

Super-parameterization in ocean modeling: Application to deep convection

Jean-Michel Campin^{*}, Chris Hill, Helen Jones, John Marshall

Department of Earth, Atmospheric and Planetary Sciences, Massachusetts Institute of Technology, Cambridge, MA 02139, USA

ARTICLE INFO

Article history:

Received 29 December 2009

Received in revised form 14 October 2010

Accepted 15 October 2010

Available online 23 October 2010

Keywords:

Multi-scale modeling

Ocean modeling

Deep convection

ABSTRACT

We explore the efficacy of “super parameterization” (**SP**) in ocean modeling in which local 2-d non-hydrostatic plume-resolving fine-grained (**FG**) models are embedded at each vertical column of a coarse-grained (**CG**) hydrostatic model. A general multi-scale algorithm is described in which tendencies from the **FG** models are projected onto the **CG** model which in turn constrains the average state of the **FG** models, coupling the two models together. The approach is tested in the context of models of open ocean deep convection and compared with a pure hydrostatic, coarse resolution model using convective adjustment (**HYD**) and a full 3-d non-hydrostatic plume-resolving simulation (**NH**). The **SP** model is found to be greatly superior to **HYD** at much less computational cost than the fully non-hydrostatic calculation.

© 2010 Elsevier Ltd. All rights reserved.

1. Introduction

Turbulent mixing plays a central role in setting the stratification of the upper ocean both in open basins and in coastal areas. Models used to parameterize turbulence in the surface mixed-layer are based on “boundary layer” representations (e.g., Mellor and Yamada, 1974; Price et al., 1986; Large et al., 1997) and perform best if the stratification is not too strong and the flow remains highly turbulent. Models used to parameterize turbulence below the surface mixed layer are based on “wave-wave interaction models” (e.g., Müller et al., 1986) and work well in weakly turbulent environments, such as the ocean thermocline. However it is quite clear that such models are not appropriate in the near field of energetic forcing such as just below the mixed-layer and the benthic boundary layer. Contrary to the traditional notion of the mixed-layer base as a boundary between quiescent and turbulent regions, turbulent mixing does not immediately drop to the small interior values at the base of the mixed layer. Instead, there is a transition layer across which mixing rates decay with depth from high values at the surface to extremely low values in the interior. The penetration depth is typically a few tens of meters, but can occasionally extend down a hundred meters or more. Such subtleties are extraordinarily difficult—perhaps impossible—to capture using conventional turbulence and mixing models.

Here we explore a new approach to the parameterization of subgridscale processes in ocean models which offers a route forward on from 1-d representations. It attempts to resolve, rather than parameterize, small-scale processes. Fine-grid non-hydro-

static models (**FG**) are embedded into a coarse-grid hydrostatic model (**CG**). Rather than employ a one-dimensional (1-d) parameterization of small-scale processes, the **CG** model includes tendencies from an array of **FG** models running at each horizontal grid-point of the **CG**, as sketched in Fig. 1. The **FG** models attempt to resolve, rather than parameterize, the major part of the turbulent mixing processes.¹ The coupling between **FG** and **CG** is two-way—the **FGs** receive information about the large-scale shear and temperature/salinity (θ/S) environment from the **CG**, compute momentum and θ/S tendencies by integrating forward **FG** submodels, and then return the tendencies to the **CG**. In this way we obviate the need for a conventional 1-d closure.

Our approach is motivated by the belief that in addition to traditional 1-d boundary layer approaches to the parameterization of turbulence in the surface mixed-layer (e.g., Kraus and Turner, 1967; Mellor and Yamada, 1982; Price et al., 1986; Large et al., 1994; Nurser, 1996) it is important to explore alternative routes that take advantage of modern massively parallel computers, permitting aspects of small-scale motions to be resolved rather than their transfer properties represented parametrically. However, a brute-force approach in which plume-resolving resolution is employed everywhere cannot yet be fully realized because of limitations in computational resources. Instead, here we experiment with high-resolution local sub-models that, initially at least, are run as vertical 2-d slices at each horizontal grid column of the **CG**.

In order to develop an appropriate algorithmic approach to the embedding of non-hydrostatic submodels in a hydrostatic

¹ A very fine resolution would be necessary if the **FG** model were to resolve the full spectra of turbulence down to the Kolmogorov scale. Thus turbulent viscosity and mixing coefficients employed in the **FG** remain typically much greater than molecular values. However, for simplicity and to contrast with the **CG** capability, we refer to the **FG** model as “resolving small scale processes”.

^{*} Corresponding author. Tel.: +1 617 253 0098; fax: +1 617 253 4464.

E-mail addresses: jmc@ocean.mit.edu (J.-M. Campin), cnh@mit.edu (C. Hill), helen@ocean.mit.edu (H. Jones), jmarsh@mit.edu (J. Marshall).

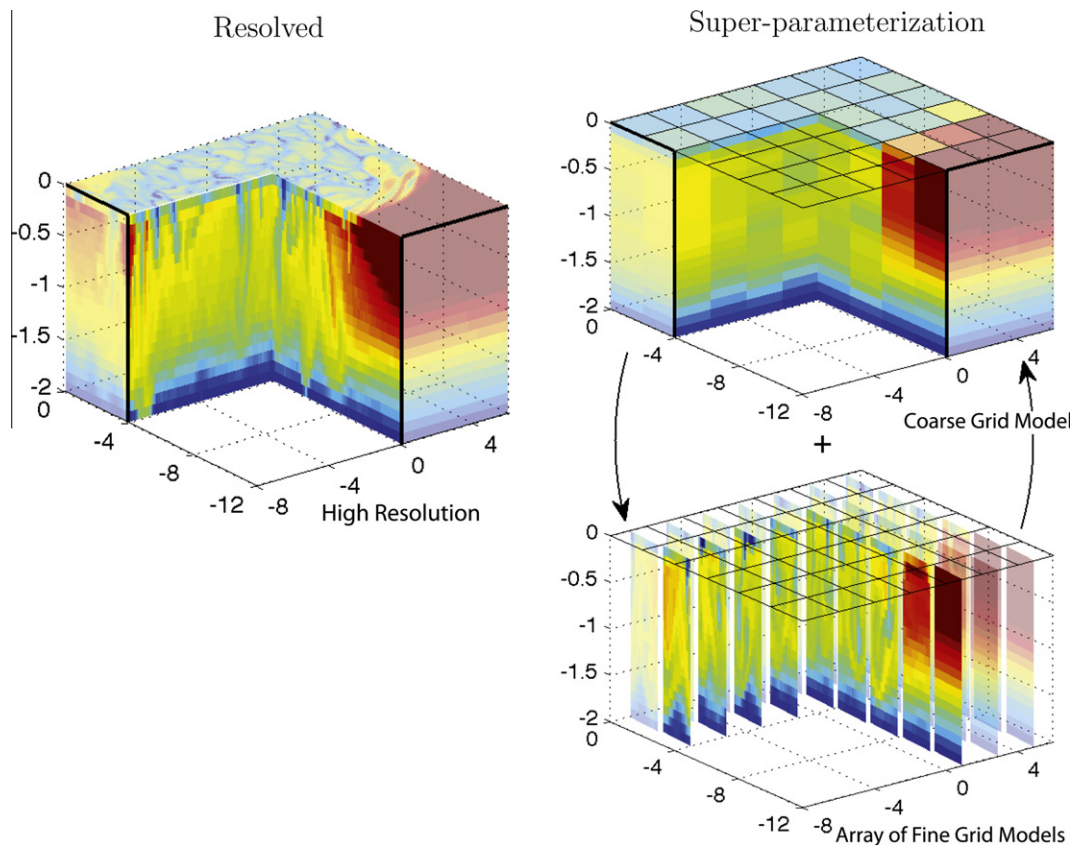


Fig. 1. 3-d view of the temperature field (red is warm, blue is cold) in two simulations of chimney convection similar to Jones and Marshall, 1993. Left side: from a high resolution simulation which resolves small scale plume processes. Right side: from a super-parameterized model in which a coarse-grained (CG) large-scale model (top right panel) representing balanced motion is integrated forward with embedded fine-grained (FG) (bottom right panel) running at each column of the large-scale grid. The FG is non-hydrostatic and attempts to resolve the small-scale processes. The FG's and the CG are integrated forward together and exchange information following the algorithm set out in Section 3.

large-scale model, we focus on the interplay of plumes and baroclinic instability in the context of open ocean deep convection. We employ the idealized configuration introduced by Jones and Marshall (1993). It is shown that the use of non-hydrostatic FG models embedded in a hydrostatic CG model is able to capture key aspects of the evolving flow at a computational cost which is orders of magnitude smaller than the 3-d plume resolving model over the entire domain. In particular the fidelity of the solutions is clearly superior to that obtained using a convective adjustment scheme.

The approach explored here is new to oceanography but has been, and is being vigorously pursued in atmospheric modeling, where it goes under the name of “super-parameterization” (SP=CG + FG). Parameterization of convective cloud processes by overset grid (Meakin, 1999) explicit models is an active area of research following Grabowski (2001). Khairoutdinov and Randall (2001) explored the impact of introducing such a scheme into the NCAR Community Climate System Model. Subsequently several authors have reported on research that expands on these ideas in global, meso-scale and idealized atmospheric models (see for example Khairoutdinov et al., 2005; Wyant et al., 2006; Tao et al., 2009; Grabowski, 2006; Majda, 2007; Khairoutdinov et al., 2008). Recently Grabowski and collaborators have also examined multiscale approaches to cloud droplet growth in the presence of turbulence (Wang et al., 2005). Application of horizontal overset mesh ideas to coupling of land-surface models to atmospheric models has been explored by Molod et al. (2004). Subsequent work has looked at overset vertical meshes for atmospheric boundary layer physics (Molod, 2009). Other research (Freitas et al.,

2006a,b) has used the super-parameterization concept more directly to tackle the interaction of land-surface fires and the atmosphere, again using overset methods.

Our paper is set out as follows. In Section 2 we describe the numerical simulation of open ocean deep convection that provides our test bed to explore SP. In Section 3 we introduce the numerical scheme used to couple CG and FG models. Section 4 evaluates the fidelity of our SP approach. In Section 5 we discuss and conclude.

