

Design and implementation of the infrastructure of HadGEM3: the next-generation Met Office climate modelling system

H. T. Hewitt¹, D. Copsy¹, I. D. Culverwell¹, C. M. Harris¹, R. S. R. Hill¹, A. B. Keen¹, A. J. McLaren¹, and E. C. Hunke²

¹Met Office, Hadley Centre, Exeter, Devon, UK

²T-3 Fluid Dynamics and Solid Mechanics Group, Los Alamos National Laboratory, Los Alamos, NM, USA

Received: 5 October 2010 – Published in Geosci. Model Dev. Discuss.: 27 October 2010

Revised: 11 March 2011 – Accepted: 18 March 2011 – Published: 4 April 2011

Abstract. This paper describes the development of a technically robust climate modelling system, HadGEM3, which couples the Met Office Unified Model atmosphere component, the NEMO ocean model and the Los Alamos sea ice model (CICE) using the OASIS coupler. Details of the coupling and technical solutions of the physical model (HadGEM3-AO) are documented, in addition to a description of the configurations of the individual submodels. The paper demonstrates that the implementation of the model has resulted in accurate conservation of heat and freshwater across the model components. The model performance in early versions of this climate model is briefly described to demonstrate that the results are scientifically credible. HadGEM3-AO is the basis for a number of modelling efforts outside of the Met Office, both within the UK and internationally. This documentation of the HadGEM3-AO system provides a detailed reference for developers of HadGEM3-based climate configurations.

1 Introduction

Driven by the consensus that climate change is now taking place (IPCC, 2007), there is an increasing demand for climate predictions on regional scales. This is leading to a move towards a family of climate models with options for both increased complexity and resolution (Martin et al., 2011). Climate models at high resolution need to be particularly robust, as shorter grid-lengths mean that models are closer to the limits of the stability criteria. This can be a particular problem around the North Pole in ocean and sea ice models – convergence of grids towards the pole can lead to a requirement

to heavily filter fields to avoid grid point noise and maintain model stability (Johns et al., 2006). Technical constraints on filtering have been found to make the filtering proportionately more expensive, and to limit scalability, which is an increasingly important consideration at higher resolution.

Previous Hadley Centre Global Environmental Models (HadGEM1, HadGEM2) have improved formulation and resolution in the Unified Model (UM) atmosphere component (HadGEM1, Johns et al., 2006) and increased complexity in the carbon cycle (HadGEM2, Collins et al., 2008). The focus of this paper is the technical development of the HadGEM3 climate model which is intended to be the next-generation climate model. Although HadGEM3 has been initially built at a resolution comparable to HadGEM1, the intention is that the model structure will form the basis for running models at significantly higher resolution as well as at lower resolution with increased complexity (as in the Quantifying Uncertainty in the Earth SysTem [QUEST] Model [QESM]). This will establish a HadGEM3 family of models which have the same underlying structure and physics. In this paper, we discuss only the physical model (hereafter referred to as HadGEM3-AO). In future, this will form the basis for the Earth System model (HadGEM3-ES).

The requirement for high resolution climate models means that it is no longer practical to maintain ocean and sea ice components which have a geographic North Pole. The grid convergence of HadGEM1 required not just filtering to maintain stability but also the insertion of a polar island. While a higher resolution version of this model was built as part of the HiGEM project (Shaffrey et al., 2009), the fidelity of the simulation of ocean and sea ice in the Arctic under such approaches is questionable. The approach taken here in HadGEM3-AO is to introduce new ocean and sea ice components that are formulated with generalised orthogonal coordinates and are therefore able to employ a tripolar grid (e.g. Murray, 1996).



Correspondence to: H. Hewitt
(helene.hewitt@metoffice.gov.uk)

Introducing new ocean and sea ice components into the climate model has enabled us to integrate our model development efforts with the international modelling community by using community models – NEMO (Madec, 2008) for the ocean and CICE (Hunke and Lipscomb, 2008) for the sea ice. The use of common code structures means that implementations of new physics can be shared but still allows different modelling centres to use different dynamics and parameterizations, maintaining model diversity.

The HadGEM3 model uses the communication model OASIS3 (Valcke, 2006) to enable coupling between the UM atmosphere, NEMO ocean and CICE sea ice. The use of OASIS means that a number of different algorithms for coupling are available and we have been able to investigate their impact on conservation. In addition, the introduction of the new component models means that we have also been able to develop a hierarchy of models to aid diagnosis of coupled model errors; forced sea ice models and forced ocean-ice models are now available in addition to pre-existing AMIP style atmosphere models (i.e. atmosphere-only models with prescribed sea surface temperature (SST) fields).

The development of HadGEM3-AO represents a departure from our usual mode of model development (HadCM3, HadGEM1 and HadGEM2) where we have worked towards a single fixed version of the model with specified physics (schemes and parameters) and a defined resolution. In previous model developments, model versions prior to the final definition have not been utilised for any applications. Our strategy with HadGEM3-AO is twofold: firstly, the intention is to develop a model family whereby the model is available at a range of resolutions and secondly, model development will continue in cycles with updated versions of the model released in a similar manner to the operational weather forecast model development cycle. The version of the model will be identified by a resolution and a revision number. For these reasons, it will be possible to use HadGEM3-AO in applications such as seasonal forecasting (Arribas et al., 2011) where the modelling system is updated on a regular basis following an operational cycle. However, we have not reached a stage in the development cycle of HadGEM3-AO where we are ready to commit to a version that can be used for long climate runs as required by the Intergovernmental Panel on Climate Change (IPCC).

In this paper, where scientific details about the HadGEM3-AO model are referred to, they are documenting HadGEM3-AO revision 1.1 (r1.1¹), while technical details are specific to implementation on the NEC SX8 supercomputer (as used

¹ Revision r1.1 originated at UM7.1 but has been upgraded to UM7.2, UM7.4 and UM7.6. The model is described in the Unified Model User Interface [UMUI] as job name ahzze. This job is essentially the version of the model assessed in March 2009 as part of the HadGEM3 development project (with minor changes to the cloud inhomogeneity and ocean vertical mixing scheme to make it more suitable for seasonal use).

at the Met Office until early 2009). In particular, note that r1.1 refers to a revision within the framework and is not, in principle, fixed to any particular resolution. For the purposes of this paper we will show results from a model run at a resolution comparable to HadGEM1.


In this paper, Sect. 2 introduces the coupling of different components of the climate system and describes the three scientific submodels (UM atmosphere, NEMO ocean and CICE sea ice) as well as the inter-component communication model (OASIS). In Sect. 3, we describe the exchange of fields between different submodels while Sect. 4 discusses conservation of heat, freshwater and momentum. Section. 5 describes the model performance in HadGEM3-AO r1.1. Finally, Sect. 6 summarises the results and discusses any outstanding issues to be addressed.

2 Coupling atmosphere, ocean and sea ice models

HadGEM3-AO couples together three submodels (UM atmosphere, NEMO and CICE) to form the climate modelling system. To simulate the climate, the submodels should exchange heat, freshwater, salt and momentum, in a conserving manner. The heat exchanges between the atmosphere, ocean and sea ice models (i.e., the interfaces between individual submodels rather than within the submodels themselves) are illustrated in Fig. 1a. The components of the heat exchange between the atmosphere and ocean are solar, longwave and turbulent (sensible and latent) fluxes. Snowfall landing in the ocean, and subsequently melting, adds an additional heat flux to the ocean. Sea ice exchanges heat with the atmosphere at its top surface which can cause the surface to melt. At the bottom surface of the sea ice, heat is exchanged with the ocean causing ice growth or melt. Heat is also passed from the sea ice to the ocean when frazil ice forms in the ocean. Freshwater exchanges (illustrated in Fig. 1b) involve rain and snow from the atmosphere providing freshwater to the ocean model or the sea ice model where ice is present. The ocean submodel uses a linear free surface and therefore freshwater fluxes are converted to salt fluxes for the ocean. River input supplies freshwater to the ocean. Evaporation causes freshwater loss from the ocean; sublimation/condensation processes cause a loss/gain of freshwater from the sea ice component. The formation or melting of sea ice creates fluxes of freshwater and salt between the ocean and sea ice. Exchanges of momentum are not shown schematically but consist of the surface wind stress forcing the ocean and sea ice models. The ocean and sea ice are further coupled through an ocean-ice stress created by the difference in the surface ocean and ice velocities.

As discussed earlier, the intention of the HadGEM3-AO development is that rather than a specified model developed with a fixed horizontal and vertical resolution and prescribed parameter settings in each submodel, the model will be developed as a family of models with a range of resolutions

a) Heat exchanges



b) Freshwater exchanges




Fig. 1. Schematic diagrams showing the exchanges between the atmosphere (which includes land surface), sea ice and ocean models in HadGEM3 of: (a) heat; and (b) freshwater. Arrows indicate the direction of information flow in the model rather than the direction of fluxes. ⁽¹⁾ Evaporation/sublimation on sea ice will affect the sea surface height albeit only by a small amount. This effect is currently not included in HadGEM3 but will be added in future developments. ⁽²⁾ The fraction of snowfall that falls on the sea ice will have an immediate impact on the sea surface height and is added to this ocean flux. The snow landing on the ice will not affect the sea surface salinity until it melts, at which point the snow melt is included in the ice-ocean freshwater flux term. ⁽³⁾ The fraction of rainfall that falls on the sea ice is assumed to drain immediately to the leads and therefore simultaneously impacts both salinity and sea surface height. The rainfall fraction is added directly to the flux affecting the sea surface height. It is also passed to the sea ice model where it is added into the ice-ocean freshwater flux which then contributes to the flux affecting sea surface salinity.

existing for each submodel. For that reason, the model components are described only briefly and in generic terms here and the reader is referred to appendices for the details of the submodel configurations which were used to build HadGEM3-AO r1.1.

The HadGEM3-A r1.1 atmosphere submodel (Appendix A) is a particular configuration of the Met Office's Unified Model (UM; Davies et al., 2005). The UM em-

ployes a semi-Lagrangian advection scheme, complex parameterizations of a number of processes (including radiation, convection, boundary layer mixing and cloud microphysics) as well as aerosol effects and feedbacks². The version of HadGEM3-A described in Sect. 5 has a horizontal resolution of 1.25° latitude by 1.875° longitude with 38 vertical levels. As the atmosphere submodel is run on a regular grid it suffers problems associated with the convergence of the meridians in polar regions. However, the semi-implicit, semi-lagrangian nature of the dynamical core, combined with Fourier filtering at high latitudes (Davies et al., 2005) permits the use of timesteps that violate the CFL criterion (where the Courant number exceeds 1 as described by Staniforth and Côté, 1991). HadGEM3-A r1.1 includes coastal tiling where grid boxes on the atmosphere grid can contain part land and part ocean, improving the representation of the coastline.

The land surface hydrology scheme used is the Met Office Surface Exchange Scheme (MOSES-II, Cox et al., 1999) which is the same as that used by HadGEM1 and HadGEM2 but with updated soil properties and soil hydraulics. The river routing scheme used is the Total Runoff Integrating Pathways scheme (TRIP; Oki and Sud, 1998) also largely unchanged from HadGEM1 and HadGEM2. Both the land surface hydrology scheme and the river routing scheme run as subcomponents of the atmospheric model as they are included in the same executable and code base as the rest of the atmospheric code. Hence all future references to the coupling between the atmospheric model and the ocean and sea ice models include the land surface and rivers as part of the atmospheric model.

The ocean submodel (Appendix B) is based on the “OPA” part of the NEMO ocean framework (Madec, 2008). Some of the motivations for using NEMO as the HadGEM3-O ocean submodel have already been described in the introduction; the option to use a tripolar grid eliminates many of the numerical problems associated with more traditional grids and integrating our model development efforts with an international community should be beneficial in maintaining state-of-the-art parameterizations and numerical schemes. As for the atmosphere, the ocean submodel is discretised on a C grid (Arakawa and Lamb, 1977). NEMO has developed a family of tripolar horizontal grids (Madec, 2008) called ORCA. Each configuration in the ORCA family has a regular Mercator grid in the Southern Hemisphere, and a non-regular grid in the Northern Hemisphere, in which the North pole singularity has been replaced by a line between points in Canada and Siberia. Some ORCA configurations, including ORCA1, use latitudinal grid refinement in the tropics. The version of HadGEM3-O described in Sect. 5 employs the nominal 1 degree resolution horizontal grid “ORCA1” (see Fig. 2 and Appendices B and C for further details) and 42 vertical levels.

The sea ice submodel (Appendix D) is based on the Los Alamos sea ice model (CICE) (Hunke and Lipscomb, 2008)

² In HadGEM3-AO aerosol fluxes are not passed across the coupler.




Fig. 2. ORCA1 grid. Every third gridline is plotted. The contoured field is the grid anisotropy, $\Delta x:\Delta y$. The tropical grid refinement is clear, as is the quasi-uniform grid in the Arctic. The land sea mask in the Arctic (and elsewhere) has been smoothed to inhibit the build-up of sea ice in narrow channels.

which is a state-of-the-art sea ice model suitable for use at the resolutions currently employed for global climate modelling. Previous climate models HadGEM1 and HadGEM2 used components of the CICE model embedded within the ocean model. HadGEM3-AO uses CICE itself as a subcomponent of the coupled system. Using CICE maintained the sea ice model science as we moved from HadGEM2 to HadGEM3. Surface sea ice temperature, atmosphere to ice fluxes and the conductive heat flux through the ice are calculated in the atmosphere component (as in HadGEM1, McLaren et al., 2006) while the remaining calculations (ice growth and melt, dynamics and ridging in thickness categories) are carried out by the CICE submodel on the ORCA grid at the same resolution used by the ocean component. The ocean and sea ice subcomponents always need to run on the same grids to ensure that fluxes of heat and freshwater can be accurately maintained between both submodels.

One key difference between NEMO and CICE is that CICE employs a B grid whereas NEMO uses a C grid; while the tracer points of the NEMO and CICE grids align, the velocity points do not. CICE has both components of velocity at the “corners” of the grid cells while NEMO has velocity points staggered on the midpoints of the appropriate edges (Arakawa and Lamb, 1977). This means that single grid-block wide channels allow transport in NEMO, but are impenetrable to flow in CICE. For this reason, channels in the

land-sea mask (further details in Appendix C) are at least two grid cells wide in order to deter the potential build-up of sea ice. We have not experienced any problems associated with the different velocity grids in NEMO and CICE.

The need for efficient load balancing, the fact that NEMO and CICE employ the same scalar grid and the fact that the model was developed on a vector machine employing relatively few CPUs led to the decision that NEMO and CICE, although existing as separate codes, should be run as a single executable in HadGEM3-AO; therefore coupling between NEMO and CICE does not use OASIS. This is relatively easily implemented and allows the ocean and sea ice codes to be maintained separately with the ability to run as standalone models. This aspect of code management could also be applied in the future to other codes such as atmosphere and land surface.

OASIS3 (Valcke, 2006) is a coupler which is designed to coordinate the exchange of fields between models. Details of the OASIS3 configuration are given in Appendix E. At the time of developing the climate model OASIS3 was a non-parallel code (i.e. all coupling was performed through a single process) which, in principle, could have presented a performance bottleneck. In practise this was not the case given the small number of processors involved on the NEC SX8 and the low coupling frequency employed. The parallel OASIS4 (Valcke and Redler, 2006) potentially provides a more scalable option suited to use on supercomputers with many processors but on the timescales of HadGEM3-AO development, OASIS4 was considered unlikely to offer the required functionality and robustness.

The various submodels of the coupled system were compiled using the same programming environment, compiler versions, MPI (Message Passing Interface) version and netCDF libraries (see Appendix F for further details). Figure 3 shows the sequence of events during a coupled model run. The three components start and execute simultaneously with the only implied inter-component communication being the points at which each component receives incoming regridded coupling data. At this point, if the required data is unavailable in OASIS3 then the receiving component is forced to wait until the data is available. It is therefore desirable that the system is well load balanced so that each scientific component passes data to OASIS3 in advance of the receiving component requesting data from the coupler.

3 Fields exchanged between submodels

This section gives a detailed description of the variables exchanged between the UM atmosphere, NEMO and CICE in HadGEM3-AO. First, in Fig. 4, we show how atmosphere fields are passed to NEMO and CICE, and how these are used in combination with sea ice fields to calculate fluxes for NEMO. Then in Fig. 5 we show how NEMO and CICE fields are combined and passed back to the UM atmosphere.

