A suite of ocean-only test cases

Stephen.Griffies@noaa.gov
NOAA/GFDL
Princeton, USA
May 14, 2012

Abstract

This document presents a suite of MOM test cases that employ just the ocean-solo driver. These tests range from a two-dimensional torus to a global ocean model. Many of the tests exercise options for numerical integrity, yet which are not meant for physically relevant simulations. So the user should generally *not* take one of these tests without scrutinizing the configuration.

Contents

1 General comments on the test cases

Test cases in this document are offered as a means to explore various numerical and physical options, thus allowing the user to verify that the code is performing in a manner consistent with that at GFDL. This form of verification is critical as one adopts the code for his or her particular research purposes. The test cases also provide a sense for some of the options available in the code, though by no means are all options exercised in the test cases.

1.1 Regression Testing for Computational Integrity

Output from the test cases provided with the MOM distribution is based on short integrations that verify the computational integrity of a simulation. The associated runscripts are provided that allow the user to rerun the regressions. These *regression test suites* (RTS) aim to verify that the following identities hold, with precision maintained to all computational bits:

- Stopping and then restarting the integration will not change answers. That is, we insist on the identity

$$\text{X day integration} = \text{X/2 day integration} + \text{X/2 day integration}. \tag{1}$$

  This test verifies that all the relevant fields are properly stored in the restart files, and that no spurious reinitialization step is performed during the beginning of the second leg.

- Changing the number of computer processors will not change the answers. This test examines whether the code is properly written for parallel machines. Its satisfaction requires that all message passing be correctly performed so that accessed halo points are filled with their proper values. In the early days of parallel computing with MOM, this test was very tedious to satisfy, since our experience was based on seriel computing. Now, after some years of experience, it is generally straightforward to code in a manner that ensures answers do not change when processor counts change.

Satisfaction of these two tests is critical to maintain computational integrity of the code.

It is important to note that the runscripts for the regressions are distributed with the following settings:

- `diag_step` is set to a small number, such as 1 or 12, in various diagnostic modules. This setting means that various numerical diagnostics are run at a very high frequency. These diagnostics can be expensive. It is therefore strongly recommended that the setting for `diag_step` be increased to a much larger number when running experiments for long periods of simulation time. Otherwise, the model will be unreasonably slow.
- The diagnostic tables are setup to output netCDF diagnostics at a very high frequency, such as daily. Again, this high frequency is unreasonable when running simulations for long periods of time. The output produced will be enormous, and the model will run at a much slower pace. Additionally, a large number of diagnostics are included in the `diag_table`, many of which may not be of interest to the user. Therefore, prior running an experiment, it is important to edit the `diag_table` to refine the desired output.

1.2 Perspectives on the test cases

The material in this document provides a rough guide to the various ocean-solo test cases. We present a sprinkling of model output to allow interested users to run simulations at their institution to verify that the code has been properly ported. Many details of the experiments are omitted, with examination of the supplied runscripts providing more details. Furthermore, a full accounting of the test cases, both their design and simulation characteristics, is beyond the scope of this document. Indeed, a full discussion would constitute a research paper. We thus present a taste, with further details readily found by diving into the model, running experiments, and performing analysis.

Some test cases are based on research experiments conducted at GFDL and elsewhere. They may thus serve as useful starting points for research using MOM. It is nonetheless critical that the user *not* blindly assume that a test case is precisely appropriate for a particular research project. Instead, one is strongly encouraged to scrutinize each option in a test case before concluding that it is relevant.

As there are many options in MOM, it is not feasible to exercise all options with only a few test cases. Hence, some tests are distributed with more options enabled than scientifically appropriate. Conversely, many options are not fully exemplified by the test cases. Omitted options include the *experimental* options sprinkled through MOM, with these options *not* supported for general use. The developers are aware of the limitations in the test cases, but choose to release the incomplete suite of tests in hopes that *something* is preferable to *nothing*.

2 Torus test case

The torus test case `mom4_torus` consists of a flat bottom square domain on the F-plane with doubly periodic boundary conditions. It provides an idealized test bed for implementing various numerical schemes, with applications to tracer advection frequently exercised at GFDL. The flow field can be specified to have uniform horizontal velocity. For the illustration considered here, we consider a zonal flow with $(u, v) = (0.25 \mathrm{m\,s}^{-1}, 0)$.

There are various options in

`mom4/ocean_tracers/ocean_passive_mod`

for specifying the initial profile of idealized passive tracers. We consider two profiles shown in the top panel of Figure 1. The first profile is a square wave or pulse, and the second is a smooth Gaussian bump. The

initial tracer concentrations all live within the range [0,1]. Numerical solutions that fall outside this range constitute spurious unphysical results arising from errors in the advection scheme. A perfect advection scheme would advect the profiles without alteration.

The tracer profiles are advected with the following Courant number

$$\begin{aligned} C &= \frac{u\,\Delta t}{\Delta x} \\ &= \frac{0.25\mathrm{m\,s^{-1}} \times 10800\mathrm{s}}{1.04\times 10^{5}\mathrm{m}} \\ &\approx 0.5 \end{aligned} \tag{2}$$

using the staggered time stepping scheme standard in MOM. For the centred 2nd order and 4th order advection schemes, however, stability requires leap-frog time stepping, in which case the Courant number is ≈ 1 since the time step for advection is $2\Delta t$. We consider the following advection schemes available with MOM4p1:

- **2nd**: The second order centred differences scheme has been the traditional scheme employed in earlier versions of MOM. It is available only with the leap-frog version of MOM, since it is unstable with a forward time step used with the forward-backward method.
- **4th**: The fourth order centred scheme also is available only for the leap-frog version of MOM4p1. This scheme assumes the grid is uniformly spaced (in metres), and so is less than fourth order accurate when the grid is stretched, in either the horizontal or vertical.
- **Upwind**: This is a first order upwind scheme, which is highly diffusive yet monotonic.
- **Quicker**: The quicker scheme is third order upwind biased and based on the work of Leonard (1979). Holland et al. (1998) and Pacanowski and Griffies (1999) discuss implementations in ocean climate models. This scheme does not have flux limiters, so it is not monotonic.
- **MDPPM**: The piece-wise parabolic method with flux limiters.
- **Super B**: Multi-dimensional third order upwind biased approach of Hundsdorfer and Trompert (1994), with Super-B flux limiters.[1]
- **Sweby**: Multi-dimensional third order upwind biased approach of Hundsdorfer and Trompert (1994), with flux limiters of Sweby (1984).[2]
- **PSOM unlimited**: The second order moment scheme of Prather (1986).
- **PSOM limited**: The second order moment scheme of Prather (1986) with flux limiters from Merryfield and Holloway (2003).

After 100 days, the tracer concentration profiles have evolved to those in the second, third, and fourth panels of Figure 1. There are clear differences between the various schemes.

The 2nd and 4th order schemes exhibit nontrivial extrema for the square pulse. Extrema arise from the dispersion errors intrinsic to these schemes. The errors are especially large when advecting profiles with strong gradients, such as the square pulse. In contrast, they perform much better for the Gaussian pulse, due to the pulse's much more smooth initial profile. The first order upwind scheme produces no extrema, yet it is highly dissipative, with the square pulse nearly as damped as the Gaussian pulse after 100 days. The Quicker scheme is somewhat of a compromise between the upwind scheme and the 2nd and 4th order schemes. It contains diffusion intrinsic to the algorithm, which smooths the profile to help maintain a more reasonable level of boundedness. However, its performance for the square pulse remains unsatisfying, as there remain some nontrivial unphysical extrema.

[1] This scheme was ported to MOM by Alistair Adcroft, based on his implementation in the MITgcm. The online documentation of the MITgcm at `http://mitgcm.org` contains useful discussions and details about this advection scheme.

[2] This scheme was ported to MOM by Alistair Adcroft, based on his implementation in the MITgcm. The online documentation of the MITgcm at `http://mitgcm.org` contains useful discussions and details about this advection scheme.

In the third panel, the Sweby, Super B, and MDPPM schemes all show similar behaviour to each other. In particular, each maintains the tracer within its physical bounds, since these schemes are flux limited. More detailed analysis reveals that MDPPM is a bit less dissipative than the other two schemes, thus allowing for a slightly better maintenance of the square pulse shape. It is interesting that each scheme converts the initially smooth Gaussian pulse into a more square feature over time. Such is a property common to many advection schemes with flux limiters.

The fourth panel shows results for the limited and unlimited PSOM schemes. The unlimited PSOM scheme produces extrema for the square pulse, but with far smaller amplitude than Quicker, 2nd, or 4th order advection. Even so, it does a wonderful job maintaining the shape of the square pulse. The limited PSOM scheme also maintains the square pulse, but clips the extrema thus retaining tracer concentrations within their physically relevant bounds.

Based on performance in the square pulse test, the limited PSOM scheme appears to be the best if we insist on maintaining boundedness of the tracer, with the unlimited PSOM scheme preferable if slight extrema are allowed. Use of the PSOM schemes come at the price of adding 10 extra three-dimensional arrays for each tracer. This added memory may preclude PSOM for some models, depending on computer platform and number of tracers. One approach to reducing the memory overhead is to employ the PSOM scheme for the active tracers (temperature and salinity), and another scheme, such as MDPPM, for the passive tracers.

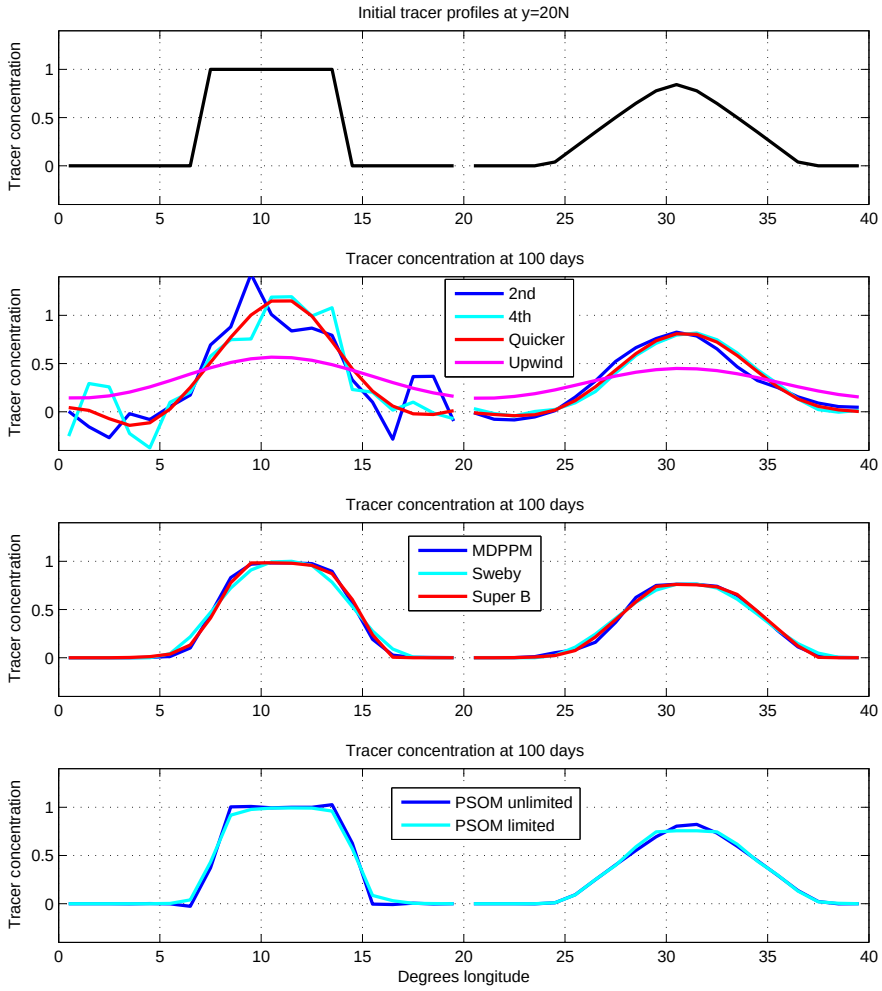

Figure 1: Top panel: Initial passive tracer profile for the torus test case. The two profiles are artificially offset in the zonal direction for purposes of clarity in presentation. Shown here is a plot through $y = 20°N$. Second panel: passive tracer profiles after 100 days of integration using 2nd, 4th, Quicker, and Upwind advection schemes. Third panel: Results using MDPPM, Sweby, and Super B. Fourth panel: Results from the limited and unlimited PSOM scheme.

