 6h, using linear interpolation between the forcing time steps.

A schematic of this approach is given in Figure 5.7. Black dots represent an initial condition from reanalysis data, available at every forcing time step (gray vertical lines). Using the parallel, atmosphere-first, and ocean-first algorithms, we produce three new initial conditions 48h later (colorful lines and dots). In general, these result in a very different model state than the reanalysis data.

For each of the  $3 \cdot 104 = 312$  initial conditions, we run two day-forecasts all four implemented coupling schemes: the three non-iterative coupling algorithms and Schwarz waveform relaxation. This gives  $4 \cdot 312 = 1,248$  two-day simulations.

To determine whether the Schwarz iterations have converged, we use both runtime criteria, Equation (4.3) and Equation (4.4). An iteration is defined to have converged only if both criteria are met for all coupling fields. For each initial condition, we consider the converged SWR solution to be the "most correct" one. This does not mean that it is the closest to observations. But it is the one closest to the solution of the discretized coupling problem (which would be attained by exchanging a higher order interpolant and not just an average of the coupling fields).

## 5.2.2. Results with the NWP Setup

### Convergence of Schwarz Iterations

Out of 312 two-day forecasts with SWR as a coupling scheme, 295 converged according to both runtime convergence criteria introduced in Section 4.3 (94.6%). Figure 5.8 shows the

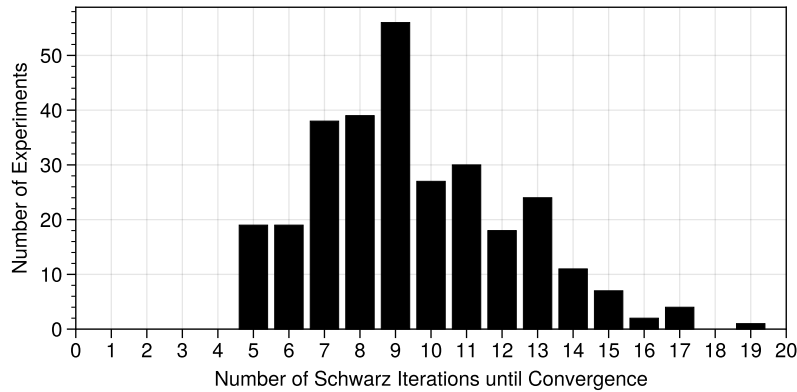


Figure 5.8.: Number of Schwarz iterations necessary to reach convergence in the NWP case study. The 17 non-converged simulations are not included.

distribution of necessary iterations to convergence for these 295 simulations. In all cases, less than 20 iterations were necessary for convergence despite the large Schwarz window size of 48h.<sup>5</sup> In 171 forecasts, approximately 55% of cases, less than ten iterations were sufficient. This confirms the benefits of using a runtime criterion to determine convergence, instead of enforcing a constant number of Schwarz iterations.

In the 17 cases where the SCM did not converge, small oscillations formed after at most 20 Schwarz iterations. OpenIFS and NEMO alternated between two model states which are physically close together (e.g., with less than 0.01°C difference in sea surface temperature). We observed oscillations with a period of four iterations and, in one case, oscillations with a period of two iterations. We return to this topic in Section 6.1 but do not investigate it further here. In the following plots, we have not excluded the cases with oscillations, since the differences between the two model states are always relatively small from a physical point of view. We have always used iteration 39 in these cases.

### Ensemble Spread

Past studies have implied that Schwarz waveform relaxation reduces the ensemble spread in atmosphere-ocean coupling (Connors & Ganis, 2011; Lemarié et al., 2014). To find out if this is the same in the AOSCM, we compare the initial and final spread of the atmospheric state.<sup>6</sup> To quantify model spread, we use atmospheric temperature  $T$  as a prognostic variable representing the current model state.<sup>7</sup> We define the spread as the vertical sum of the standard deviation  $\sigma(T)$  across the three available initial conditions:

<sup>5</sup>By contrast, Marti et al. (2021) used 50 iterations in each coupling period, yielding significantly larger computational cost.

<sup>6</sup>Ocean quantities cannot be taken into account because we did not perturb the ocean initial condition.

<sup>7</sup>As  $T$  is given in Kelvin, this means we will not run into sign changes or values close to 0.

Coupling Scheme	Smallest Spread Ratio Count
Atmosphere-first	23/104 (22%)
Ocean-first	29/104 (28%)
Parallel	26/104 (25%)
Schwarz Waveform Relaxation	26/104 (25%)

Table 5.2.: Fraction of forecasts where a given coupling scheme has a lower  $r_{\text{spread}}$  than the other three coupling schemes.

$$\text{spread}(t) = \sum_{i=1}^{60} \sigma(T_i(t)) \quad (5.1)$$

The initial spread ranges from  $4 \cdot 10^{-2} \text{°C}$  to  $10 \text{°C}$ , which might indicate that our method of creating perturbed initial data produces rather inconsistent results. For each start date and coupling scheme, we can also compute the final spread,  $\text{spread}(t_{\text{end}})$ . Combining the two, we define the spread ratio

$$r_{\text{spread}} = \frac{\text{spread}(t_{\text{end}})}{\text{spread}(t_{\text{start}})}. \quad (5.2)$$

Overall, the AOSCM model spread increases over the course of the two day forecast, with a median spread ratio between 1.9 and 2.3 per coupling scheme. However, sometimes the model spread decreases over the course of the simulation, i.e.,  $r_{\text{spread}} < 1$ . Such a behavior would be a rare exception for a 3D general circulation model. In our case study, this happens in 57 out of 416 cases (13.7%). A potential reason for this is that the externally prescribed dynamics constrain the AOSCM and prevent free model drift. Other times,  $r_{\text{spread}} \gg 1$ , with maximum values between 29 and 43 attained on July 12.

We want to see whether Schwarz waveform relaxation leads to a significantly smaller model spread compared to the non-iterative coupling algorithms. To this end, we check for each forecast date which coupling algorithm leads to the smallest spread ratio. The results are given in Table 5.2. What stands out is that they are almost uniformly distributed, with every coupling scheme performing "best" in about 25% of cases. This suggests that the coupling scheme does not have a strong influence on model spread in the EC-Earth AOSCM, as opposed to previous studies done by Connors and Ganis (2011) and Lemarié et al. (2014). More analysis of the data strengthens this hypothesis: By all measures we considered, the model spread seems to be mostly influenced by the initial and forcing data, while the coupling algorithm plays a minor role. Schwarz waveform relaxation does not perform significantly better than the non-iterative coupling algorithms—but neither does any other algorithm.

**Which Non-Iterative Coupling Method is the Best?**

Finally, we want to use the experiment results to determine which non-iterative coupling algorithm performs the best for the 312 initial conditions we produced as described in Section 5.2.1. We consider a coupling algorithm to give “the best” result if its 48h forecast is closer to the converged SWR solution than the other two coupling methods. The previous study by Marti et al. (2021) claims that the sequential atmosphere-first algorithm will give significantly better results than the ocean-first algorithm, with the parallel scheme performing the worst. The results of the control experiment in Section 5.1 suggest that we will reach a similar conclusion when looking at sea surface temperature, while a more nuanced result could be obtained for the atmospheric quantities. Recall that in the control experiment, the parallel algorithm outperformed the two sequential coupling algorithms for the first three days of the simulation when looking at atmospheric temperature and moisture.

We select three variables to measure the difference between forecasts: sea surface temperature (SST) as a representative quantity for the ocean, and atmospheric temperature  $T$  and moisture  $q$  in the boundary layer. To simplify the analysis, we do not recompute the size of the boundary layer for each experiment but instead select the lowest ten model levels of OpenIFS, which span from surface pressure down to  $p = 913 \pm 5$  hPa ( $z \approx 900$ m). For all three variables, we compute the  $\|\cdot\|_2$ -norm of the difference between the result obtained with the non-iterative coupling algorithm and converged SWR:

$$e_i(SST) = \|SST_i - SST_{\text{SWR}}\|_2 = |SST_i - SST_{\text{SWR}}| \quad (5.3)$$

$$e_i(T) = \|T_i - T_{\text{SWR}}\|_2 \quad (5.4)$$

$$e_i(q) = \|q_i - q_{\text{SWR}}\|_2, \quad (5.5)$$

where  $i \in \{\text{atm, oce, par}\}$  denotes the coupling scheme.