2. Target application: simulation of open-ocean deep convection

We have developed the super-parameterization algorithm in the context of the plume-resolving “chimney” problem of Jones and Marshall (1993). This has been the focus of many parameterization attempts reviewed in Marshall and Schott (1999). Fig. 2a shows the development of open-ocean deep convection in a very high resolution (100 m in the horizontal and vertical) non-hydrostatic (NH) model which simultaneously resolves convective plumes and large-scale geostrophically balanced motion. The domain is doubly periodic and has a horizontal scale of $40 \text{ km} \times 40 \text{ km}$ with uniform depth of 2 km. An 800 Wm^{-2} heat loss is applied at the surface over a centered circular region with radius of 10 km. In all cases, the fluid starts at rest with a uniform stratification of Brunt-Väisälä frequency $N = 3 \times 10^{-4} \text{ s}^{-1}$. In all calculations presented here a linear equation of state is used with a thermal expansion coefficient of $2 \times 10^{-4} \text{ K}^{-1}$. There is no salt. A background constant Coriolis parameter of $f = 10^{-4} \text{ s}^{-1}$ is used.

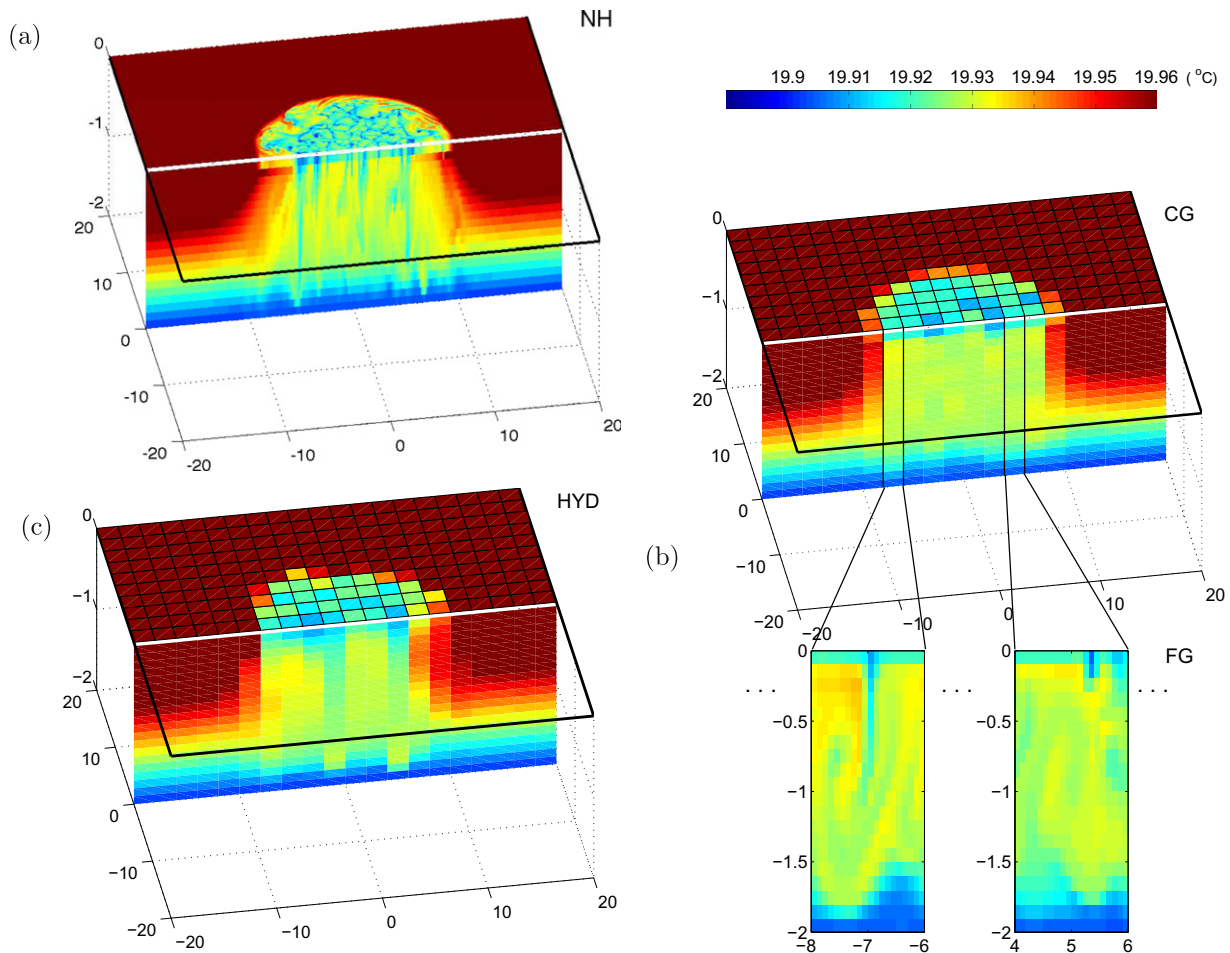


Fig. 2. Overview of the temperature field in the three models considered in this paper represented in a 3-d cut out view. Panel (a) (top left) shows the temperature field in a high-resolution non-hydrostatic simulation of chimney convection by Jones and Marshall (1993). Cooling of a weakly stratified ocean at rest over a disc (shown in the half plan view of the surface temperature field) generates convection (as seen in the vertical section of temperature). The axes are labelled in km. Since the horizontal resolution is 100 m, the 400×400 horizontal domain has a dimension of (40×40) km. Panel (b) (right side) shows the same fields from the **CG** model run at a horizontal resolution of (2×2) km with embedded **FG** models at each horizontal grid point. Vertical slice from two **FG** models illustrates how the embedded non-hydrostatic **FG** models resolve plume dynamics with a resolution of (100×100) m. Panel (c) (bottom left) shows the same fields but from the coarse resolution (2×2) km hydrostatic model (**HYD**) with convective adjustment.

Simulations extend for a four day period, early in which plumes of cooled water develop and create a roughly cylindrical mixed patch of higher density fluid (see Fig. 2a). The resulting radial density gradient induces a circulation which becomes baroclinically unstable, generating geostrophically balanced eddies that are shed from the convection region, carry the mixed dense water away from the cooling patch, and bring stratified fluid from the periphery. Dynamics on the scale of the plumes are non-hydrostatic and require higher resolution than the geostrophically balanced deformation scale eddies that emerge later in the simulation.

Because of this relatively clear scale separation between plumes and larger scale balanced motions, this chimney convection problem constitutes a good test case for a multi-grid approach where the two different scales are resolved by two model components at different grid resolution: The “super-parameterized” model, **SP**, introduces a two-dimensional (vertical slice) plume-resolving model, **FG**, in each grid column of a 2×2 km coarse-grid hydrostatic model, **CG**, which cover the full domain and resolves the geostrophic eddy scale. The **CG** and **FG** exchange information as described in Section 3. A time-step of the **SP = CG + FG** system involves stepping forward the hydrostatic **CG** equations and the non-hydrostatic **FG** equations. However, because the **FG** model is two-dimensional and local to a grid column it is computationally cheaper than the full **NH** model (see below).

With **NH** as our reference, we now wish to evaluate the difference between (i) a conventional hydrostatic model (**HYD**) using the same resolution as the **CG** component but making use of a convective adjustment scheme (which mixes unstable water through an enhanced vertical diffusivity—see Klinger et al., 1996) to represent the effects of plume dynamics and (ii) the super-parameterized model **SP = CG + FG**; in this latter model, **FG** takes the place of the convective adjustment scheme of **HYD**.

In the calculations presented here, key mesh parameters of the three model setups are given in Table 1. In order to simplify the

Table 1

The three model configurations used in this study: coarse resolution hydrostatic model (**HYD**), high resolution non-hydrostatic model (**NH**) and super-parameterization model (**SP = CG + FG**).

	HYD	SP = CG + FG	NH
Horizontal resolution	2 km	CG: 2 km	100 m
Horizontal grid dimension	20×20	FG: 100 m CG: 20×20 (FG: 20×1) $\times 400$	400×400
Vertical resolution	100 m	CG: 100 m FG: 100 m	100 m
Total number of grid cells	8000	168,000	3,200,000

comparison between our simulations several parameters are identical across models. It is important to realize, however, that our super-parameterization algorithm does not require this. For example, all models used here have the same vertical resolution and use the same time-steps. In the configuration chosen the **NH** model requires 400 times the number of grid cells as the **HYD** model due to its plume resolving horizontal resolution. The **SP = CG + FG** scheme falls between the two, with 21 times the number of grid cells as **HYD**.

Results from the three models are shown in Fig. 2. Fig. 2a shows the high-resolution non-hydrostatic reference simulation (**NH**). Cooling of an ocean at rest over a disc (seen in the half plan view of the surface temperature field) generates convection (as seen in the vertical section of temperature). Fig. 2c shows the same fields from the coarse resolution hydrostatic **HYD** model with convective adjustment. Fig. 2b shows the temperature field from the two components of the **SP** model. In the upper panel, the **CG** with embedded **FG** models at each horizontal grid point is shown. In the lower panel, plumes resolved in two of those embedded 2-d plume models can be seen.

A test of the super-parameterization approach is the degree to which the large-scale evolution of the solution, as revealed in Fig. 2a, is better captured in **SP = CG + FG** (Fig. 2b) than **HYD** with a conventional parameterization (Fig. 2c). This will be evaluated in Section 4.

3. FG ↔ CG algorithm

3.1. Overview

The idea behind super-parameterization is that the coarse-grained model (the **CG**) carries information about the large-scale dynamics, represented by variables with the subscript ‘c’. Thus, for example, **CG** might be a primitive equation model running on a grid with a horizontal spacing of a few kilometers. The dynamics of small-scale motions—the scales we wish to parameterize—is represented by local fine scale models (the **FGs**), represented by variables with subscript ‘f’, and run at each vertical grid column of the **CG**. Like many vertical (1-d) parameterizations, a scale separation is assumed between the **CG** resolved motions and sub-grid scale (SGS) motions that **FG** intends to resolve. There is no limitation on the coarseness of the **CG** model and a 2-d **FG** slice need not necessarily extend to a full **CG** grid-cell size, as long as it produces reliable SGS averages.