Table 1a. Details of combining the atmosphere fields in Fig. 4. These manipulations are in the UM subroutine *OASIS3_GETA20*

Operations: combining atmosphere fields	
$T1_AT = AO_3 + AO_4$	Total snowfall rate = Large-scale + Convective rates
$T2_AT = AO_5 + AO_6$	Total rainfall rate = Large-scale + Convective rates
$T3_AT = - (LF + LC) \times AO_7$	Latent heat flux into ice = $- (\text{latent heat of fusion} + \text{condensation}) \times \text{sublimation rate}$
$T4_AT = AO_11 + AO_12 - AO_13 - LC \times AO_8$	Total surface heat flux = solar + longwave + sensible + latent heat fluxes
$T8_AT = AO_9 / (1 - \text{land frac in grid box})$	Runoff into ocean = Atmosphere grid-box-mean runoff / fraction of sea in grid box

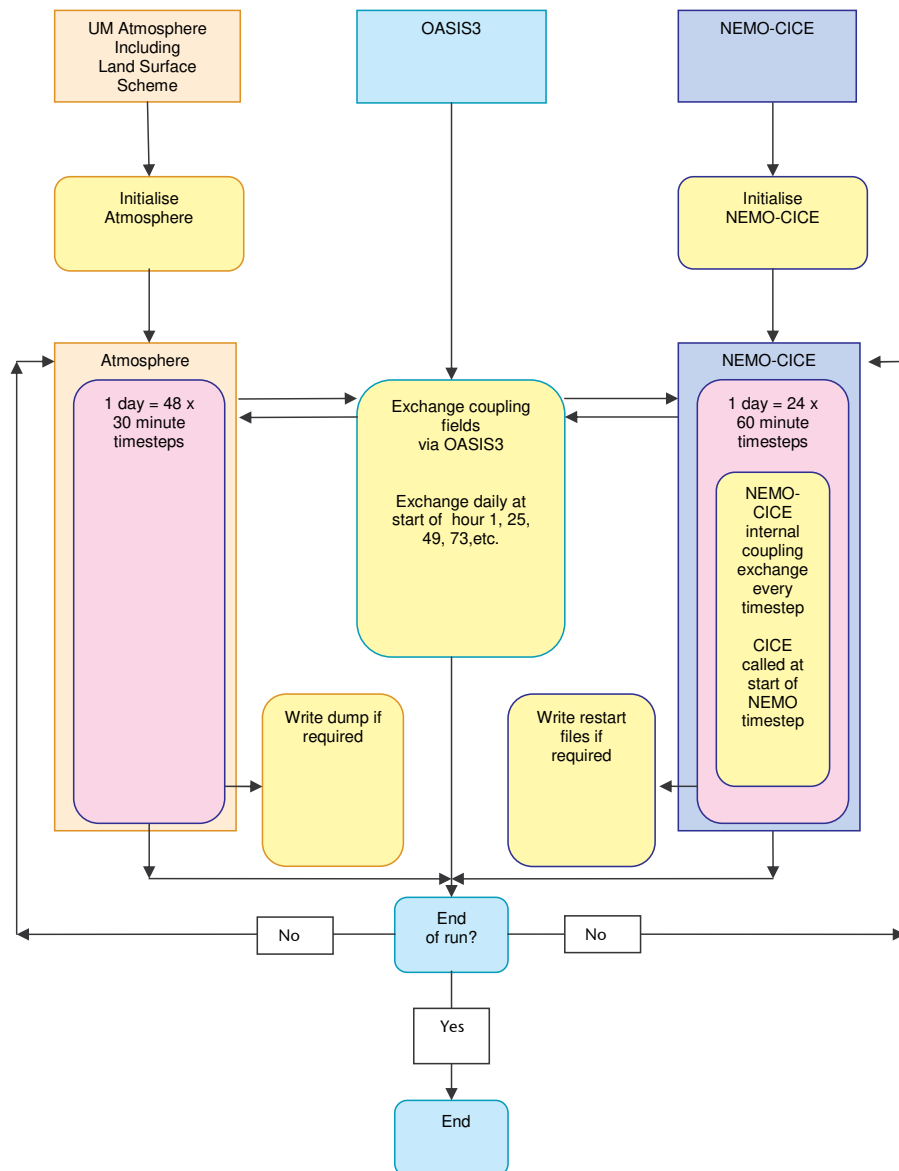


Fig. 3. Schematic representation of the sequence of events in coupled models highlighting the critical positioning of the coupling exchanges and creation of model restart and dump files in relation to each component.

Table 1b. Details of regridding the atmosphere fields onto the ocean grid in Fig. 4. These manipulations are carried out by OASIS itself. “n” is the ice thickness category index.

Regridding: atmosphere to ocean grid	
AO_1 → T5_NU ^a	Transfer x comp of wind stress to ocean U points
AO_2 → T6_NV ^a	Transfer y comp of wind stress to ocean V points
T1_AT → T1_NT	Transfer total snowfall rate to ocean/ice T points
T2_AT → T2_NT	Transfer total rainfall rate to ocean/ice T points
T3_AT → T3_NT	Transfer latent heat flux into ice to ocean/ice T points
AO_8 → T7_NT	Transfer evaporation to ocean T points
T8_AT → NI_7	Transfer runoff into ocean T grid and reverse sign
AO_10 → T9_NT	Transfer solar penetrative radiation to ocean T points
T4_AT → T4_NT	Transfer total surface heat flux to ocean T points
AO_14(n) → CI_6(n)	Transfer category “botmelt” to ocean/ice T points
AO_15(n) → T10_NT(n)	Transfer category “topmelt” to ocean/ice T points

^a Vectors are rotated to NEMO model grid after call to OASIS in NEMO subroutine *SBC_CPLIN*

Table 1c. Detail of regridding the ice fields onto the ocean U and V points in Fig. 4. These manipulations are in the NEMO subroutine *CICE_SBC_OUT*.

Regridding: ice to ocean grid	
CO_11 → T11_NU	Transfer aggregate ice concentration to ocean U points
CO_11 → T11_NV	Transfer aggregate ice concentration to ocean V points
CO_6 → T12_NU	Transfer x-comp. of GBM ocean-ice stress to ocean U points
CO_7 → T13_NV	Transfer y-comp. of GBM ocean-ice stress to ocean V points

In the model described in this paper NEMO and CICE are coupled at every timestep (this is when the operations below the dashed line in Figs. 4 and 5 are performed) and the ocean-ice system is coupled to the UM atmosphere every 24 h (which corresponds to 24 ocean-ice timesteps and 48 atmosphere timesteps). The coupling frequency between the submodels will change in future HadGEM3-AO developments.

In the figures below, input and output fields are referenced as XY_n where X is (A, N, C) for (Atmosphere, NEMO, CICE), Y is (I, O) for (Input, Output) and n indicates the number of the field. Temporary variables used during the coupling are referenced as Tn_XZ where n is the number, X is the model, and Z is (T, U, V) according to the grid points on which these fields are found (see Tables 1g and 2f).

3.1 Atmosphere to Ice to Ocean coupling (Fig. 4)

Atmosphere fields averaged over the coupling period are prepared for passing to OASIS3 through the manipulations detailed in Table 1a (and performed in the *OASIS3_PUTA20* subroutine within the UM code). These are largely straightforward operations which combine fields or convert units prior to the OASIS3 interpolation. Note that the ice heat

fluxes (AO_14 and AO_15) and ice sublimation rate (AO_7) calculated in the atmosphere model are “per unit grid-box area”³ rather than “per unit ice area” (i.e. they are multiplied by the ice concentration in the atmosphere model). This is to ensure energy conservation between submodels as the ice area changes in CICE during the coupling period. Similarly, the evaporation (AO_8), solar, longwave and sensible heat fluxes over the open ocean (AO_10, AO_11, AO_12, AO_13) are multiplied by the leads fraction when diagnosed in the atmosphere code. The runoff field (AO_9) has to be scaled by the fraction of land in the atmosphere grid box (see Table 1a) to account correctly for coastal tiling which allows atmosphere grid boxes to be part land and part ocean.

Once the fields have been regridded (from the atmosphere to the ocean grid) by OASIS3 and vector fields have been rotated within NEMO, the operations in Table 1d and f provide the correct fields for driving the NEMO and CICE models. In both sets of operations ice concentrations from the previous CICE timestep (CO_11) are used to ensure consistency when atmosphere fluxes and stresses are partitioned between ocean and sea ice. The ocean-ice stresses from CICE (CO_6 and CO_7) are grid-box means (i.e. they have been multiplied by ice fraction within CICE) so can be used directly as

³ At points with coastal tiling, the “grid box area” is the non-land area of the grid box.

Table 1d. Details of partitioning the atmospheric stresses and fluxes for the ice model in Fig. 4. These manipulations are in the NEMO subroutine *CICE_SBC_IN*. “n” is the ice thickness category index.

Operation: partitioning of stresses and fluxes for ice model; freeze-melt potential calculation ^a	
$T14_NU = T11_NU \times T5_NU$	<i>x</i> comp: Wind stress used by ice model = ice concentration \times GBM wind stress
$T15_NV = T11_NV \times T6_NV$	<i>y</i> comp: Wind stress used by ice model = ice concentration \times GBM wind stress
$CL_3 = CO_11 \times T1_NT$	Snowfall on ice = ice concentration \times GBM snowfall
$CL_4 = CO_11 \times T2_NT$	Rainfall on ice = ice concentration \times GBM rainfall
$CL_5(n) = T3_NT \times CO_1(n)/CO_11$ ^b	Category latent heat = GBM latent heat \times category ice concentration/aggregate ice concentration
$CL_7(n) = T10_NT(n) + CL_6$	Category surface heat flux = category “topmelt” + conductive flux through ice
$CL_8 = rau0 \times rcp \times e3t(1) \times (Tocnfrz - CL_9)/(2 \times dt)$ ^c	Freeze-melt potential = ocean density \times ocean heat capacity \times ocean surface layer thickness \times (freezing temperature – SST)/ (2 \times CICE timestep)

^a The ice concentration (CO_11) used here and in Table 1f must be the same and is the value at the end of the previous CICE timestep.

^b In the latent heat calculation, if CO_11 = 0 (no ice), CL_5(1)=T3_NT and CL_5(2:5)=0.0. In CICE, this will then be passed to the ocean to ensure energy/water conservation.

^c The freeze-melt potential is applied over a leapfrog step in NEMO, hence the division by 2 \times dt.

Table 1e. Details of regridding the ocean fields onto the ice U/V points in Fig. 4. These manipulations are in the NEMO subroutine *CICE_SBC_IN*.

Regridding: ocean to ice grid	
$T14_NU \rightarrow CL_1$	Transfer <i>x</i> component of wind stress to ice U/V points
$T14_NV \rightarrow CL_2$	Transfer <i>y</i> component of wind stress to ice U/V points

components of NI_1 and NI_2 once they have been shifted from U/V points on the CICE grid to U and V points on the NEMO grid.

The formation of frazil ice in CICE is determined by the difference between the SST and the freezing point temperature using the “freeze-melt potential” (CL_8) as defined in Table 1d. When the freeze-melt potential is positive (i.e. SST is below the freezing point of -1.8 °C) then all this potential is used to form frazil ice in CICE and so the SST is set to freezing by adding this amount of energy to the total heat flux into NEMO (NI_6). However when the freeze-melt potential is negative (i.e. SST is above freezing) it is not added into NI_6; in this case CICE computes an ocean-to-ice heat flux which melts existing ice and is limited by the negative freeze-melt potential value. The CICE computed flux is then added to the net heat flux (CO_10) output from the ice model. The calculation of freeze-melt potential currently takes place immediately before CICE is called and uses the SST from the previous NEMO timestep because it will then be applied over a leapfrog step in NEMO. The timing and details of

the freeze-melt potential calculation would have to be carefully considered if the NEMO-CICE coupling frequency was changed. The subsurface ocean is not checked for temperatures which are below freezing.

The net heat flux from CICE (CO_10) includes the ice-ocean heat flux at the base of the ice, the lateral heat flux at the sides of the ice and also any atmosphere-ice heat fluxes required for conservation when ice has melted during a coupling period with the atmosphere. For example, as noted in Table 1d, the latent heat flux (T3_NT) is partitioned between different ice categories unless the aggregate ice concentration is now zero; in this case all latent heat is put into category 1 and will be returned by CICE to the ocean to ensure conservation.

Two freshwater fluxes are passed into NEMO: a freshwater flux for the free surface calculation (NI_3, emp [evaporation minus precipitation] in NEMO code) and a freshwater equivalent flux for the surface salinity calculation (NI_4, emps). Melting and formation of ice does not affect the integrated column pressure in the ocean and so we take a simplifying approach of disallowing mass exchange between ice and ocean (see e.g. Schmidt et al., 2004). Hence the freshwater flux for the free surface calculation does not include the freshwater flux from the ice model. Snow and rainfall landing on the ice is immediately assumed to have an impact on the free surface (i.e. NI_3 includes all precipitation whether or not it lands in the open ocean). Sublimation/condensation at the sea ice surface (computed from the latent heat flux CL_5) does not currently impact on the free surface calculation which is unrealistic but will be included

Table 1f. Details of combining the atmosphere and ice fields to create surface fields for the ocean in Fig. 4. These manipulations are in the NEMO subroutine *CICE_SBC_OUT*.

Operations: combining atmos and ice fields to create surface fields for ocean	
$NL.1 = (1 - T11_NU) \times T5_NU - T12_NU$	x comp: Surface stress = (1-ice concentration) \times GBM wind stress – GBM ocean-ice stress
$NL.2 = (1 - T11_NV) \times T6_NV + T13_NV$	y comp: Surface stress = (1-ice concentration) \times GBM wind stress – GBM ocean-ice stress
$NL.3 = T7_NT - T1_NT - T2_NT + NL.7$	Freshwater flux for free surface calculation = evaporation – snowfall – rainfall + runoff
$NL.4 = T7_NT - (1 - CO.11) \times T1_NT - (1 - CO.11) \times T2_NT + NL.7 + CO.9/CL.10 - CO.8$	Freshwater equivalent flux for salinity calculation = evaporation – (1-ice concentration) \times snowfall – (1-ice concentration) \times rainfall + runoff + salt flux from ice/surface salinity – freshwater flux from ice
$NL.5 = T9_NT^a$	Total solar penetrative flux = open-ocean flux
$NL.6 = T4_NT - T9_NT + CO.10 - (1 - CO.11) \times T1_NT \times l_{fus} + H(CL.8) \times CL.8$	Total non-penetrative heat flux = heat flux from atmos – open-ocean solar penetrative flux + heat flux from ice model + (1-ice concentration) \times snowfall \times latent heat of fusion ^b + freeze-melt potential ^c

^a Currently solar radiation is only absorbed at the surface of the sea ice in the atmosphere model. Future developments will allow solar radiation to penetrate into the ice, which may then penetrate through to the ocean adding an extra term to the solar penetrative flux for NEMO.

^b This term is heat associated with snowfall melting in ocean.

^c If it is positive.

Table 1g. Table of temporary variables as used in the coupling operations described in Tables 1a, d and f.

Temporary variables	
T1	Total snowfall rate
T2	Total rainfall rate
T3	Latent heat flux into ice
T4	Total surface heat flux
T5	x-component of wind stress
T6	y-component of wind stress
T7	Evaporation
T8	Runoff
T9	Solar penetrative radiation
T10	Category “topmelt”
T11	Aggregate ice concentration (from end of previous timestep)
T12	x-component of ocean-ice stress
T13	y-component of ocean-ice stress
T14	x-component of wind stress used in ice model
T15	y-component of wind stress used in ice model
T16	Freeze-melt potential

in future revisions. The freshwater equivalent flux for salinity includes all the freshwater fluxes from the atmosphere and CICE, together with the CICE salt flux converted to a freshwater equivalent. Rainfall over sea ice is assumed to drain into the leads immediately in CICE and is included in the freshwater flux output from CICE (CO.8).

NEMO requires two surface heat fluxes: the total solar penetrative flux (NL.5, qsr in NEMO code) and a total non-penetrative flux (NL.6, qns). Currently, the total solar penetrative flux is simply the open-ocean field passed from the atmosphere. All solar absorption for sea ice is assumed to occur at the surface in the sea ice thermodynamics in the atmosphere model, hence, there is no solar flux passed through the ice to the ocean. The total non-penetrative flux consists of the remaining open-ocean fluxes from the atmosphere, fluxes from the CICE and the heat required to melt snow falling in the ocean.

3.2 Ocean to Ice to Atmosphere Coupling (Fig. 5)

The manipulations and calculations shown in Fig. 5 are generally more straightforward than those in Fig. 4. Within NEMO the operations are largely confined to transformations between T, U and V points as and when required for NEMO and CICE. Every time coupling to the atmosphere is required, NEMO and CICE surface currents are combined according to ice concentration as detailed in Table 2c. The vector components are then rotated from the model grid to the geographic grid. All NEMO-CICE fields passed to OASIS3 for regridding are instantaneous quantities at the coupling timestep and have not been meaned over the coupling period.

Once coupling fields have been received by the atmosphere *OASIS3_GETO2A* subroutine there are various

Table 1h. Table of constants as used in model components and coupling code relating to Fig. 4.