3 Symmetric box test case

The symmetric box test case `mom4_symmetric_box` consists of a flat bottom domain which is symmetric across the equator. It has solid walls and is forced with zonally constant forcing. The initial conditions are constant temperature and salinity. The surface restoring fields are show in Figure 2. There is zero wind stress applied. Density is determined by the realistic equation of state of Jackett et al. (2006), and the prognostic temperature variable is the conservative temperature of McDougall (2003).

The aim of this test case is to examine numerical methods with respect symmetry across the equator. Small discretization errors which are not symmetric will accumulate and become visible in flow features. This test provides a useful means to measure this accumulation. Consequently, there is a relatively large suite of physical parameterizations enabled, many of which are incompatible with one another. The purpose of enabling so many options is to thoroughly examine symmetry properties of the schemes.

Figure 3 shows the surface height and zonally averaged temperature and zonal velocity after 1000 days of simulation. The surface height shows signs of asymmetry, whereas the zonally averaged temperature and zonal velocity are reasonably symmetric.

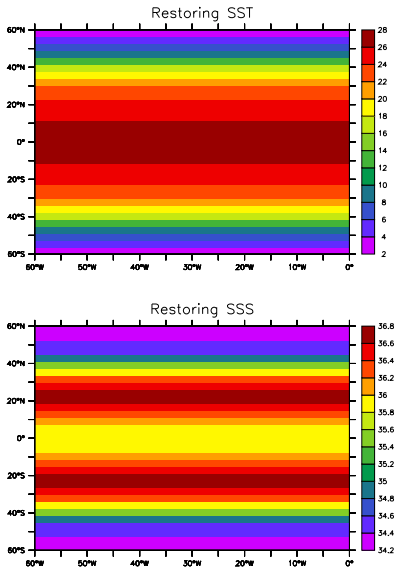

Figure 2: Restoring temperature and salinity for the symmetric box test case.

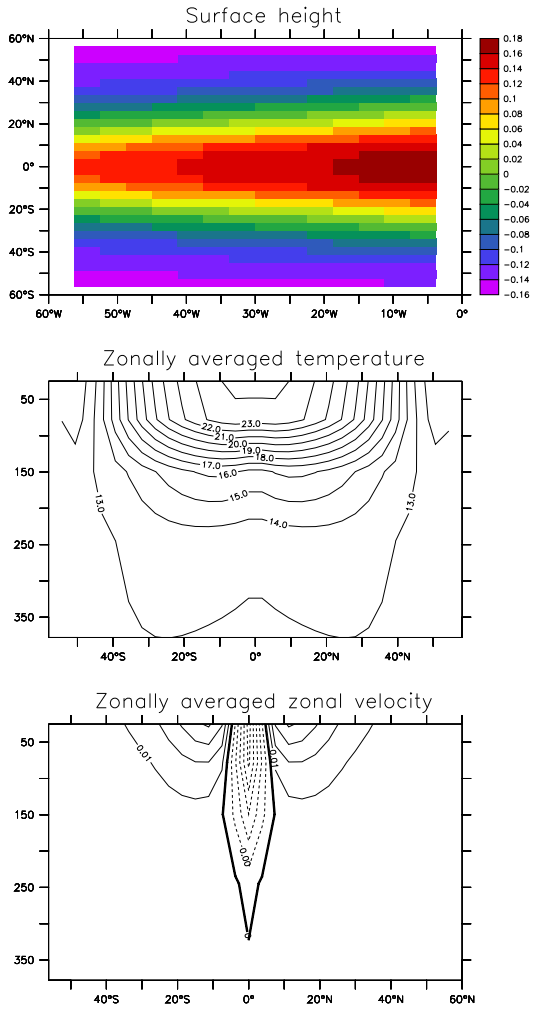

Figure 3: Snapshots of simulation features after 1000 days. Top panel: surface height in metres; Middle panel: zonally averaged temperature; bottom panel: zonally averaged zonal velocity.

4 Wind driven gyre test case

The test case `mom4_gyre` consists of a flat bottom middle latitude domain driven by a zonal wind stress. The temperature is initially stratified in the vertical (Figure 4), with density a linear function of temperature and independent of pressure and salinity. We employ the modified pressure coordinate

$$p^* = p_b^o \left(\frac{p - p_a}{p_b - p_a} \right) \tag{3}$$

for the vertical coordinate, with the applied pressure $p_a = 0$, and the static reference bottom pressure

$$p_b^0 = g \int_{-H}^{0} dz \, \rho^{\text{init}} \tag{4}$$

determined by the initial density profile ρ^{init}.

This test case has been used at GFDL for examining the spurious mixing properties of various advection schemes. For this purpose, we initialize to unity a passive tracer on each of the 50 vertical levels, with zero value off the chosen level (Figure 4). As the initial temperature profile is independent of horizontal position, the passive tracers are initialized on surfaces of constant density. We use the MOM4p1 MDPPM scheme for the test shown here.

During spinup of the circulation (Figure 5), passive tracers are advected around the gyre. Horizontal gradients on a temperature surface arise from inhomogeneous stretching and compression of the surface as a function of the level of eddy activity. The near surface has an Ekman layer, where vertical overturning occurs, as parameterized with convective adjustment. This overturning is in response to the Ekman suction.

Beneath the Ekman layer, spreading of passive tracer across surfaces of constant temperature (Figure 6) occurs for two reasons:

- nonzero vertical diffusivity, set here to $10^{-5} \text{m}^2 \text{s}^{-1}$,
- spurious mixing from numerical advection.

Quantifying levels of spurious mixing is straightforward. We do so by matching the evolution of passive tracer in density space to the evolution achieved with pure vertical diffusion, with a diffusivity fit to match that from the spurious mixing. This approach is a complement to the more complex sorting approach employed by Griffies et al. (2000) (see Griffies (2012)).

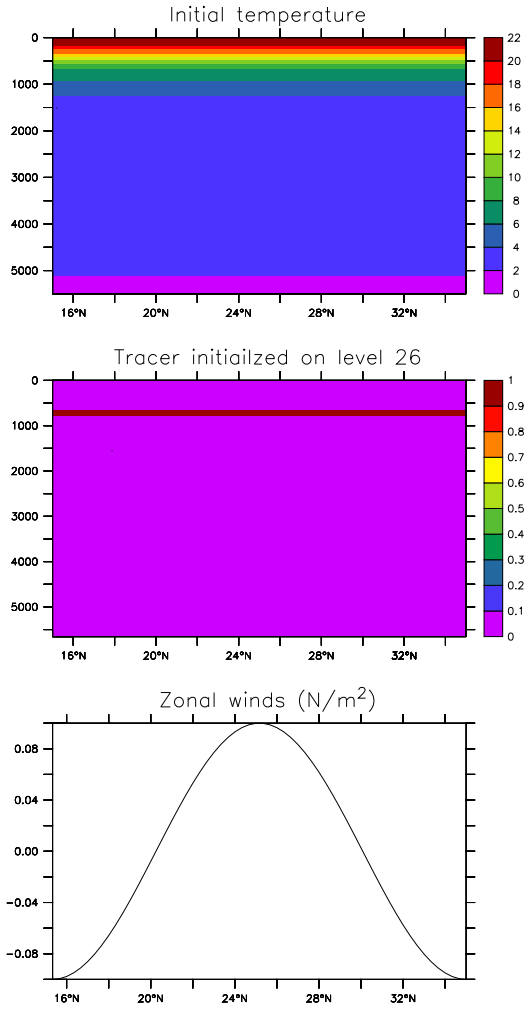

Figure 4: Upper panel: Section illustrating the initial temperature for the gyre experiment, with no horizontal variation in the initial temperature field. Middle panel: Illustration of an initial passive tracer, placed in this case on model level 26. Lower panel: The zonal wind stress applied to the gyre model, which spins up a two-gyre circulation within a few years.

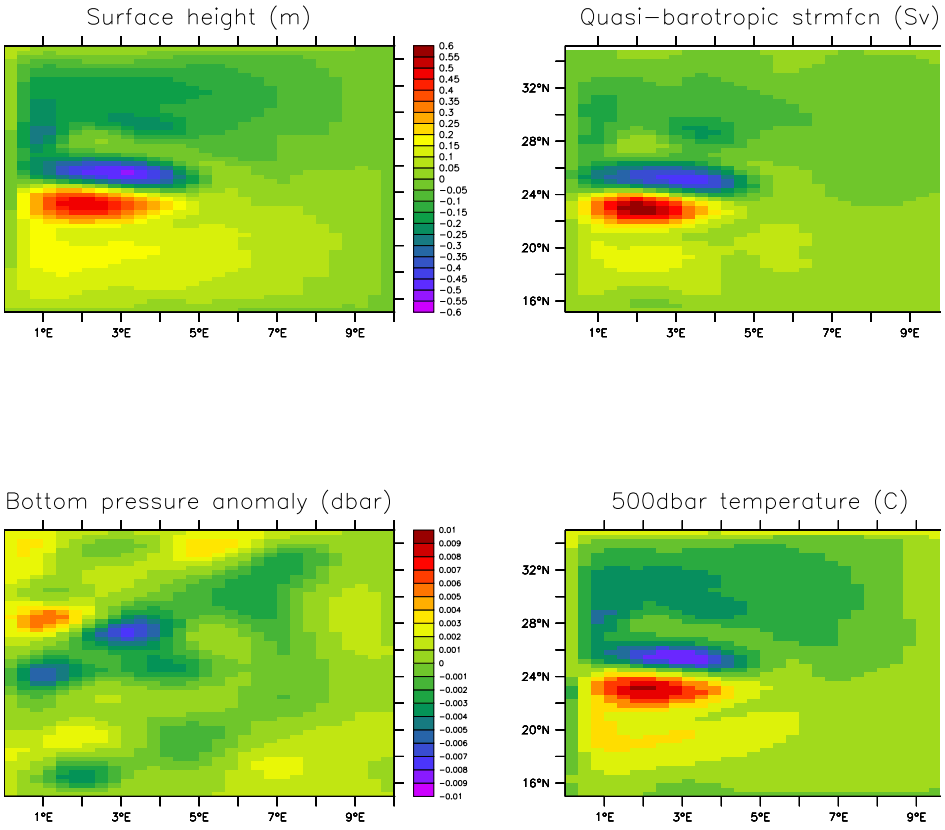

Figure 5: Simulation characteristics at 12 months. Shown here are monthly means for the following fields. Upper left panel: surface height (metre); Lower left panel: bottom pressure anomaly (dbar) $p_b - p_b^0$; upper right panel: quasi-barotropic streamfunction; Lower right: Temperature at 500dbar depth. Note the very small bottom pressure anomalies, relative to the surface height undulations, is consistent with the circulation being largely isolated to the upper ocean.

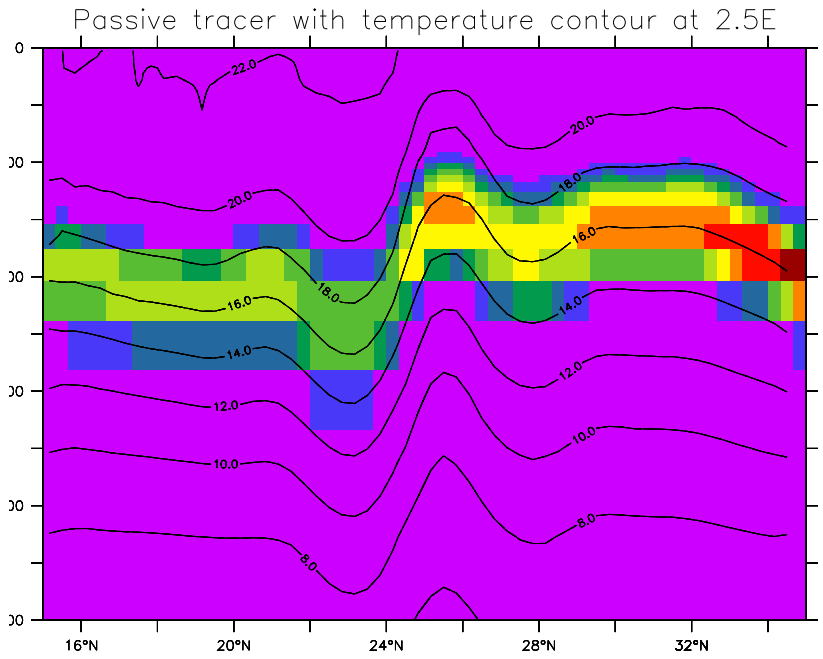

Figure 6: Monthly averaged passive tracer at 2.5E, overlaid with contours of constant temperature. Notice how the tracer undulates with the undulating temperature surfaces. Also, the tracer has spread across surfaces of constant density, with this spread associated with the nonzero vertical diffusivity and the spurious mixing from numerical advection errors.