As for the ensemble spread, we can count how often a coupling scheme has the lowest error for each variable, i.e.,  $i \equiv \arg \min_i e_i$ . This gives a ranking of non-iterative coupling schemes, which we summarize in Table 5.3. If the coupling algorithm does not have a strong influence on the performance, we would expect an even distribution of about 104/312 “wins” per coupling algorithm. The ranking in the table confirms the results of Marti et al. (2021), particularly when looking at SST: The atmosphere-first algorithm outperforms the other two coupling algorithms in more than 2/3 of the cases. The parallel coupling scheme has the worst score but is not far behind the ocean-first coupling algorithm. This falls in line with the hypothesis that the phase error present in these two methods is responsible for most of the error in the sea surface temperature.

For atmospheric  $T$  and  $q$  we see identical results. Here, we obtain the same ranking as for the ocean, but the distribution is less skewed: The sequential atmosphere-first method produces the closest results to Schwarz waveform relaxation in 42% of cases, while the parallel algorithm “wins” in 24% of cases (almost twice as often as for the sea surface

Sea Surface Temperature		Atmospheric $T, q$ in Boundary Layer	
1. Atmosphere-first	(212/312, 68%)	1. Atmosphere-first	(132/312, 42%)
2. Ocean-first	(57/312, 18%)	2. Ocean-first	(106/312, 34%)
3. Parallel	(43/312, 14%)	3. Parallel	(74/312, 24%)

Table 5.3.: Ranking of non-iterative coupling schemes in terms of how often they produce the two-day forecast closest to the converged SWR result.

temperature). This suggests that the control experiment was an outlier simulation in terms of coupling algorithm performance—but not as rare as one would expect from the results previously found by Marti et al. (2021).

The ranking in Table 5.3 does not give any information on the size of  $e_i$ . The maximum errors across all  $3 \cdot 312 = 936$  non-iterative experiments are significant considering the short duration of the simulation:

$$\begin{aligned}
 e_{\max}(SST) &= 0.27^\circ\text{C} \\
 e_{\max}(T) &= 5.19^\circ\text{C} \\
 e_{\max}(q) &= 3.25 \text{ g kg}^{-1}
 \end{aligned}$$

However, such large errors come up rarely while the vast majority of experiments is considerably closer to the converged SWR result. To illustrate this, we compute a weighted error  $e/e_{\max}$ , group the result into bins, and count how often a coupling scheme appears in each error range. The results are given in Figure 5.9, which is by design similar to Figure 5 of Marti et al. (2021). In that paper, they used the maximum difference in SST between two subsequent coupling windows, but this value is less easy to obtain in our experiment. Figure 5.9 confirms the ranking in Table 5.3. We see that the sequential atmosphere-first algorithm strongly outperforms the other two coupling schemes for sea surface temperature, whereas the difference is less clear for atmospheric quantities. What stands out in Figure 5.9 is that in most experiments, the error stays below 10% of the respective  $e_{\max}$ .

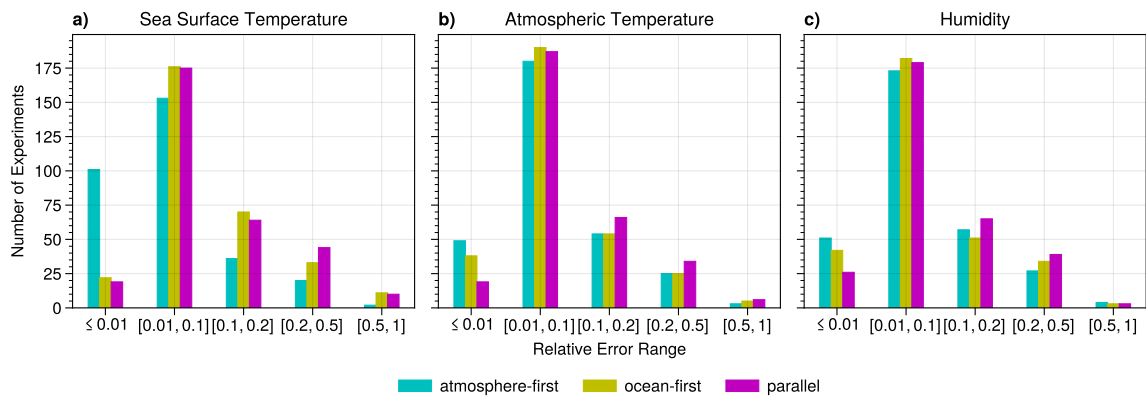


Figure 5.9.: Weighted error  $e/e_{\max}$  for the three observed quantities, grouped by coupling scheme.

## 6. Discussion

In the last chapter, we have seen various effects of switching the coupling scheme in the EC-Earth AOSCM. We will now go on to discuss the implications of the numerical results in Section 6.1. Section 6.2 reviews the benefits and limitations of the implementation choices we made in order to support SWR in the single column model.

### 6.1. Discussion of Numerical Results

#### 6.1.1. Impact of Coupling Schemes on Numerical Results

We carried out two separate numerical experiments with the EC-Earth AOSCM, a single four-day control experiment, as well as a collection of 1,248 two-day forecasts. We placed the AOSCM at the PAPA station, an ice-free location in the Pacific which has been thoroughly studied in the context of Earth system model development (Hartung et al., 2018). For both experiments, we used realistic time step sizes and restricted the simulation lengths to prevent model drift. We saw in the control experiment and the NWP case study that the coupling method can lead to significant differences in the numerical solution.

The ocean results are in line with previous research by Marti et al. (2021), Valcke (2021), and Voldoire et al. (2022). Differences in ocean output are most visible in sea surface temperature, where we see a phase shift depending on the coupling algorithm (Figure 5.1a). This phase shift can be explained from the difference in coupling lag, as described in Section 2.2: In the parallel and ocean-first algorithms, NEMO sees atmospheric fluxes from the coupling window  $[t_{n-1}, t_n]$  during the integration from  $t_n$  to  $t_{n+1}$ . This means that the ocean sees a shifted diurnal cycle depending on the coupling scheme, causing a numerical error. When using the sequential atmosphere-first algorithm or SWR, the lag is eliminated. This result carries over to the NWP case study, where we clearly see that the sequential atmosphere-first algorithm outperforms the other two non-iterative coupling schemes in Figure 5.8 and Table 5.3. The atmosphere-first algorithm gives the best result in more than two thirds of cases and the error magnitude with respect to SWR is significantly lower than what is observed with the parallel and ocean-first algorithms. The differences in ocean salinity and ocean currents are smaller but they show a similar pattern as sea surface temperature: the ocean-first and parallel algorithm deviate the most, indicating that the coupling lag is responsible for most of the behavior. Since the sea surface temperature shows the clearest signal and is the only ocean variable which directly affects the atmosphere, it appears to be a good measure for NEMO's response to changes in interface conditions.

In contrast to Marti et al. (2021), we studied atmospheric variables in more detail: In general, numerical differences are larger in the atmospheric output, which is unsurprising seeing as the ocean is a significantly slower physical system. We saw that wind speeds react very little relative to their strength and seem to be mostly driven by the externally provided atmospheric dynamics (horizontal advection, vertical advection, and geostrophic wind), cf. Figure 5.1b. In the control experiment, atmospheric temperature and moisture reacted very sensitively to changes in the coupling scheme. As opposed to Marti et al. (2021), we cannot conclude that atmosphere-first outperforms the other two non-iterative coupling algorithms in the control experiment. While the overall ranking of non-iterative coupling schemes in the NWP case study agrees with previous findings, the atmosphere-first algorithm does not outperform the parallel and ocean-first algorithms as clearly as for NEMO output. While sea surface temperature might be a sufficient measure to determine SWR convergence in models like the AOSCM and IPSL-CM6<sup>1</sup>, it is not sufficient to draw general conclusions about the performance of a coupling scheme for the atmosphere-ocean problem.