Constant	Value	NEMO (location)	CICE (location)	UM (location)
Latent heat of condensation	$2.501 \times 10^6 \text{ J kg}^{-1}$			LC (<i>c_lheat.h</i>)
Latent heat of fusion	$0.334 \times 10^6 \text{ J kg}^{-1}$	lfus (<i>physct</i>)		LF (<i>c_lheat.h</i>)
Ocean density	1026 kg m^{-3}	rau0 (<i>physct</i>)		
Ocean heat capacity	$4000 \text{ J kg}^{-1} \text{ K}^{-1}$	rcp (<i>physct</i>)		
Surface ocean layer depth	10 m	e3t(1) (<i>dom_oce</i>)		
Freezing temperature	$-1.8 \text{ }^\circ\text{C}$		Tocnfrz (<i>ice_constants</i>)	
CICE timestep	1 h		dt (<i>ice_calendar</i>)	
Heaviside function $H(x)$	$1 (x > 0), 0 (x < 0)$			

Table 2a. Details of regridding the ocean fields onto the ice U/V points in Fig. 5. These manipulations are in the NEMO subroutine *CICE_SBC_IN*.

Regridding: ocean to ice grid	
NO.1 → CI.11	Transfer x component of ocean surface current to ice U/V points
NO.2 → CI.12	Transfer y component of ocean surface current to ice U/V points

transformations required as shown in Table 2e. The most significant of these are the imposition of a minimum ice category concentration (2×10^{-4} is used in HadGEM3-AO r1.1), and also a minimum local ice thickness (currently 0.01 m) which is applied so as to retain the same ice volume. These minimum values are necessary for stability in the atmosphere model.

4 Conservative interpolation of fluxes between the atmosphere and ocean or sea ice

4.1 Theory

Conservation of fluxes as they are transferred between sub-models on different grids (e.g. a regular grid in the atmosphere and ORCA grid for the ocean and sea ice) has been a primary concern in the development of HadGEM3-AO. Heat and freshwater flux conservation are necessary for the long-term stability and integrity of coupled climate integrations.

HadGEM3-AO performs the coupling with OASIS3, whose interpolation element is the SCRIP interpolation code (Jones, 1999). SCRIP provides the means to interpolate fluxes in various ways (bilinear, first- and second-order conservative, etc) between grids of general polygons on the sphere. Under first-order conservative interpolation, which is used to remap all scalar fields from atmosphere to ocean in HadGEM3-AO, fluxes $\{f_j\}$ on a source grid g are mapped to fluxes $\{F_J\}$ on a target grid G , according to the linear remapping (Jones, 1999):

$$F_J = \sum_j W_{Jj} f_j, \quad (1)$$

where the matrix element W_{Jj} is calculated according to the “FRACAREA” normalisation option in SCRIP. The weights W_{Jj} only depends on the geometry and masking of the source and target grids. The non-zero weights between given source and target grids are calculated and stored (as netCDF weights files) by OASIS3 (and SCRIP). Interpolation then amounts to the simple matrix multiplication in Eq. (1). A standalone version of SCRIP allows remapping matrices to be generated and tested offline, before being used directly in OASIS3.

A second-order conservative interpolation scheme is available in SCRIP. This assumes a piecewise linear variation of source flux across each source cell. (First-order interpolation assumes piecewise constant source fluxes.) Unfortunately, such schemes can generate undershoots and overshoots near sharp gradients in the source fields. This makes them unsuitable for interpolating fields that must be positive everywhere, such as short wave radiation and (especially) river runoff. We therefore use first order conservative interpolation, which maintains positivity, when mapping fields from the atmosphere to the ocean. For ocean to atmosphere interpolation, we use second order conservative interpolation but trim the sea ice fraction to be between 0 and 1. The resulting weak non-conservation in this direction is not found to cause problems in practice.

The remapping of vector fields (wind stress and ocean/ice surface current) between regular and tripolar grids is handled by bilinear interpolation, as was the case for earlier Hadley Centre coupled models (Johns et al., 2006). Conservation of momentum across the coupler has not been addressed – not least because of the significant technical difficulties of interpolating vector fields between arbitrary grids. We do

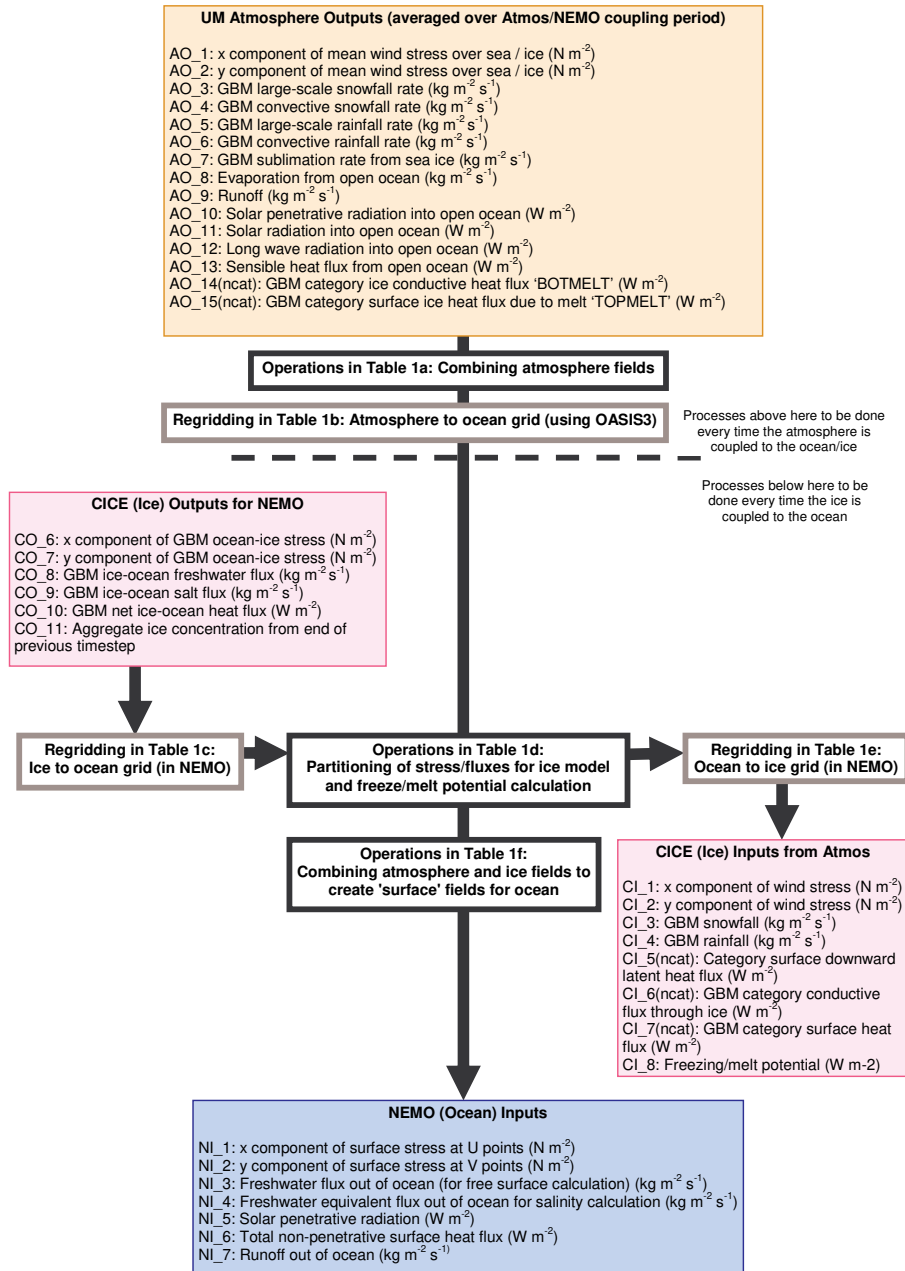


Fig. 4. Flow diagram showing the coupling of the 3 submodels in HadGEM3 in the direction atmosphere to sea ice to ocean. Tables 1a–f detail the operations and regridding shown in the figure. Temporary fields are listed in Table 1g and constants used in the coupling are shown in Table 1h. All fields are two dimensional horizontal fields unless marked “(ncat)”, in which case they have a third dimension of the number of ice thickness categories (currently 5). GBM stands for Grid-Box-Mean. Input and output fields are referenced as XY_n where X is (A, N, C) for (Atmosphere, NEMO, CICE), Y is (I,O) for (Input, Output) and n indicates the number of the field. Temporary variables used during the coupling are referenced as Tn_{XZ} where n is the number, X is the model, and Z is (T, U, V) according to the grid points on which these fields are found (see Table 1g).

not however believe that there are any serious issues arising from the non-conservation of momentum (since this was also used in HadGEM1 and HadGEM2).

As a result of these and other considerations, the remapping options in Table 3 were chosen in HadGEM3-AO.

The NEMO and CICE models are on the same scalar grid in HadGEM3-AO, but their equations are differenced on the Arakawa C grid and B grid respectively. The components of vector fields therefore appear at different locations, even though (by design) scalar fields coincide. In HadGEM3-AO,

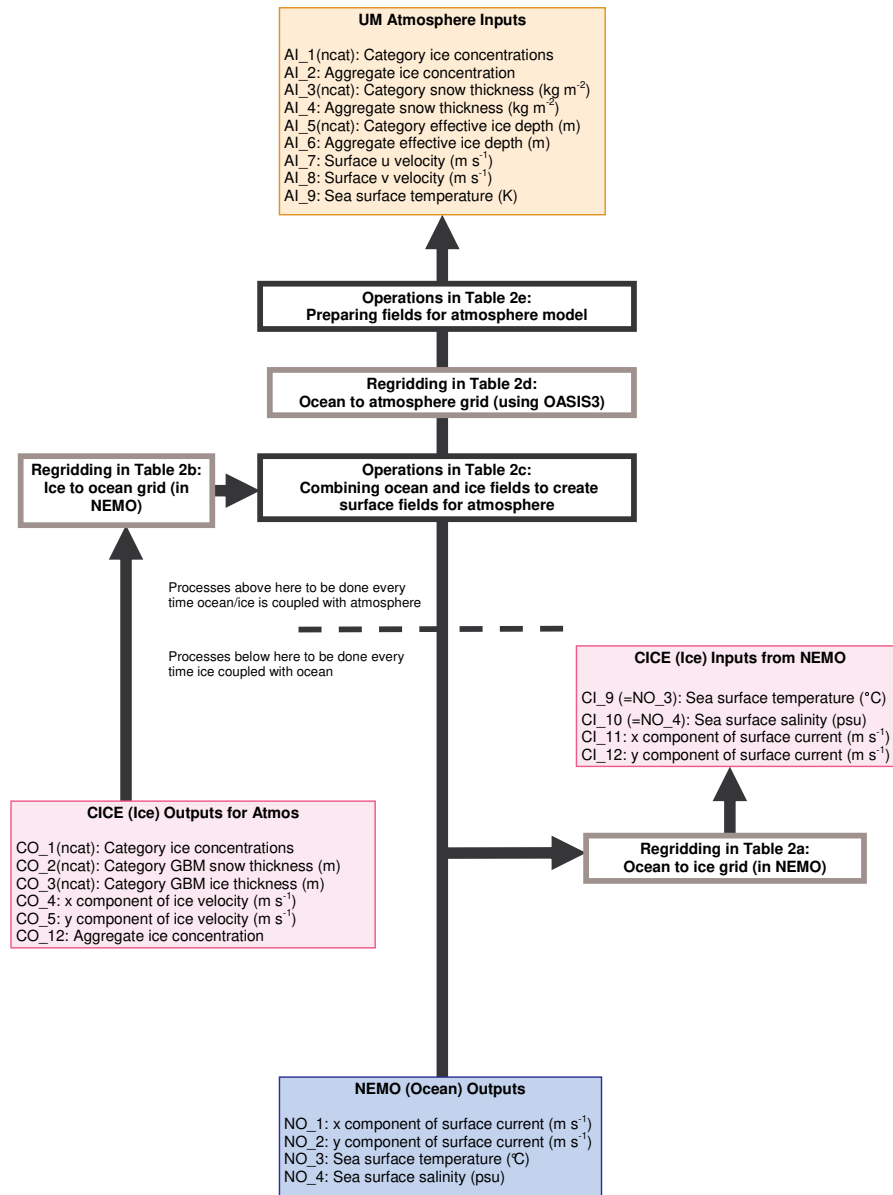


Fig. 5. Flow diagram showing the coupling of the 3 submodels in HadGEM3 in the direction ocean to sea ice to atmosphere. Tables 2a–e detail the operations and regridding shown in the figure. Temporary fields are listed in Table 2f and constants used in the coupling are shown in Table 2g. Details of annotation are given in the caption for Fig. 4.

Table 2b. Details of regridding the ice fields onto the ocean U and V points in Fig. 5. These manipulations are in the NEMO subroutine *CICE_SBC_OUT*.

Regridding: ice to ocean grid	
CO_12 → T17_NU	Transfer aggregate ice concentration to ocean U points
CO_12 → T17_NV	Transfer aggregate ice concentration to ocean V points
CO_4 → T18_NU	Transfer x-component of ice velocity to ocean U points
CO_5 → T19_NV	Transfer y-component of ice velocity to ocean V points

Table 2c. Details of combining the ocean and ice fields to create surface fields for the atmosphere in Fig. 5. These manipulations are in the NEMO subroutine *CICE_SBC_HADGAM*.

Operations: combining ocean and ice fields to create surface fields for atmosphere	
$T20_NU = (1 - T17_NU) \times NO_1$ $+ T17_NU \times T18_NU$	x-comp: Surface velocity = $(1 - \text{ice concentration}) \times \text{ocean velocity}$ $+ \text{ice concentration} \times \text{ice velocity}$
$T21_NV = (1 - T17_NV) \times NO_2$ $+ T17_NV \times T19_NV$	y-comp: Surface velocity = $(1 - \text{ice concentration}) \times \text{ocean velocity}$ $+ \text{ice concentration} \times \text{ice velocity}$

Table 2d. Details of regridding the ocean/ice fields onto the atmosphere grid in Fig. 5. These manipulations are carried out by OASIS itself. “n” is the ice thickness category index.

Regridding: ocean to atmosphere grid	
$NO_3 \rightarrow AI9^a$	Transfer sea surface temperature to atmosphere T points
$CO_1(n) \rightarrow AI_1(n)$	Transfer category ice concentrations to atmosphere T points
$CO_2(n) \rightarrow T22_AT(n)$	Transfer category GBM snow thickness to atmosphere T points
$CO_3(n) \rightarrow T23_AT(n)$	Transfer category GBM ice thickness to atmosphere T points
$T20_NU \rightarrow AI_7^b$	Transfer x-component of surface velocity to atmosphere T points
$T21_NV \rightarrow AI_8^b$	Transfer y-component of surface velocity to atmosphere T points

^a SST is converted to Kelvin in *SBC_CPL_OUT* immediately before call to send field to OASIS

^b Vectors are rotated to the geographical grid before call to OASIS in NEMO subroutine *SBC_CPL_OUT*

vectors passed between NEMO and CICE are regridded in NEMO coupling routines, by simply averaging the values at the two nearest neighbours on the other grid.

The OASIS3 coupler requires separate netCDF format files containing grid and land-sea mask definitions. The grid file contains cell latitudes and longitudes pertaining to each type of grid involved in coupling. That is, separate sets of coordinates are required for each of the T, U and V points on both the atmosphere and ocean grids. In addition, sets of four grid-cell corner coordinates are required for each grid cell on each grid type. Similarly the land-sea mask file requires a separate land-sea mask for each grid type employed – six in total. These required netCDF files may be generated at run time via OASIS3 library calls embedded in the model code. However, in practice, performance overheads and difficulty with sequencing of polar row cell corners meant that it was preferable to generate these files offline and link them to the model at run time. This approach reduced the run time overhead and gave the most accurate and reliable grid definitions and land-sea masks.

It is possible to allow OASIS3 to generate the remapping weights files at run time by using the grid and mask information. In practice this turned out to be very time-consuming on the NEC with each weights file taking approximately 12 min to be generated. With six such files needed for a coupled run, it was clearly unacceptable to spend over an hour in generating the weights before being able to run the model. Therefore, following the same approach used for the grid files, the weights files were generated offline with SCRIP and linked

to the model at run-time. This device also allowed us to test the weights files in isolation, and make manual corrections where necessary.

4.2 Results

In general, the budgets of fluxes calculated with the FRACAREA normalisation of Eq. (1) would be computed according to

$$\sum_J \Omega_J M_J F_J A_J = \sum_j \omega_j m_j f_j a_j, \quad (2)$$

where ω_j (and Ω_J) are the unmasked fractions of grid G on grid g (and grid g on grid G respectively).

In HadGEM3-AO, however, as in earlier Met Office coupled models (Johns et al., 2006), the ocean model land-sea mask *defines* the atmosphere model land-sea mask, in such a way that valid ocean model grid cells are always fully overlain by sea points of the atmosphere model. The ocean grid fraction is therefore always unity, while the atmospheric equivalent is the usual sea fraction α . Budgeting of fluxes therefore amounts to multiplying atmospheric fields by α before integrating. This is the case whether the mapping is from atmosphere to ocean or ocean to atmosphere.

The results of interpolating conservatively like this are exemplified by Fig. 6a which is a time series of various terms in the global heat budget of the HadGEM3-AO r1.1 run. Shown in red is the heat flux from the atmosphere and sea ice to the ocean; shown in blue is the heat flux as seen by the ocean model. (As explained above, atmosphere fields

Table 2e. Details of preparing the fields for the atmosphere model in Fig. 5. These manipulations are in the UM subroutine *OASIS3_GETA20*. “n” is the ice thickness category index.