5 Box sector test case

The box test case `mom4_box` consists of a flat bottom northern hemisphere sector domain with thermohaline forcing at the upper boundary (Figure 7). There is zero wind stress. The initial salinity is constant, and the initial temperature has a nontrivial zonally symmetric thermocline structure (Figure 7). Density is a linear function of temperature, and is independent of salinity and pressure.

The domain and grid are of a modest size ($24(x)\times35(y)\times18(z)$), thus allowing this test to be run on most any computer. The resulting circulation is driven by density gradients set up by the thermohaline forcing. It exhibits an overturning circulation, with sinking in the north and rising throughout the domain, which is reminscient of many similar idealized simulations run in the 1980s and 1990s.

We illustrate the behaviour of this test with the four following vertical coordinates.

- Geopotential vertical coordinate, as in MOM4.0;
- The depth based vertical coordinate

$$z^* = H\left(\frac{z-\eta}{H+\eta}\right) \tag{5}$$

- The pressure coordinate;
- The pressure based vertical coordinate

$$p^* = p_{\rm b}^o\left(\frac{p-p_{\rm a}}{p_{\rm b}-p_{\rm a}}\right), \tag{6}$$

  with $p_{\rm b}^o$ the initial bottom pressure, $p_{\rm b}$ the bottom pressure, and $p_{\rm a}$ the applied pressure at the ocean surface, which is set to zero for this suite of tests.

Figures 8, 9, and 10 illustrate the zonal averaged temperature, salinity, and ideal age tracer, each averaged over years 901-1000. The differences in the simulations is very minor, especially when recognizing the different vertical axes extents for the pressure versus depth based vertical coordinates. Figure 11 shows the meridional overturning streamfunction for the simulations, again revealing very minor differences.

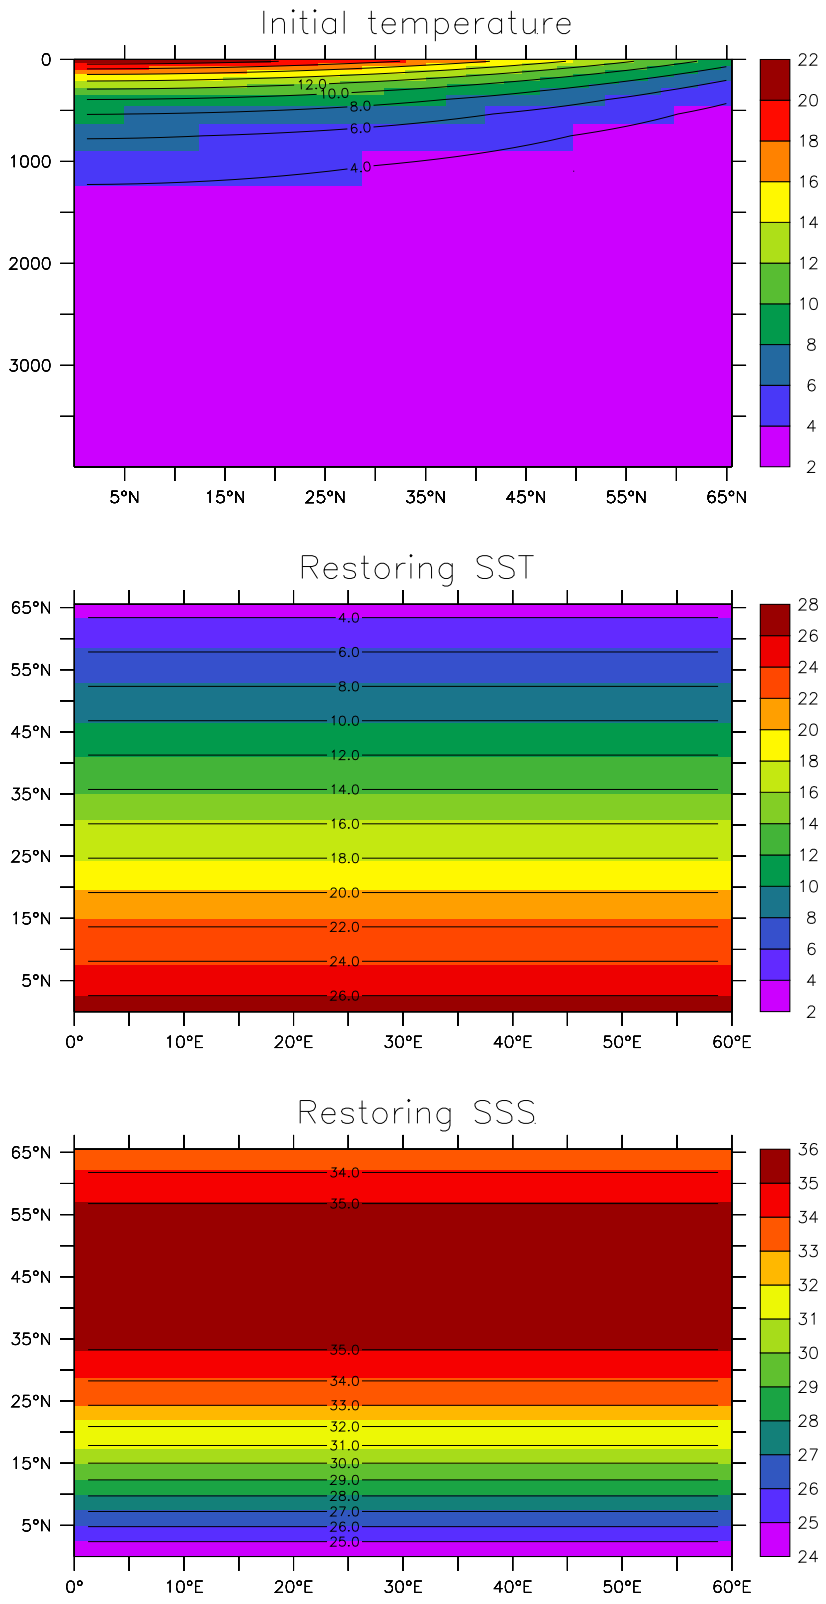

Figure 7: Initial and boundary conditions for the box test case. Upper panel: initial temperature. Middle panel: SST used for computing the restoring flux of heat. Lower panel: SSS used for computing restoring flux of salt or fresh water.

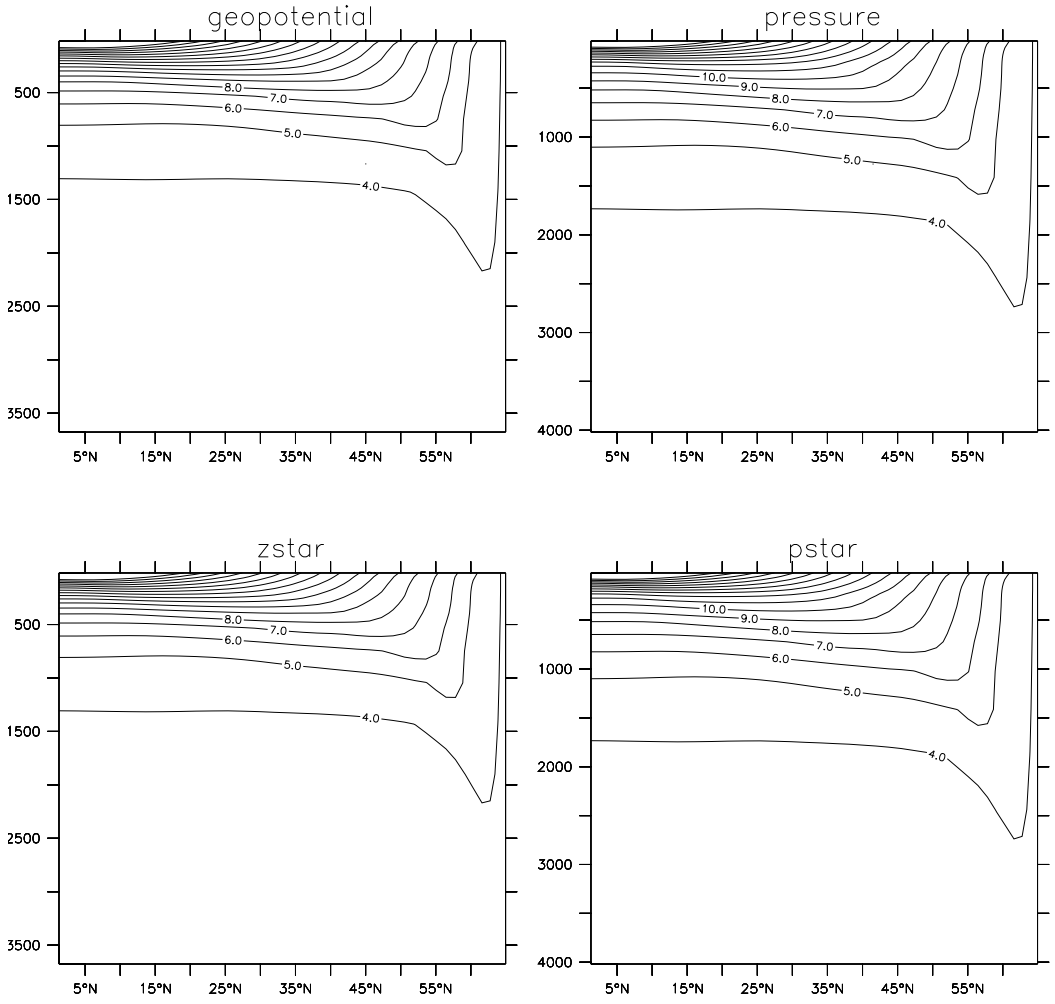

Figure 8: Zonal averaged temperature time averaged over the years 901-1000 in the box test case. Shown here are results for the four vertical coordinates considered in this chapter. Note that for the depth based coordinates, the vertical refers to metres, whereas for the pressure based it is dbars. This accounts for the slight offset in the vertical extent of the domains.