The algorithm employed can be decomposed into four steps:

1. integrate each **FG** forward to compute tendencies on the fine grid
2. average **FG** tendencies to the coarse grid ‘c’
3. integrate **CG** forward incorporating coarse grid averaged tendencies from **FG**
4. adjust state variables ($\mathbf{v}, \theta, \dots$) of each **FG** model to make them consistent with the corresponding coarse grained vertical profile.

An adjustment is used in step 4 to ensure that fine-scale variables interpolated to the coarse model are the same as coarse-grained variables, thus keeping the ‘c’ and ‘f’ variables consistent with one-another. This is similar to the “gridalt” technology described in Molod (2009). Both momentum and tracer variables are treated in the super-parameterization. Note that different time-steps can be taken in **CG** and **FG** models. However, in our application here, only a modest speed-up would be achieved given the load balance of computation between **CG** and **FG** models (see Section 3.3.2).

A super-parameterization approach becomes computationally feasible because the embedded 2-d non-hydrostatic models can be run very efficiently, and so one can afford to integrate a large

array of such models, for example at each horizontal grid-point of the **CG**. One could also contemplate running small 3-d non-hydrostatic submodels, computational resources permitting. We will see that in the present application super-parameterization increases the CPU cost of simulations by a factor of less than one hundred, but can make efficient use of massively parallel computers. In addition, super-parameterization makes it possible for an ocean model to converge to a fully 3-d non-hydrostatic model as the horizontal grid spacing of the model is decreased.

3.2. Implementation details

We now describe one particular implementation of **SP** that we have used to explore the potential of this approach in ocean modeling. The **HYD** and **NH** model are standard configurations of the MITgcm (Marshall et al., 1997a,b, 1998). Here we focus on outlining the approach used in the **SP (= CG + FG)** model. All models are based on configurations of the MITgcm software. Each model steps forward prognostic equations for potential temperature (θ), two horizontal components of velocity (u and v) and solves a vertically integrated implicit equation for the surface elevation field (η). In addition the non-hydrostatic model, **NH**, and non-hydrostatic sub-model, **FG**, step forward a prognostic equation for vertical velocity (w) and solve for a non-hydrostatic pressure field (P_{nh}). Hereafter the vector notation $\mathbf{v} = \mathbf{u} + w\hat{\mathbf{z}} = u\hat{\mathbf{x}} + v\hat{\mathbf{y}} + w\hat{\mathbf{z}}$ is used to represent the three component velocity $(\tilde{u}, \tilde{v}, \tilde{w})$ along three orthogonal axes $\tilde{\mathbf{x}}, \tilde{\mathbf{y}}$ and $\tilde{\mathbf{z}}$.

3.2.1. Tracer equations

We formulate **SP** as a coupled system in which two sets of equations are stepped forward, one set for variables $(\theta_c, u_c, v_c, \eta_c)$ in **CG** and one set for variables $(\theta_f, u_f, v_f, \eta_f, (P_{nh})_f)$ in the **FG** sub-models. The fine-grid 2-d (x, z) model **FG** is configured in a doubly periodic domain on an f-plane and it is assumed that all gradients in the y-direction vanishes. In our scheme the coupling between **CG** and **FG** occurs through the prognostic equations for θ, u, v . For example, the temperature equation has the form

$$\text{Fine : } \frac{\partial \theta_f}{\partial t} = -\mathbf{v}_f \cdot \nabla \theta_f \quad (1)$$

$$\text{Coarse : } \frac{\partial \theta_c}{\partial t} = -\mathbf{v}_c \cdot \nabla \theta_c + F_\theta^{SGS} \quad (2)$$

The term F_θ^{SGS} represents sub-grid scale (SGS) forcing effects that are calculated from **FG** components (1) and then averaged to the corresponding **CG** column:

$$F_\theta^{SGS} = \left[\frac{\partial \theta_f}{\partial t} \right]_c \quad (3)$$

In addition, each **FG** sub-model is subject to the constraint:

$$[\theta_f]_c = \theta_c(z) \quad (4)$$

where the $[\]_c$ operator is defined as the horizontal average over the small **FG** domain and maps to the corresponding **CG** water column in which it is embedded. In our calculations (see Section 4) the surface forcing is directly applied to each **FG** sub-model and is either prescribed or computed from local surface conditions, depending on which tracer is considered (temperature or passive tracer). The average surface forcing is transmitted to the **CG** component through the **SP** model mapping operator $[\]_c$, by F_θ^{SGS} as defined in Eq. (3).

In addition to the prognostic Eq. (1), the constraint (4) is applied at the beginning of a new time-step by adjusting the **FG** field θ_f to a new value θ_f^* , to ensure that the mean vertical profiles in the **FG** sub-models match the vertical profile of the corresponding **CG** water column:

$$\theta_f^*(x) = \text{Fct}(\theta_f(x), [\theta_f]_c, \theta_c) \quad (5)$$

where “Fct” represents an appropriate mapping function and x is the horizontal coordinate in **FG**. Use of the simplest mapping function

$$\theta_f^*(x) = \theta_f(x) - [\theta_f]_c + \theta_c$$

can produce unphysical extrema in the **FG** temperature field θ_f^* , which lie outside the range of the original temperature bounds of both **FG** and **CG** solutions. Instead, a linear mapping that prevents false extrema is used:

$$\theta_f^*(x) = A + (\theta_f(x) - A) \times \frac{\theta_c - A}{[\theta_f]_c - A}$$

with $A = \min(\theta_c, \min_x \theta_f)$ if $\theta_c < [\theta_f]_c$
and $A = \max(\theta_c, \max_x \theta_f)$ if $\theta_c > [\theta_f]_c$

In the algorithm described by Grabowski (2001), the coupling between **CG** and **FG** appears as a relaxation term in Eqs. (1) and (2) (his Eqs. 3a and 3b). This becomes similar to the present formulation (our Eqs. (3) and (4)) once the relaxation time-scale is set equal to the model time step, as discussed in Khairoutdinov et al. (2005).

3.2.2. Momentum equations and orientation of 2-d FG models

The algorithm that applies to the momentum equations is similar to that described above for temperature, but there are several differences. Since only one horizontal dimension (x_f) is represented in the **FG** model, the horizontal momentum equation is much simpler in the cross plane direction (y_f) (where there is zero gradient) than along x_f direction. However, the orientation of the **FG** x -axis does not need to coincide with the **CG** grid axes, and can be selected in a physically sensible way. In order to capture effects of the large scale vertical shear present in the **CG** solution, the orientation of the **FG** model is allowed to evolve and to align along the direction of maximum vertical shear. In practice, the orientation is relaxed towards a target direction, α_{tg} , defined as follows. If α is the angle relative to the **CG** x -axis, the function:

$$[\text{RMS}(\alpha)]^2 = \frac{1}{H} \int_H [(u_c - \bar{u}_c^z) \cos(\alpha) + (v_c - \bar{v}_c^z) \sin(\alpha)]^2 dz$$

reaches a maximum for $\alpha = \alpha_{tg}$, corresponding to the target orientation α_{tg} . Here $\mathbf{u}_c = (u_c, v_c)$ is the horizontal **CG** velocity and $\langle \cdot \rangle^z = \frac{1}{H} \int_H (\cdot) dz$ represents the vertical averaging operator over the total depth “ H ”. The **FG** orientation vector V_f which defines \hat{x}_f is then relaxed towards the target orientation vector V_{tg} on a time scale τ_{rot} :

$$\frac{\partial}{\partial t} V_f = (V_{tg} - V_f) / \tau_{rot} \quad \text{where}$$

$$V_{tg} = \text{RMS}(\alpha_{tg}) \left[\cos(\alpha_{tg}) \hat{x}_c + \sin(\alpha_{tg}) \hat{y}_c \right] \quad (6)$$

Note that both the magnitude $|V_f|$ (same units as V_{tg} , in m/s) and the direction of the vector V_f evolve with time following the target orientation vector V_{tg} .

Using a relaxation time scale τ_{rot} shorter than the typical time-scale of coarse-grid flow adjustment (e.g., an inertial period, in our case ~ 17 h) allows **FG** models to track the direction of **CG** large vertical shear. A time scale much longer than the model time step (in our case $\Delta t = 60$ s) prevents the occurrence of sudden changes of orientation due to numerical noise in the **CG** model. A time scale $\tau_{rot} = 1$ h was found to be satisfactory in **SP** simulations. Fig. 3 plots the **FG** orientation vector, V_f , field after four days of simulation. We see that it aligns along the “rim current” associated with the strong thermal wind shear on the periphery of the dense mixed patch where sea level is depressed, as discussed in Jones and Marshall (1993).