Operations: preparing fields for atmosphere model ^a	
$AI_1(n) = \min(AI_1(n), T23_AT(n)/0.01)$	If local ice thickness is less than 1 cm, reshape ice so local thickness is 1 cm without changing volume
If $AI_1(n) < 2.0 \times 10^{-4}$, then $AI_1(n) = T23_AT(n) = T22_AT(n) = 0.0$	Impose minimum ice concentration of 2×10^{-4}
$AI_2 = \sum_n AI_1(n)$	Aggregate ice concentration = sum of all category concentrations
$T24_AT(n) = T23_AT(n) + T22_AT(n) \times (kappai/kappas)$	GBM effective ice thickness = GBM ice thickness + GBM snow thickness \times (ice conductivity/snow conductivity)
$T25_AT = \sum_n T24_AT(n)$	Aggregate GBM effective ice thickness = sum of all category values
$AI_5(n) = T24_AT(n)/AI_1(n)$	Local effective ice thickness = GBM effective ice thickness/ice concentration
$AI_6 = T25_AT/AI_2$	Aggregate local effective ice thickness = Aggregate GBM effective ice thickness/aggregate ice concentration
$T26_AT = \sum_n T22_AT(n)$	Aggregate GBM snow thickness = Sum of all category values
$AI_3(n) = T22_AT(n) \times rhosnow/AI_1(n)$	Local snow thickness (in kg m^{-2}) = GBM snow thickness (in m) \times snow density/ice concentration
$AI_4 = T26_AT \times rhosnow/AI_2$	Aggregate local snow thickness (in kg m^{-2}) = Aggregate GBM snow thickness (in m) \times snow density/aggregate ice concentration

^a TSTAR (global surface temperature field) and TSTAR.SICE (area weighted average surface temperature for ocean and sea ice) in the UM are also altered in this subroutine to take account of the latest SST.

Table 2f. Table of temporary variables as used in the coupling operations described in Tables 2c and e.

Temporary variables	
T17	Aggregate ice concentration
T18	x-component of ice velocity
T19	y-component of ice velocity
T20	x-component of surface velocity
T21	y-component of surface velocity
T22	Category GBM snow thickness
T23	Category GBM ice thickness
T24	Category GBM effective ice thickness
T25	Aggregate GBM effective ice thickness
T26	Aggregate GBM snow thickness

need to be multiplied by the sea fraction α when reckoning flux budgets.) If the heat flux coupling were truly conservative, these would agree closely, and they clearly do: to within $2 \times 10^{-3} \text{ W m}^{-2}$ on average. Note that a steady heat flux imbalance of this magnitude would cause a mean ocean temperature drift of about $6 \times 10^{-4} \text{ K}$ per century, which is smaller than we would expect to see in any medium range climate

Table 2g. Table of constants as used in model components and coupling code relating to Fig. 5.

Constant	Value	UM name (location)
Ice thermal conductivity	$2.09 \text{ W m}^{-1} \text{ K}^{-1}$	kappai (<i>c.kappai.h</i>)
Snow thermal conductivity	$0.31 \text{ W m}^{-1} \text{ K}^{-1}$	kappas (<i>c.kappai.h</i>)
Snow density	330 kg m^{-3}	rhosnow (<i>c.kappai.h</i>)

change simulation. Figure 6b is the analogous plot for the freshwater flux in the same HadGEM3-AO r1.1 run. Again, freshwater is conserved quite closely across the coupler: the global, annual mean freshwater fluxes on the atmosphere (red) and ocean (blue) sides of the coupler agree to within 10^{-3} Sv on average. A steady freshwater flux of this magnitude would equate to a mean sea level rise of $9 \times 10^{-3} \text{ m}$ per century, which again is smaller than we would expect to see in any medium range climate change simulation. These results demonstrate that the heat and freshwater budgets are closed to a high degree of accuracy with a small “coupler error” due to the numerics of regridding.

Also shown are estimates of the rate of change of the ocean heat content, as derived from the ocean temperature

Table 3. HadGEM3 interpolation method choices.

Remapping		Method
Source	Destination	
Atmosphere T points	NEMO T points	First order conservative
Atmosphere U, V points	NEMO U, V points	Bilinear
NEMO T points	Atmosphere T points	Second order conservative
NEMO U, V points	Atmosphere U, V points	Bilinear

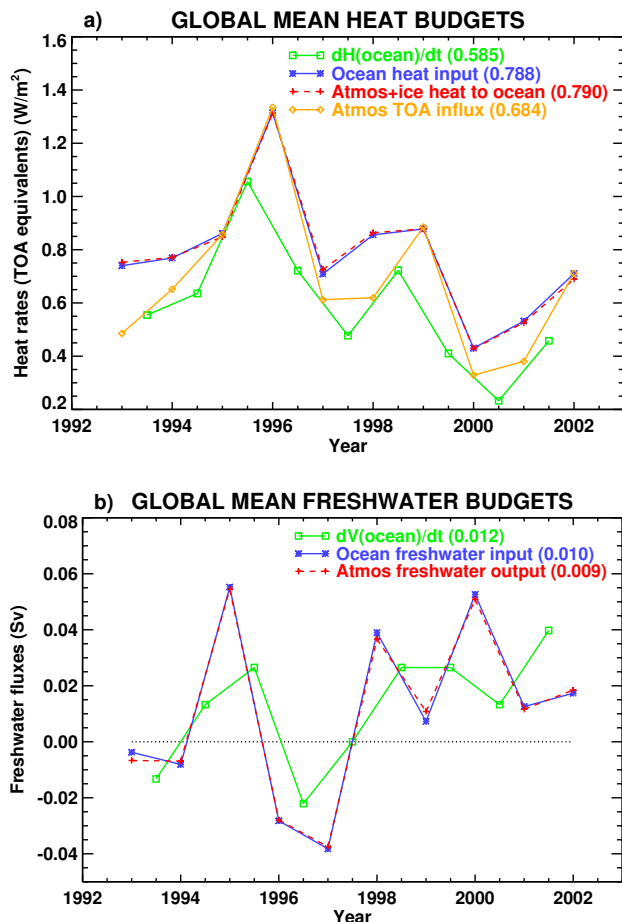


Fig. 6. Time series of global mean flux budgets. Overall time averages in parentheses. (a) Heat into ocean compared to heat leaving atmosphere and sea ice. Also shown are net top of atmosphere fluxes, and estimated rate of change of ocean heat content (from annual mean fields). (b) Freshwater into ocean and out of atmosphere, and estimated rate of change of ocean volume (from annual mean fields).

(Fig. 6a), and an estimate of the rate of change of the ocean freshwater content, as derived from the sea surface height (Fig. 6b). The difference between the variation of integrated heat/freshwater contents (green) and that of the driving fluxes (blue) represents an “ocean conservation error”. The ocean

conservation error is estimated as approximately 0.2 W m^{-2} (when calculated as the equivalent top of atmosphere (TOA) flux) for heat and 0.002 Sv for freshwater. The heat flux error is an order of magnitude larger than the “coupler error” discussed in the previous paragraph. The freshwater error and some of the heat flux error is due to the fluxes and rate of change being evaluated at different times of the year using annual means (for reasons of practical availability)⁴. However, the systematic heat flux bias is probably largely due to anomalous advective fluxes associated with the linear free surface. The error we see here is consistent with short tests using HadGEM3-AO and work in the HiGEM model (I. Stevens, personal communication, 2007). On-going work to investigate the bias associated with the linear free surface may motivate the use of a non-linear free surface in future versions of HadGEM3-AO.

5 Model performance

In this section we discuss both the technical and the scientific performance of HadGEM3-AO. Technically, the performance is highly dependent on the exact supercomputer being used and this is discussed below. Scientifically, it is possible to demonstrate that, based on a 30 year run, the results of HadGEM3-AO r1.1 are credible.

The particular integration of HadGEM3-AO investigated in this section is forced by greenhouse gas concentrations, ozone concentrations and aerosol emissions representative of the 1980s. All other atmospheric variables are initialised from an analysis of the real atmosphere on 1 September 2008 (as used by the Met Office weather forecasting centre). The ocean is initialised at rest, but with temperature and salinity fields interpolated to 1 September from a monthly climatology produced using years 2004–2008 of the EN3 analysis (Ingleby and Huddleston, 2007). The sea ice initial conditions are September mean fields of ice concentration and

⁴The fluxes and the heat contents are annual means (1 Jan – 30 Dec) and are 2nd order accurate estimates of the fluxes at 1 Jul. The dH/dt calculated from these is therefore a (2nd order) estimate of the value at 1 Jan. However, the model has not yet at equilibrium, so there is a 1st order time truncation difference between them.

thickness from a HadGEM1 run averaged over the period 1985–2005. Obviously there is a mismatch between the prescribed radiative forcings and the initial conditions used by the model. Unfortunately, while less than ideal, such inconsistencies are common place in climate modelling. For example, “pre-industrial control experiments” are often initialised with present day conditions but then forced with pre-industrial radiative forcing.

The r1.1 coupled model described in this section represents a very early version of the HadGEM3-AO prototype and, in particular, has not been “tuned” (that is to say, there has been little exploration of parameter space in the various model components). Here we briefly examine some aspects of the atmosphere, ocean and sea ice performance in this version and where appropriate we have compared with the results from a decadal climatology of HadGEM1 with comparable radiative forcing to HadGEM3 r1.1.

5.1 Technical performance

A key issue when running a climate model is to be able to run the model as fast as possible while maintaining efficient use of the available computing resources. In the case of running the HadGEM3-AO configuration on the SX8 supercomputer, load imbalance between atmosphere and the NEMO-CICE component led to a loss of performance. Timing data from running in a 7-1-1 configuration (i.e. seven processors devoted to the atmosphere model, one processor to the OASIS3 coupler and one to the NEMO-CICE model) showed that typically the ocean model timestep was completed well in advance of the atmosphere. The actual work performed by OASIS3 was significantly less than that performed by the component models. The consequence of this was that, for the majority of time, the coupler was simply idling rather than performing useful work. In principle, it may be possible to run the coupler and NEMO-CICE component on the same processor, thus freeing up an extra processor for use by the atmosphere. However, fundamental rearrangement of the code and control systems would be required to effectively allow OASIS3 and NEMO-CICE to run sequentially while allowing the atmosphere to continue concurrently. Such a code rearrangement would be highly architecture-dependent and thus far has not been explored.

5.2 Atmospheric model performance

The precipitation climatology from HadGEM3-AO r1.1, for both December to February (DJF) and June to August (JJA), is shown in Fig. 7a and b. Comparing these plots with GPCP (Adler et al., 2003) observations (Fig. 7c and d) shows that the model generates fairly realistic precipitation patterns, with an Inter-Tropical Convergence Zone (ITCZ) stretching across the Pacific and Atlantic Oceans in the same locations as that observed and local precipitation maxima at mid-

latitudes located off the eastern coast of the USA and to the east of Japan.

In Fig. 8a and b the precipitation biases in the model are compared with observations. In the Pacific the model has excessive rainfall in the ITCZ and SPCZ (South Pacific Convergence Zone) regions to the north and south of the equator and not enough rainfall on the equator. Rainfall on the equator is closely coupled to sea surface temperatures (SSTs) on the equator in what is known as the equatorial Pacific cold tongue region. SSTs in this region are too cold in the model (see Sect. 5.3 – ocean model performance), suppressing precipitation on the equator and allowing more rain to fall to the north and south.

Excessive precipitation is also seen in the equatorial Indian Ocean with an associated dry bias over India in JJA. Precipitation in these regions is highly sensitive to the model convection scheme. Inroads have been made improving Indian JJA precipitation by delaying convection so that less rain falls over the Indian Ocean and more rain falls downstream as the air passes over India. This work is ongoing but some positive improvements can already be seen comparing the HadGEM3-AO bias with that of HadGEM1 (Fig. 8c and d) particularly in DJF.

Figures 9a and b show the 850 hPa horizontal wind climatology from HadGEM3-AO r1.1. Comparing these plots with ERA40 reanalysis (Fig. 8c and d) shows that the main features of the circulation are reproduced including:

- The midlatitude westerly jets at 45°S in the North Pacific and the North Atlantic;
- Low-level easterlies crossing the equatorial Pacific and Atlantic Oceans;
- The Indian JJA monsoon jet crossing the western equatorial Indian Ocean and turning eastwards to head over India.

Figure 10a and b show the 850 hPa horizontal wind biases when compared to ERA40 reanalysis. These wind biases are relatively small (when compared to precipitation biases) and the only bias of concern is in the equatorial Pacific where the easterly winds are too strong, particularly in DJF. This is partly responsible for the excessively cold tongue seen in the sea surface. However, compared with the biases seen in HadGEM1 (Fig. 10c) the excessive easterlies in the western Pacific in DJF are reduced in HadGEM3 r1.1. In addition, HadGEM1 has large biases in the Southern Ocean in JJA (Fig. 10d) which are much improved in HadGEM3 r1.1 (Fig. 10b).

Oceanic surface wind stresses relative to the Scatterometer Climatology of Ocean Winds (SCOW) show similar biases to the 850 hPa wind biases although the biases on the equator have reduced and the biases in the Southern Ocean have increased a little.

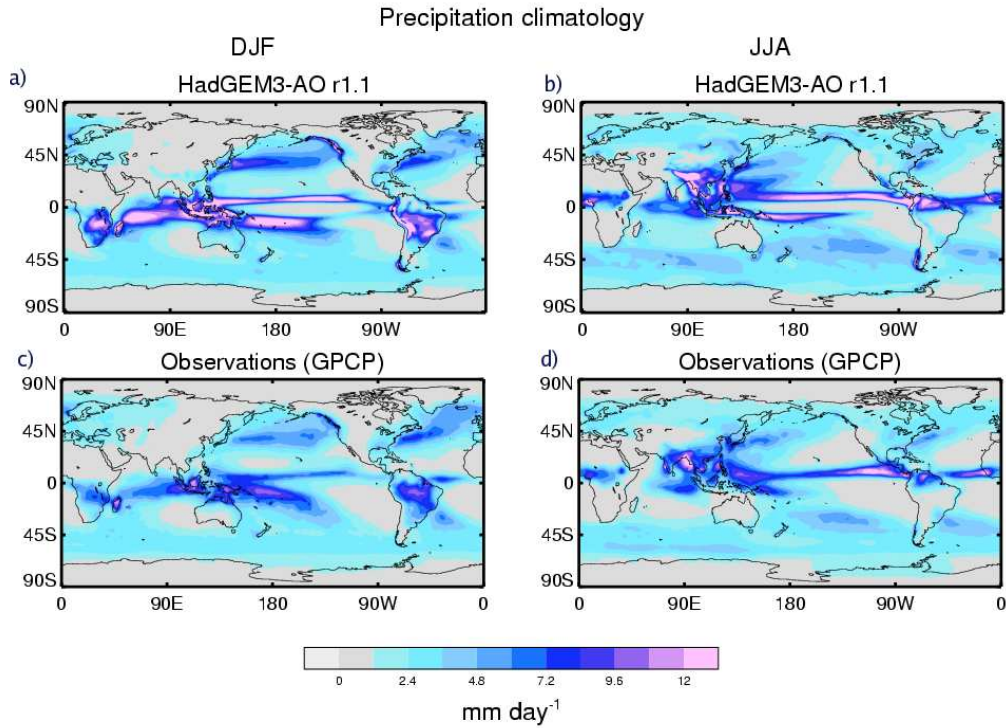


Fig. 7. Precipitation climatology (in mm day⁻¹) of a thirty year period of HadGEM3 r1.1 (plots **a–b**) and observations from the Global Precipitation Climatology Project (GPCP, Adler et al., 2003) covering the period 1979 to 1998 (plots **c–d**) for the seasons December to February (plots **a** and **c**) and June to August (plots **b** and **d**).




Fig. 8. Biases in precipitation (relative to GPCP climatology shown in Fig. 7) for HadGEM1 (**a–b**) and HadGEM3 r1.1 (**c–d**) for the seasons December to February (plots **a** and **c**) and June to August (plots **b** and **d**).




Fig. 9. Horizontal wind climatology (in m s^{-1}) of HadGEM3 r1.1 (plots **a–b**), reanalysis from the ECMWF 40 year reanalysis (ERA40, Uppala et al., 2005) covering the period 1979 to 1998 (plots **c–d**) for the seasons December to February (plots **a** and **c**) and June to August (plots **b** and **d**). The colours show windspeed and arrows show wind direction.