The large differences in atmospheric temperature and moisture in the control experiment are not directly explainable by the comparatively small changes in sea surface temperature. Instead, physical parameterizations need to be taken into account for studying the atmosphere-ocean coupling problem properly. In our case, we offer a mechanistic explanation for the feedback at play in the OpenIFS-SCM: Small SST changes, depending on the coupling algorithm, can trigger different boundary layer parameterizations in the atmosphere. These are directly linked to the convective mass flux term used by the EDMF scheme for vertical turbulence parameterization. The convective term influences the vertical stratification of the atmosphere (Figure 5.4 vs. Figure 5.5) and can affect the radiation scheme inside OpenIFS (Figure 5.2). The boundary layer types in OpenIFS are determined using sharp thresholds and purely based on the instantaneous values of atmospheric quantities. Figure 5.3 shows how sensitively the parameterization can react to changes in the coupling algorithm and across Schwarz iterations. This, and the large changes in temperature and moisture it leads to, indicates potential for future work to reduce unstable and perhaps unrealistic model behavior in certain situations. In previous analytical work studying the convergence of SWR in atmosphere-ocean coupling, vertical turbulence was only approximated with a diffusive term (e.g., Clement et al., 2020; Connors & Ganis, 2011; Lemarié et al., 2013a, 2013b). The experiments in this thesis indicate that the mass flux term can have significant impact on the behavior of the model.

### 6.1.2. Schwarz Waveform Relaxation Convergence

The NWP case study and the control experiment showed that Schwarz waveform relaxation can converge with a relatively low number of iterations for multi-day simulations in the AOSCM. In all simulations where SWR converged, less than 20 iterations were suf-

---

<sup>1</sup>see the discussion in Section 4.3



ficient, and a majority of the experiments converged in less than ten iterations. This is a notable result, considering the fact that we did not use any acceleration methods: Marti et al. (2021) used 50 iterations per Schwarz window, while the Schwarz window size was equal to the coupling period in their case.

However, there are physical situations where SWR does not converge. Instead, oscillations with a period of four or two iterations developed, with a comparatively small amplitude. This happened in 17 out of the 312 two-day forecasts presented in Section 5.2. Such a model behavior was previously observed, though never fully explained, in the results of Marti et al. (2021) and Valcke (2021). It is unclear whether the two states are two possible solutions to the discrete coupling problem, or if the SWR method fails to converge to a stable solution in between the states. Investigating these cases further could reveal whether specific physical conditions or parameterizations are responsible for such oscillations. It remains an open question whether series acceleration methods (e.g., Gatzhammer, 2014) or adjustments to the coupling conditions (as suggested in Gross et al., 2018) are a better path to address this issue.

As seen in Figure 5.6 and Figure 5.8, the runtime convergence criteria we introduced in Section 4.3 are a valuable addition to the Schwarz algorithm in atmosphere-ocean coupling. Our approach of looking at multiple, already exposed, coupling variables seems to be robust and applicable to different models and implementations. For future experiments, a stricter SST threshold would be adequate when using the criterion in Equation (4.4), which is most comparable to convergence criteria used by other SWR implementations such as the one by R uth et al. (2021).

### 6.1.3. Limitations of the Numerical Experiments

The conclusions we can draw from the numerical experiments in Chapter 5 are limited by the available input files. Particularly, for technical reasons, we only had atmospheric initial and forcing data available for a single month at the PAPA station. This input file contains six-hourly data, which OpenIFS interpolates linearly in every model time step, which is “likely not a good assumption for the temporal evolution of the forcing fields” (Hartung et al., 2018, p. 4133). Repeating similar experiments with atmospheric data of higher frequency, at different (ice-free) locations, and during more than one season, is advisable to strengthen the robustness of our claims.

Using the GLORYS12V1 reanalysis product allowed us to easily obtain input data for NEMO which is globally available, time-continuous, and requires little preprocessing. However, we note that this data set only contains daily averages, which might impair the quality of the results.

We have not compared any of our results to observation data. In a first step, this is also not necessary: The primary questions of interest were how far standard coupling algorithms deviate from the converged SWR solution. To answer this, no observations have to be taken into account. For future model development however, comparison with observation data is important: Weather and climate models are used for real-world predictions

and cannot be considered in isolation from measurements. As noted by Gross et al. (2018, p. 3515), “when a novel coupling scheme is implemented in a tuned model, the solution is likely to be worse for the new coupling method if the model is then not retuned, even if the new coupling scheme would lead to a superior solution in the absence of tuning.” This should be taken into account if one combines coupling algorithm comparisons with parameter-tuning studies like the one by Deppenmeier et al. (2020).

The EC-Earth AOSCM uses realistic physical parameterization and is, in many ways, closer to a 3D general circulation model than simplified models such as the ones studied by, e.g., Clement et al. (2020), Connors and Doland (2019), and They et al. (2022). However, we stress that results from a single column model cannot be directly compared with comprehensive general circulation models. One example where this was very prominent in our experiments was the topic of ensemble spread (cf. Section 5.2): Since the atmosphere in the AOSCM is forced in every model state by the (linearly interpolated) large-scale dynamical forcing, it does not evolve as freely as typical NWP models. For this reason, we have to conclude that it is not a suitable tool to study the impact of coupling schemes on ensemble spread.

## 6.2. Discussion of the Implementation

As part of this thesis, the EC-Earth AOSCM was extended to allow users to switch between the parallel, sequential atmosphere-first and sequential ocean-first coupling schemes. To support this feature in a coupled model using OASIS3-MCT is straightforward and we have provided the technical details in Appendix A.4. Furthermore, it is now possible to do additive Schwarz waveform relaxation with the AOSCM, using a similar approach as developed by Valcke (2021) as part of the COCOA project.<sup>2</sup> The implementation is based on a Python wrapper which replaces the standard way of running the AOSCM (adapting an XML file and manually running a shell script). To enable future research with different coupling configurations, both within and outside of the EC-Earth AOSCM, the implementation follows good practice to be maintainable: It is thoroughly documented and tutorials for EC-Earth users are available. The code makes use of modern Python features, well-tested open-source libraries, and is written in a modular way.

Overall, our implementation treats OpenIFS and NEMO as black boxes. It is based on fundamental concepts of the OASIS coupling software, which is used in various climate models (Valcke, 2013). The outer Python layer separates the Schwarz algorithm from the rest of the model, as opposed to the SWR implementation by Marti et al. (2021). This separation of responsibilities made a working prototype of SWR for the AOSCM realizable in a matter of weeks. Switching the underlying version of OpenIFS or placing the column at a location with sea ice does not require any changes in the Python wrapper. As soon

---

<sup>2</sup>Additionally, a multiplicative Schwarz method can be replicated with this algorithm, as described in Appendix A.6, but it introduces a lot of duplicate computation. This is acceptable in the case of the AOSCM, which has a low computational cost, but not feasible for 3D coupled general circulation models.

as coupling variables are modified or the horizontal model grids are adapted, adjustments will be necessary. However, such changes will in general affect the coupling setup, i.e., OASIS configuration settings. Thus we conclude that our SWR implementation does not significantly decrease the maintainability and extensibility of the AOSCM. Since we only introduce an outer layer, the model is still usable in the same manner as before: No existing user is required to modify their setup to use the default coupling configuration. Because of these benefits, we propose that SWR logic should generally be implemented separate from the model components; either directly inside the coupling software, as done by R uth et al. (2021), or in a separate component such as the Python wrapper we use.

Conceptually, each Schwarz iteration is a full AOSCM run in our implementation. This creates an overhead of setting up the model structure before every iteration and renaming the output directory afterwards. OASIS itself neither has the control, nor the capability, to have a model component repeat a coupling period. However, this functionality is a crucial part of Schwarz waveform relaxation (cf. Figure 2.5). The only way to support it without heavily modifying the model components and OASIS was to move this step beyond a single AOSCM run.

As a consequence of the previous choice, we so far only support a Schwarz window size equal to the full simulation length of the AOSCM. Larger Schwarz window sizes generally increase the number of iterations for SWR to converge (Gander & Halpern, 2007). This limits the feasible simulation length, as multi-week simulations can easily require upwards of a hundred iterations. A first step to improve this would be to split up a simulation into multiple sub-intervals (e.g., 24h). The output of the previous interval could be used to restart the model. This solution has two issues: First of all, we once again introduce an overhead of setting up the AOSCM for each Schwarz iteration *and* Schwarz window. One could argue that this is somewhat counteracted by the fact that less iterations will be necessary with shorter Schwarz windows, reducing the overall amount of computational effort. The second issue is that the EC-Earth AOSCM, specifically the OpenIFS-SCM, is not properly restartable: While a large amount of variables is saved after a model run, some internal atmospheric variables are not available from model output. Therefore, in our implementation, a simulation split up into multiple Schwarz windows is not guaranteed to give the same results as one where the Schwarz window size is equal to the full simulation time.