The remaining aspect of the momentum coupling closely follows the tracer algorithm: the averaging operator $[\cdot]_c$ takes into account the orientation (α) of the **FG** model:

$$[\mathbf{u}_f]_c = \mathbf{u}_c(z) \quad (7)$$

which can be written:

$$[\cos \alpha u_f - \sin \alpha v_f]_c = u_c(z)$$

$$[\sin \alpha u_f + \cos \alpha v_f]_c = v_c(z)$$

The SGS forcing term is computed from the advective tendency only:

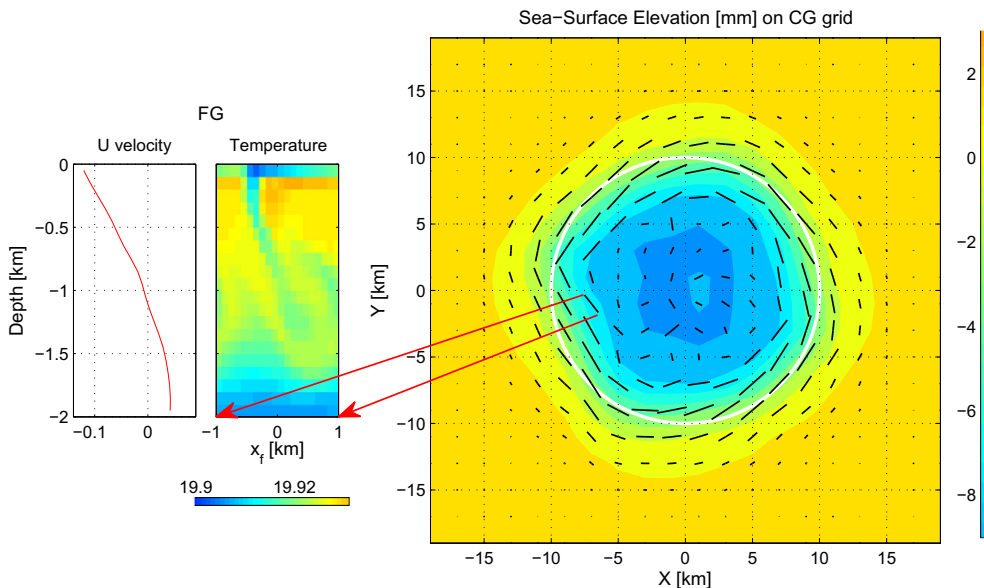


Fig. 3. Right panel: Sea surface height (in mm) from the **CG** model and **FG** orientation vector, V_f , defined in Eq. (6) after 4 days of simulation. The ocean is cooled within the white circle with a heat loss of 800 W/m^2 . The left panel shows results from one **FG** model at the location indicated by the red arrows: The along section velocity mean vertical profile (m/s) from the **CG** is represented next to temperature field in x_f - z plane, at the same time.

$$F_{\mathbf{u}}^{\text{SGS}} = \left[-\tilde{\mathbf{v}}_f \cdot \nabla \tilde{\mathbf{v}}_f \right]_c \quad (8)$$

since the other (linear) terms are explicitly represented in the **CG** model. The momentum equations on the “f” and “c” grids are:

$$\text{Fine : } \frac{\partial \tilde{\mathbf{v}}_f}{\partial t} = -\tilde{\mathbf{v}}_f \cdot \nabla \tilde{\mathbf{v}}_f - 2 \tilde{\Omega} \times \tilde{\mathbf{v}}_f - \frac{1}{\rho_0} \nabla (P_h + P_{nh})_f + D_f \quad (9)$$

$$\text{Coarse : } \frac{\partial \tilde{\mathbf{u}}_c}{\partial t} = -\tilde{\mathbf{v}}_c \cdot \nabla \tilde{\mathbf{u}}_c - 2 \tilde{\Omega} \times \tilde{\mathbf{u}}_c - \frac{1}{\rho_0} \nabla_h (P_h)_c + D_c + F_{\mathbf{u}}^{\text{SGS}} \quad (10)$$

where D is the dissipation term and P_h is the hydrostatic pressure including the surface pressure contribution.

3.2.3. Time stepping

The coarser horizontal resolution of the **CG** model allows one to use a longer time-step than in the **FG** component. However, to simplify the comparison of the different simulations, the same time-step is used in all models (**NH**, **HYD** and **SP**) and components (**CG**, **FG**). In the present implementation, the two components are stepped forward sequentially: first the orientation of the **FG** model is updated (Eq. 6) using the current **CG** velocity field and then each of the **FG** instances is advanced in time (Eqs. 1 and 9). This allows one to compute the resulting sub-grid scale contribution (Eqs. 3 and 8) required to step forward the **CG** prognostic variables (Eqs. 2 and 10). A global budget of volume and tracer in the **CG** model shows perfect conservation.

3.3. Super-parameterization coupled computation—software design and cost

3.3.1. Software implementation

We use an acyclic graph coupling approach to link the **CG** and **FG** sub-models in **SP**. Both the **CG** and **FG** sub-models are viewed as “components” at the leaves of a two branch tree, or acyclic graph, that are distributed over multiple processors using the Earth System Modeling Framework (Hill et al., 2004; Suarez et al., 2007) software library. The arrangement is illustrated in Fig. 4.

The overall computation is supervised by a root component, r , that spans a set of one or more processes and/or threads $\{p_r\}$ and executes concurrently over p_r . Child “components”, c and f , that correspond to submodels **CG** and **FG** are controlled by the root component r . The c and f components are designed to execute concurrently, under the overall control of r , on process/thread sets $\{p_c\}$ and $\{p_f\}$. The sets $\{p_c\}$ and $\{p_f\}$ are subsets of $\{p_r\}$. In the experiments described here we define the c component on a single process that integrates forward the **CG** terms and $(n_x \times n_y)$ f components, each of which integrates forward the **FG** terms for a single grid column. The assignment of process/thread subsets $\{p_c\}$ and passing of information between the c and f components, is orchestrated by the r component. Using this approach the r component can flexibly map the computations to available compute processes/threads.

In the experiments described here we execute the coarse component on a single process. The work of the $(n_x \times n_y)$ f components is

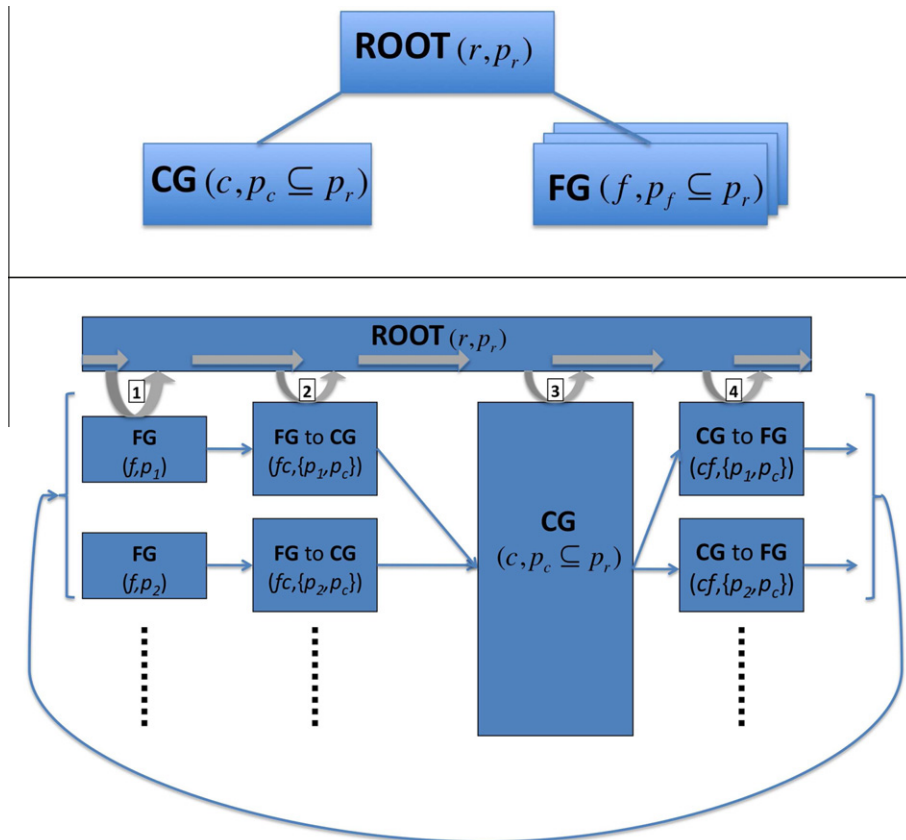


Fig. 4. Top panel: The computational elements of the SP scheme involving **CG** and **FG** models organized as components under a parent component **ROOT** using the Earth System Modeling Framework. The **ROOT** acts to schedule **CG** and **FG**s pairing in a set of computational threads/processes p_r that includes all the **CG** threads/processes p_c and all the **FG** threads/processes p_f . Bottom panel: The sequence of events in a model timestep of **SP** coordinated by **ROOT**. The numbers 1–4 correspond to the algorithm steps 1–4 described in Section 3.1. In step (1) the **FG** models are integrated forward. There are $n_x \times n_y$ independent **FG** models, one for each grid cell. These models are spread over threads/processes p_n , where $n = 1, m$ within the set p_f . In step (2) updated tendencies from the **FG** models are mapped to the **CG** model grid and transferred to the **CG** model threads/processes p_c . In step (3) the **CG** model executes on a set of threads/processes p_c . Finally, in step (4), **CG** model state information is distributed to the **FG** models. Control returns to the **ROOT** component in between each step, allowing it to coordinate **SP** model integration.




Fig. 5. Evolution of (a) global mean tracer concentration and (b) standard deviation within the convective region for the three simulations: **NH** (black line), **SP** (dash red line) and **HYD** (blue line). (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

spread over 4 processes, so that each process handles $(n_x \times n_y) / 4$ f components. Grouping f components provides a way to amortize data copy costs between c and f components. Map and inverse map functions $R: c \rightarrow f, R^{-1}: f \rightarrow c$ are defined using the ESMF regrid library. For this computation the R map is a general scatter that decimates the c index space and communicates different parts of it to different members of p_f that control different f sub-models. The R^{-1} map is a general gather that collects information from the members of p_f and passes an assembled c index space to p_c . The necessary routing and data transfer for this is handled by the ESMF library.