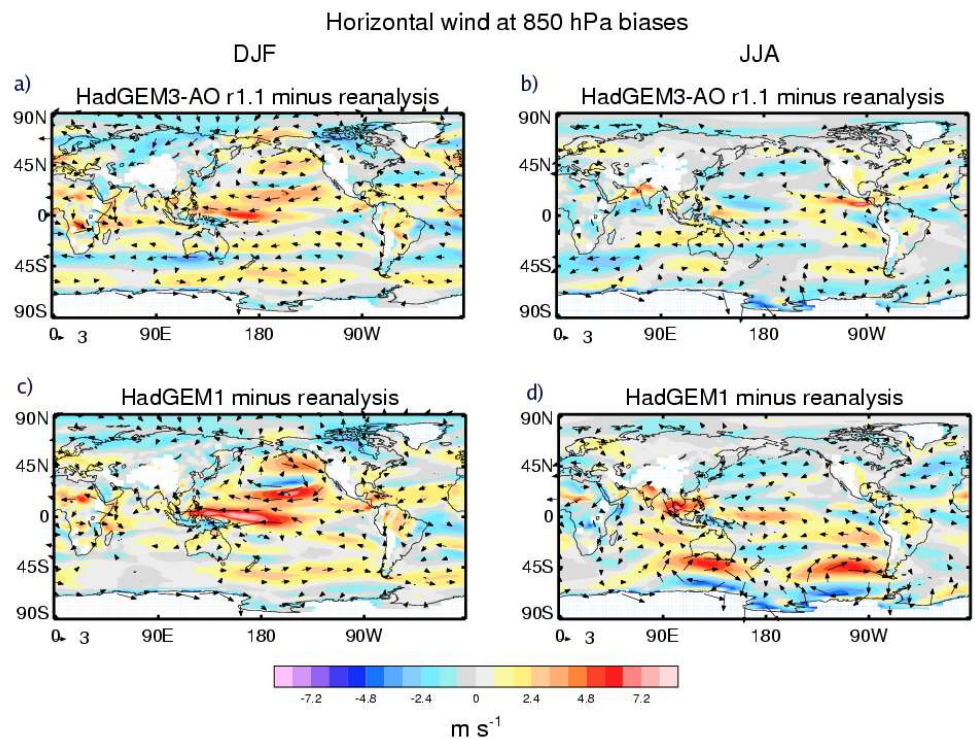


Fig. 10. Biases in horizontal wind climatology (relative to ERA40 climatology shown in Fig. 9) for HadGEM1 (**a–b**) and HadGEM3 r1.1 (**c–d**) for the seasons December to February (plots **a** and **c**) and June to August (plots **b** and **d**).

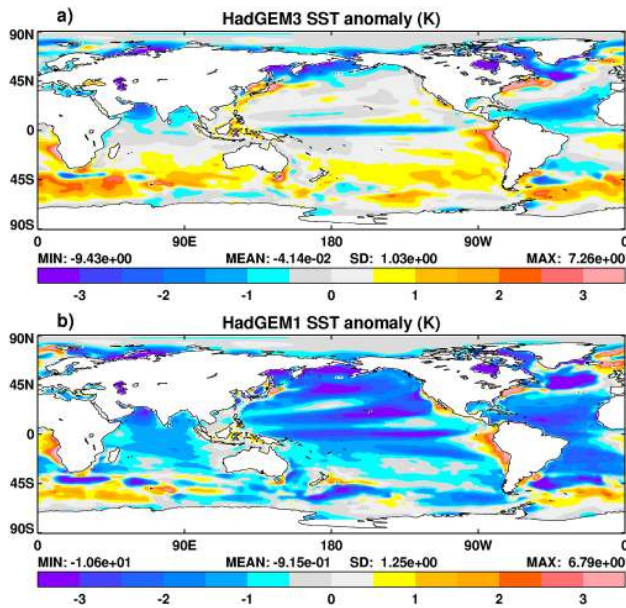


Fig. 11. Sea surface temperature (K) anomalies with respect to initial conditions for (a) HadGEM1 and (b) HadGEM3 r1.1

5.3 Ocean model performance

Figure 11a compares 10-year means (for the third decade of the integration) of sea surface temperature (SST) against the EN3 climatology used to initialise the model. Many of the SST biases are a common feature of coupled climate models and are areas of on-going work. Notable are the cold bias in the equatorial Pacific (discussed in Sect. 5.2), the biases in the North Atlantic (which are found to be significantly improved in higher resolution versions of HadGEM3-AO where the path of the North Atlantic current is much better captured) and the large warm biases in the upwelling regions off the west coasts of Africa and South America – again these are found to be improved by higher ocean resolution. The widespread warm SST bias in the Southern Ocean is seen in many climate models. Deficiencies in the modelling of clouds in this region are thought to play a role, but the ocean vertical mixing is also an important factor and coupled models using the Turbulent Kinetic Energy scheme in NEMO appear to be particularly prone to this problem at present (more details in Appendix B). The SST biases in HadGEM3 r1.1 represent an improvement over HadGEM1 (Fig. 11b) where the global mean SST bias was -0.9 K compared with HadGEM3 r1.1 where the global mean SST bias is -0.04 K.

From Fig. 12 it can be seen that freshwater and heat transport in the model (both globally and within individual basins such as the Atlantic) are consistent with the observational uncertainties for these quantities. For the northward heat transport shown in Fig. 12a and b the difference between the implied values (calculated from the surface heat fluxes assuming zero heat flow at the northern boundary) indicates where

the depth-integrated ocean heat content is changing, for example in the Southern Ocean (as can be seen in Fig. 13).

5.4 Sea ice model performance

In the Arctic, the modelled sea ice extent is greater than values calculated from the HadISST observational data set (Rayner et al., 2003) at all times of year (Fig. 14a). In the winter this is mainly due to too much ice in the North Pacific, which is consistent with cold SSTs in this region. In the summer too much ice remains in the Canadian Archipelago. The spatial distribution of ice thickness is promising – with the thickest ice banked up against the Canadian Archipelago and north of Greenland as observed – although overall the ice is rather too thin (not shown). In the Antarctic the ice extent tends to be lower than observed (Fig. 14b). The winter extent is close to observations but in summer there is too little ice and the ice is thin at all times of year. This may be a consequence of the warm SSTs in the Southern Ocean. Overall the simulation of the sea ice is good for a prototype model run, particularly in the Arctic. The simulation of Antarctic sea ice should improve if the SST bias in the Southern Ocean is reduced as further investigation into the SST bias has found it to be related to a bias in the atmospheric forcing and the vertical mixing scheme used in NEMO and not directly linked to the low sea ice extent. The importance of the SST is also apparent when we compare with HadGEM1 sea ice extents (Fig. 14a and b); while the seasonal cycle is similar in the Arctic, in the Antarctic sea ice extents are generally larger in HadGEM1 which is consistent with the colder SST bias in the Southern Ocean in HadGEM1.

5.5 Suitability for applications

Overall the HadGEM3-AO r1.1 model performs well, and in many respects matches the performance of earlier Hadley Centre climate models such as HadGEM1. Indeed, this version of the model performed sufficiently well to be used in GloSea4 (Arribas et al., 2011), the Met Office seasonal forecasting system. While model biases remain undesirable, some of the model errors may be less serious in GloSea4 because the seasonal forecast includes a bias correction method. GloSea4 updates the model version at regular intervals which means that as improved versions of HadGEM3-AO (i.e. different revisions) are made available they will be used in the seasonal forecast. A key aim of our model development program is continually to improve the predictive ability of the HadGEM3-AO model family on seasonal to decadal timescales, and regional spatial scales. The use of revision numbers enables us to use the model for scientific studies but also to continue to develop it. The model is not currently suitable for submission to IPCC reports because the parameters have not been optimised to allow long term climate stability in a control run – a necessity before forced experiments can be performed.

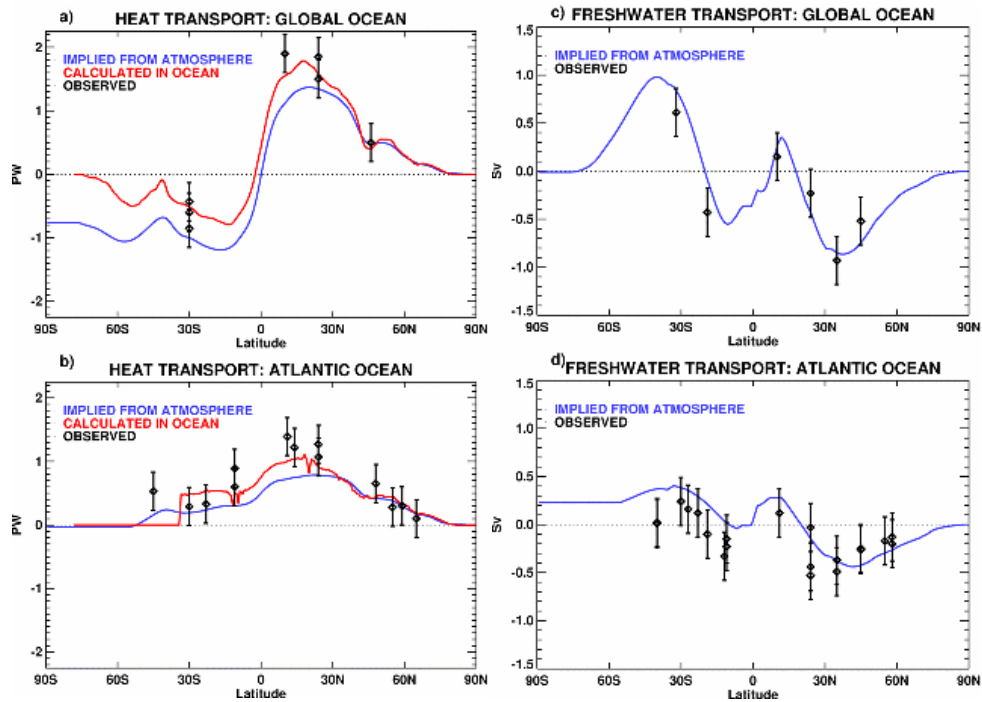


Fig. 12. Integrated northward heat and freshwater ocean transports. (a) Global heat transport, inferred from atmospheric heat fluxes, and calculated within NEMO, compared to observations. (b) Atlantic heat transport, inferred from atmospheric heat fluxes, and calculated within NEMO (including the diffusive component), compared to observations. (c) Global freshwater transport, inferred from atmospheric heat fluxes and compared to observations. (d) Atlantic freshwater transport, inferred from atmospheric heat fluxes and compared to observations.

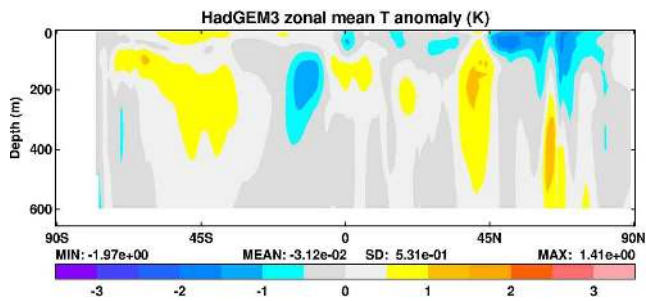


Fig. 13. Zonal mean temperature (K) anomalies with respect to initial conditions for HadGEM3 r1.1.

6 Summary and discussion

This paper has focussed on a detailed explanation of the coupling configuration developed for the HadGEM3-AO climate model and given this focus, scientific details about the component models are largely set out in appendices. This paper lays out the fields from all component models (and associated manipulations) which need to be exchanged for conservative coupling and discusses the conservative transformation of fields between different grids in the context of the OASIS coupler. The intention is that the documentation provided here should be a useful tool for model development and intercomparison exercises.

A key aspect to the coupling of various component models is that heat and freshwater should be conserved throughout the climate system. We are able to demonstrate that HadGEM3-AO can conserve heat and freshwater sufficiently accurately to be useful in practice; drifts of temperature and water volume arising from non-conservation of fluxes are much smaller than those likely to be seen in climate change integrations. In addition, the HadGEM3-AO r1.1 prototype version has a reasonable simulation of the atmospheric and oceanic circulation and sea ice distribution. This is particularly reassuring since the r1.1 version has not been subjected to any tuning to improve the simulation and includes new ocean and sea ice components as well as modifications to the atmosphere model relative to HadGEM1 (Johns et al., 2006).

This paper describes how a single executable can be built from two separate code bases; in this case, the NEMO ocean model and CICE sea ice model. This approach may also be applied to atmosphere and land submodels in the future; as described in Sect. 2, currently the MOSES-II scheme is embedded within the UM atmosphere model. However, in the future, MOSES-II will cease to be used and be replaced by the Joint UK Land Environment Simulator (JULES; Blyth et al., 2010). JULES exists as a code independent of the UM atmosphere model but it is likely that the land scheme will need to run synchronously with the atmosphere model. Producing a single executable from two code bases, as we have

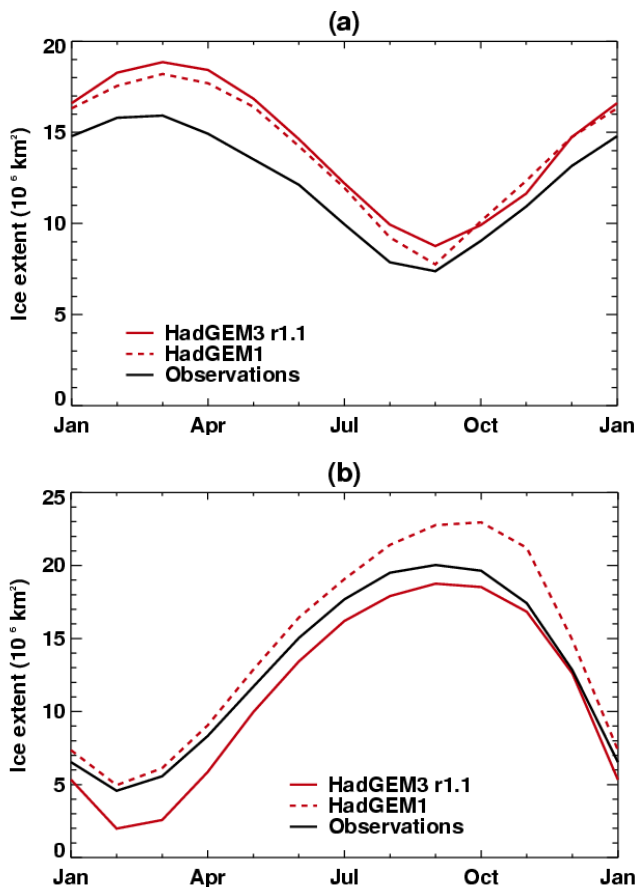


Fig. 14. Seasonal cycle of sea ice extent (10^6 km^2) for years 21 to 30 of HadGEM3 r1.1 (red line) and HadGEM1 (red dashed line) together with the HadISST observational data set values (Rayner et al., 2003) averaged over the period 1980–1999 (black line) for (a) Northern Hemisphere; and (b) Southern Hemisphere.

done with NEMO and CICE, maintains the integrity of the individual codes and requires minimal changes to the component models.

As discussed earlier, the technical details of the implementation of a coupled model can be highly dependent on the model configuration and supercomputer architecture on which the model is deployed. In particular, if the model is deployed on a scalar MPP machine (e.g. the IBM power6 now used at the Met Office), opportunities for efficient use of resources are increased. With higher numbers of processors available, the consequent greater flexibility in achieving optimal load balance reduces the need to run NEMO and CICE as a single executable (although there may be other reasons for continuing to do this). The ACCESS coupled model (T. Hirst, personal communication, 2008) is largely based on the HadGEM3-AO infrastructure but uses the AusCOM ocean model based on GFDL's MOM4. AusCOM and CICE are run as separate executables, demonstrating the viability of both approaches within the coupling infrastructure

outlined here. In addition, increasing processor numbers associated with larger distributed memory supercomputers and higher model resolutions are expected to increase the need for scalable, parallel coupling (as available from OASIS4) in order to maintain efficient performance.

Overall, this paper serves as documentation of the coupling infrastructure for users of the HadGEM3-AO climate model and as a guide for other modellers who are interested in building conservatively coupled GCMs.

Appendix A

Atmosphere science and parameter settings for HadGEM3 r1.1

HadGEM3-AO r1.1 uses the Met Office Unified Model (UM) as its atmospheric submodel. The HadGEM3-A r1.1 version of the UM is similar to the HadGEM1 version (described by Martin et al., 2006). The major differences between these models are summarised in Table A1, with the individual schemes explained in more detail in the following subsections.

A1 Dynamical core and resolution

The dynamical core of HadGEM3-A is mainly unaltered to that used in HadGEM1. It is described by Davies et al. (2005) and is nonhydrostatic, fully compressible, and uses a semi-Lagrangian advection scheme. The vertical coordinate is height based and terrain following. Each of the schemes listed in this section calculate increments (of temperature, humidity, wind, etc.) which are then combined at the end of each timestep.

The model uses an Arakawa C grid where the zonal and meridional winds lie horizontally between the density points. Vertical winds, temperatures and tracers lie vertically between the density points. This arrangement is well suited to work with the semi-Lagrangian advection scheme.

The horizontal and vertical resolution of HadGEM3-A r1.1 remains unchanged from HadGEM1. It has a horizontal resolution of 1.25° latitude by 1.875° longitude (N96), 38 vertical levels (top at 39 km – “low top”) and a 30 min timestep. Various increased vertical and horizontal resolution versions of HadGEM3-A r1.1 have also been tested including:

- An 85 vertical level model (“high top”) with an 85 km model top, a resolved stratosphere and an improved stratospheric Quasi-Biennial Oscillation;
- An N216 (0.554° latitude by 0.833° longitude) horizontal resolution model that resolves more small-scale features and modes of variability.