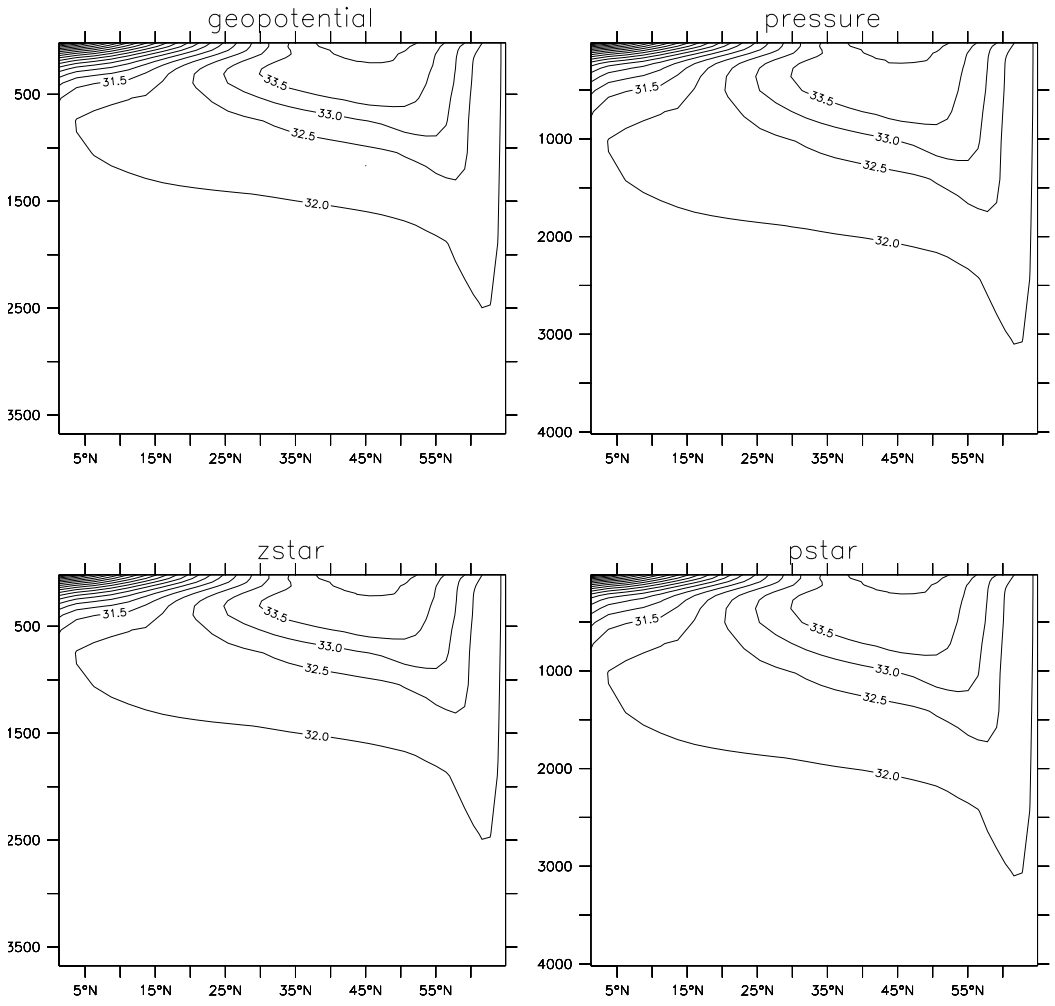

Figure 9: As in Figure 8, but for salinity.

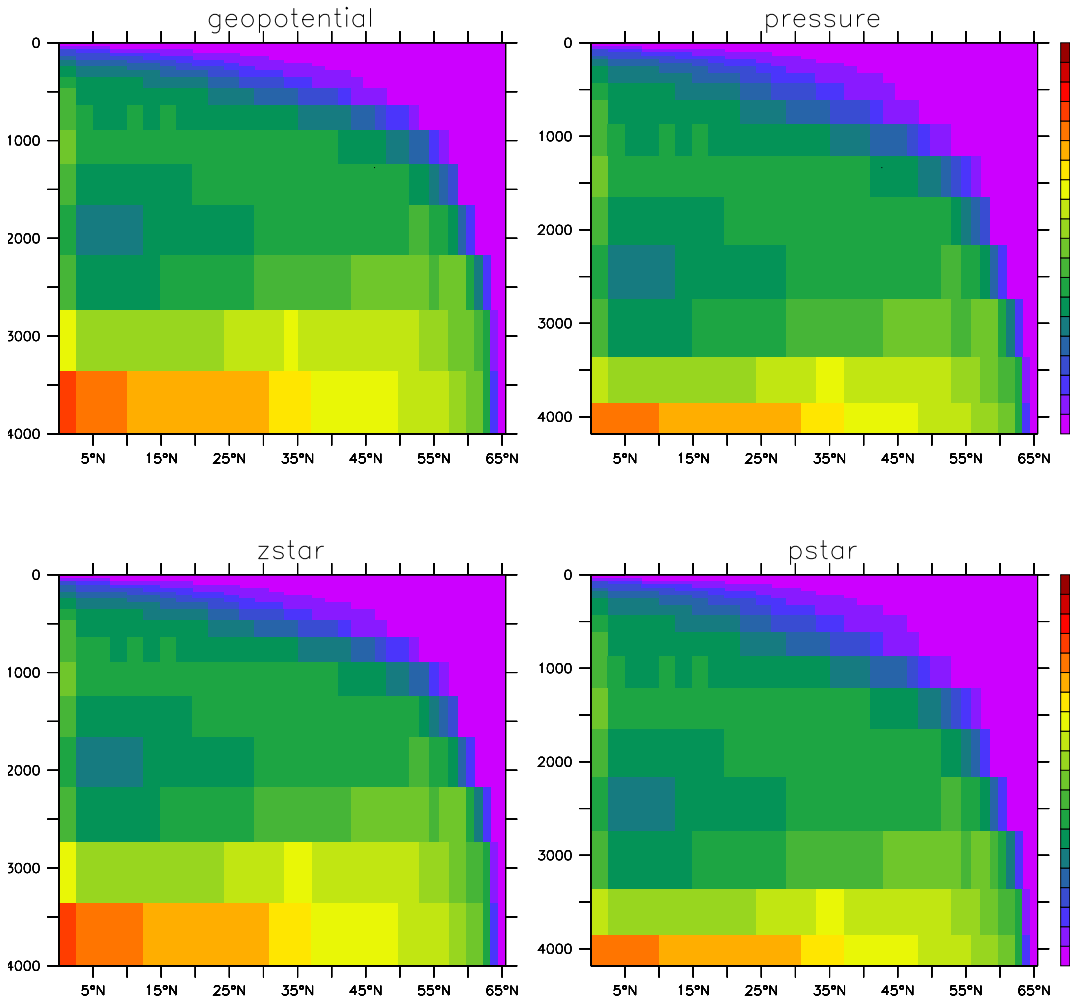

Figure 10: As in Figure 8, but for the age tracer.

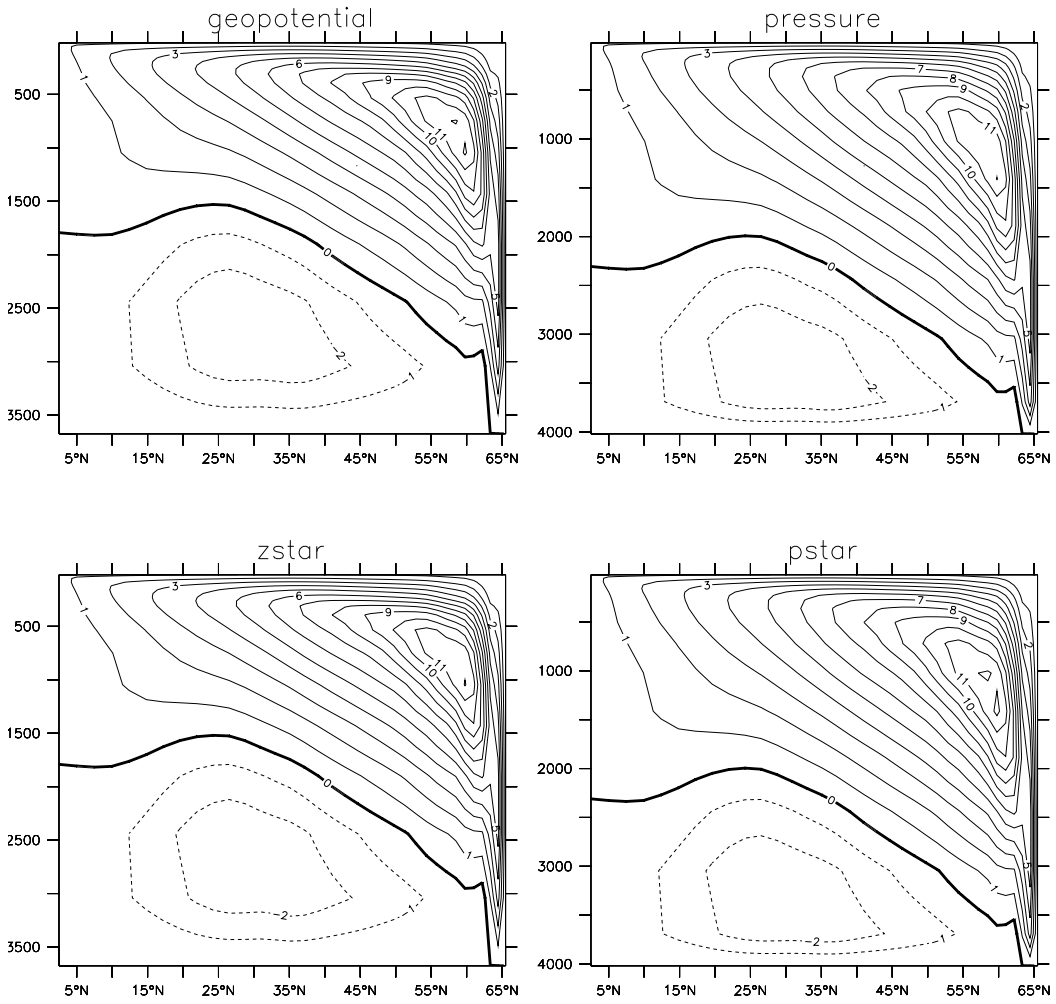

Figure 11: As in Figure 8, but for the meridional overturning streamfunction.

6 Box-channel test case

The box-channel test case `mom4_box_channel` consists of a flat bottom southern hemisphere sector domain with a channel to the south of the sector. This idealized domain is useful for certain studies of geometries analogous to that in Nature, where a southern channel feature (i.e., the Antarctic Circumpolar Current) navigates around the globe and connects to the three main basins (Atlantic, Indian, and Pacific). This geometry was constructed by Rongrong Zhao at GFDL and used by her to study elements of subgrid scale (SGS) parameterizations.

The initial conditions and surface restoring fields are show in Figure 12. There is zero wind stress applied. Density is a linear function of temperature, and is independent of salinity and pressure. Circulation arises just due to the buoyancy forcing.

The main features highlighted in this test case involve sensitivity of the deep water formation to the SGS parameterizations in the tracer equation. We consider four cases:

- No lateral SGS parameterization; only the Sweby advection scheme for temperature, PSOM for salinity, and MDPPM for age.
- Horizontal diffusion as an SGS parameterization; otherwise the same as above.
- Replace the horizontal diffusion with the neutral physics scheme with boundary layer determined as in MOM4.0.
- Replace the horizontal diffusion with the neutral physics scheme with boundary layer determined as in Ferrari et al. (2008).

Figure 13 shows the zonally averaged temperature as time averaged over years 901-1000. Note the large increase in vertical stratification as neutral physics is enabled. Indeed, the Ferrari et al. (2008) scheme appears to fully shut-down ventilation. Such is reflected in Figure 14, which shows the zonally averaged ideal age tracer.

Another feature to note is the colder deep temperatures using the Ferrari et al. (2008) boundary layer formulation. This latter result is spurious, since the initial conditions and surface restoring are all warmer than 4C. The MOM implementation of the Ferrari et al. (2008) introduces an unacceptable level of spurious extrema. Such extrema are difficult to control with rotated physics schemes, although they are generally not as egregious as seen here with the new boundary layer formulation.