3.3.2. Computational cost

Computational costs for the different simulations **HYD**, **SP** and **NH** can be understood in terms of Table 1. The **HYD** simulation requires solving the hydrostatic equations of motion on a grid of size $n_x \times n_y \times n_z$ at an associated computational cost per grid cell of $\gamma_{\text{dyn}} + \gamma_{\text{conv}}$, where γ_{dyn} is the cost of hydrostatic dynamics and γ_{conv} is the cost of the convective adjustment scheme. The total cost of a time-step, C_{HYD} is then

$$\frac{C_{\text{HYD}}}{n_z} = \gamma_{\text{dyn}} N_h + \gamma_{\text{conv}} N_h \quad (11)$$

where $N_h = n_x \times n_y$. The **SP** simulation cost per time-step can be similarly expressed as

$$\frac{C_{\text{SP}}}{n_z} = \gamma_{\text{dyn}} N_h + \gamma_{\text{nh}} N_h s \quad (12)$$

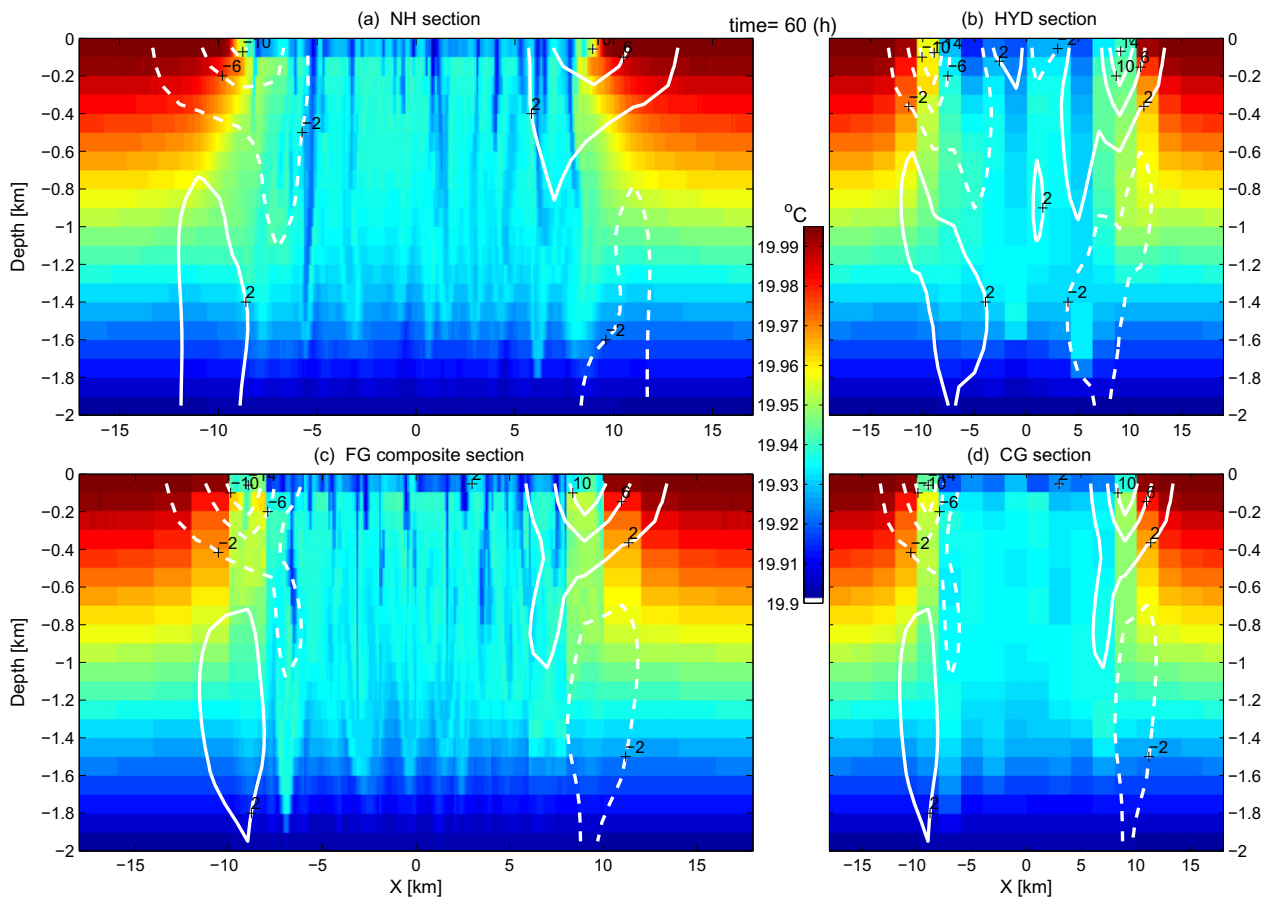


Fig. 6. Temperature distribution in one radial vertical section after 60 h of simulation, for the three cases: (a) Top left: **NH**; (b) Top right: **HYD** with convective adjustment; (c) Bottom left: **SP** showing a composite of all FG instances along the radial section; (d) Bottom right: same as bottom left but averaged on the **CG** grid. The velocity component normal to the corresponding section is represented by white isolines with labels in cm/s.

where γ_{nh} is the cost per grid cell of the non-hydrostatic algorithm and s is the scale ratio of the embedded model resolution to the coarse model resolution. In the present application (Table 1), $s = 2 \text{ km}/100 \text{ m} = 20$. The cost of the **NH** model is then

$$\frac{C_{NH}}{n_z} = \gamma_{nh} N_h s^2 \quad (13)$$

The non-hydrostatic algorithm involves all the hydrostatic terms in γ_{dyn} and some additional terms including an elliptic problem for the time dependent three-dimensional pressure field (solved iteratively), so that the ratio $\delta = \gamma_{nh}/\gamma_{dyn}$ is found to be between 1.6 and 4. In the experiments reported here it is also true that $\gamma_{conv} \ll \gamma_{dyn}$ – the experiments use a simple elevated vertical diffusivity to convectively mix water parcels that are found to be statically unstable with respect to a common reference level (Klinger et al., 1996). We therefore see that the ratio of compute cost of **HYD:SP:NH** is, to first order, given by $1: \delta s: \delta s^2$. From Table 1, $s = 20$, so that given the same timestep is used in all experiments and using the middle range value $\delta = 2.5$ (for this problem) for the non-hydrostatic to hydrostatic computation ratio, the computational cost **HYD:SP:NH** is 1:50:1000 per timestep. This is roughly the ratio of computer time used in the present calculations when running single CPU benchmarks.

We can offset some of the increased computational cost by exploiting parallelism in the **SP** and **NH** simulations that is greater than in the **HYD** simulations. Execution time is a function of both computational cost and the degree to which concurrent computations can be executed efficiently in parallel. In all of **HYD**, **SP** and **NH** there is abundant data parallelism (of the order of the number

of grid cells). However, in practice, process to process and/or thread to thread synchronization, communication and computer resource contention mean that, here, the **HYD** calculation is performed as a single process. In contrast, the **SP** computation is executed as a single process/thread **CG** dynamics computation that executes concurrently with multiple two-dimensional **FG**_{2d} computations. The full, three-dimensional, **NH** computation is also executed in parallel. This can bring the execution time ratios to less than the computational cost ratios (by utilizing more compute resources). In this way the **SP** computation wall-clock time could, with the right hardware mix, be brought closer to the pure **HYD** time. In other words, writing the wall-clock time ratio for **HYD:SP:NH** as $1:C_{SP}:C_{NH}$ the extra parallelism allows $C_{SP} < \delta s$ and $C_{NH} < \delta s^2$. We have not fully explored this optimization at this stage.

4. Results and analysis

For the purposes of comparing the different algorithms we treat the **NH** as our reference solution. To aid in the comparison, a passive tracer has been added in the three models, and evolves according to the same algorithm as that for temperature. The initial tracer distribution is zero everywhere except a concentration of unity is prescribed in the surface level. Moreover, in contrast to temperature, a strong restoring of the tracer to a concentration of unity is applied at the surface. The **HYD** model uses a simple convective adjustment in which vertical diffusivity is set to a large value to mix water parcels that are found to be statically unstable following (Klinger et al., 1996). The magnitude of the enhanced convective

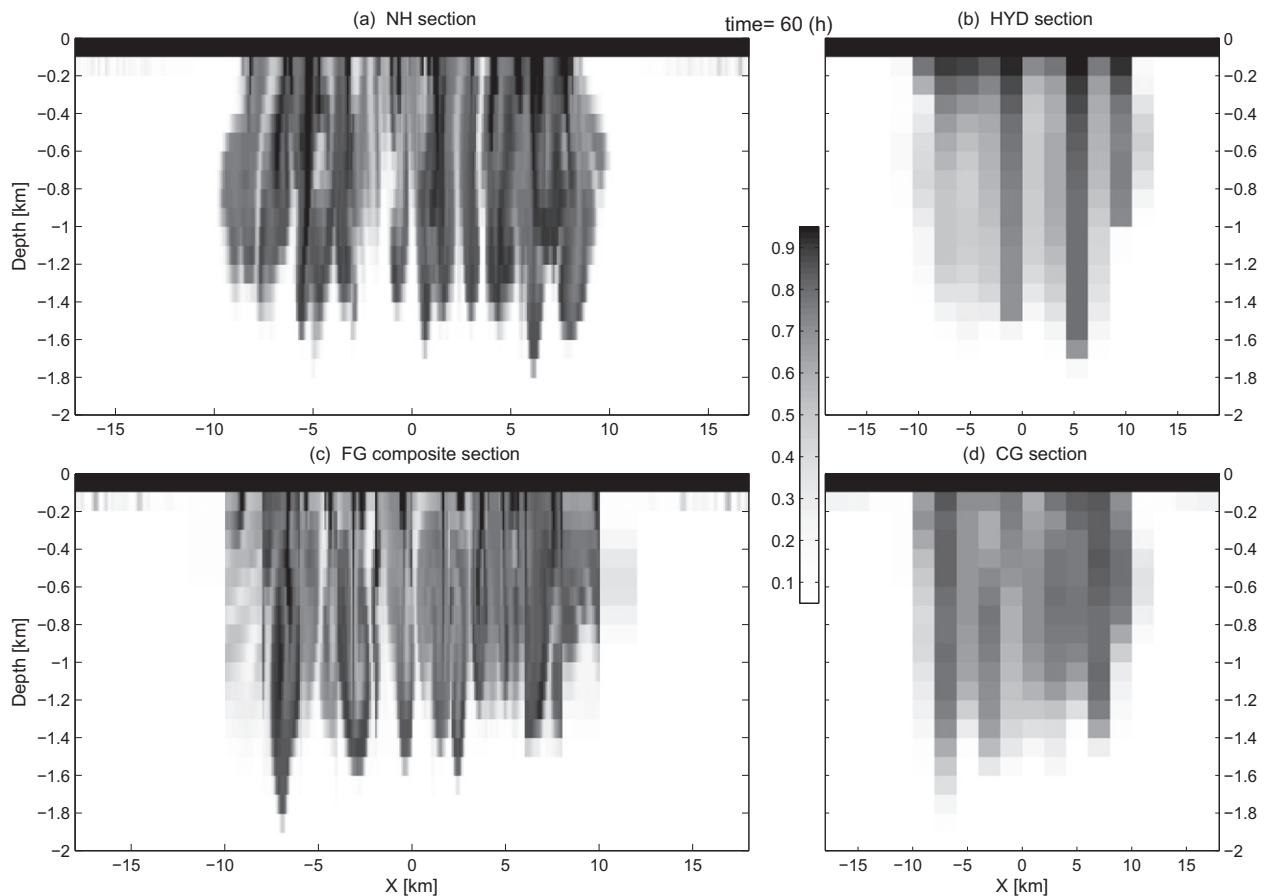


Fig. 7. Tracer distribution in one radial vertical section after 60 h of simulation, for the three cases: (a) Top left: **NH**; (b) Top right: **HYD** with convective adjustment; (c) Bottom left: **SP** showing a composite of all **FG** instances along the radial section; (d) Bottom right: same as bottom left but averaged on the **CG** grid.