Results from higher resolution models are beyond the scope of this paper, but form part of the development of the HadGEM3 model family.

Table A1. Main differences between the atmospheric submodels in HadGEM1 and HadGEM3 r1.1.

	HadGEM1	HadGEM3 r1.1
Convection	Mass flux scheme; 1 h relative humidity based CAPE closure; All mass detrained at top of convection.	Mass flux scheme; 2 h vertical wind based CAPE closure; Smoothed adaptive detrainment throughout convective events.
Orography	Global Land-Based 1 km Base Elevation (GLOBE) data at 1' resolution.	GLOBE data at 30'' resolution including increased subgrid orography.
Hydrology and soil properties	MOSES-II hydrology and surface exchange scheme; Clapp and Hornberger soil hydraulics; Wilson and Henderson-Sellers (WHS) soils.	MOSES-II hydrology and surface exchange scheme; Van Genuchten soil hydraulics; International Geosphere- Biosphere Programme (IGBP) soils.
Clouds	Diagnostic cloud scheme: cloud properties are calculated from temperature and humidity each timestep and then discarded.	Prognostic Cloud fraction and Prognostic Condensate scheme (PC2): cloud properties are retained from one timestep to the next and advected with the wind.
Aerosols	Sulphates Soot Biomass Dust (from WHS soils) Sea salt	Sulphates Soot Biomass Dust (from IGBP soils) Sea salt Biogenic (climatology only)

A2 Radiation and ozone

The majority of the radiation code remains unaltered from HadGEM1 which uses the Edwards and Slingo (1996) radiation scheme. The most notable radiation code change in HadGEM3-A r1.1 is improved pressure and temperature scaling (designed by W. Zhong, personal communication, 2007). This code improves longwave fluxes travelling through the stratosphere where pressures are low, making considerable improvements to the “high top” model.

The HadGEM3-A r1.1 radiation scheme also includes a scaling to account for cloud inhomogeneity. This effect increases transmissivity in cloudy grid boxes, increasing the amount of shortwave radiation reaching the surface. This code change is a tuning rather than a permanent solution and was included to reduce cold sea surface temperature biases. Future releases of HadGEM3-A may not include these cloud inhomogeneity changes.

The ozone distribution also remains unaltered from HadGEM1, using the SPARC data set (Randel and Wu, 2007). The “low top” model involves a scheme that moves the ozone to follow the tropopause and conserves column mass, while the “high top” model simply imposes a fixed global field.

A3 Aerosols

HadGEM3-A r1.1 models many different aerosol species as listed at the bottom of Table A1. It differs from HadGEM1 in that the dust scheme has been tuned to model dust generated from IGBP (International Geosphere-Biosphere Program) soils. This improved dust scheme increases the amount of atmospheric dust in the model, reducing the amount of shortwave reaching the surface over the tropical Atlantic, the Sahara, the Arabian Sea and southern Asia. All these dust changes improve the model when compared to observations. All of the aerosol schemes in HadGEM3-A r1.1 are interactive (aerosols are emitted, advected and deposited) apart from the new biogenic aerosol scheme where climatological concentrations are used for their radiative effects.

A4 Land surface processes and hydrology

HadGEM3-A r1.1 uses a new method proposed by van Genuchten (1980) for calculating soil hydraulic conductivity. This reduces runoff, producing wetter soils with more evaporation in regions with high precipitation like the Amazon and equatorial Africa. Increased evaporation also leads to surface cooling and colder surface temperatures. This is particularly apparent in Europe where an existing warm bias is removed in Boreal summer and a slight cold bias is introduced in Boreal winter.

The orography used in HadGEM3-A r1.1 is the 30" GLOBE data set that provides orographic heights at about 1 km horizontal resolution. The Raymond (1988) filter has been used to process the data from the raw 30" dataset using a filter length scale of 6 km. Envelope orography was not used as the UM uses a flow blocking parameterization that includes the effects of sub grid scale ridges onto the grid scale flow. Converting the GLOBE data set for use in the model's N96 horizontal grid involved diagnosing a large amount of small-scale orography as sub-grid-scale, and the effects of this on the model were parameterized by the orographic gravity wave drag scheme. This scheme had to be retuned from that used in HadGEM1 which used a coarser orography data set. More information on the generation of the model orography and the flow blocking parameterization can be found in Webster et al. (2003).

A5 Boundary layer turbulent mixing

The boundary layer scheme of HadGEM3-A r1.1 is that of Lock et al. (2000) and is similar to that used in HadGEM1 with the following changes:

- A new non-local stress formulation has been included, following Brown and Grant (1997). In addition, revisions have been made to both the implementation of the entrainment parametrization and the diagnosis of the vertical extent of the diffusion profiles to make these more accurate and robust;
- Improvements have been made to surface scalar transfer over the sea to bring its dependence on windspeed more in line with recent observations. At the same time, the effects of salinity in reducing the saturation vapour pressure (previously neglected in the model) have been included. Details are given in Edwards (2007);
- HadGEM3-A uses an unconditionally stable, non-oscillatory, boundary layer solver as described in Wood et al. (2007). This reduces noise in regions that have stable boundary layers and removes two-timestep-oscillations found to be an issue in HadGEM1.

HadGEM1 also suffered with a lack of rain over the Maritime Continent (Indonesia). This was improved in HadGEM3-AO with the use of a "buddy" scheme for coastal grid points which uses an average windspeed over neighbouring sea points to split the level-1 windspeed into separate land and sea contributions. This enhances the wind-speed over the sea part of the grid box giving improved scalar fluxes there. The rain over the Maritime Continent is much improved with this new scheme.

A6 Convection

The convection scheme of HadGEM3-A r1.1 is similar to that in HadGEM1 but with the following changes:

- The HadGEM3-A r1.1 convection scheme uses smoothed adaptive detrainment which better represents the detrainment of the convective plume into its environment. The old scheme detrained all of the plume at the cloud top where the parcel becomes neutrally buoyant. The new scheme allows air to be detrained out of the parcel's ascent, maintaining parcel buoyancy so that the parcel can reach higher up in the atmosphere. It is thought that this better represents a cluster of convective clouds where some clouds are relatively shallow and some are relatively tall;
- HadGEM3-A r1.1 uses a 2-h CAPE closure scheme based on vertical windspeed. The CAPE closure scheme specifies the timescale (in this case two hours) over which the convection scheme will try and dissipate the convectively available potential energy (CAPE). The convection scheme does this by adjusting the updraught mass flux into the base of the convective column. In certain conditions the CAPE closure timescale is not short enough to prevent the build up of excessive moisture in a column, known as a grid point storm. Grid point storms are diagnosed by testing whether the vertical windspeed exceeds 0.3 m s^{-1} and when this happens the CAPE timescale is reduced to less than two hours;
- in HadGEM3-A r1.1 the critical water content for precipitation is a function of cloud depth so that shallower clouds have to contain more water before they start to generate precipitation. This is an improvement to HadGEM1 when precipitation only occurred if shallow and deep clouds exceeded a critical water content for precipitation (fixed at 1 g kg^{-1}) and if shallow clouds exceeded a critical (land/sea-dependent) depth;
- the convection scheme also includes revised parcel perturbations for shallow convection that help make vertical fluxes more consistent between the boundary layer and convection schemes.

A7 Cloud and precipitation microphysics

The cloud scheme in HadGEM3-A r1.1 is substantially different to the diagnostic scheme used in HadGEM1. It uses the Prognostic Cloud Prognostic Condensate (PC2) scheme described by Wilson et al. (2008). The PC2 cloud scheme retains the following prognostic variables from one timestep to the next:

- Liquid water cloud fraction (as a fraction of the grid box)
- Ice cloud fraction (as a fraction of the grid box)
- Liquid water condensate amount (kg of liquid water per kg of air)
- Ice condensate amount (kg of ice per kg of air)

The following schemes are allowed to modify these four variables:

- Shortwave radiation
- Longwave radiation
- Boundary layer
- Cloud erosion (cloud mixes with surrounding air)
- Large-scale precipitation (including phase changes)
- Convection
- Advection
- Pressure change

Thus it is possible to slowly evolve cloud throughout the model run, with cloud being generated, increased, decreased and destroyed by all possible physical processes. The net effect of PC2 (compared to HadGEM1's cloud scheme) is for cloud to be retained for longer, increasing (and improving) cloud cover and precipitation over many tropical land masses and the Maritime Continent.

Appendix B

Ocean science and parameter settings for HadGEM3 r1.1

Although we largely use “NEMO” as a short-hand for the ocean component of HadGEM3-AO, formally NEMO is a framework of ‘ocean related engines’ including OPA (for ocean dynamics and thermodynamics), LIM (for sea ice dynamics and thermodynamics) and TOP (for biogeochemistry). HadGEM3-O only makes use of the OPA part – this is an ocean model originally based on OPA8.1 (Madec et al., 1998) but which has undergone continual development within the NEMO framework. The model does not include any biogeochemistry at present and the LIM sea ice code from NEMO is not used because CICE has been chosen as an alternative sea ice model. It is planned however that the interface to CICE will become part of the standard NEMO code in the near future.

The NEMO code is written in FORTRAN 90 (making use of some of the newer features including modules and interfaces) and C preprocessing is used (in combination with namelist logicals) to determine which parts of the code are

used in a particular integration. NEMO is designed to have a transparent structure with a high degree of internal modularity. This makes the code easier to understand and helps in the development process. Coding standards (Madec, 2006) with fixed conventions for the naming of variables, logicals etc are also useful in keeping the code easily readable. In general the code is designed not to have too many machine customisations although it is optimised for vector computers (like the NEC SX8 on which the coupled model described has been run). When running on many processors the domain can be decomposed using MPI although this was not in fact necessary for the SX8 integrations because the atmosphere model is much slower than the ocean (see details in Sect. 5.1).

B1 Horizontal, temporal and vertical discretisation

The equations in NEMO are written in an orthogonal curvilinear coordinate system in which the grid spacing can vary spatially. It is this feature which makes it possible to use horizontal grids (like the tripolar ORCA grids) which are not essentially latitude-longitude grids because the so-called ‘scale factors’ can vary spatially. The set of standard NEMO configurations does not include a 1 degree resolution tripolar grid in the ORCA family - the details of the construction of this grid (which includes enhanced resolution in the tropics) are available in Appendix C.

As mentioned in Sect. 2, NEMO uses an Arakawa C grid (Arakawa and Lamb, 1977) for the spatial discretisation of the equations (whereas CICE uses a B grid). NEMO has an Asselin time filter (Asselin, 1972) to control noise in its leapfrog timestepping scheme – again this differs from CICE where a forward timestep is used.

Various vertical coordinates are available as options within NEMO (including z -, s - and σ -coordinates) but to date all HadGEM3-O configurations have used the z -coordinate option. There are 42 z -levels, whose thicknesses vary smoothly from 10 m at the surface to 300 m at depth. Partial steps – bottom-layer grid cells whose lower boundary is the specified model bathymetry rather than the next available grid-cell boundary – are used for improved representation of bathymetry, especially where it has a small slope (Madec, 2008). For stability reasons it was found necessary to use a partial step thickness no less than the minimum of 20 m and one tenth of the full cell thickness (this is set using parameters `e3zps_min` and `e3zps_rat` in the namelist). The minimum allowed ocean depth is 30 m.

B2 ORCA1 configuration and scientific choices

At the time of development, a one degree configuration does not form part of the standard NEMO release, in addition to the construction of the ORCA1 grid and the vertical level set, there were many decisions to be made about the scientific parameters and numerical schemes to be used. This was

Table B1. The C preprocessor keys used for NEMO in HadGEM3r1.1.

cpp key	Purpose
key_vectopt_loop	Enables collapse of internal loops for vector machine
key_vectopt_memory	Reduces memory requirement at expense of computational time (only used in TKE routine)
key_dynspgflt	Surface pressure gradient activated with filtered free surface
key_diaev	Eddy-induced velocities included as diagnostics
key_diaht	Computation of various heat contents and thermocline depths activated
key_ldfslp	Calculation of isopycnal slopes activated (required for isopycnal diffusion)
key_traldf_eiv	Eddy-induced (Gent-McWilliams) advection activated
key_traldf_c2d	2-D spatial variation of lateral (isopycnal) diffusivity and eddy (Gent-McWilliams) coefficient
key_dynldf_c3d	3-D spatial variation of coefficient for lateral viscosity
key_dstatem	Required for reading in initial state 3-D temperature field
key_dtasal	Required for reading in initial state 3-D salinity field
key_trabbc	Bottom temperature boundary condition activated
key_zdftke	Turbulent Kinetic Energy (TKE) vertical mixing activated
key_zdfddm	Double diffusive mixing activated
key_mpp_mpi	Allows massively parallel processing using MPI (although NEMO actually ran on a single processor)
key_orca_r1	ORCA1 configuration
key_orca1_upsadv	Extra upstream advection activated at various river outflows
key_oasis3	Coupling to UM atmosphere using OASIS3
key_cpl_rootexchg	Passing all data to/from OASIS3 using a single NEMO processor (in fact NEMO ran on a single processor anyway)
key_cice	Using CICE rather than LIM ice model

largely a case of expert judgement (with considerable advice from G. Madec, personal communications, 2006) in combination with various sensitivity and stability tests.

Key choices were to use a 1 h timestep, a horizontal eddy viscosity coefficient ($ahm0$) of $10^4 \text{ m}^2 \text{ s}^{-1}$ and a lateral eddy diffusivity coefficient ($aht0$) of $10^3 \text{ m}^2 \text{ s}^{-1}$. $aht0$ is scaled with the grid “scale factors” such that values decrease to approximately $500 \text{ m}^2 \text{ s}^{-1}$ in regions such as the Southern Ocean. The tracer diffusion is along isopycnals and there is a Gent-McWilliams (Gent and McWilliams, 1990) term applied in the advective formulation with the NEMO default of a spatially and temporally varying coefficient (calculated from the local growth rate of baroclinic instability). As in the ORCA2 configuration the horizontal (Laplacian) viscosity reduces to $aht0$ in the tropics, except in the Western boundaries where it relaxes back to $ahm0$ according to scaling factors bilinearly interpolated from the ORCA2 values. The value of ahm has a vertical dependence such that it is an order of magnitude greater at depth.

Second-order centred tracer advection is used except at river outflow points where a mixed upstream-centred scheme is used to avoid problems with negative salinities. The relevant coefficients were interpolated from ORCA2 and then slightly modified to be more appropriate for the runoff distribution generated by the UM atmosphere.

Surface boundary layer and interior vertical mixing of both tracers and momentum is treated consistently by a Turbulent Kinetic Energy (TKE) scheme (Gaspar et al., 1990; Blanke and Delecluse, 1993). There are many tuneable aspects of this scheme as implemented in NEMO – the values of various parameters currently in use can be seen in the supplementary material. The background vertical viscosity is $1.2 \times 10^{-4} \text{ m}^2 \text{ s}^{-1}$ and the background tracer diffusivity increases linearly with depth (Kraus, 1990) from a surface value of $1.2 \times 10^{-5} \text{ m}^2 \text{ s}^{-1}$.

For the dynamics the NEMO vector-form momentum advection scheme is used along with the enstrophy conserving scheme for the vorticity term.

The equation of state used is Jackett and McDougall (1995) and double diffusive mixing is activated with a maximum vertical mixing on salinity of $1 \times 10^{-4} \text{ m}^2 \text{ s}^{-1}$ and a heat-salt buoyancy flux ratio of 1.6. Convection is modelled using an enhanced vertical diffusion (for both tracers and momentum) of $100 \text{ m}^2 \text{ s}^{-1}$.

The shortwave radiation from the atmosphere is distributed in the surface layers of the ocean using parameters consistent with water of type I (Jerlov, 1968). However the less penetrating fraction (with an exponential decay depth of 0.35 m) is reduced from the standard 0.58 (Paulson and Simpson, 1977) to 0.26; this attempts to take account of the 200–689 nm wavelength band for the UM “blue” radiation

compared to the normally assumed 400–1000 nm band. At the ocean floor there is a linear bottom friction coefficient of 4×10^{-4} and a spatially varying geothermal heat flux (interpolated from the ORCA2 field inferred from Stein and Stein, 1992). The model does not have a bottom boundary layer scheme or any other special treatment to improve North Atlantic Deep Water formation.

For the model configuration described an implicit linear free-surface was used with a preconditioned conjugate gradient (pcg) solver. The absolute precision required for convergence is 1×10^{-6} with a maximum of 2000 iterations.

For reference, the complete NEMO namelist is given in the supplementary material, and the associated cpp (C pre-processor) keys are given in Table B1.