In the algorithm implemented as part of this thesis, the model components read in a single value of each coupling variable, representing the average over the interval  $[t_n, t_{n+1}]$ . It would also be possible for each model to read in values more frequently, even in a multirate setting (where NEMO and OpenIFS use different time step sizes). For example, the Python wrapper could interpolate the coupling variables during a coupling period to provide good approximations for the target component. However, this would not change the fact that each model component considers the coupling variable to be constant during a single time step. For true higher order waveform relaxation, the time integration method *inside* OpenIFS/NEMO should have access to the interpolant, not a constant field (R uth et al., 2021). In this sense, "true" higher-order SWR is not achievable for the AOSCM

without significant changes to the model code. A next step could nevertheless be to investigate which components would have to be adapted and how. Furthermore, it is an open question whether an intermediate step towards higher order SWR, such as the one illustrated above, already has measurable impacts on the numerical solution produced by the Schwarz method.

As a final discussion point, we consider potential pitfalls when reusing our approach for a 3D coupled general circulation model. It is a drawback of our implementation that we manually remap the coupling variables between iterations. OASIS does not support remapping fields if they are simply written out to or read in from a file. This is not a big issue in the single column model, where the remapping operation consists of converting a  $1 \times 1$  to a  $3 \times 3$  array and vice versa. But for 3D models, this task is computationally and mathematically involved.<sup>3</sup> To use our implementation for a 3D model, the Python wrapper should not be responsible for regridding the coupling fields. Instead, OASIS should take care of this task. This could be achieved by extending the OASIS feature set (i.e., supporting spatial operations for OASIS output and input fields). Such a solution would be ideal from the perspective of our implementation, since it takes away responsibility from the outer layer. Alternatively, it might be possible to call the remapping functionality outside of direct communication between models, which would introduce an additional step in the outer layer. If this is the case, no development work inside OASIS3-MCT would be necessary.

As indicated by the results of Marti et al. (2021), we expect a slower convergence speed of Schwarz iterations for 3D simulations or when more than two model components are coupled. This might raise the need for acceleration methods when using Schwarz waveform relaxation. Options for this, according to existing literature, appear to be (Aitken) underrelaxation or interface Quasi-Newton methods (Gatzhammer, 2014; Keyes et al., 2013; R uth et al., 2021). To our knowledge, these have thus far not been studied in the context of the atmosphere-ocean coupling problem.

---

<sup>3</sup>In fact, ensuring that regridding is done accurately, reliably, and efficiently, is one of the reasons why coupling software such as OASIS was developed in the first place (Valcke, 2013, Section 2).

## 7. Conclusion and Outlook

This master’s thesis set out to investigate Schwarz waveform relaxation (SWR) in the EC-Earth coupled atmosphere-ocean single column model (AOSCM). The goal was to make the AOSCM a tool for comparing coupling algorithms and their impact on numerical solutions of the atmosphere-ocean coupling problem. To this end, we added a simple way for users to switch between three non-iterative coupling schemes: the parallel, sequential atmosphere-first, and sequential ocean-first algorithm. The first two are most commonly used in operational climate models. Furthermore, we created an outer layer for the AOSCM to support Schwarz waveform relaxation with a Schwarz window size equal to the simulation length. All of these are optional runtime features which do not affect the default behavior of the AOSCM. The implementation treats the model components (the 1D versions of OpenIFS and NEMO) as black boxes and is therefore minimally invasive.

We have carried out two sets of numerical experiments with the AOSCM, both at the PAPA station in the Northern Pacific Ocean during July 2014. We investigated the model behavior for a single four-day simulation, studying various output variables for both the atmosphere and the ocean. This is in contrast to previous numerical studies on atmosphere-ocean coupling, where the focus was exclusively on the results seen in sea surface temperature. We also compared 1,248 two-day simulations to obtain statistical insight into the behavior of the AOSCM for time scales representative in numerical weather prediction.

Our experiments show that changing the coupling algorithm can have a significant impact on the numerical solution, although the differences are small in the majority of simulations at this location and time of year. We have identified the same phase shift in ocean variables as previous research by Marti et al. (2021) and Voltaire et al. (2022). The control experiment revealed that the physical parameterizations in OpenIFS can react very sensitively to small changes in interface boundary conditions. In this specific example, the behavior seems to be related to the boundary layer type in OpenIFS, which is chosen based on the instantaneous state of the atmosphere. Depending on the selected type, different parameterizations for vertical turbulence are chosen for the boundary layer. Particularly, the convective mass flux term in the turbulence closure directly affects the atmospheric stratification, which can cause large differences in temperature and moisture. Such a decision-tree-based parameterization is difficult, if not impossible, to study theoretically.

We used the AOSCM to attempt a ranking of the non-iterative coupling algorithms in terms of how close they are to the converged SWR result, similar to Marti et al. (2021): The sequential atmosphere-first algorithm performs substantially better than the sequential ocean-first algorithm, with the parallel coupling scheme coming in third place. This result is in line with previous research, particularly when considering sea surface temperature.

For atmospheric variables, the three algorithms lie considerably closer together and the conclusions are less obvious.

As part of the implementation, we introduced runtime convergence criteria for Schwarz iterations. These are model-agnostic and combine previous work on multiphysics coupling in general and atmosphere-ocean coupling in particular. The numerical results in Chapter 5 indicate that runtime convergence checks help reduce the computational effort of Schwarz waveform relaxation in atmosphere-ocean coupling: The SWR algorithm converged in at most 19 iterations in 95% of the experiments we conducted in this thesis, while more than half of the experiments required less than ten iterations. In the remaining simulations, the model developed small oscillations between two states.

An important part of this thesis was the consistent definition of the coupling problem solved by the AOSCM in Chapter 3. This merges previous theoretical work with the technical documentations of OpenIFS and NEMO. We explicitly include the physical parameterizations in the primitive equations, which makes the required interface boundary conditions easily identifiable. Our overview can serve as a baseline to explain AOSCM behavior, derive new idealized models, or compare the EC-Earth AOSCM with other coupled atmosphere-ocean models.

As discussed in Chapter 6 we propose to confirm our numerical results with a larger set of locations, time spans, and atmospheric forcing data of higher frequency. One could also repeat the same simulations with another coupled SCM like CNRM-CM6-1D, where a prototype for SWR experiments has been developed as part of the COCOA project (Valcke, 2021; Voltaire et al., 2022). Such a comparison can indicate whether our observations are specific to the physical parameterizations in the EC-Earth AOSCM. The coupling problem formulation and the numerical experiments in this thesis are restricted to an ice-free scenario, although NEMO contains a sea ice model. Including the sea ice-coupling conditions would be an interesting next step, as the atmosphere-ocean-sea ice problem has so far only been investigated with a simplified sea ice model (Lozano, 2022).

Furthermore, we suggest to investigate the cases where the SWR algorithm does not converge. Research questions could be which physical situations cause this behavior and whether series acceleration methods or adapted interface boundary conditions are a better path to address this issue.

Finally, various directions exist to continue the work on algorithmic aspects of SWR in atmosphere-ocean coupling. Obvious improvements to our implementation include adding support for shorter Schwarz window sizes and using OASIS3-MCT instead of the Python wrapper for regriding. Both steps are necessary in order to reuse our approach for, e.g., full complexity models. An open research topic is the potential of higher-order Schwarz waveform relaxation in the context of atmosphere-ocean coupling. The possible benefits in terms of numerical results, but also the implementation challenges of properly supporting it in Earth system models, have not been investigated as of now. The EC-Earth AOSCM is a promising candidate to examine both aspects: It is based on the same code as coupled general circulation models while the computational cost is low and issues with non-matching grids are removed.