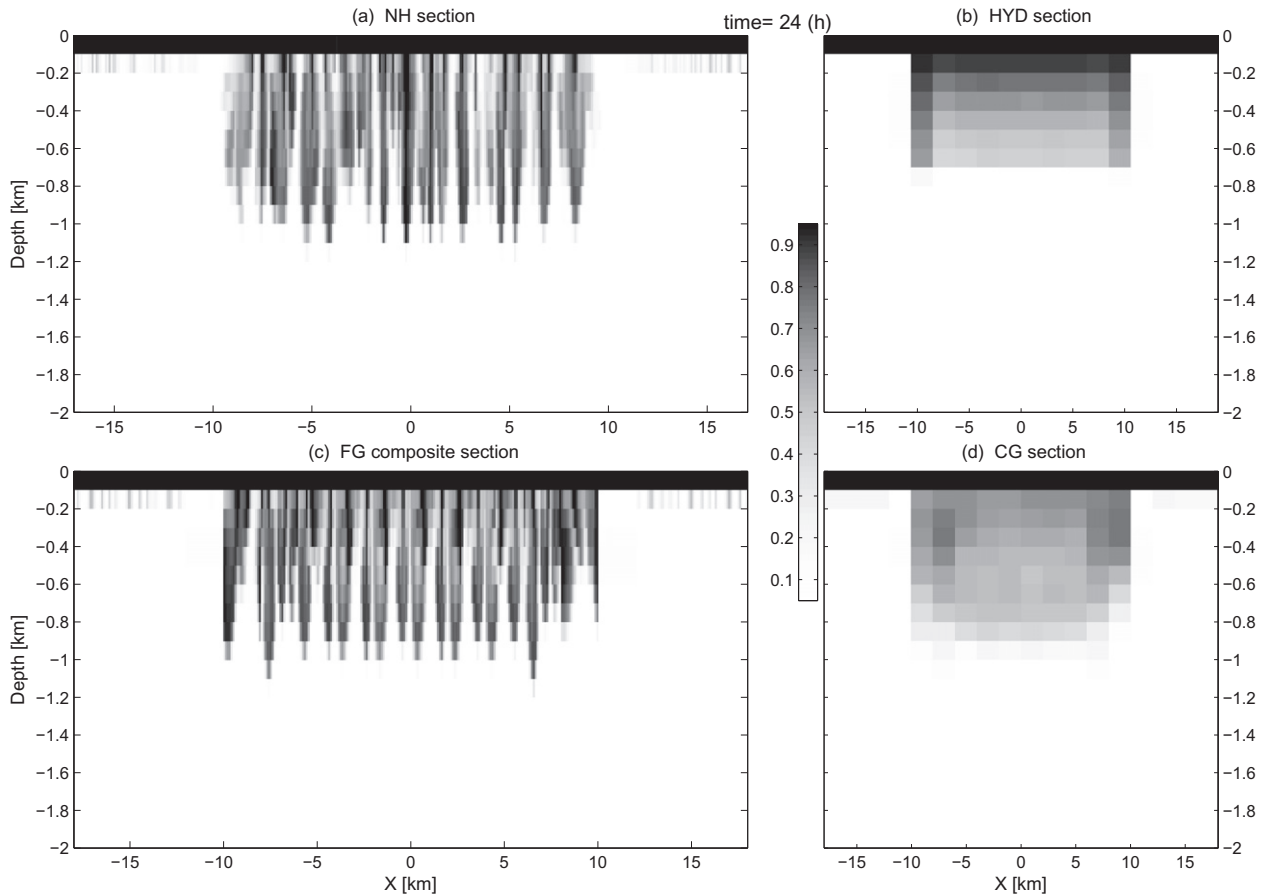


Fig. 8. Tracer distribution along the section shown in Fig. 7 after 24 h of simulation.

mixing in **HYD** has been adjusted to $K_{conv} = 2.5 \text{ m}^2/\text{s}$ to ensure that the integrated tracer content matches the **NH** simulation after 60 h of simulation (see Fig. 5a). We are interested in the degree to which **SP** captures the large-scale evolution of the temperature, currents and tracer distributions of **NH** relative to that of **HYD**.

The evolution of tracer integrated properties shown on Fig. 5 provides an overview answer. The total tracer content in **SP** is barely distinguishable from the reference **NH** simulation (Fig. 5a) whereas the **HYD** model deviates by as much as 10% early on in the simulation (soon after 1 day). The super-parameterization **SP** is also in better agreement (although not perfect) with **NH** than **HYD** regarding the tracer distribution inside the convective patch, as seen from the evolution of the tracer variance inside the cooling region (Fig. 5b).

The following instantaneous vertical sections (Figs. 6–8) give an insight into the three model simulations. After 60 h of cooling at the surface, a mixed patch of dense water occupies the central part of the domain, as seen from Fig. 6 for all three simulations (**NH**, **HYP** and **SP**) with, around the patch, cyclonic circulation (anti clock-wise) near the surface and of opposite sign at great depth. The effect of mixing down surface water properties is obvious from the tracer sections (Fig. 7) either as a consequence of convective adjustment in **HYD** (Fig. 7b) or as a result of explicitly resolved 3-d convective plumes in **NH** (Fig. 7a) or just 2-d plumes in **FG** (Fig. 7c). The convective plumes look quite similar (e.g., reaching similar depth, around 1.5 km) in the **NH** model (Figs. 6a and 7a) and in the composite sections of **FGs** (Figs. 6c and 7c) despite the expected discontinuity at each **FG** model boundary. The over-all similarity also extends to the **HYD** temperature and tracer section

(Figs. 6b and 7b) when compared to **CG** coarse resolution section (Figs. 6d and 7d) as far as mixing depth is concerned. However, early on the **HYD** model is the outlier. For example after 1 day, convection is shallower and more uniform in **HYD** than in the **NH** and **SP**, as is evident from the tracer section shown in Fig. 8. Note that in **HYD** the convective adjustment time scale ($\sim 1/K_{conv}$) is kept constant; see (Send and Marshall, 1995) for a discussion of how this might be varied.

Qualitatively, the **SP** model captures several features of the **NH** simulation which **HYD** poorly represents:

- The evolution of the convective plumes in the **NH** and **FG** simulations is rather similar, both in term of timing and form, as seen in temperature and tracer sections (panels a and c from Figs. 6–8). In particular, since the 2-d **FG** models feel the large scale vertical shear present in the **CG** simulation (see left panel of Fig. 3), they capture the tilt of the plumes observed in the **NH** simulation (e.g., on Figs. 6a and 7a the cold and tracer enriched plume around $x = -5.5 \text{ km}$ is tilted towards the left as it descends).
- Convective plumes can penetrate deeper than their neutrally buoyant level, burrowing into the stratified (colder) water below. This is a short-lived non-hydrostatic process which might have non-negligible effects (e.g., Wang, 2006) even though lighter water generally “rebounds” soon after (seen on animations not shown here). For this reason, it is often ignored in oceanic convection schemes (e.g., Paluszkiwicz and Romea, 1997, p. 113) but is generally accounted for in atmospheric convection (e.g., Boers, 1989). The presence at depth of water