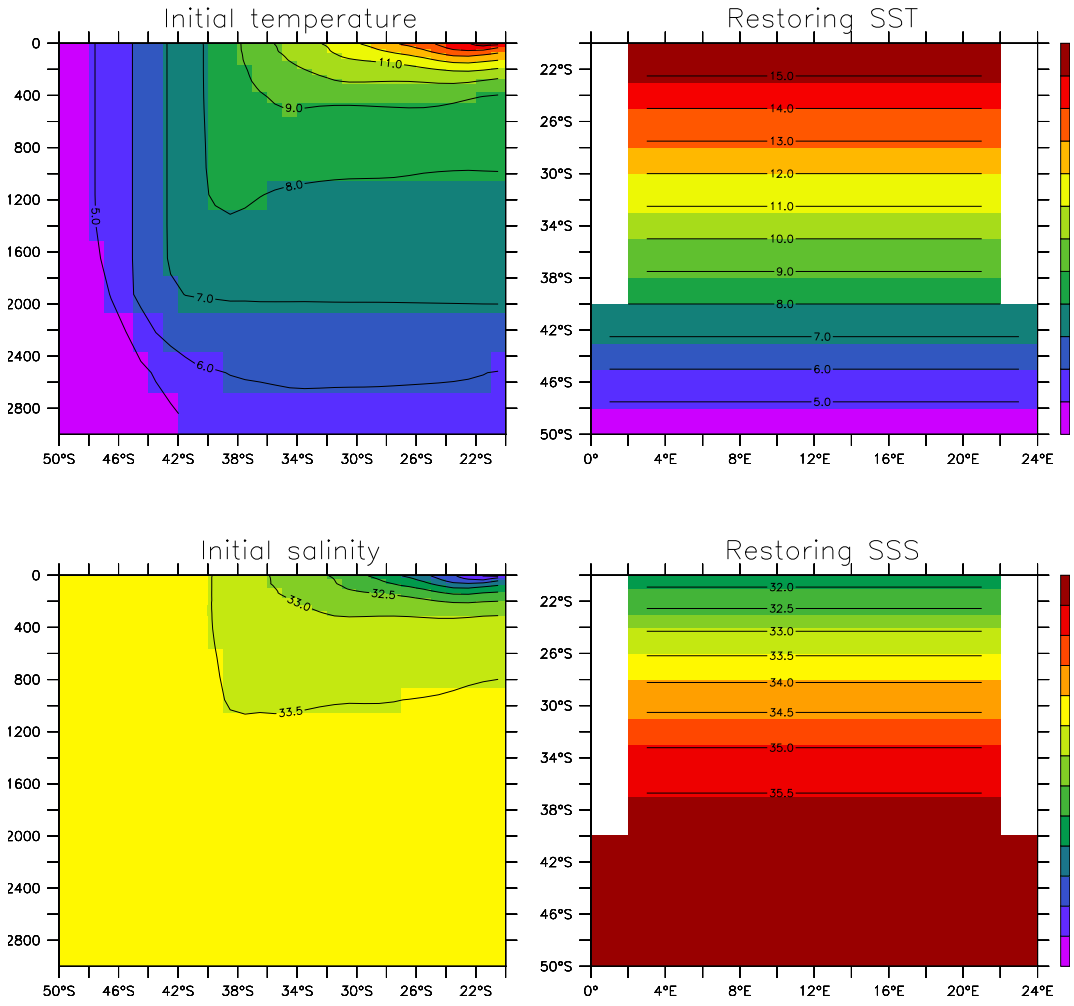

Figure 12: Initial and boundary conditions for the box test case. Upper panel: initial temperature. Middle panel: SST used for computing the restoring flux of heat. Lower panel: SSS used for computing restoring flux of salt or fresh water.

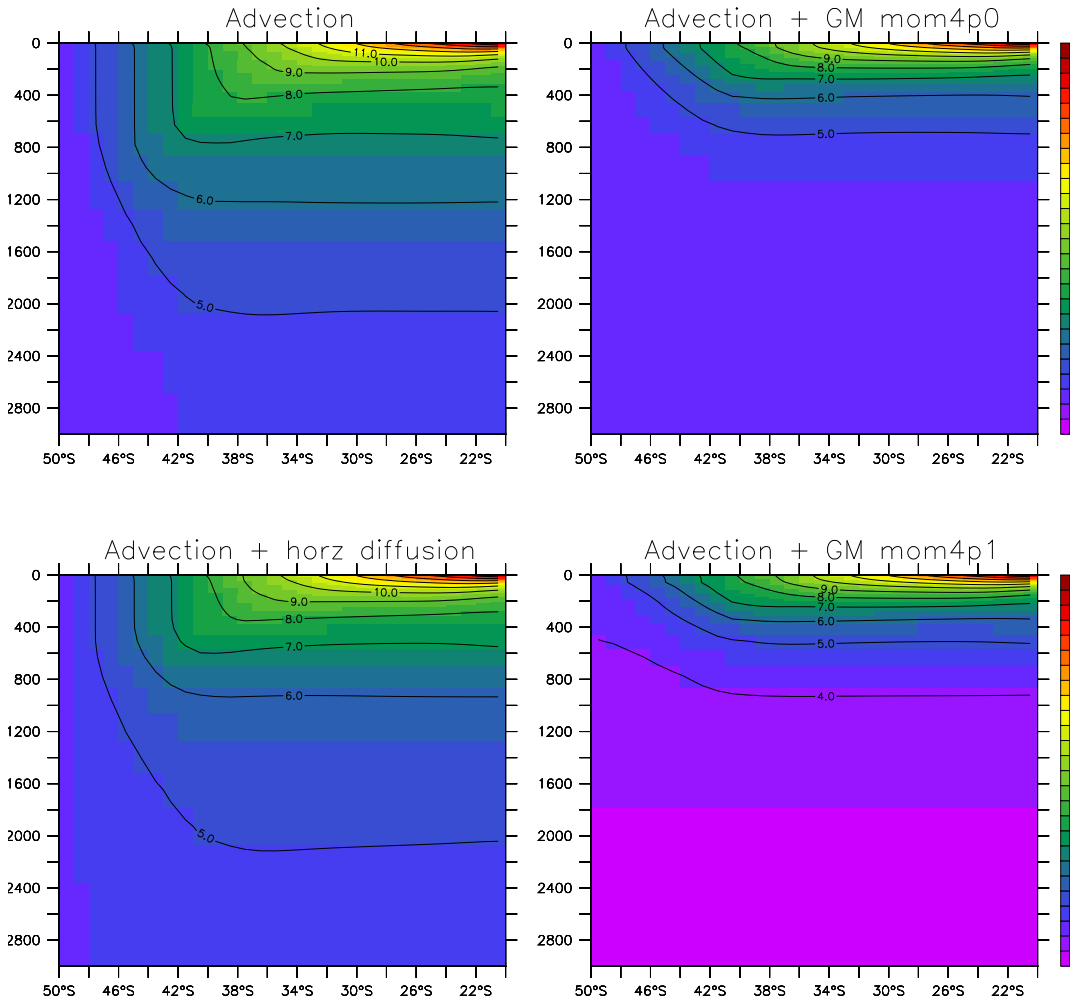

Figure 13: Time mean and zonally averaged temperature from years 901-1000 for the box channel test case. Upper left panel: results with advection yet without lateral SGS parameterization; lower left panel: advection plus horizontal diffusion; upper right: advection plus MOM4p0 neutral physics; lower right: advection plus MOM4p0 neutral physics.

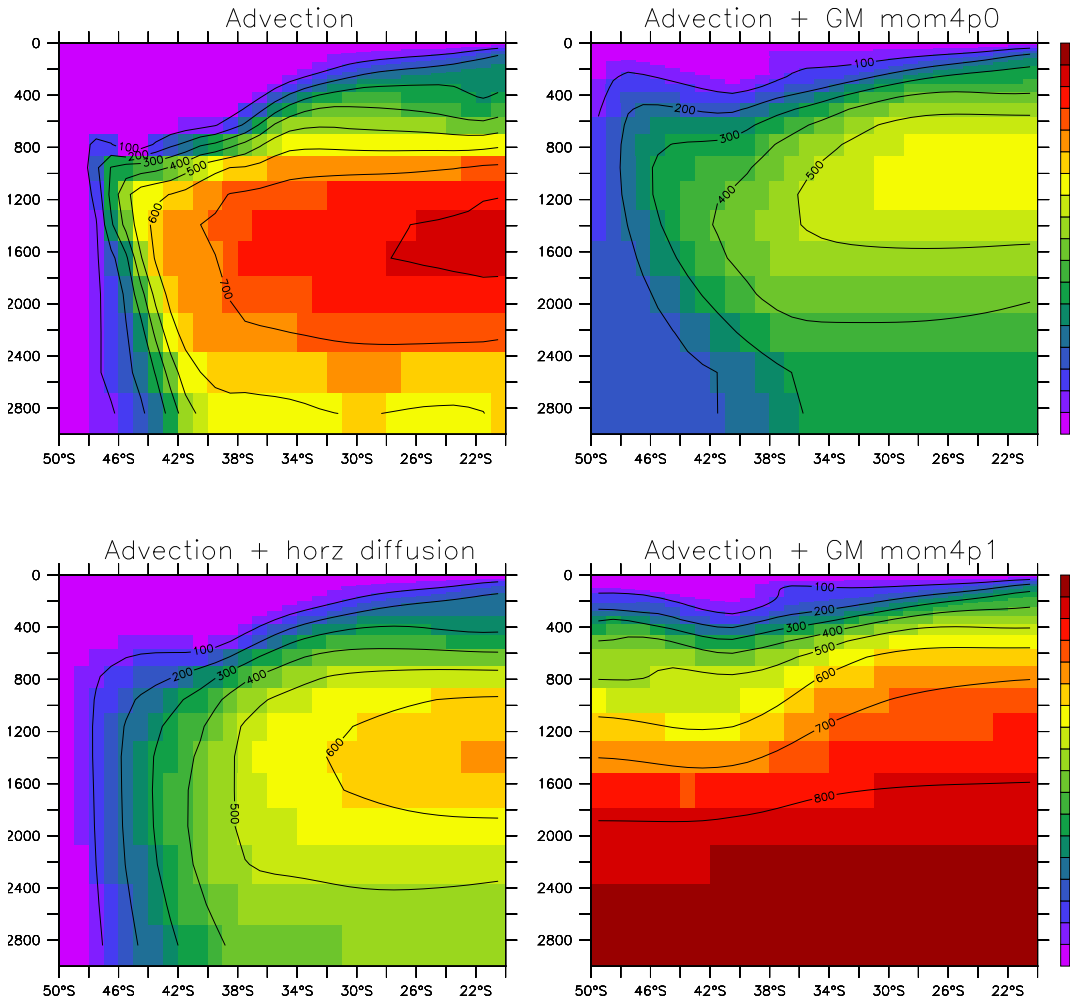

Figure 14: Time mean and zonally averaged age from years 901-1000 for the box channel test case. Upper left panel: results with advection yet without lateral SGS parameterization; lower left panel: advection plus horizontal diffusion; upper right: advection plus MOM4p0 neutral physics; lower right: advection plus MOM4p0 neutral physics.

7 Bowl test case

The bowl test case `mom4_bowl` is based on a configuration built for studies of overflow processes by Winton et al. (1998). It consists of a shallow shelf in the northern portion of the domain, with a deeper bowl to the south (Figure 15). The initial conditions place cold dense water on the shelf (Figure 16). Density is a linear function of temperature, and is independent of salinity and pressure.

Integration proceeds with no surface forcing. Hence, this is a relaxation or adjustment experiment, with no steady state realized. The initial pressure forces set up acceleration which causes the water to move off the shelf and into the deeper portion of the domain. The numerical and physical processes active in this relaxation experiment are described by Winton et al. (1998). In particular, the test case is useful to examine the sensitivity of the deep water formed to various numerical and physical options. It can be used in tandem with the test case `mom4_dome` described in Section 8.

Level models are notoriously poor at representing overflow processes (Winton et al., 1998). Various approaches have thus been proposed to resolve, or reduce, the problems. We test the following four configurations here:

- Terrain following pressure based vertical coordinate

$$\sigma^{(p)} = \left(\frac{p - p_a}{p_b - p_a} \right), \tag{7}$$

  with p_b the bottom pressure, and p_a the applied pressure at the ocean surface, which is set to zero for this suite of tests. For this vertical coordinate, we transport tracer just by advection plus some vertical diffusion.

- Quasi-horizontal pressure based vertical coordinate

$$p^* = p_b^o \left(\frac{p - p_a}{p_b - p_a} \right), \tag{8}$$

  with p_b^o the initial bottom pressure. In this test, we also just employ advection plus vertical diffusion.

- In this test, we use p^* with the addition of nontrivial sigma diffusion in the bottom cell (see Griffies (2012)).

- In this test, we use p^* with the addition of nontrivial sigma diffusion in the bottom cell, as well as the *overexchange* scheme (see Griffies (2012)).

Figures 17 and 18 show the salinity (a passive tracer) and temperature at the botom of the domain, after one year as averaged over the last five days of the year. Both of the cases without any added subgrid scale methods show advection transporting the tracers southwards along the western boundary. The terrain following $\sigma^{(p)}$ case brings the tracer further into the abyss than the p^* case. When adding sigma diffusion or overexchange to the p^* cases, the bottom signal is greatly diffusved, as expected since these parameterizations are diffusive. They also cause the tracer to bleed into the abyss more uniformly, rather than following the path along the wester wall.