# Appendix





# A. Appendix

## A.1. Reynolds Averaging

The central lines of thought in this section are mainly taken from Vallis (2017, Chapter 11), Smits (2009), and Olbers et al. (2012, Chapters 11, 12). The processes in the Earth system fundamentally span "the entire range of spatial and temporal scales" (Gross et al., 2018, p. 3506). In Earth system models, on the other hand, the primitive equations are discretized in space and time, and not all scales are resolved in the discretization. A way to make this apparent is Reynolds averaging, here studied generally for a conserved quantity  $\phi$ . We can view the discrete values of  $\phi$  produced by the model as averages with respect to the discretization, i.e., they represent a mean in space and time which we denote by  $\overline{\phi}$  and call the *Reynolds average* of  $\phi$ . The fluctuations of  $\phi$  which are not resolved in the discretization are referred to as  $\phi'$ , leading to the *Reynolds decomposition*:

$$\phi = \overline{\phi} + \phi'. \quad (\text{A.1})$$

Note that, in the rest of this thesis and in literature, we drop the overline of the conserved quantity, i.e.,  $\phi = \overline{\phi}$ , thus pretending that this separation of resolved and unresolved scales does not actually exist.

We assume that the following identities hold (for a discussion of these see, e.g., Olbers et al. (2012)):

$$\overline{\overline{\phi}} = \overline{\phi} \quad (\text{A.2})$$

$$\overline{\phi'} = 0 \quad (\text{A.3})$$

$$\overline{\phi\psi} = \overline{\phi}\overline{\psi} \quad (\text{A.4})$$

$$\overline{\phi + \psi} = \overline{\phi} + \overline{\psi} \quad (\text{A.5})$$

$$\overline{\frac{\partial}{\partial \gamma} \phi} = \frac{\partial}{\partial \gamma} \overline{\phi}, \quad (\text{A.6})$$

with  $\phi$  and  $\psi$  conserved quantities and  $\gamma$  an independent variable.

We will now consider the incompressible Navier-Stokes equations in Cartesian coordinates and, particularly, the conservation of zonal momentum. By taking the Reynolds average of the momentum conservation equation and using the identities in Equations (A.2)

to (A.6), one obtains (cf. Vallis, 2017):

$$\frac{\partial \bar{u}}{\partial t} + (\bar{\mathbf{u}} \cdot \nabla) \bar{u} = -\frac{\partial \bar{p}}{\partial x} - \nabla \cdot \overline{\mathbf{u}'u'}. \quad (\text{A.7})$$

Recall that  $\mathbf{u} = (u, v, w)^T$  is the velocity vector and  $p$  denotes pressure. The final term is referred to as the Reynolds stresses, which "represent the effects of eddies on the mean flow" (Vallis, 2017, p. 414). Similar terms are encountered in the conservation equations for temperature, salinity, and humidity, assuming again that the flow is incompressible. The second order interaction of the eddy fluctuations  $\overline{u'u'}$ ,  $\overline{v'u'}$ ,  $\overline{w'u'}$  in the Reynolds stress term in turn depend on third order eddy interactions  $\overline{u'u'u'}$ , etc. Thus, as long as not all temporal and spatial scales are resolved in the Navier-Stokes equations, there will be more unknowns than equations, i.e., one encounters a *closure problem* related to turbulence: Is there a way to represent the higher order fluctuations solely in terms of lower order fluctuations? In particular, is it possible to represent the Reynolds stress terms using only mean flow quantities? Assuming that this is possible is referred to as the *turbulent closure hypothesis* and the resulting approximations are given in Section 3.1.

"The effect of the small-scale turbulence in the stratified ocean and atmosphere, with its small aspect ratio, is predominantly given by vertical processes." (Olbers et al., 2012, p. 390) For this reason, only the Reynolds stresses interacting with the vertical velocity fluctuations  $w'$  is kept in the conservation equations, i.e., they reduce to:

$$\nabla \cdot \overline{\mathbf{u}'\phi'} = \frac{\partial}{\partial z} \overline{w'\phi'}. \quad (\text{A.8})$$

On a final note: In this derivation, we have always assumed incompressibility of the fluid in the derivation of the Reynolds stress terms. While this is also the case in the primitive equations for NEMO (cf. Section 3.2), the equations solved by OpenIFS are not using this assumption. But with the hydrostatic approximation and when using pressure coordinates, the continuity equation simplifies in a way that allows a reformulation similar to Equation (A.7) Vallis (2017, p. 2.6.2), with the vertical velocity in pressure coordinates  $w$  instead of  $w$ .

## A.2. Classification of AOSCM Interface Conditions

If one only considers the type of data which NEMO and OpenIFS exchange during coupling (cf. Figure 3.1), one can conclude that NEMO sends the value of prognostic variables, whereas OpenIFS sends back tendencies, i.e., derivatives of prognostic variables. This might suggest that ocean-atmosphere coupling is realized with a Dirichlet-Neumann type coupling condition. However, this is not the case, as is apparent when looking at the interface conditions in Equations (3.24) to (3.29). As stated in Chapter 2, the condition enforced is a continuity of air-sea fluxes at the interface.

However, the interface operator  $\mathcal{C}$  applied to a conserved variable  $\phi$  is not the normal

derivative  $\partial_z \phi$  but instead  $K_\phi \partial_z \phi$ . Since  $K_\phi$  depends on  $\phi$ , we cannot simply treat this as a multiplication by a constant factor. The matter is further complicated by the fact that the fluxes  $\mathcal{J}$  are computed using nonlinear functions.

It is possible to express the interface conditions as a Robin-like condition in special cases such as the sensible heat flux, using Equation (3.20):

$$K_T^a \frac{\partial T^a}{\partial z} \Big|_{z=0} = C_H |U(z_1)| (T^a(z_1) - T^a|_{z=0}), \quad (\text{A.9})$$

which we can reformulate as:

$$\left( C_H |U(z_1)| T^a + K_T^a \frac{\partial T^a}{\partial z} \right) \Big|_{z=0} = C_H |U(z_1)| T^a(z_1). \quad (\text{A.10})$$

Note that this is not possible for, e.g., the horizontal velocities.

Because most of the interface condition treatment happens inside OpenIFS, this more nuanced view is hidden when only considering the data sent between the models. For an accurate numerical analysis of the coupling problem, it is however essential to consider the full complexity. We point to past studies such as Clement et al. (2020), Connors and Ganis (2011), and They et al. (2022) for examples in this regard.

### A.3. Fundamentals of the OASIS Coupler

This section is mainly based on the OASIS3-MCT user guide (Valcke et al., 2015). The *namcouple* is an input file read by OASIS and contains runtime information about the experiment. Particularly, the second section provides information for all coupling and I/O fields used during a simulation, e.g.: the coupling period ( $\Delta t$  in Section 2.2), the coupling field type, the LAG parameter, as well as spatial and temporal transformations OASIS should apply. Each of these fields has to be declared at compile time by model components, using calls to OASIS library functions. However, not all coupling fields declared in the model component have to be part of the *namcouple*.

As a temporal transformation, we always use AVERAGE, thus computing the mean of a field over the coupling period. To map the horizontal grids onto each other, the MAPPING transformation can be used.

#### A.3.1. Coupling Field Types in OASIS

There are four coupling field types important to understand the functionality of OASIS we make use of in this work:

1. EXPORTED: This type of coupling field is sent from one component to another, with potential spatial and temporal transformations applied by OASIS. In a normal coupled simulation, this would be the default status for most fields.

2. `EXPOUT`: A version of `EXPORTED` particularly suited for debugging cases. OASIS sends the coupling field from the source to the target component and applies transformations. Additionally, it writes out two debug files: One by the source component, before sending the field, and one by the target component, after the data was received and all transformations were applied.
3. `OUTPUT`: The field is not sent to the target grid but written out into a file. No spatial transformations are possible when writing out the file.
4. `INPUT`: The field is not received from a source component but read from an input file. No transformations (neither temporal or spatial) can be applied in this case and the `LAG` parameter is not supported.

### A.3.2. The LAG Parameter

A coupling field is sent from the source to the target component when the `oasis_put()` and `oasis_get()` calls match, as determined by OASIS. In the default case, data would be exchanged when the current time step of the source and target component is equal *and* is an integer multiple of the coupling period, i.e.,

$$t_{\text{source}} = t_n = t_{\text{target}}. \quad (\text{A.11})$$

Recall that we defined  $t_n$  as a coupling time step in Section 2.2. The `LAG` parameter shifts the `oasis_put()` calls in time relative to the `oasis_get()` calls:

$$t_{\text{source}} + \text{LAG} = t_n = t_{\text{target}}. \quad (\text{A.12})$$

If  $\text{LAG} = 0$ , we recover the default behavior.