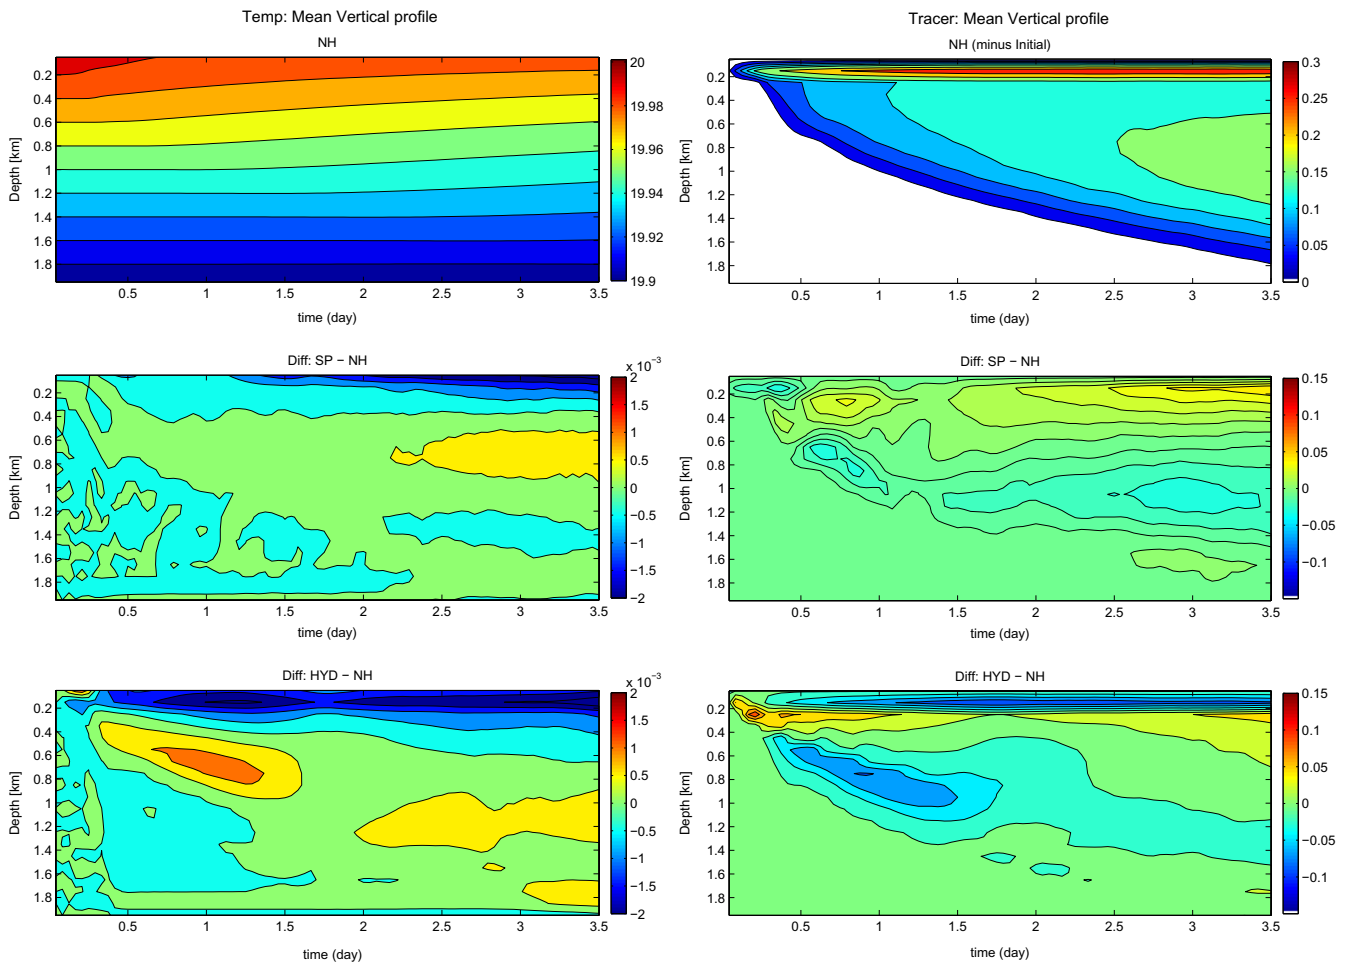


Fig. 9. Evolution of the spatial mean vertical mean profile of Temperature (left) and surface passive tracer (right), with time (in days) as horizontal axis and depth (in km) as vertical axis. Top: **NH** simulation; middle: Difference **SP** minus **NH**; bottom: Difference **HYD** minus **NH**.

warmer than the initial background (seen on Fig. 6a around $x = +6$ km and on Fig. 6c around $x = -7$ km) can be interpreted as a signature of penetrating plumes. It is interesting to note that the convective adjustment solution also shows some warm incursion at depth (Fig. 6b, around $x = +5$ km) which could result from large scale advection.

- The convective plumes overturn the water column and bring cold and tracer depleted water upward, as seen on Fig. 7 in both **NH** and **FG** sections. As a consequence, the resulting coarse-grid average vertical profile is not always monotonic, tracer content being sometimes lower near the surface with patches of higher concentration at depth (Fig. 7d). Use of a diffusive convective adjustment scheme, in contrast, maintains decreasing tracer concentration with depth in each column at all times (see Fig. 7b). This feature has motivated development of non-local convective parameterization schemes (e.g., KPP, Large et al., 1994) but is a natural feature of **SP**.
- In the convective patch, the level just below the surface remains warmer than above and immediately below, in both **NH** and **SP** simulations (Fig. 6), indicating a reduced vertical exchange at the first vertical interface. This is consistent with turbulence scaling where vertical velocity decreases toward zero at the surface. By contrast, no inversion develops in our **HYD** simulation and the sub-surface level is significantly colder (see also Fig. 9). However, these differences may be an artifact of the low vertical resolution near the surface and are expected to become smaller with increasing vertical resolution.

- The horizontal variance inside the convective patch is the most significant difference between **HYD** and either **SP** or **NH** simulations (as seen from Figs. 6–8 and as discussed below from Fig. 10). The nature of the convective adjustment (on–off, depending on a single threshold criteria) is likely to be responsible for too high a sensitivity and may result, after 60 h, the large horizontal variability observed within the patch (Figs. 6b and 7b) which feeds back on the dynamics (note the noisy velocity field at the center of the patch on Fig. 6b which is absent from both **SP** or **NH** results). In contrast, after only 1 day of simulation, the convective adjustment solution **HYD** is horizontally uniform inside the patch (Fig. 8b) except just at the edge where mixing is intensified, but none of the two other simulations (**SP**, **NH**) shows similar contrast. Stochastic convective parameterizations (e.g., Lin and Neelin, 2003) attempt to address this issue.

The time evolution of the mean vertical profiles of temperature and passive tracer (Fig. 9) allows one to quantify the differences between the three simulations. In both cases, **HYD** is much further away from our reference true solution **NH** than is **SP**, especially during the first 2 days of simulation. At this time the convective patch in **HYD** is too shallow and cold with rather too high a concentration of tracer in the upper part of the patch, and too warm with too low a concentration beneath. There are also noticeable differences in the initial response to buoyancy loss: the convective adjustment in **HYD** begins to mix immediately whereas the plume

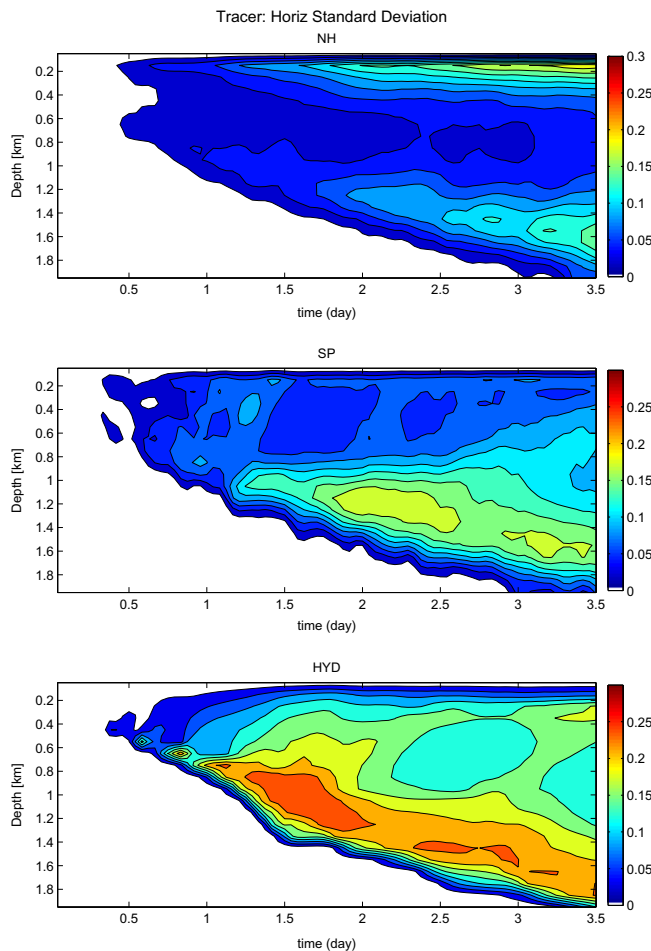


Fig. 10. Evolution of the standard deviation of tracer at each vertical level inside the convective region (where surface cooling is applied), with time (in days) as the horizontal axis and depth (in km) as the vertical axis. All statistics are computed from the average coarse grid (ignoring sub-grid/fine grid variance). (a) Top: **NH** simulation; (b) middle: **SP** simulation; (c) bottom: **HYD** simulation.

dynamics spin up gradually resulting in significant tracer differences during the first 6 h (Fig. 5a). More elaborate mixing parameterizations than those used in **HYD** attempt to include additional prognostic variables (e.g., Mellor and Yamada, 1982) and are likely to better capture such subtleties in the time evolution of deepening convection.

The spatial variability observed inside the convective patch (Fig. 5b) and in particular the horizontal component represented on Fig. 10, reveal significant differences between the three simulations, especially in the case of the **HYD** model, which consistently overestimates the variability at depth; as mentioned earlier, dynamical feedbacks on the convective adjustment scheme may amplify the problem. The **SP** simulation is much closer to the reference experiment **NH**, but also experiences some difficulties in reproducing the horizontal tracer variability (e.g., the near surface maximum in Fig. 10a and a too early development of the deep maximum, Fig. 10b). The use of 2-d **FG** models and an imperfect scale separation are likely contributors.

5. Conclusions

We have described a multi-scale, super-parameterization approach to the representation of open-ocean deep convection in models. By coupling a two-dimensional, plume resolving, non-hydrostatic sub-model at each grid column of a coarser resolution

hydrostatic model, the resulting **SP** model captures much of the richness of a full non-hydrostatic simulation. Features such as tilted plumes, induced by the vertical shear of horizontal currents, and transient overshoots past the level of neutral buoyancy at the base of the mixed layer, are very difficult to capture in hydrostatic, **HYD**, simulations with a parametric representation of convection. Sophisticated one-dimensional vertical mixing schemes, such as KPP, attempt to represent aspects of such processes. However, much detail has to be sacrificed, some of which may be very important for the evolution of bulk properties of the convecting boundary layer. For example, the temporal and spatial inhomogeneities in the 2-d **FG** sub-model produce mixing which is less sensitive to details of the large scale model than the single threshold convective adjustment to which we compared. As a result, the large scale distribution of temperature and tracer within the convective chimney is more horizontally homogeneous. The 2-d **FG** plume-resolving model also impacts the transient evolution of tracers and is likely to be important in rectifying biological and biogeochemical processes, as well as in modulating vertical transport of physical properties such as buoyancy and momentum.

The computational cost of the **SP** approach is significantly less than that of a full 3-d **NH** model. The independent 2-d plume models provide a rich source of parallelism that can be efficiently exploited to amortize computation and reduce the clock-time of the computation. Moreover the 2-d models could be honed to a performance peak which has not been attempted here. Thus the general approach outlined could prove beneficial to emerging petascale ocean applications which are targeting basin and global scale simulations at a few kilometer resolution. Embedding a 2-d non-hydrostatic special-purpose model in such integrations would provide a computationally tractable way to incorporate non-hydrostatic effects in global models in the relatively near future.