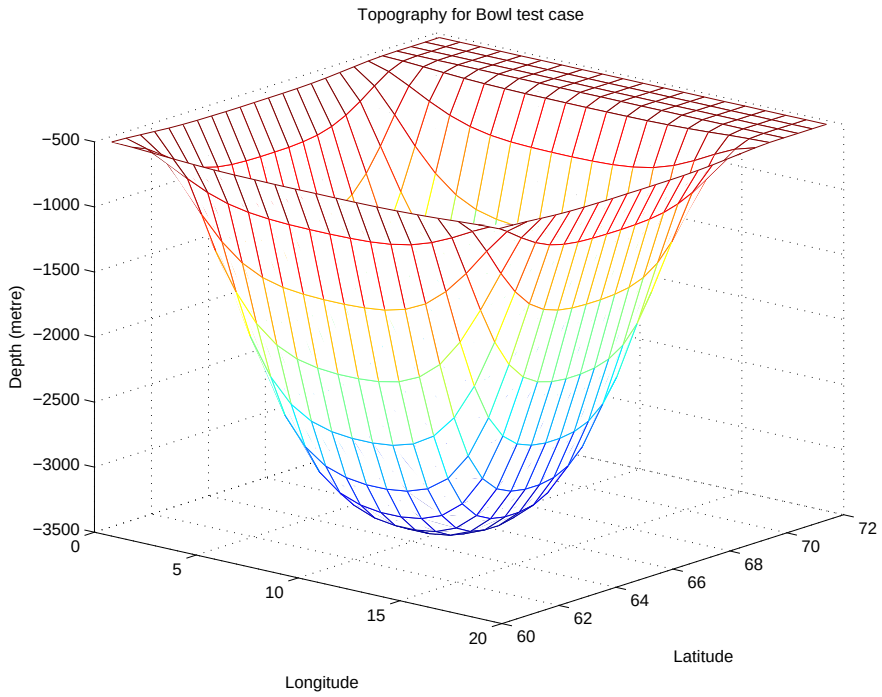

Figure 15: Geometry of the bowl test case, with shallow shelf in the north and deeper bowl to the south.

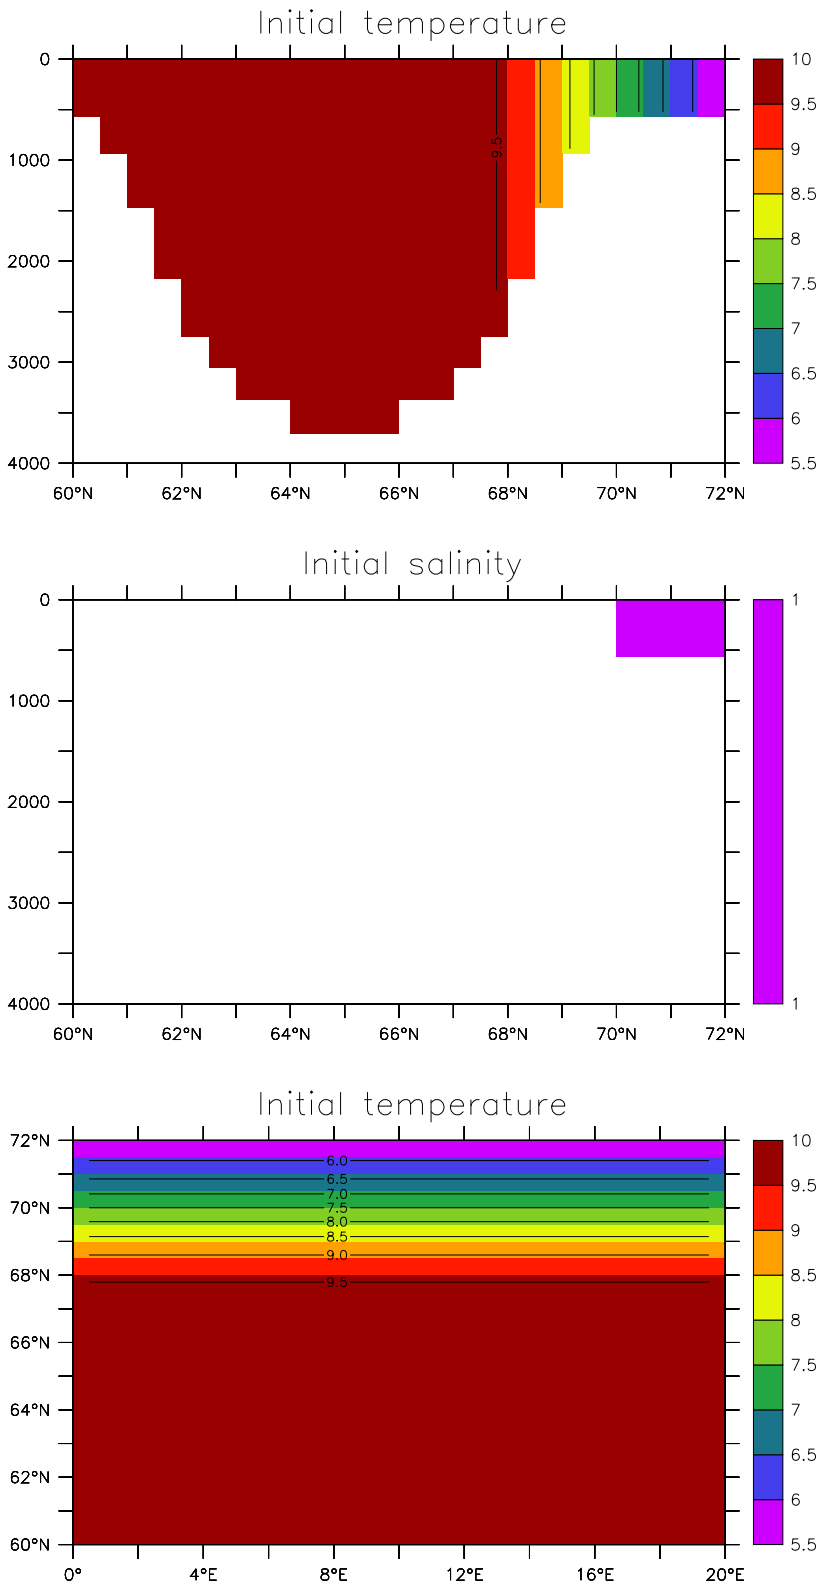

Figure 16: Initial conditions for temperature and salinity in the bowl test case. Upper panel: meridional-depth slice of the zonally symmetric initial temperature. Middle panel: Initial salinity, which has a unit value on the shelf and zero elsewhere. Lower panel: Plan view of the initial SST.

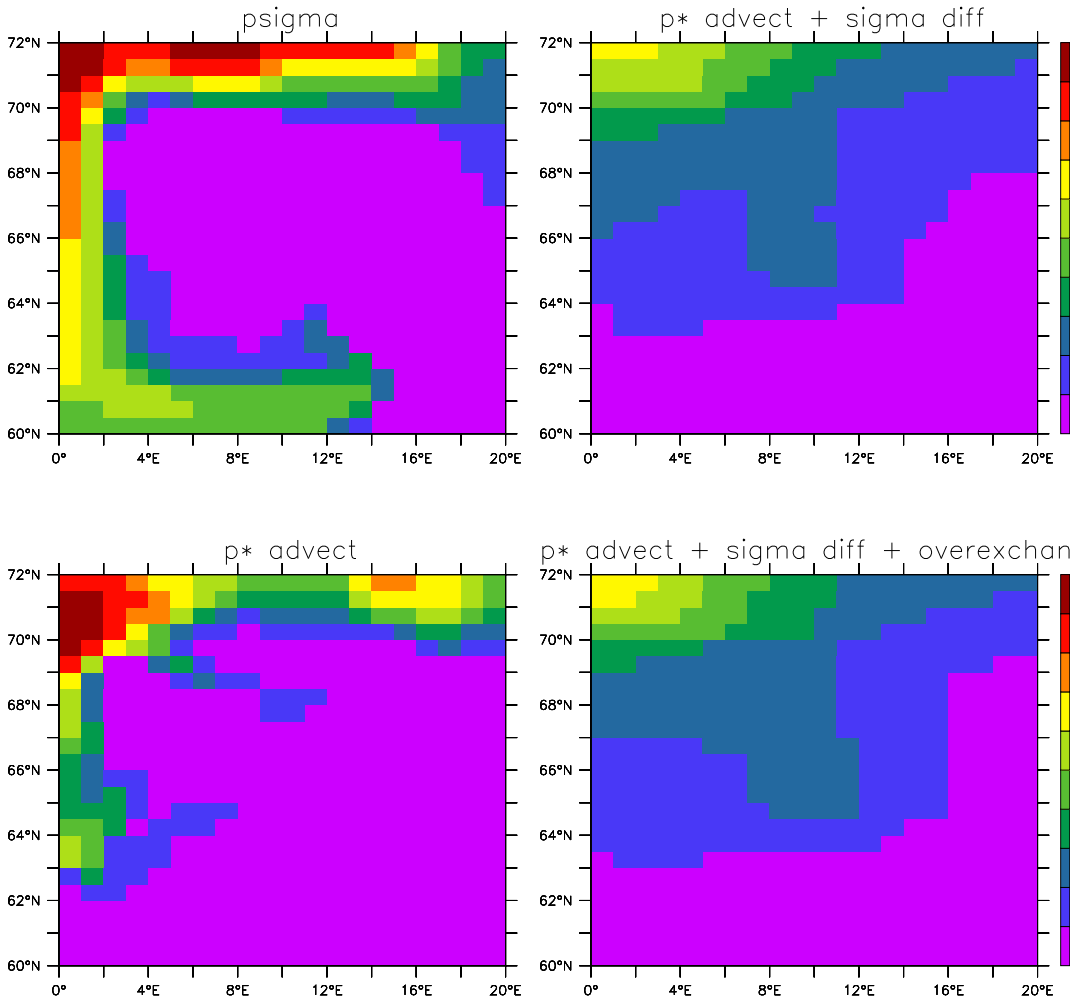

Figure 17: Salinity at the bottom of the domain after one year, as averaged over the last five days of the year. Upper left panel: results from $\sigma^{(p)}$ vertical coordinate using advection alone for the tracer transport; Lower left panel: results from p^* vertical coordinate using advection alone for the tracer transport; Upper right panel: p^* with advection and sigma diffusion applied in the bottom cell; Lower right panel: p^* with advection and sigma diffusion applied in the bottom cell and the overexchange scheme.

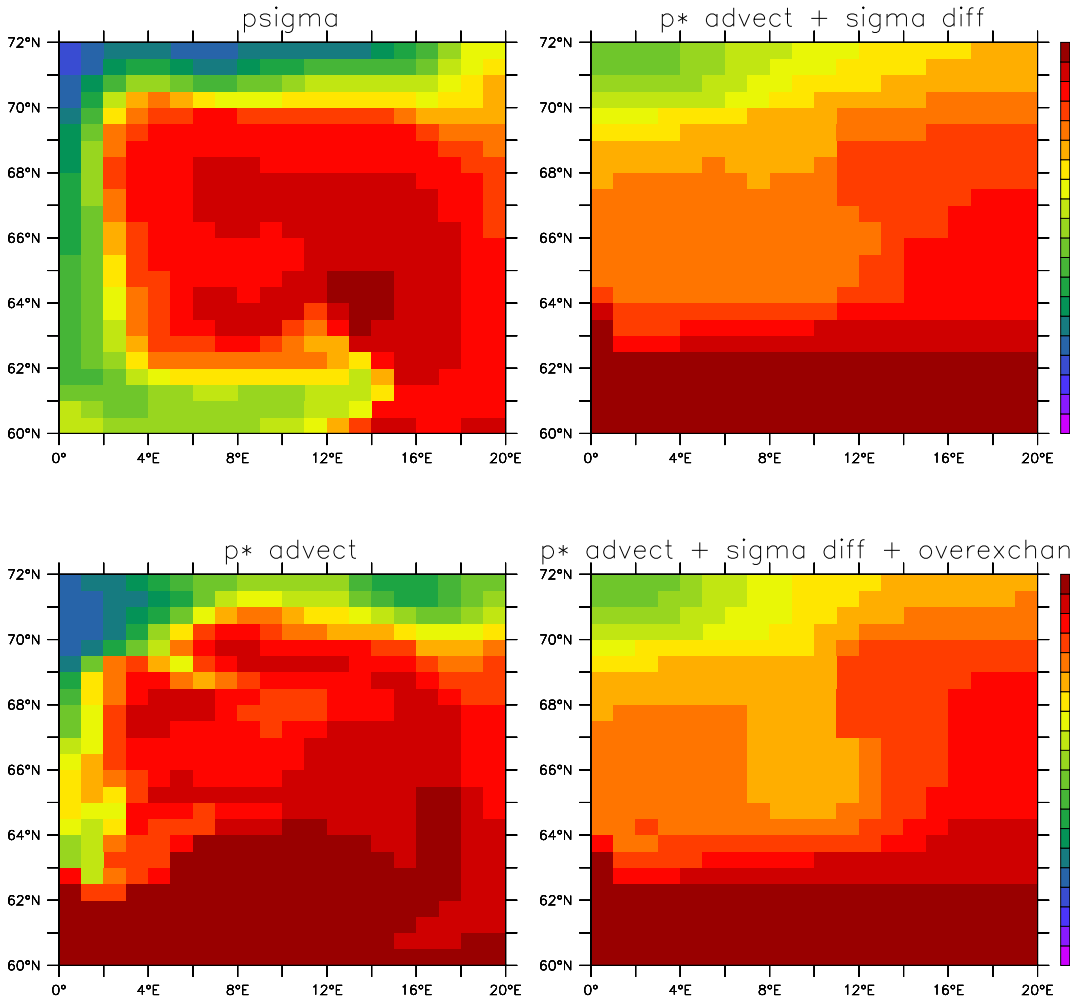

Figure 18: As in Figure 17, but for temperature.