In case  $\text{LAG} > 0$ , the target component receives the data from an *earlier* point in time,  $t_n - \text{LAG}$ . Assuming the two components take roughly the same amount of time to complete a coupling window, this usually means that the call to `oasis_put()` already happened when `oasis_get()` is called. This can prevent deadlocks and reduces waiting times, along with a potential increase in inter-model parallelism. At the first coupling time step  $t_0$ , the source model has not produced any data that can be sent. In this case, the user has to provide a so-called restart file, which gets read in and sent to the target model.

Choosing  $\text{LAG} < 0$  has the effect that the target component receives data from a later point in time as a boundary condition. This has drawbacks from a computational efficiency point of view. However, it enables sequential coupling schemes which can yield more accurate results (Marti et al., 2021).

Modifying the `LAG` parameter thus fundamentally changes the coupling logic and gives the user control over the question which boundary condition a model sees during a coupling window.

## A.4. Sequential Coupling with OASIS

In the parallel coupling case, both OpenIFS and NEMO use their respective model time step size as the `LAG` parameter. The component model that computes the coupling window  $[t_n, t_{n+1}]$  first has to use its time step size minus the coupling period as a lag. The other component, which uses the newer result in its computation of the same coupling window, keeps its model time step size as the coupling lag.

We give an example of sequential atmosphere-first coupling, assuming the following parameters: We choose a coupling period of  $\Delta t = 3600\text{s}$ , an atmosphere time step size  $\Delta t_{\text{atm}} = 900\text{s}$ , and an ocean time step size of  $\Delta t_{\text{oce}} = 1800\text{s}$ . As required, the coupling period is an integer multiple of both model time step sizes.

Fields that are sent from OpenIFS to NEMO, e.g., the wind stresses, use  $\text{LAG} = \Delta t_{\text{atm}} - \Delta t = -2700\text{s}$ . In general, the component model "going first" has  $\text{LAG} \leq 0$ . Fields sent from NEMO to OpenIFS, in our case only the sea surface temperature, have  $\text{LAG} = \Delta t_{\text{oce}} = 1800\text{s}$ .

Because the ocean-to-atmosphere fields have a positive `LAG` parameter, i.e., they are read in from a restart file in the first coupling window. On the other hand, atmosphere-to-ocean fields have a non-positive lag. Thus, they are only sent after computation by the atmospheric model, in all coupling windows including the first one.

In sequential ocean-first coupling, using the same logic and example, one obtains 900s as the lag for atmosphere-to-ocean fields and -1800s as the lag for ocean-to-atmosphere fields.

## A.5. OASIS Configuration Examples for SWR

We include two excerpts from the OASIS *namcouple* files used to configure the AOSCM in a SWR simulation. We restrict ourselves to sea surface temperature, which NEMO sends to OpenIFS. We assume the same simulation parameters as in Appendix A.3.2 and that the parallel algorithm is used in the first Schwarz iteration. The respective part of the *namcouple* takes the following form:

---

```

1 O_SSTSST A_SST 1 3600 2 rstos.nc EXPOUT
2 3 3 1 1 OC1D ASCM LAG=1800
3 R 0 R 0
4 LOCTRANS MAPPING
5     AVERAGE
6     rmp_OC1D_to_ASCM.nc

```

---

For the specific syntax, we point to the OASIS3-MCT user guide (Valcke et al., 2015, Chapter 3). These lines contain the following information:

## A. Appendix

---

- the coupling field is called `O_SSTSST` on the source component (NEMO) and `A_SST` on the target component (OpenIFS);
- the coupling period is  $\Delta t = 3600\text{s} = 1\text{h}$ ;
- the coupling lag is equal to  $\Delta t_{\text{oce}} = 1800\text{s}$ ;
- two transformations are applied by OASIS during communication: averaging over the coupling period and remapping from the  $3 \times 3$  OC1D grid to the  $1 \times 1$  ASCM grid, using the weights specified in `rmp_OC1D_to_ASCM.nc`;
- the coupling fields should be sent *and* written out to debug files, as specified by `EXPOUT`.

For a regular coupled run of the AOSCM, `EXPORTED` would also be a valid choice for the coupling field type. However, our SWR implementation requires the OASIS output files in order to reuse data from previous iterations, as explained in Section 4.2.

For every Schwarz iteration after the initial guess, a different variant of the *namcouple* is used:

---

```
1 # write out current iteration
2 O_SSTSST O_SSTSST 1 3600 1 rstos.nc OUTPUT
3 OC1D OC1D LAG=1800
4 LOCTRANS
5     AVERAGE
6 # read in from previous iteration
7 A_SST A_SST 1 3600 0 A_SST.nc INPUT
```

---

The treatment of the sea surface temperature is now split up into two different OASIS tasks: saving the values of the current iteration to an output file and loading data from the previous iteration. Note that here, neither regridding nor renaming coupling variables is done by OASIS: `OUTPUT` and `INPUT` fields do not allow for the same transformations as when data is communicated between two components. Instead, our Python wrapper takes care of this task. Time-averaging is supported for `OUTPUT` fields, which is why we can utilize it in our implementation.

## A.6. Multiplicative SWR with the AOSCM

We want to show how our implementation reduces to a multiplicative SWR method when one of the sequential coupling algorithms is used in the first Schwarz iteration. Recall that the first Schwarz iteration corresponds to a regular coupled AOSCM run, whereas the two models do not communicate directly in later iterations. As an example, we consider

the case of the sequential atmosphere-first algorithm, depicted in Figure 2.4a. For now we restrict ourselves to the case where the coupling window size is equal to the Schwarz window size.


We now omit the details about which data is sent in the sequential atmosphere-first algorithm and arrive at the following control flow: First, the atmosphere is evolved from time step  $t_n$  to  $t_{n+1}$ . Then, it sends data to the ocean model, which uses this data in its interface boundary conditions while integrating from  $t_n$  to  $t_{n+1}$ . This data dependency is visualized in the leftmost part of Figure A.1a.

After this initial iteration, the Python wrapper post-processes the OASIS output files and they are used for the interface boundary conditions of the next iteration. As explained in Section 4.2, the communication now happens *between*, not *during* model runs. Figure A.1a shows that the ocean model receives the same atmospheric data in both the first and the second Schwarz iteration. The numerical solution of a model component only depends on the values of the initial, boundary, and interface conditions. The only input data changing for model components between two Schwarz iterations are the interface conditions, which are provided by OASIS. In this example, NEMO receives the same interface data in the first two iterations and will thus produce the same numerical solution. This in turn means that NEMO provides the same coupling data to OpenIFS in iterations 2 and 3, which leads to OpenIFS producing the same numerical solution in these iterations, etc.


In Figure A.1b, we mark all duplicate communication and computation in gray to illustrate this effect. All other arrows depict the production of new numerical solutions and the communication of new interface data, as long as the Schwarz iterations have not converged. We can remove the duplicate operations from Figure A.1b and “compress” the diagram to obtain a mathematically equivalent depiction of the thus implemented algorithm in Figure A.1c. Comparing this to Figure 2.5b, we see that this corresponds to a multiplicative SWR algorithm.

In our implementation, the Schwarz window size is equal to the simulation length, which makes the explanation and visualization of the behavior more complex. However, even in this case the implemented algorithm corresponds to a variant of a multiplicative SWR method. We have numerically confirmed these observations for a single coupling window and for longer simulations. We have also seen in example simulations that the multiplicative SWR method takes about half as many iterations to converge as the additive SWR method, as expected (Gander, 2008).


Notably, the duplicate computation and communication makes this realization of a multiplicative SWR method highly inefficient. We do not advise using this approach for anything that exceeds computationally cheap simulations like the short experiments we did with the EC-Earth AOSCM.



(a) Control flow if the sequential atmosphere-first algorithm is used in the first Schwarz iteration.



(b) As in (a) but duplicate communication and computation is marked in gray.



(c) Effectively implemented algorithm, which looks identical to Figure 2.5b.

Figure A.1.: Diagram to illustrate how a sequential coupling algorithm combined with our additive SWR method reduces to multiplicative SWR. Colors and notation as in Figure 2.5.