B3 NEMO code version

The code used was essentially the NEMO 3.0 code release (Madec, 2008) but with a number of local code changes. Briefly these were as follows:

- ORCA1: A few routines required minor changes to provide an ORCA1 option (either with a hard-wired parameter consistent with the 1 degree grid or 42 vertical level set, or to replicate the behaviour in other ORCA configurations as desired – see discussion above);
- HadGEM3-O: Code was added to deal with the coupling fields to and from both the UM atmosphere and the CICE ice model (see details of the required transformations in Sect. 3.1);
- diagnostic changes: Minor code changes were included to make the netCDF output files CF-compliant (Eaton et al., 2009), to include an 18 °C isotherm depth diagnostic, to include a Kara mixed layer depth (Kara et al., 2000) diagnostic, and to make the poleward transport diagnostic calculation and control more consistent with the other diagnostics.

Appendix C

Grids, bathymetry and land-sea masks

C1 Grids used by the different submodels

The ocean and sea ice components of HadGEM3-AO are run on the ORCA family of tripolar grids. We describe the configuration which is nominally 1 degree and known as “ORCA1” but the same regridding considerations apply to other resolutions of the ORCA family. The ORCA1 grid is shown in Fig. 2. It has the following characteristics:

- A uniform resolution “Mercator” grid in the Southern Hemisphere, in which the longitudinal resolution $\Delta\lambda=1^\circ$, and the latitudinal resolution $\Delta\phi=1^\circ\cos\phi$, thus ensuring isotropic grid blocks ($\Delta x:\Delta y=1$);
- grid refinement in the tropics between $\pm 20^\circ$ N, in which $\Delta\phi$ reduces smoothly to $1/3^\circ$ at the equator, while $\Delta\lambda$ remains 1° . Such tropical grid refinements have proved useful for modelling equatorial ocean behaviour in the past (Johns et al., 2006);
- a dipolar grid north of 20° N, in which the traditional coordinate singularity at the North pole has been replaced by an arc of a Great Circle – the ‘north fold’ – between two numerical poles sited in Canada and Siberia. This device places the singular points outside the ocean model domain, and leaves a nearly uniform resolution mesh in the Arctic Ocean. It thereby removes the need for costly and deleterious high frequency filtering at high latitudes. (On a traditional uniform resolution grid, Fourier filtering would be necessary to allow the model to run in regions where the small grid spacing, arising from the convergent meridians, would otherwise demand prohibitively small timesteps.)

C2 Bathymetry and land-sea mask

The ORCA1 bathymetry has been derived from the ETOPO2 2” global bathymetry data set (ETOPO2v2, 2006) by collaborators at the National Oceanography Centre, Southampton (Coward, 2006). The resulting land-sea mask was then smoothed, principally to diminish the risk of isolated pillars of sea ice building up. In the Arctic, care has been taken to open the Kennedy Channel (between Ellesmere Island and Greenland) and the Gulf of Ob, and to remove resolvable, but potentially problematic, islands in the Canadian Archipelago. It was decided to fill in the North West Passage, even though technically it could be open at this resolution, as it was felt that the numerical risks associated with sea-ice accumulation in the narrow Passage outweighed the benefits of representing it. The land-sea mask in the Arctic can be seen in Fig. 2. In the Antarctic, ice shelves were eliminated and coasts, particularly of the Antarctic Peninsula, were smoothed.

Appendix D

Sea ice scientific choices and parameters for HadGEM3 r1.1

The CICE sea ice model solves an equation for the evolution of the ice thickness distribution (ITD) in time and space (Thorndike et al., 1975): the evolution is determined by advection, thermodynamic growth and melt, and mechanical redistribution (ridging). Full details of the model are available in the CICE user manual (Hunke and Lipscomb, 2008);

here we summarise the main features of the model used in HadGEM3, and detail the specific settings and choices for HadGEM3 r1.1.

D1 Horizontal, temporal and vertical discretisation

HadGEM3 r1.1 uses a revision (146) of CICE version 4.0 which allows a tripolar grid to be employed. We use essentially the same family of ORCA grids as the NEMO model (described more fully in Appendix C), although CICE uses an Arakawa B grid rather than a C grid and so the CICE velocity grid points are not coincident with the NEMO velocity points. The implication of this on the coupling is described in Sect. 3. The grid and land-mask definitions required by CICE are read in directly from a file, as are the initial conditions.

The sub-grid-scale ITD is modelled by dividing the ice pack at each grid point into a number of thickness categories. HadGEM3 uses 5 categories, plus an open-water category, which has been shown to be sufficient for climate modelling (Bitz et al., 2001). The lower bounds for the 5 thickness categories are: 0, 0.6, 1.4, 2.4 and 3.6 m.

HadGEM3 uses the zero-layer thermodynamic model of Semtner (1976) to calculate the growth and melt of the sea ice, with one layer of snow and one layer of ice in the vertical. This is not the standard scheme implemented in CICE, which has a multilayer ice model (Bitz and Lipscomb, 1999). We could not use the CICE multilayer thermodynamics in HadGEM3 because the UM atmosphere currently calculates the surface temperature and the conductive heat flux into the ice at sea ice points. This would not be consistent with the CICE multilayer thermodynamics scheme, which calculates these quantities itself. Planned modifications to CICE and the UM atmosphere will enable the CICE multilayer model to be used with the UM atmosphere in the future. However for HadGEM3 we adapted CICE to use the zero-layer surface and diffusive ice fluxes received from the UM atmosphere. The full zero-layer thermodynamics scheme was also added to CICE for use in standalone runs (as detailed below). CICE solves for the surface temperature and fluxes in the same way as the multilayer case but the calculations are much simpler.

D2 Thermodynamics

The sea ice albedo is calculated within the UM atmosphere, and is a function of temperature and snow cover, including a parameterization to represent the impact of melt ponds. This is the same scheme that was used in HadGEM1 (McLaren et al., 2006) with one exception of the snow-free sea ice albedo being increased from 0.57 to 0.61 following a change made in HadGEM2 (Martin et al., 2011). The UM atmosphere also calculates the surface sea ice temperature and the atmosphere-to-ice fluxes as in HadGEM1 (see McLaren et al., 2006 for details). Within CICE these fluxes (downward latent heat flux, conductive flux through the ice, and surface

heat flux) determine the rate at which the ice grows or melts in each thickness category. This calculation also uses the enthalpy of each snow and ice layer. The enthalpy is defined as the negative of the energy required to melt a unit volume of ice/snow and raise its temperature to 0°C. For the zero-layer thermodynamics used in HadGEM3, the enthalpy is simply the negative product of the density and the latent heat of fusion. The calculated thermodynamic growth or melt rates are then used in the linear remapping scheme of Lipscomb (2001) to transport the ice between thickness categories.

D3 Dynamics and ridging

The ice velocities are calculated by solving the 2-D momentum equation for the force balance per unit area in the ice pack (Hibler, 1979), including terms for wind stress, ocean stress, internal ice stress, and stresses due to Coriolis effects. The internal ice stress is calculated using the elastic viscous plastic (EVP) scheme (Hunke and Dukowicz, 2002), which assumes the ice has a viscous plastic rheology, and incorporates an elastic wave modification to improve the computational efficiency. HadGEM3 r1.1 uses the Rothrock (1975) formulation for ice strength with the standard CICE value of 17 for the parameter C_f (ratio of total energy losses to potential energy change). The sea ice is advected using the CICE incremental remapping scheme (Lipscomb and Hunke, 2004).

The mechanical redistribution (or ridging) scheme in CICE converts thinner ice to thicker ice and open water, and is applied after the advection of ice. When the ice is converging, enough ice ridges to ensure that the ice area does not exceed the grid-cell area. The scheme is based on work by Thorndike et al. (1975), Hibler (1980), Flato and Hibler (1995), and Rothrock (1975). It favours the closing of open water and ridging of the thinnest ice over the ridging of thicker ice. In HadGEM3 r1.1, the ridging participation function suggested by Lipscomb et al. (2007) is used. The ridged ice is then distributed between thickness categories assuming an exponential ITD (Lipscomb et al., 2007).

D4 Implementation of zero-layer thermodynamics in CICE for standalone runs

Note that this scheme is only used in standalone CICE runs with ‘HadGEM3’-like science (i.e. `calc_Tsfc` is set to true and `heat_capacity` to false in the CICE namelist, and there is only one snow and one ice layer).

The zero-layer thermodynamics uses the same basic equation as the multiple layer model to solve for the surface temperature and surface fluxes (as described in Hunke and Lipscomb, 2008). However, because the conductive flux through the ice is not dependent on any internal temperatures, it is much easier to solve for the surface temperature.

Table D1. The C preprocessor keys used for CICE in HadGEM3r1.1.

cpp key	Purpose
coupled	Coupled run
ncdf	NetCDF format options available for input and output files
CICE_IN_NEMO	CICE is run in the NEMO environment. CICE is called from the NEMO surface module which also supplies the coupling fields to and from CICE
ORCA_GRID	Controls reading in grid, land masks and forcing data on the ORCA family of grids.
key_oasis3	Coupling uses OASIS3
REPRODUCIBLE	Ensures global sums bit compare for parallel model runs with different grid decompositions

At the surface, if the surface temperature (T_{sf}) is below 0°C , the conductive flux from the top surface to the bottom of the ice (F_{ct}) must balance the net surface flux from the atmosphere to ice (F_0),

$$F_0 = F_{ct}, \quad (\text{D1})$$

and both fluxes are evaluated at time $m + 1$ (m being the current timestep). The surface fluxes are computed in the same way as the CICE multilayer model. In the zero-layer model, any solar energy calculated to be absorbed internally within the ice is added to the surface absorption. A fraction of the shortwave radiation is still able to pass through the ice to the ocean, in contrast to the coupled configuration.

The conductive flux in the zero-layer model is

$$F_{ct} = \frac{k_i}{h_{\text{eff}}} (T_{sf} - T_f), \quad (\text{D2})$$

where k_i is the conductivity of ice ($2.0 \text{ W m}^{-1} \text{ K}^{-1}$), T_f is the freezing temperature (i.e. the temperature of the base of the ice). h_{eff} is the effective ice thickness defined as $h_i + \frac{k_s}{k_i} h_s$ (with h_i and h_s being the ice and snow thicknesses, and k_s the conductivity of snow, $0.31 \text{ W m}^{-1} \text{ K}^{-1}$.)

As in the multiple layer code, we make the linear approximation

$$F_0^{m+1} = F_0^* + \left(\frac{dF_0}{dT_{sf}} \right)^* (T_{sf}^{m+1} - T_{sf}^*), \quad (\text{D3})$$

where T_{sf}^* is the surface temperature from the most recent iteration, and F_0^* and $(dF_0/dT_{sf})^*$ are functions of T_{sf}^* . (The outgoing longwave, sensible and latent heat fluxes are the components of F_0^* that are dependent on T_{sf}^* .) Substituting Eqs. (D2) and (D3) into (D1) and rearranging we get

$$\left[\left(\frac{dF_0}{dT_{sf}} \right)^* - \frac{k_i}{h_{\text{eff}}} \right] T_{sf}^{m+1} = \left(\frac{dF_0}{dT_{sf}} \right)^* T_{sf}^* - F_0^* - \frac{k_i}{h_{\text{eff}}} T_f. \quad (\text{D4})$$

This equation is solved by iteration and the conditions on T_{sf} for an acceptable solution are those of the multilayer model.

D5 CICE settings used for HadGEM3 r1.1

The key settings for HadGEM3 are highlighted here, and we include the CICE namelist used for HadGEM3 r1.1. For coupling with the UM atmosphere, `heat_capacity` and `calc_Tsfc` are both set to false. This means that zero-layer thermodynamics are used and that CICE does not calculate any surface fluxes or the surface ice temperature. Note that setting `calc_Tsfc` to false also means that the albedo settings in the CICE namelist are irrelevant as the albedo is not calculated by CICE. Wind stresses are passed from the UM atmosphere rather than being calculated in CICE, so `calc_strair` is set to false. A constant value for the freezing point of sea water is used (-1.8°C), by setting `Tfrzpt=constant`. This is required for consistency with the UM atmosphere ice thermodynamics. The `ns_boundary_type` is set to `tripole` for the ORCA1 grid, indicating a tripolar grid with the “north fold” occurring along velocity points. The alternative tripole setting `tripoleT` is used for ORCA2 and ORCA025 grids where the north fold occurs along temperature points. The CICE cpp keys used in HadGEM3 r1.1 are shown in Table D1. The CICE namelist used in HadGEM3 r1.1, which has been edited to detail the scientific options only, can be viewed in the supplementary material.

Appendix E

OASIS3 build

HadGEM3 r1.1 employed version 2.5 of the OASIS3 coupler, dated 28 December 2007. It was compiled for use on the NEC SX8 using MPI-SX version 151 and `sxf90` FORTRAN compiler version 360. For reasons of portability, MPI-1 was employed. The build was linked to the `netcdf` version 3.6.0-p1 library. The optional OASIS3 cpp key “`DEBUG`” was typically used throughout the debugging and development phases to obtain additional debugging output. The extended debugging option, activated via the cpp key “`_DEBUG`” was not generally employed following the discovery that it tended

to cause false deadlock detection in cases where remapping weights were being generated during the run. The standard OASIS3 configuration and make process was modified slightly to allow the “PRISM_HOME” environment variable used in the make process to be specified once only, in the top-level make.inc file. Each new build of OASIS3 was initially tested by employing the standard “toy” models supplied with the distribution, before attempting to deploy it in a full coupled model system employing the UM atmosphere and NEMO-CICE models.

E1 Management of namcouple and name_table.txt files

The namcouple file, used to supply the OASIS3 coupler with all the necessary information needed to run, contained all the key word, resource, and field operation definitions required by OASIS3. The bulk of the information contained in these files needed to be defined before the model run commenced. However, the Unified Model control scripts were adapted to allow for certain items, such as the processor configuration for each model component, executable names and model run length to be defined and automatically generated based on job submission parameters.

The name_table.txt file was only used for informational purposes in defining names of coupling fields to be output in any netCDF debugging files generated by OASIS3. Since the HadGEM3 model contained a number of fields not in the original file, new descriptions and field index numbers were added to this file. These were principally to support multi-category ice fields associated with coupling to the CICE model.

E2 The OASIS3 SPAWN option

The OASIS3 “SPAWN” option, whereby OASIS3 could be allowed to create and destroy model component subprocesses as and when required, was not employed since it offered no advantage in the Met Office batch queuing system. Its reliance on MPI-2 functionality additionally rendered it not strictly portable at the time of development and therefore unsuitable for use in operational systems.

E3 OASIS restart files

It was essential for traceability reasons that models could be stopped and restarted at arbitrary times, producing exactly the same results as if they had run continuously for an equivalent period (“bit reproducibility”). OASIS3 provides such a mechanism via use of its own form of netCDF restart file. However, since coupled model file management already involved numerous files of various types and formats, then avoidance of the need to manage a further file type was considered desirable. Each component model dump contained, or could be modified to contain, sufficient information to achieve this bit reproducibility by rendering restart data files self-consistent, i.e. each could be used as the basis

to restart the appropriate component without recourse to OASIS3 restart files and without the attendant data management issues.

Appendix F

Building HadGEM3 configuration

F1 Compilation of UM, NEMO and CICE components

The UM is compiled using automatic promotion of all integer and real values to 64 bits unless explicitly typed. With OASIS3 typically requiring 32 bit integers, the routines interfacing to the OASIS PSMILE library calls needed to be excluded from the compiler option to promote all variables to 64 bits. This was achieved by explicit typing and use of the NEC sxf90 “-dw” compiler option to employ 4 byte numeric storage unless otherwise stated.

GCOM (General COMmunications) is an interface library used by the UM which provides an interface or “wrapper” for MPI calls within the code. The presence of the GCOM wrapper around MPI communications in the UM caused some difficulty. The MPI calls within the OASIS3 PSMILE routines caused conflicts with the MPI calls in GCOM and the atmosphere code. To avoid this, GCOM was rebuilt using an MPL FORTRAN layer around the MPI calls.

GCOM worked on the assumption that it was employed within a model component which had no need to communicate with any external components. Thus, it employed the MPI_COMM_WORLD MPI communicator directly for all its operations. This conflicted with the need for OASIS3 to be able to define MPI sub-communicators for each model component. Hence modifications were made to GCOM, allowing it to use communicators provided by OASIS3 in place of MPI_COMM_WORLD.

The NEMO-CICE component is, by default, compiled to promote only real variables to 64 bits. Additionally, NEMO and CICE generally use explicit typing of most variables and were not therefore reliant on blanket UM-style compiler options to promote variables to a particular level of precision. Hence there were fewer issues about matching variable types in calls to PSMILE routines from within NEMO.

F2 Unified Model User Interface (UMUI) control

The UM atmosphere is based on code which, in the main, uses dynamic memory allocation. Hence once an atmosphere component has been compiled, the number and configuration of processors used to run it can be altered without the need to recompile. The same is not true of the NEMO and CICE components. These use static allocation, with model dimensions defined at compilation time. Hence any change in processor arrangement required modifications to the source code followed by recompilation. This procedure was automated through the use of numeric values for relevant parameters

supplied via `cpp` keys. These changes (with the aid of environment variables and specially developed scripts embedded in the UMUI) defined and updated critical information, such as the processor configuration, prior to compilation.