8 DOME test case

The experiment mom4_dome is based on the idealized overflow test established for the project Dynamics of Overflow, Mixing, and Entrainment (DOME). Details of the experiment are provided in various published papers, such as Legg et al. (2006).

As with the bowl test case mom4_bowl discussed in Section 7, the DOME test case provides a means to test various overflow parameterizations. In both tests, the simulations highlight the ability, or inability, of the simulated flow to overcome the geostrophic balance, which aims to align the flow along isobaths. In coarsely resolved level models, spuriously large levels of ambient waters tend to be entrained to the dense pulse, thus compromising the ability of the simulation to form dense deep water, and to feel the effects of the bottom (Winton et al., 1998). The main difference between mom4_dome and mom4_bowl is that mom4_dome allows for a steady state to be achieved, since dense water continues to be injected from the embayment. The mom4_bowl test, in contrast, is an adjustment experiment, with no boundary forcing considered.

Coarse resolution level models are notoriously poor at representing overflow processes (Winton et al., 1998). Various approaches have thus been proposed to resolve, or to reduce, the problem (see Griffies (2012)). For the simulations discussed here, we employ the quasi-horizontal coordinate p^*. Any of the other quasi-horizontal coordinates, such as z, z^*, or p would show similar results. In contrast, the terrain following coordinates $\sigma^{(z)}$ or $\sigma^{(p)}$ show much more downslope flow, due to their ability to better represent the interaction between the flow and the terrain (not shown).

The configuration consists of a shelf in the north with a shallow embayment. We use 1/2° grid resolution in both latitude and longitude, which yields two zonal grid points in the embayment. The vertical grid is the same 50 levels used in the global configuration used in Griffies et al. (2005).

Temperature is initialized with a zonally symmetric profile shown in Figure 19. Cold water in the north lives at the bottom of the embayment. It is injected southward by a transport imposed at the embayment's northern boundary. Salinity is initialized to zero everywhere, but it is given a value of unity for water that is injected from the embayment. Density is a linear function of temperature, with no dependence on salinity and pressure. Salinity thus provides a passive marker for injected water that is transported southward and into the abyss. Both temperature and salinity are damped to their initial conditions at the eastern and western boundaries. This damping, applied through a sponge condition, eliminates the Kelvin wave signals. Finally, for tracer advection, we use the Prather (1986) scheme, employed without flux limiters.

We test the following five cases in this chapter:

- No lateral or downslope parameterization; just advective transport;
- Advection plus sigma diffusion acting in the bottom grid cell;
- Advection plus sigma diffusion acting in the bottom grid cell plus the overflow scheme of Campin and Goosse (1999) (see Griffies (2012));
- Advection plus sigma diffusion acting in the bottom grid cell plus the overexchange scheme (see Griffies (2012));
- Advection plus sigma diffusion acting in the bottom grid cell plus the mixdownslope scheme (see Griffies (2012));

Four months of integration is sufficient to highlight differences between various overflow options available in MOM4p1. Figure 20 shows salinity at the bottom of the domain, averaged over the last five days of the four month experiment. The case with pure advection (plus some vertical diffusion) shows little downslope flow, as well as some extrema (values less than zero) arising from the absence of flux limiters on the advection scheme. Adding sigma diffusion allows for salinity to penetrate further south. There are also no tracer extrema, likely due to the smoothing of the otherwise strong gradients near the bottom. Adding the overflow scheme of Campin and Goosse (1999) reduces the depth where salinity penetrates. Apparently it is acting in a manner that handicaps the penetration of dense water southwards, perhaps due to over dilution. In contrast, the *overexchange* scheme allows for the salinity to penetrate further southward and hence deeper. For this test, we applied the *overexchange* scheme in four adjacent grid cells. Finally, the *mixdownslope* also allows for penetration further southward, although with a local extrema a few points away from

the embayment, perhaps due to the nonlocal manner of transport. The penetration of salinity southward is reflected in the zonally averaged temperature shown in Figure 21. Consistent with the salinity tracer, the densest waters arise from the case with the *overexchange* and *mixdownslope* schemes. The different density profiles give rise to differences in flow characteristics. Figure 22 shows the barotropic quasi-streamfunction (see Griffies (2012))

$$\psi(x,y) = -\int_{y_{\text{south}}}^{y} \mathrm{d}y' \, U(x,y'), \tag{9}$$

where

$$U = \int_{-H}^{\eta} \mathrm{d}z \, u \tag{10}$$

is the vertically integrated zonal transport.[3] The vertically integrated transport is far stronger in the cases with dense shelf water penetrating into the deep.

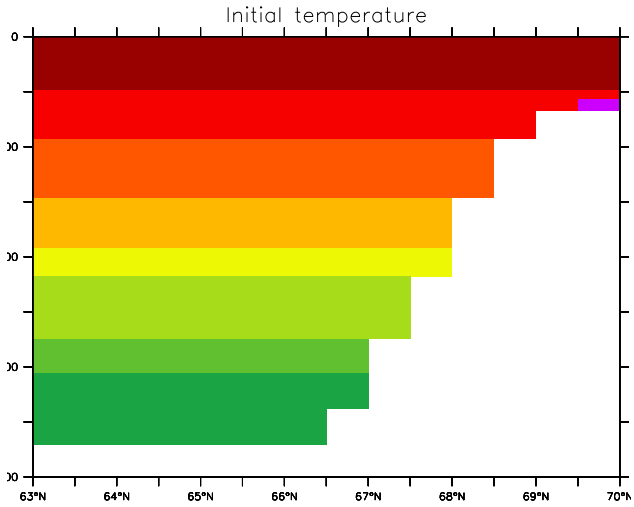

Figure 19: Initial temperature for the DOME test case. The cold water in the far north is at the bottom of the embayment. Southward flow is injected from this embayment, allowing for the cold and dense water to leave the embayment and enter the abyss.

[3] In a rigid lid model, or a free surface model having reached a steady state in the absence of surface boundary fluxes of mass, then ψ is a vertically integrated streamfunction. For the present case, with a transient free surface model, ψ is only a direct correlate to the true vertically integrated transport, hence the qualifier *quasi*.

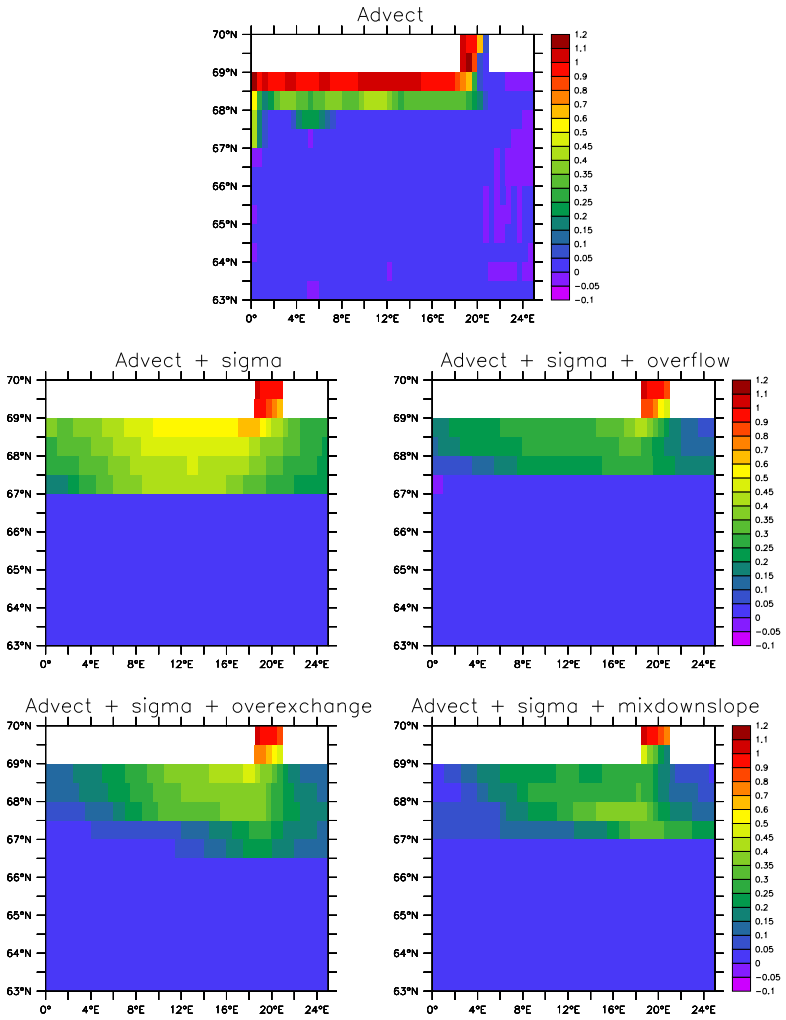

Figure 20: Bottom salinity averaged over the last five days of a four month integration. Top panel: no lateral or terrain following parameterization. Middle left: advection plus sigma diffusion applied in the bottom grid cell. Middle right: advection plus sigma diffusion plus the overflow scheme of Campin and Goosse (1999) (Section **??**). Bottom left: advection plus sigma diffusion plus the *overexchange* scheme (Section **??**) Bottom right: advection plus sigma diffusion plus the *mixdownslope* scheme (Section **??**).

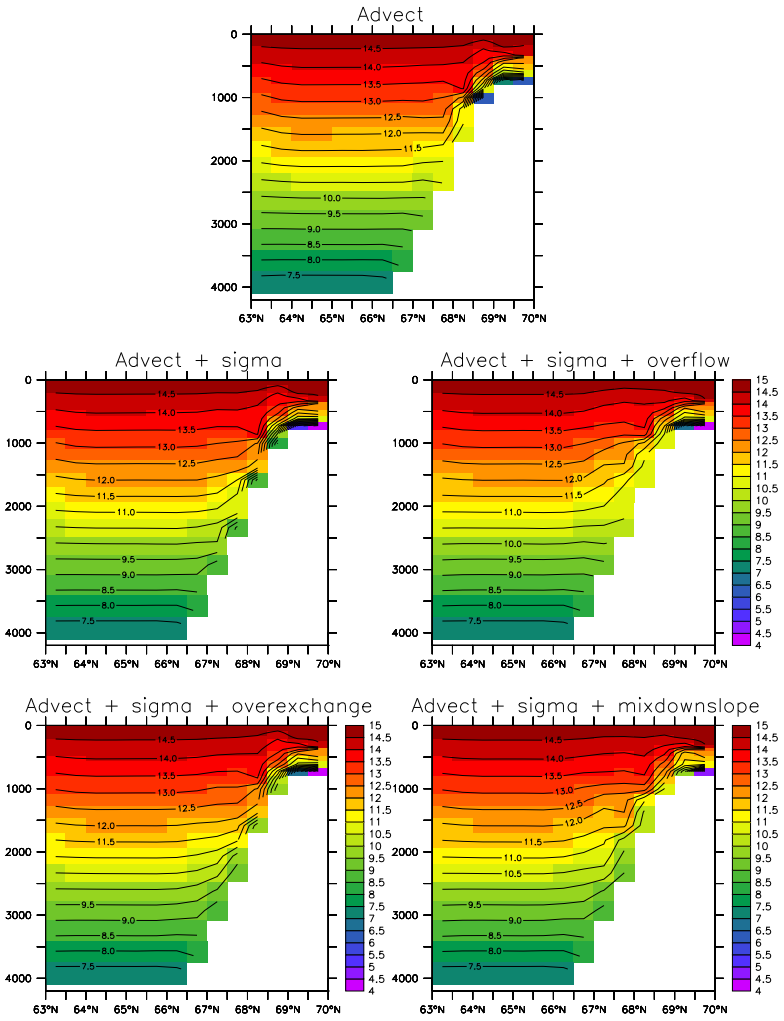

Figure 21: Zonally averaged temperature for the DOME test case, with panels corresponding to those in Figure 20.