# Bibliography

- American Meteorological Society. (2022). *Glossary of Meteorology*.
- Blayo, E., Cherel, D., & Rousseau, A. (2016). Towards Optimized Schwarz Methods for the Navier–Stokes Equations. *Journal of Scientific Computing*, 66(1), 275–295. <https://doi.org/10.1007/s10915-015-0020-9>
- Blayo, E., Rousseau, A., & Tayachi, M. (2017). Boundary conditions and Schwarz waveform relaxation method for linear viscous Shallow Water equations in hydrodynamics. *The SMAI journal of computational mathematics*, 3, 117–137. <https://doi.org/10.5802/smai-jcm.22>
- Clement, S., Lemarié, F., & Blayo, E. (2020). Discrete analysis of Schwarz Waveform Relaxation for a simplified air-sea coupling problem with nonlinear transmission conditions. *DD26 - 26th International Domain Decomposition Conference*, 1–8. Retrieved November 25, 2022, from <https://hal.archives-ouvertes.fr/hal-03434816>
- Clement, S., Lemarié, F., & Blayo, E. (2022). Discrete analysis of Schwarz waveform relaxation for a diffusion reaction problem with discontinuous coefficients. *The SMAI Journal of computational mathematics*, 8, 99–124. <https://doi.org/10.5802/smai-jcm.81>
- Connors, J. M., & Doland, R. D. (2019). Stability of Two Conservative, High-Order Fluid-Fluid Coupling Methods. *Advances in Applied Mathematics and Mechanics*, 11(6), 1287–1338. <https://doi.org/10.4208/aamm.OA-2018-0212>
- Connors, J. M., & Ganis, B. (2011). Stability of algorithms for a two domain natural convection problem and observed model uncertainty. *Computational Geosciences*, 15(3), 509–527. <https://doi.org/10.1007/s10596-010-9219-x>
- Connors, J. M., & Miloua, A. (2011). Partitioned time discretization for parallel solution of coupled ODE systems. *BIT Numerical Mathematics*, 51(2), 253–273. <https://doi.org/10.1007/s10543-010-0295-z>
- Couvelard, X., Lemarié, F., Samson, G., Redelsperger, J.-L., Ardhuin, F., Benshila, R., & Madec, G. (2020). Development of a two-way-coupled ocean–wave model: Assessment on a global NEMO(v3.6)–WW3(v6.02) coupled configuration. *Geoscientific Model Development*, 13(7), 3067–3090. <https://doi.org/10.5194/gmd-13-3067-2020>
- Craig, A., Valcke, S., & Coquart, L. (2017). Development and performance of a new version of the OASIS coupler, OASIS3-MCT.3.0. *Geoscientific Model Development*, 10(9), 3297–3308. <https://doi.org/10.5194/gmd-10-3297-2017>
- Dee, D. P., Uppala, S. M., Simmons, A. J., Berrisford, P., Poli, P., Kobayashi, S., Andrae, U., Balmaseda, M. A., Balsamo, G., Bauer, P., Bechtold, P., Beljaars, A. C. M., van de Berg, L., Bidlot, J., Bormann, N., Delsol, C., Dragani, R., Fuentes, M., Geer, A. J.,

- ... Vitart, F. (2011). The ERA-Interim reanalysis: Configuration and performance of the data assimilation system. *Quarterly Journal of the Royal Meteorological Society*, 137(656), 553–597. <https://doi.org/10.1002/qj.828>
- Deppenmeier, A.-L., Haarsma, R. J., van Heerwaarden, C., & Hazeleger, W. (2020). The Southeastern Tropical Atlantic SST Bias Investigated with a Coupled Atmosphere–Ocean Single-Column Model at a PIRATA Mooring Site. *Journal of Climate*, 33(14), 6255–6271. <https://doi.org/10.1175/JCLI-D-19-0608.1>
- ECMWF. (2014). *IFS Documentation CY40R1*. ECMWF. <https://doi.org/10.21957/f56vvey1x>
- ECMWF. (2021). *IFS Documentation CY47R3*. ECMWF. <https://doi.org/10.21957/b18qsx663>
- Fairall, C. W., Bradley, E. F., Godfrey, J. S., Wick, G. A., Edson, J. B., & Young, G. S. (1996). Cool-skin and warm-layer effects on sea surface temperature. *Journal of Geophysical Research: Oceans*, 101(C1), 1295–1308. <https://doi.org/10.1029/95JC03190>
- Fitch, A. C. (2022). Improving stratocumulus cloud turbulence and entrainment parametrizations in OpenIFS. *Quarterly Journal of the Royal Meteorological Society*, 148(745), 1782–1804. <https://doi.org/10.1002/qj.4278>
- Gander, M. J. (2008). Schwarz methods over the course of time. *Electronic Transactions on Numerical Analysis*, 31, 228–255.
- Gander, M. J. (2015). Waveform Relaxation. In B. Engquist (Ed.), *Encyclopedia of Applied and Computational Mathematics* (pp. 1549–1552). Springer.
- Gander, M. J., & Halpern, L. (2007). Optimized Schwarz Waveform Relaxation Methods for Advection Reaction Diffusion Problems. *SIAM Journal on Numerical Analysis*, 45(2), 666–697.
- Gander, M. J., Halpern, L., & Nataf, F. (1999). Optimal convergence for overlapping and non-overlapping Schwarz waveform relaxation, 27–36.
- Gatzhammer, B. (2014). *Efficient and Flexible Partitioned Simulation of Fluid-Structure Interactions* (Doctoral dissertation). Technical University of Munich. Munich, Germany.
- Gross, M., Wan, H., Rasch, P. J., Caldwell, P. M., Williamson, D. L., Klocke, D., Jablonowski, C., Thatcher, D. R., Wood, N., Cullen, M., Beare, B., Willett, M., Lemarié, F., Blayo, E., Malardel, S., Termonia, P., Gassmann, A., Lauritzen, P. H., Johansen, H., ... Leung, R. (2018). Physics–Dynamics Coupling in Weather, Climate, and Earth System Models: Challenges and Recent Progress. *Monthly Weather Review*, 146(11), 3505–3544. <https://doi.org/10.1175/MWR-D-17-0345.1>
- Hartung, K., Svensson, G., Struthers, H., Deppenmeier, A.-L., & Hazeleger, W. (2018). An EC-Earth coupled atmosphere–ocean single-column model (AOSCM.v1\_EC-Earth3) for studying coupled marine and polar processes. *Geoscientific Model Development*, 11(10), 4117–4137. <https://doi.org/10.5194/gmd-11-4117-2018>
- Hourdin, F., Mauritsen, T., Gettelman, A., Golaz, J.-C., Balaji, V., Duan, Q., Folini, D., Ji, D., Klocke, D., Qian, Y., Rauser, F., Rio, C., Tomassini, L., Watanabe, M., & Williamson, D. (2017). The Art and Science of Climate Model Tuning. *Bulletin of the American Meteorological Society*, 98(3), 589–602. <https://doi.org/10.1175/BAMS-D-15-00135.1>