The numerical recipe we have outlined could also be applied to the representation of other physical processes where there is a relatively clean separation of scales and where approximate and uncertain parameterizations are currently employed. For example, various parameterizations have been devised to capture the sub-grid scale sinking of dense, salty water during ice formation. A two-dimensional, non-hydrostatic model could provide an alternative rooted in fundamental principles and obviate the need for the specification of tunable vertical depth scales used in such parameterizations (e.g., Nguyen et al., 2009).

More generally the approach provides a route to the coupling of many different types of overset models. It would be possible, for example, to couple a non-eddy ocean circulation model to high resolution quasi-geostrophic submodels to explicitly represent eddy potential vorticity fluxes associated with mesoscale and sub-mesoscale variability rather than parameterize them. The stacked quasi-geostrophic models of the kind considered in Smith and Marshall (2009), could be run at each horizontal grid cell of a course-resolution ocean model. Eddy potential vorticity fluxes would be passed to the course-grid model where they would appear as a source in the residual momentum equation, as outlined in Ferreira et al. (2005) and Ferreira and Marshall, 2006. In return the coarse-grid model would provide the large-scale stratification and horizontal velocity profiles required by the subgrid model.

In the near term we are beginning to examine the inclusion of momentum forcing into the testbed problem we have outlined, and focus the analysis on the vertical flux of momentum as well as the resolved and sub-grid scale kinetic energy budget (see, e.g., the discussion of energetics in Hughes et al., 2009). This will provide an opportunity to clarify the optimal choice of resolution for both **CG** and **FG** components as well as the domain size of the fine-grid model. The next step will be to include 2-d embedded plume models into a realistic regional configuration which, with a larger computation domain, could provide a more precise and

more practical computing efficiency analysis. This plume superparameterized implementation is a precursor to the more elaborate applications outline above.

Acknowledgements

The authors thank anonymous reviewers for constructive comments that improved the manuscript. This work was supported by a grant from ONR and NASA (ECCO-2, MAR Program).

References

- Boers, R., 1989. A parameterization of the depth of the entrainment zone. *Journal of Applied Meteorology* 28, 107–111.
- Ferreira, D., Marshall, J., 2006. Formulation and implementation of a residual-mean ocean circulation model. *Ocean Modelling* 13, 86–107.
- Ferreira, D., Marshall, J., Heimbach, P., 2005. Estimating eddy stresses by fitting dynamics to observations using a residual-mean ocean circulation model. *Journal of Physical Oceanography* 35, 1891–1910.
- Freitas, S.R., Longo, K., Fazenda, A., Rodrigues, L.F., 2006a. Using the superparameterization concept to include the sub-grid plume-rise of vegetation fires in low resolution atmospheric chemistry-transport models. In: *Proceedings of 8 ICSHMO*, pp. 109–113.
- Freitas, S.R., Longo, K.M., Andreae, M.O., 2006b. Impact of including the plume rise of vegetation fires in numerical simulations of associated atmospheric pollutants. *Geophysical Research Letters* 33, L17808.
- Grabowski, W.W., 2001. Coupling cloud processes with the large-scale dynamics using the Cloud-Resolving Convection Parameterization (CRCP). *Journal of the Atmospheric Sciences* 58, 978–997.
- Grabowski, W., 2006. Comment on Preliminary tests of multiscale modeling with a two-dimensional framework: Sensitivity to coupling methods. *Monthly Weather Review* 134, 2021–2026.
- Hill, C., DeLuca, C., Balaji, V., Suarez, M., Da Silva, A., 2004. The architecture of the Earth System Modeling Framework. *Computing in Science and Engineering* 6 (1), 18–28.
- Hughes, G.O., Hogg, A.M., Griffiths, R.W., 2009. Available potential energy and irreversible mixing in the meridional overturning circulation. *Journal of Physical Oceanography* 39, 3130–3146.
- Jones, H., Marshall, J., 1993. Convection with rotation in a neutral ocean a study of open-ocean deep convection. *Journal of Physical Oceanography* 23, 1009–1039.
- Khairoutdinov, M., Randall, D., 2001. A cloud resolving model as a cloud parameterization in the NCAR Community Climate System Model. *Geophysical Research Letters* 28 (18), 3617–3620.
- Khairoutdinov, M., Randall, D., DeMott, C., 2005. Simulations of the atmospheric general circulation using a cloud-resolving model. *Journal of the Atmospheric Sciences* 62, 2136–2154.
- Khairoutdinov, M., DeMott, C., Randall, D.A., 2008. Evaluation of the simulated interannual and subseasonal variability in an AMIP-style simulation using the CSU Multiscale Modeling Framework. *Journal of Climate* 21, 413–431.
- Klinger, B.A., Marshall, J., Send, U., 1996. Representation of convective plumes by vertical adjustment. *Journal of Geophysical Research* 101 (C8), 18,175–18,182.
- Kraus, E., Turner, J., 1967. A one-dimensional model of the seasonal thermocline II: the general theory and its consequences. *Tellus* 19, 98–105.
- Large, W., McWilliams, J., Doney, S., 1994. Oceanic vertical mixing: a review and a model with nonlocal boundary layer parameterization. *Reviews of Geophysics* 32, 363–403.
- Large, W., Danabasoglu, G., Doney, S., McWilliams, J., 1997. Sensitivity to surface forcing and boundary layer mixing in a global ocean model: annual-mean climatology. *Journal of Physical Oceanography* 27, 2418–2447.
- Lin, J.W.-B., Neelin, J.D., 2003. Toward stochastic deep convective parameterization in general circulation models. *Geophysical Research Letters* 30, 1162.
- Majda, A., 2007. Multiscale models with moisture and systematic strategies for superparameterization. *Journal of the Atmospheric Sciences* 64, 2726–2734.
- Marshall, J., Schott, F., 1999. Open ocean deep convection: observations, models and theory. *Reviews of Geophysics* 37, 1–64.
- Marshall, J., Adcroft, A., Hill, C., Perelman, L., Heisey, C., 1997a. A finite-volume, incompressible Navier Stokes model for studies of the ocean on parallel computers. *Journal of Geophysical Research* 102, 5753–5766.
- Marshall, J., Hill, C., Perelman, L., Adcroft, A., 1997b. Hydrostatic, quasi-hydrostatic, and nonhydrostatic ocean modeling. *Journal of Geophysical Research* 102, 5733–5752.
- Marshall, J., Jones, H., Hill, C., 1998. Efficient ocean modeling using non-hydrostatic algorithms. *Journal of Marine Systems* 18, 115–134.
- Meakin, R.L., 1999. Composite overset structured grids. In: Weatherill, N.P., Soni, B.K., Thompson, J.F. (Eds.), *Handbook of Grid Generation*. CRC Press.
- Mellor, G.L., Yamada, T., 1974. A hierarchy of turbulent closure models for planetary boundary layers. *Journal of the Atmospheric Sciences* 31 (4), 1791–1806.
- Mellor, G.L., Yamada, T., 1982. Development of a turbulence closure model for geophysical fluid problems. *Reviews of Geophysics* 20 (4), 851–875.
- Molod, A., 2009. Running GCM physics and dynamics on different grids: algorithm and tests. *Tellus A* 61A (3), 361–395.
- Molod, A., Salmun, H., Waugh, D., 2004. The impact on a GCM climate of an extended mosaic technique for the land-atmosphere. *Journal of Climate* 17 (20), 3877–3891.
- Müller, P., Holloway, G., Henyey, F., Pomphrey, N., 1986. Nonlinear interactions among gravity waves. *Reviews of Geophysics* 24 (4), 493–536.
- Nguyen, A., Menemenlis, D., Kwok, R., 2009. Improved modeling of the arctic halocline with a sub-grid-scale brine rejection parameterization. *Journal of Geophysical Research* 114, C11014.
- Nursler, A.J.G., 1996. A review of models and observations of the oceanic mixed layer. Internal document No 14, Southampton Oceanography Centre, Southampton, UK, 247 pp.
- Paluszkiwicz, T., Romea, R., 1997. A one-dimensional model for the parameterization of deep convection in the ocean. *Dynamics of Atmospheres and Oceans* 26, 95–130.
- Price, J., Weller, R., Pinkel, R., 1986. Diurnal cycling: observations and models of the upper ocean response to diurnal heating. *Journal of Geophysical Research-Oceans* 91 (C7), 8411–8427.
- Send, U., Marshall, J., 1995. Integral effects of deep convection. *Journal of Physical Oceanography* 25, 855–872.
- Smith, K.S., Marshall, J., 2009. Evidence for enhanced eddy mixing at mid-depth in the southern ocean. *Journal of Physical Oceanography* 39, 50–69.
- Suarez, M., Trayanov, A., Hill, C., Schopf, P., Vikhliav, Y., 2007. MAPL: a high-level programming paradigm to support more rapid and robust encoding of hierarchical trees of interacting high-performance components. In: *CompFrame'07: Proceedings of the 2007 symposium on Component and Framework Technology in High-Performance and Scientific Computing*. ACM, New York, NY, USA, pp. 11–20.
- Tao, W., Chern, J., Atlas, R., Randall, D., 2009. A multiscale modeling system: developments, applications, and critical issues. *Bulletin of the American Meteorological Society*, 515–534.
- Wang, D., 2006. Effects of the earth's rotation on convection: Turbulent statistics, scaling laws and lagrangian diffusion. *Dynamics of Atmospheres and Oceans* 41, 103–120.
- Wang, L., Ayala, O., Kasprzak, S., Grabowski, W., 2005. Collision efficiency of hydrodynamically interacting cloud droplets in turbulent atmosphere. *Journal of the Atmospheric Sciences* 62, 2433–2450.
- Wyant, M., Khairoutdinov, M., Bretherton, C., 2006. Climate sensitivity and cloud response of a GCM with a superparameterization. *Geophysical Research Letters* 33, 4.