Additionally, various scripts and modifications were developed to manage automatic resubmission, continuation runs and data archiving. These all centred around the existing UM control system, but introduced new concepts to deal with multiple component models employing various file types not previously catered for by the UM system.

F3 Debugging and visualisation of exchanged fields

Field definitions specified in the controlling OASIS namcouple file employed the EXPOUT option for coupling fields in all development and debugging cases. This option indicates that the relevant coupling field is output to a separate, OASIS3-managed, netCDF file as well as being involved in the usual coupling exchanges. These files were used to facilitate debugging and allow visualisation of fields sent and received during coupling exchanges.

For performance reasons, during operational runs, the EXPORTED option was employed in namcouple files, preventing output of coupling fields to separate netCDF files, thus reducing any OASIS3 runtime costs and minimising demands on disk space.

Supplementary material related to this article is available online at:

<http://www.geosci-model-dev.net/4/223/2011/gmd-4-223-2011-supplement.pdf>.

Acknowledgements. We thank all colleagues at the Met Office Hadley Centre who have contributed towards the development of the model; Gervan Madec and other members of the NEMO System Team for their help with setting up the ORCA1 configuration; Andrew Coward for construction of the original ORCA1 ocean bathymetry; William Lipscomb for his help and advice with using CICE; Adrian Lock, Rachel Stratton, Cyril Morcrette and James Manners for their help describing the atmospheric sub-components; Rene Redler and Sophie Valcke for advice in setting up OASIS and Kristian Mogensen for useful discussions on the practicalities of OASIS use; and Jonathan Gregory for his help in extending the tripolar boundary condition in CICE and discussions surrounding CF-compliance of NEMO output. This work was supported by the Joint DECC and Defra Integrated Climate Programme – DECC/Defra (GA01101).

Edited by: D. Lunt

Reviewed by: two anonymous referees

References

- Adler, R. F., Huffman, G. J., Chang, A., Ferraro, R., Xie, P., Janowiak, J., Rudolf, B., Schneider, U., Curtis, S., Bolvin, D., Gruber, A., Susskind, J., and Arkin, P.: The Version 2 Global Precipitation Climatology Project (GPCP) Monthly Precipitation Analysis (1979–Present), *J. Hydrometeorol.*, 4, 1147–1167, 2003.
- Arakawa, A. and Lamb, V. R.: Computational design of the basic dynamical process of the UCLA general circulation model, *Methods in Computational Physics*, 17, Academic Press, New York, 173–265, 1977.
- Arribas, A., Glover, M., Maidens, A., Peterson, K., Gordon, M., MacLachlan, C., Cusack, S., and Scaife, A.: The GloSea4 ensemble prediction system for seasonal forecasting, *Mon. Weather Rev.*, in press, 2011.
- Asselin, R.: Frequency filter for time integrations, *Mon. Weather Rev.*, 100, 487–490, 1972.
- Bitz, C., Holland, M., Weaver, A., and Eby, M.: Simulating the ice-thickness distribution in a coupled climate model, *J. Geophys. Res.*, 106, 2441–2463, 2001.
- Bitz, C. M. and Lipscomb, W. H.: An energy-conserving thermodynamic model of sea ice, *J. Geophys. Res.*, 104 (C7), 15669–15677, 1999.
- Blanke, B. and Delecluse, P.: Variability of the tropical Atlantic ocean simulated by a general circulation model with two different mixed-layer physics, *J. Phys. Oceanogr.*, 23, 1363–1388, 1993.
- Blyth, E., Clark, D. B., Ellis, R., Huntingford, C., Los, S., Pryor, M., Best, M., and Sitch, S.: A comprehensive set of benchmark tests for a land surface model of simultaneous fluxes of water and carbon at both the global and seasonal scale, *Geosci. Model Dev. Discuss.*, 3, 1829–1859, doi:10.5194/gmdd-3-1829-2010, 2010.
- Brown, A. R. and Grant, A. L. M.: Non-local mixing of momentum in the convective boundary layer, *Bound. Layer Meteorol.*, 84, 1, 1–22, 1997.
- Collins, W. J., Bellouin, N., Doutriaux-Boucher, M., Gedney, N., Hinton, T., Jones, C. D., Liddicoat, S., Martin, G., O'Connor, F., Rae, J., Senior, C., Totterdell, I., Woodward, S., Reichler, T., and Kim, J.: Evaluation of HadGEM2 model, Hadley Centre Technical Note, 74, 44 pp., 2008.
- Coward, A. C.: The ORCA1 model bathymetry, <http://www.noc.soton.ac.uk/nemo/?page=bathyimages>, National Oceanography Centre, Southampton, UK, 2006.
- Cox, P. M., Betts, R. A., Bunton, C. B., Essery, R. L. H., Rowntree, P. R., and Smith, J.: The impact of new land surface physics on the GCM simulation of climate and climate sensitivity, *Clim. Dynam.*, 15, 183–203, 1999.
- Davies, T., Cullen, M. J. P., Malcolm, A. J., Mawson, M. H., Staniforth, A., White, A. A., and Wood, N.: A new dynamical core for the Met Office's global and regional modelling of the atmosphere, *Q. J. Roy. Meteor. Soc.*, 131, 1759–1782, 2005.
- Eaton, B., Gregory, J., Drach, R., Taylor, K., and Hankin, S.: NetCDF Climate and Forecast (CF) Metadata Conventions, <http://cf-pcmdi.llnl.gov/documents/cf-conventions/1.4/cf-conventions-multi.html>, 2009.
- Edwards, J. M.: Oceanic Latent Heat Fluxes: Consistency with the atmospheric hydrological and energy cycles and general circulation modelling, *J. Geophys. Res.*, 112, D06115, doi:10.1029/2006JD007324, 2007.
- Edwards, J. M. and Slingo, A.: Studies with a flexible new radia-

- tion code. I: Choosing a configuration for a large-scale model, Q. J. Roy. Meteor. Soc., 122, 689–719, 1996. ETOPO2v2: 2-Minute Gridded Global Relief Data (ETOPO2v2), <http://www.ngdc.noaa.gov/mgg/fliers/01mgg04.html>, World Data Center for Geophysics & Marine Geology, Boulder, USA, 2006.
- Flato, G. and Hibler, W.: Ridging and strength in modeling the thickness distribution of Arctic sea ice, J. Geophys. Res., 100, 18611–18626, 1995.
- Gaspar, P., Gregoris Y., and Lefevre, J.-M.: A simple eddy kinetic energy model for simulations of the oceanic vertical mixing. Tests at station Papa and long-term upper ocean study site, J. Geophys. Res., 95(C9), 16179–16193, 1990.
- Gent, P. R. and McWilliams, J. C.: Isopycnal mixing in ocean circulation models, J. Phys. Oceanogr., 20(1), 150–155, 1990.
- Hibler, W.: A dynamical thermodynamic sea ice model, J. Phys. Oceanogr., 9, 817–846, 1979.
- Hibler, W.: Modeling a variable thickness sea ice cover, Mon. Weather Rev., 108, 1943–1973, 1980.
- Hunke, E. C. and Dukowicz, J. K.: The Elastic-Viscous-Plastic Sea Ice Dynamics Model in General Orthogonal Curvilinear Coordinates on a Sphere—Incorporation of Metric Terms, Mon. Weather Rev., 130, 1848–1865, 2002.
- Hunke, E. C. and Lipscomb, W. H.: CICE: the Los Alamos sea ice model documentation and software user's manual, Version 4.0, LA-CC-06-012, Los Alamos National Laboratory, N.M., 2008.
- Ingleby, B. and Huddleston, M.: Quality control of ocean temperature and salinity profiles – historical and real-time data, J. Mar. Syst., 65, 158–175, 2007.
- IPCC: Climate Change 2007: The Physical Science Basis, in: Contribution of Working Group I to the Fourth Assessment Report of the Intergovernmental Panel on Climate Change, edited by: Solomon, S., Qin, D., Manning, M., Chen, Z., Marquis, M., Averyt, K. B., Tignor, M., and Miller, H. L., Cambridge University Press, Cambridge, United Kingdom and New York, NY, USA, 996 pp., 2007.
- Jackett, D. R. and McDougall, T. J.: Minimal adjustment of hydrographic data to achieve static stability, J. Atmos. Ocean. Tech., 12, 381–389, 1995.
- Jerlov, N. G.: Optical Oceanography, Elsevier, 194 pp., 1968.
- Johns, T. C., Durman, C. F., Banks, H. T., Roberts, M. J., McLaren, A. J., Ridley, J. K., Senior, C. A., Williams, K. D., Jones, A., Rickard, G. J., Cusack, S., Ingram, W. J., Crucifix, M., Sexton, D. M. H., Joshi, M. M., Dong, B. W., Spencer, H., Hill, R. S. R., Gregory, J. M., Keen, A. B., Pardaens, A. K., Lowe, J. A., Boda-Salcedo, A., Stark, S., and Searl, Y.: The new Hadley Centre climate model HadGEM1: Evaluation of coupled simulations in comparison to previous models, J. Climate, 19 (7), 1327–1353, 2006.
- Jones, P. W.: First and Second-Order Conservative Remapping Schemes for Grids in Spherical Coordinates, Mon. Weather Rev., 127, 2204–2210, 1999.
- Kara, A. B., Rochford, P. A., and Hurlburt, H. E.: An optimal definition for ocean mixed layer depth, J. Geophys. Res., 105 (7), 16803–16821, 2000.
- Kraus, E. B.: Diapycnal mixing, in: Climate Ocean Interaction, edited by: Schlesinger, M. E., North Atlantic Treaty Organisation, Kluwer Academic Publishers, Dordrecht, 269–293, 1990.
- Lipscomb, W.: Remapping the thickness distribution in sea ice models, J. Geophys. Res., 106, 13989–14000, 2001.
- Lipscomb, W. H. and Hunke, E. H.: Modeling sea ice transport using incremental remapping, Mon. Weather Rev., 132, 1341–1354, 2004.
- Lipscomb, W. H., Hunke, E. C., Maslowski, W., and Jakacki, J.: Improving ridging schemes for high-resolution sea ice models, J. Geophys. Res., 112, C03S91, doi:10.1029/2005JC003355, 2007.
- Lock, A. P., Brown, A. R., Bush, M. R., Martin, G. M., and Smith, R. N. B.: A new boundary layer mixing scheme. Part I: Scheme description and SCM tests, Mon. Weather Rev., 128, 3187–3199, 2000.
- Madec, G.: NEMO coding rules, http://www.nemo-ocean.eu/content/download/250/1629/file/coding_rules.OPA9.pdf, 2006.
- Madec, G.: NEMO ocean engine, Note du Pole de modélisation, Institut Pierre-Simon Laplace (IPSL), France, No 27, ISSN No 1288-1619, 2008.
- Madec, G., Delecluse, P., Imbard, M., and Levy, C.: OPA 8.1 ocean general circulation model – reference manual, Note du Pole de modélisation, Institut Pierre-Simon Laplace (IPSL), France, 11, 1998.
- Martin, G. M., Ringer, M. A., Pope, V. D., Jones, A., Dearden, C., and Hinton, T. J.: The Physical Properties of the Atmosphere in the New Hadley Centre Global Environmental Model (HadGEM1). Part I: Model Description and Global Climatology, J. Climate, 19, 1274–1301, 2006.
- Martin, G. M., Bellouin, N., Collins, W. J., Culverwell, I. D., Halloran, P., Hardiman, S., Hinton, T. J., Jones, C. D., McLaren, A., O'Connor, F., Rodriguez, J., Woodward, S., et al.: The HadGEM2 family of Met Office Unified Model climate configurations, Geosci. Model Dev. Discuss., in press, 2011.
- McLaren, A. J., Banks, H. T., Durman, C. F., Gregory, J. M., Johns, T. C., Keen, A. K., Ridley, J. K., Roberts, M. J., Lipscomb, W. H., Connolley, W. M., and Laxon, S. W.: Evaluation of the sea ice simulation in a new coupled atmosphere-ocean climate model (HadGEM1), J. Geophys. Res., 111, C12014, doi:10.1029/2005JC003033, 2006.
- Murray, R. J.: Explicit generation of orthogonal grids for ocean models, J. Comp. Phys., 126, 251–273, 1996.
- Oki, T. and Sud, Y. C.: Design of Total Runoff Integrating Pathways (TRIP)—A Global River Channel Network, Earth Interactions, 2, 1998.
- Paulson, C. A. and Simpson, J. J.: Irradiance measurements in the upper ocean, J. Phys. Oceanogr., 7, 952–956, 1977.
- Randel, W. J. and Wu, F.: A stratospheric ozone profile data set for 1979–2005: Variability, trends, and comparisons with column ozone data, J. Geophys. Res., 112, D06313, doi:10.1029/2006JD007339, 2007.
- Raymond, W. H.: High-order low-pass implicit tangent filters for use in finite area calculations, Mon. Weather Rev., 116, 2132–2141, 1988.
- Rayner, N. A., Parker, D. E., Horton, E. B., Folland, C. K., Alexander, L. V., Rowell, D. P., Kent, E. C., and Kaplan, A.: Global analysis of sea surface temperature, sea ice, and night marine air temperature since the late nineteenth century, J. Geophys. Res., 108, D14, 4407, doi:10.1029/2002JD002670, 2003.
- Rothrock, D.: The energetics of the plastic deformation of pack ice by ridging, J. Geophys. Res., 80, 4514–4519, 1975.
- Schmidt, G. A., Bitz, C. M., Mikolajewicz, U., and Tremblay, L. B.: Ice-ocean boundary conditions for coupled models, Ocean Modell., 7 (1–2), 59–74, 2004.

- Semtner, A. J.: A model for the thermodynamic growth of sea ice in numerical investigations of climate, *J. Phys. Oceanogr.*, 6, 379–389, 1976.
- Shaffrey, L. C., Stevens, I., Norton, W. A., Roberts, M. J., Vidale, P. L., Harle, J. D., Jrrar, A., Stevens, D. P., Woodage, M. ., Demory, M. E., Donners, J., Clark, D. B., Clayton, A., Cole, J. W., Wilson, S. S., Connolley, W. M., Davies, T. M., Iwi, A. M., Johns, T. C., King, J. C., New, A. L., Slingo, J. M., Slingo, A., Steenman-Clark, L., and Martin, G. M.: UK-HiGEM: The New UK High Resolution Global Environment Model. Model description and basic evaluation, *J. Climate*, 22, 1861–1896, 2009.
- Staniforth, A. and Côté, J.: Semi-Lagrangian integration schemes for atmospheric models—A review, *Mon. Weather Rev.*, 119, 2209–2223, 1991.
- Stein, C. A. and Stein, S.: A model for the global variation in oceanic depth and heat flow with lithospheric age, *Nature*, 359, 123–129, 1992.
- van Genuchten, M. Th.: A closed-form equation for predicting the hydraulic conductivity of unsaturated soils, *Soil Sci. Soc. Am. J.*, 44, 892–898, 1980.
- Thorndike, A., Rothrock, D., Maykut, G., and Colony, R.: The thickness distribution of sea ice, *J. Geophys. Res.*, 80, 4501–4513, 1975.
- Uppala, S. M., Källberg, P. W., Simmons, A. J., Andrae, U., Da Costa Bechtold, V., Fiorino, M., Gibson, J. K., Haseler, J., Hernandez, A., Kelly, G. A., Li, X., Onogi, K., Saarinen, S., Sokka, N., Allan, R. P., Andersson, E., Arpe, K., Balmaseda, M. A., Beljaars, A. C. M., Van De Berg, L., Bidlot, J., Bormann, N., Caires, S., Chevallier, F., Dethof, A., Dragosavac, M., Fisher, M., Fuentes, M., Hagemann, S., Hólm, E., Hoskins, B. J., Isaksen, I., Janssen, P. A. E. M., Jenne, R., McNally, A. P., Mahfouf, J.-F., Morcrette, J.-J., Rayner, N. A., Saunders, R. W., Simon, P., Sterl, A., Trenberth, K. E., Untch, A., Vasiljevic, D., Viterbo, P., and Woollen, J.: The ERA-40 re-analysis, *Q. J. Roy. Meteor. Soc.*, 131, 2961–3012, 2005.
- Valcke, S.: OASIS3 User Guide (prism.2-5). PRISM Support Initiative No 3, 68 pp., 2006.
- Valcke, S. and Redler, R.: OASIS4 User Guide (OASIS4.0.2), PRISM Support Initiative Report No 4, 64 pp., 2006.
- Webster, S., Brown, A. R., Cameron, D. R., and Jones, C. P.: Improving the representation of orography in the Met Office Unified Model, *Q. J. Roy. Meteorol. Soc.*, 129, 1989–2010, 2003.
- Wilson, D. R., Bushell, A. C., Kerr-Munslow, A. M., Price, J. D., and Morcrette, C. J.: PC2: A prognostic cloud fraction and condensation scheme. I: Scheme description, *Q. J. Roy. Meteor. Soc.*, 134, 2093–2107, 2008.
- Wood, N., Diamantakis, M., and Staniforth, A.: A monotonically damping, second-order accurate, unconditionally stable, numerical scheme for diffusion, *Q. J. Roy. Meteor. Soc.*, 133, 1559–1573, 2007.