- Huang, R. X. (1993). Real Freshwater Flux as a Natural Boundary Condition for the Salinity Balance and Thermohaline Circulation Forced by Evaporation and Precipitation. *Journal of Physical Oceanography*, 23(11), 2428–2446. [https://doi.org/10.1175/1520-0485\(1993\)023<2428:RFFAAN>2.0.CO;2](https://doi.org/10.1175/1520-0485(1993)023<2428:RFFAAN>2.0.CO;2)
- Keyes, D. E., McInnes, L. C., Woodward, C., Gropp, W., Myra, E., Pernice, M., Bell, J., Brown, J., Clo, A., Connors, J., Constantinescu, E., Estep, D., Evans, K., Farhat, C., Hakim, A., Hammond, G., Hansen, G., Hill, J., Isaac, T., ... Wohlmuth, B. (2013). Multiphysics Simulations: Challenges and Opportunities. *The International Journal of High Performance Computing Applications*, 27(1), 4–83. <https://doi.org/10.1177/1094342012468181>
- Lauritzen, P. H., Kevlahan, N. K.-R., Toniazzo, T., Eldred, C., Dubos, T., Gassmann, A., Larson, V. E., Jablonowski, C., Guba, O., Shipway, B., Harrop, B. E., Lemarié, F., Tailleux, R., Herrington, A. R., Large, W., Rasch, P. J., Donahue, A. S., Wan, H., Conley, A. J., & Bacmeister, J. T. (2022). Reconciling and improving formulations for thermodynamics and conservation principles in Earth System Models (ESMs). *Journal of Advances in Modeling Earth Systems*. <https://doi.org/10.1029/2022MS003117>
- Lemarié, F., Blayo, E., & Debreu, L. (2015). Analysis of Ocean-atmosphere Coupling Algorithms: Consistency and Stability. *Procedia Computer Science*, 51, 2066–2075. <https://doi.org/10.1016/j.procs.2015.05.473>
- Lemarié, F., Debreu, L., & Blayo, E. (2013a). Toward an Optimized Global-in-Time Schwarz Algorithm for Diffusion Equations with Discontinuous and Spatially Variable Coefficients, Part 1: The Constant Coefficients Case. *Electronic Transactions on Numerical Analysis*, 40, 148–169. Retrieved April 1, 2022, from <https://hal.archives-ouvertes.fr/hal-00661977>
- Lemarié, F., Debreu, L., & Blayo, E. (2013b). Toward an Optimized Global-in-Time Schwarz Algorithm for Diffusion Equations with Discontinuous and Spatially Variable Coefficients, Part 2: The Variable Coefficients Case. *Electronic Transactions on Numerical Analysis*, 40, 170–186. Retrieved April 1, 2022, from <https://hal.archives-ouvertes.fr/hal-00661978>
- Lemarié, F., Marchesiello, P., Debreu, L., & Blayo, E. (2014). Sensitivity of Ocean-Atmosphere Coupled Models to the Coupling Method: Example of Tropical Cyclone Erica.
- Lions, P.-L. (1988). On the Schwarz Alternating Method. I. *First International Symposium on Domain Decomposition Methods for Partial Differential Equations*, 1, 1–42.
- Lozano, P. (2022, July). *Analysis and Optimization of Schwarz Algorithms for Ocean-Sea Ice-Atmosphere Coupling* (MS Thesis). Université Grenoble Alpes. Grenoble, France.
- Madec, G., Bourdallé-Badie, R., Bouttier, P.-A., Bricaud, C., Bruciaferri, D., Calvert, D., Chanut, J., Clementi, E., Coward, A., Delrosso, D., Ethé, C., Flavoni, S., Graham, T., Harle, J., Iovino, D., Lea, D., Lévy, C., Lovato, T., Martin, N., ... Vancoppenolle, M. (2017). NEMO ocean engine. <https://doi.org/10.5281/zenodo.1472492>
- Maher, P., Gerber, E. P., Medeiros, B., Merlis, T. M., Sherwood, S., Sheshadri, A., Sobel, A. H., Vallis, G. K., Voigt, A., & Zurita-Gotor, P. (2019). Model Hierarchies for Un-

- derstanding Atmospheric Circulation. *Reviews of Geophysics*, 57(2), 250–280. <https://doi.org/10.1029/2018RG000607>
- Marti, O., Nguyen, S., Braconnot, P., Valcke, S., Lemarié, F., & Blayo, E. (2021). A Schwarz iterative method to evaluate ocean–atmosphere coupling schemes: Implementation and diagnostics in IPSL-CM6-SW-VLR. *Geoscientific Model Development*, 14(5), 2959–2975. <https://doi.org/10.5194/gmd-14-2959-2021>
- Mehl, M., Uekermann, B., Bijl, H., Blom, D., Gatzhammer, B., & van Zuijlen, A. (2016). Parallel coupling numerics for partitioned fluid–structure interaction simulations. *Computers & Mathematics with Applications*, 71(4), 869–891. <https://doi.org/10.1016/j.camwa.2015.12.025>
- Monin, A. S., & Obukhov, A. M. (1954). Basic laws of turbulent mixing in the surface layer of the atmosphere. *Contrib. Geophys. Inst. Acad. Sci. USSR*, 151(163).
- Olbers, D., Willebrand, J., & Eden, C. (2012). *Ocean Dynamics*. Springer Berlin Heidelberg. <https://doi.org/10.1007/978-3-642-23450-7>
- Reffray, G., Bourdalle-Badie, R., & Calone, C. (2015). Modelling turbulent vertical mixing sensitivity using a 1-D version of NEMO. *Geoscientific Model Development*, 8(1), 69–86. <https://doi.org/10.5194/gmd-8-69-2015>
- Renault, L., Lemarié, F., & Arsouze, T. (2019). On the implementation and consequences of the oceanic currents feedback in ocean–atmosphere coupled models. *Ocean Modelling*, 141, 101423. <https://doi.org/10.1016/j.ocemod.2019.101423>
- Rüth, B., Uekermann, B., Mehl, M., Birken, P., Monge, A., & Bungartz, H.-J. (2021). Quasi-Newton Waveform Iteration for Partitioned Surface-Coupled Multiphysics Applications. *International Journal for Numerical Methods in Engineering*, 122(19), 5236–5257. <https://doi.org/10.1002/nme.6443>
- Rüth, B., Uekermann, B., Mehl, M., & Bungartz, H.-J. (2018). Time stepping algorithms for partitioned multi-scale multi-physics in preCICE. *6th European Conference on Computational Mechanics (Solids, Structures and Coupled Problems)*, 2678–2689.
- Schlichting, H., & Gersten, K. (2017). *Boundary-Layer Theory*. Springer Berlin Heidelberg. <https://doi.org/10.1007/978-3-662-52919-5>
- Schwarz, H. A. (1870). *Ueber einen Grenzübergang durch alternirendes Verfahren*. Zürcher u. Furrer.
- Siebesma, A. P., Soares, P. M. M., & Teixeira, J. (2007). A Combined Eddy-Diffusivity Mass-Flux Approach for the Convective Boundary Layer. *Journal of the Atmospheric Sciences*, 64(4), 1230–1248. <https://doi.org/10.1175/JAS3888.1>
- Simmons, A. J., & Burridge, D. M. (1981). An Energy and Angular-Momentum Conserving Vertical Finite-Difference Scheme and Hybrid Vertical Coordinates. *Monthly Weather Review*, 109(4), 758–766. [https://doi.org/10.1175/1520-0493\(1981\)109\(0758:AEAAMC\)2.0.CO;2](https://doi.org/10.1175/1520-0493(1981)109(0758:AEAAMC)2.0.CO;2)
- Smits, A. J. (2009). *Viscous Flows and Turbulence* [Lecture Notes]. Retrieved November 22, 2022, from <http://profs.sci.univr.it/~zuccher/downloads/FD-MAE553-Smits.pdf>
- Streffing, J., Sidorenko, D., Semmler, T., Zampieri, L., Scholz, P., Andrés-Martínez, M., Koldunov, N., Rackow, T., Kjellsson, J., Goessling, H., Athanase, M., Wang, Q.,

- 
- Hegewald, J., Sein, D. V., Mu, L., Fladrich, U., Barbi, D., Gierz, P., Danilov, S., ... Jung, T. (2022). AWI-CM3 coupled climate model: Description and evaluation experiments for a prototype post-CMIP6 model. *Geoscientific Model Development*, 15(16), 6399–6427. <https://doi.org/10.5194/gmd-15-6399-2022>
- Thery, S., Pelletier, C., Lemarié, F., & Blayo, E. (2022). Analysis of Schwarz waveform relaxation for the coupled Ekman boundary layer problem with continuously variable coefficients. *Numerical Algorithms*, 89(3), 1145–1181. <https://doi.org/10.1007/s11075-021-01149-y>
- Valcke, S. (2013). The OASIS3 coupler: A European climate modelling community software. *Geoscientific Model Development*, 6(2), 373–388. <https://doi.org/10.5194/gmd-6-373-2013>
- Valcke, S. (2021). *Schwarz iterations for ocean-atmosphere interface coherency in CNRM-CM6-1D*.
- Valcke, S., Craig, A., & Coquart, L. (2015). OASIS3-MCT User Guide.
- Vallis, G. K. (2017). *Atmospheric and Oceanic Fluid Dynamics: Fundamentals and Large-Scale Circulation* (2nd ed.). Cambridge University Press. <https://doi.org/10.1017/9781107588417>
- Voldoire, A., Roehrig, R., Giordani, H., Waldman, R., Zhang, Y., Xie, S., & Bouin, M.-N. (2022). Assessment of the sea surface temperature diurnal cycle in CNRM-CM6-1 based on its 1D coupled configuration. *Geoscientific Model Development*, 15(8), 3347–3370. <https://doi.org/10.5194/gmd-15-3347-2022>