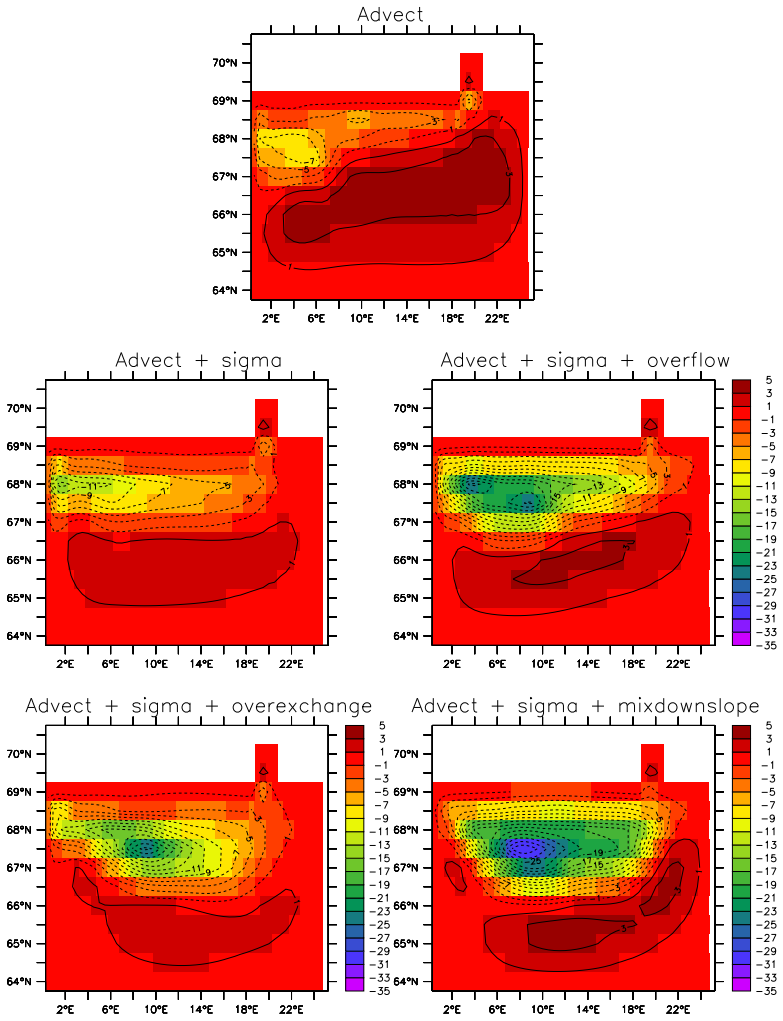

Figure 22: Quasi-streamfunction (Sv) for the vertically integrated transport in the four DOME simulations, with panels corresponding to those in Figure 20.

9 Indian Ocean domain with sponge boundaries

The test case `mom4_iom` consists of a realistic regional model of the Indian Ocean. The lateral regional boundaries are damped to climatology using sponges. This model was originally developed as part of a course on large scale ocean modelling in Bangalore, India during October 2004. The model remains a useful starting point for certain Indian Ocean research efforts using MOM. It is essentially the MOM4p1 implementation of Version 1.0 of that model configuration. The model is forced with a full suite of realistic atmospheric products, and employs a relatively modern suite of physical parameterizations.

One point of departure from the standard MOM4.0 configuration is to employ the depth based vertical coordinate

$$z^* = H\,(z-\eta)/(H+\eta). \tag{11}$$

This coordinate is useful in cases where it is desired to refine the vertical resolution in the upper ocean. Such enhanced vertical resolution may be of use for the Indian Ocean, where important vertical gradients exist especially near river mouths.

Figure 23 shows the SST, SSS, and surface height fields for the simulation after only two days of integration. As this model is not actively run at GFDL, we have not run it ourselves for much longer periods. Researchers interested in using this model configuration should consult with the MOM4 user community to be able to access the experience of others in running this model.

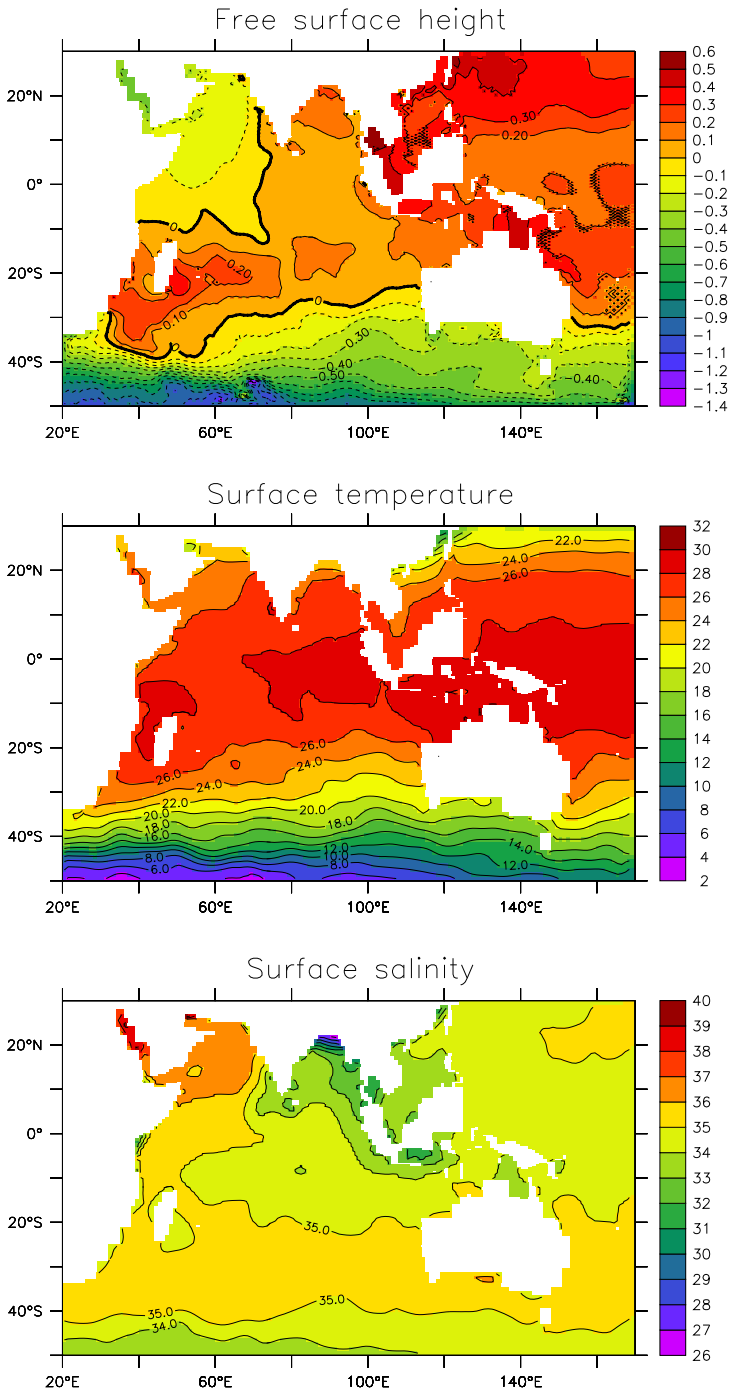

Figure 23: Daily averaged surface fields after two days integration of the Indian Ocean Model. Top panel: free surface height (metre); middle panel: surface temperature (C); bottom panel: surface salinity (psu).

10 CSIRO Mark 3.5 test

The test case `mom4_mk3p5` consists of a realistic global model with spherical coordinates. This model was developed at the CSIRO Marine and Atmospheric Research in Aspendale, Australia. It is routinely used as the ocean component for the Australian contribution to the IPCC report.

It is notable that this model was originally developed for a rigid lid algorithm in MOM1. For that algorithm, a heavy dose of polar filtering was employed. The present configuration in MOM, employing a free surface, removes all polar filtering, yet it accepts a *longer* time step than the polar filtered rigid lid configuration. This result is indicative of the often paradoxical results found with polar filtering. We strongly recommend that all polar filtering be removed from global simulations. Indeed, although a polar filtering module remains as part of MOM, it is *not supported* by GFDL scientists. The main reason is that is breaks local conservation of tracer, thus leading to spurious surface fluxes and poor high latitude simulations.

References

Campin, J.-M., Goosse, H., 1999. Parameterization of density-driven downsloping flow for a coarse-resolution ocean model in *z*-coordinate. Tellus 51A, 412–430.

Ferrari, R., McWilliams, J., Canuto, V., Dubovikov, M., 2008. Parameterization of eddy fluxes near oceanic boundaries. Journal of Climate 21, 2770–2789.

Griffies, S. M., 2012. Elements of the Modular Ocean Model (MOM): 2012 release (GFDL Ocean Group Technical Report No. 7). NOAA/Geophysical Fluid Dynamics Laboratory, Princeton, USA, 670 + xvi pages.

Griffies, S. M., Gnanadesikan, A., Dixon, K. W., Dunne, J. P., Gerdes, R., Harrison, M. J., Rosati, A., Russell, J., Samuels, B. L., Spelman, M. J., Winton, M., Zhang, R., 2005. Formulation of an ocean model for global climate simulations. Ocean Science 1, 45–79.

Griffies, S. M., Pacanowski, R. C., Hallberg, R. W., 2000. Spurious diapycnal mixing associated with advection in a *z*-coordinate ocean model. Monthly Weather Review 128, 538–564.

Holland, W. R., Chow, J. C., Bryan, F. O., 1998. Application of a third-order upwind scheme in the ncar ocean model. Journal of Climate 11, 1487–1493.

Hundsdorfer, W., Trompert, R., 1994. Method of lines and direct discretization: a comparison for linear advection. Applied Numerical Mathematics, 469–490.

Jackett, D. R., McDougall, T. J., Feistel, R., Wright, D. G., Griffies, S. M., 2006. Algorithms for density, potential temperature, conservative temperature, and freezing temperature of seawater. Journal of Atmospheric and Oceanic Technology 23, 1709–1728.

Legg, S., Hallberg, R., Girton, J., 2006. Comparison of entrainment in overflows simulated by z-coordinate, isopycnal and non-hydrostatic models. Ocean Modelling 11, 69–97.

Leonard, B. P., 1979. A stable and accurate convective modelling procedure based on quadratic upstream interpolation. Computer Methods in Applied Mechanics and Engineering 19, 59–98.

McDougall, T. J., 2003. Potential enthalpy: a conservative oceanic variable for evaluating heat content and heat fluxes. Journal of Physical Oceanography 33, 945–963.

Merryfield, W. J., Holloway, G., 2003. Application of an accurate advection algorithm to sea-ice modelling. Ocean Modelling 5, 1–15.

Pacanowski, R. C., Griffies, S. M., 1999. The MOM3 Manual. NOAA/Geophysical Fluid Dynamics Laboratory, Princeton, USA, 680 pp.

Prather, M., 1986. Numerical advection by conservation of second-order moments. Journal of Geophysical Research 91, 6671–6681.

Sweby, P., 1984. High-resolution schemes using flux limiters for hyperbolic conservation-laws. SIAM Journal of Numerical Analysis 21, 995–1011.

Winton, M., Hallberg, R., Gnanadesikan, A., 1998. Simulation of density-driven frictional downslope flow in z-coordinate ocean models. Journal of Physical Oceanography 28, 2163–2174.