Journal of Computational Physics 228 (2009) 8665–8692

Contents lists available at ScienceDirect

Journal of Computational Physics

journal homepage: www.elsevier.com/locate/jcp

High-order regridding–remapping schemes for continuous isopycnal and generalized coordinates in ocean models ☆

Laurent White [a,*], Alistair Adcroft [a], Robert Hallberg [b]

[a] Princeton University, Program in Atmospheric and Oceanic Sciences, 201 Forrestal Road, Princeton, NJ 08540, USA
[b] NOAA/Geophysical Fluid Dynamics Laboratory, Princeton University, 201 Forrestal Road, Princeton, NJ 08540, USA

ARTICLE INFO

Article history:
Received 28 April 2009
Received in revised form 19 August 2009
Accepted 22 August 2009
Available online 28 August 2009

Keywords:
Generalized vertical coordinates
Hybrid vertical coordinates
Remapping schemes
Ocean modeling
Continuous isopycnal coordinates

ABSTRACT

A hierarchy of high-order regridding–remapping schemes for use in generalized vertical coordinate ocean models is presented. The proposed regridding–remapping framework is successfully used in a series of idealized one-dimensional numerical experiments as well as two-dimensional internal wave and overflow test cases. The model is capable of replicating *z*-, sigma- and isopycnal-coordinate results, among others. Particular emphasis is placed on the design of a continuous isopycnal framework, which is a more general alternative to the layered isopycnal paradigm. Continuous isopycnal coordinates use target interface densities to define layers. In contrast to traditional layered isopycnal models, in which along-layer density gradients vanish, general coordinate approaches must deal with extra terms. For example, the calculation of pressure gradient force is more complicated and must be evaluated carefully. High-order reconstructions within boundary cells are crucial for obtaining sensible results and for reducing spurious diffusion near boundaries. Vertical advection is implicitly embedded in the remapping step and directly benefits from high-order schemes. Volume and all tracers are conserved to machine precision, which is a necessary ingredient for long-term ocean climate modeling. This hybrid vertical coordinate model provides the framework to easily capture the impact of different coordinate systems on dynamics.

1. Introduction

It is quite common to categorize ocean models according to the type of coordinates used in the vertical. Geopotential- or *z*-coordinate models use a grid for which the vertical increment at a given level does not vary horizontally (except where partial steps are used) and are particularly well suited to specifying resolution in the surface boundary layer. Terrain-following- or σ-coordinate models stretch and shrink the vertical grid in order to conform to the bottom topography and are widely used for coastal applications. Isopycnal- or ρ-coordinate models use a grid defined in terms of layers of constant potential density[1] and are ideal for representing the adiabatic nature of the ocean interior. Detailed discussions on these model categories may be found elsewhere [9,23].

Each one of these coordinates may be severely deficient in representing certain key physical processes. Two approaches have been actively pursued to remedy this issue: (1) alter and improve the subgrid-scale parameterizations to counteract issues arising from the choice of coordinates; or (2) move away from the single-coordinate paradigm and towards hybrid

☆ Revision submitted to Journal of Computational Physics on 19 August 2009.
* Corresponding author. Present address: Exxon Mobil Research and Engineering, 1545 Route 22 East, Annandale, NJ 08801, USA. Tel.: +1 908 730 2265; fax: +1 908 730 3031.
E-mail address: laurent.white@exxonmobil.com (L. White).
[1] Throughout this manuscript, density actually refers to potential density with respect to a reference pressure, usually 20 MPa in global ocean models.

0021-9991/$ - see front matter © 2009 Elsevier Inc. All rights reserved.
doi:10.1016/j.jcp.2009.08.016

representations (i.e., generalized coordinates), which are combinations of two or more vertical grid types within the same framework [16,3,8,21,2,15,23]. Due to the dynamical nature of the ocean, these hybrid coordinate systems are adapted in the course of the simulation, which is often implemented via Arbitrary Lagrangian–Eulerian (ALE) algorithms [7].

A successful hybrid coordinate system is contingent on both the regridding and remapping steps. Regridding is concerned with optimally locating the new vertical grid. The remapping step acts to remap all variables from the old grid onto the new grid. Improving the accuracy of remapping is a major research issue in hybrid coordinate ocean models. Additional complications arise when the model's regridding variable, such as density, depends on salinity and temperature via a nonlinear equation of state. When remapping of the state variables occurs, a new density profile is obtained via the equation of state. Given a set of target densities, a new grid can then be determined. If the regridding and remapping schemes are inaccurate, this coupled problem is not guaranteed to converge and the vertical grid is at risk of drifting away from any sensible state (especially when the remapping is overly diffusive). One way to circumvent this problem is to remap only one of the variables (either temperature or salinity) and, given the layer density, diagnose the variable that was not remapped [15]. However, this approach results in non-conservation of the variable that is not remapped, which is unacceptable for long-term climate simulations. The issue of building a consistent regridding scheme has yet to be fully addressed. High-order remapping schemes have already been explored by the authors [24] and the current paper extends our previous work to provide an effective and consistent hybrid coordinate framework. A consistent regridding–remapping framework is defined as one in which a motionless state is preserved by virtue of ensuring compatibility between the representation of topography, the initial conditions, the calculation of the horizontal pressure gradient, the equation of state and the remapping and regridding schemes. These compatibility conditions will be addressed in detail.

One of the objectives of building a general coordinate framework is to enable the comparison of different vertical coordinate systems within a single framework. Though such intercomparison exercises have taken place in the past [4,25], they generally involve different models, which, apart from the vertical grid, differ in many other aspects. This limitation raises the question as to whether differences in model solutions are only caused by different vertical grids and stresses the need for a single framework in which to evaluate the impact of coordinate choice.

Two directions of improvement of the generalized coordinate paradigm are investigated. First, we explore the regridding step in detail and explain how to design a continuous isopycnal framework using high-order interpolating techniques to determine the location of given target interface densities. Second, we seek to improve the reconstruction within boundary cells for both the regridding and remapping steps by no longer resorting to piecewise constants (as was done by White and Adcroft [24]), when deemed appropriate. It is shown that using high-order extrapolation at the boundaries is critical for obtaining sensible results.

2. Regridding framework

Regridding–remapping algorithms involve a regridding step, whereby a new grid is generated based on some criteria, and a remapping step, whereby the variables are remapped from the old grid onto the new grid (Fig. 1). It is generally required that the remapping be both conservative and monotonic in the sense that no new extrema should be created nor existing ones amplified. This constraint is particularly important in applications where boundedness of some variables must be guaranteed or when non-monotonicity would trigger convective adjustments [9]. High-order remapping schemes were studied by White and Adcroft [24]. We now concentrate on the regridding framework.

If a coordinate is cast in a functional form of independent variables (e.g., geopotential or terrain-following) then regridding is relatively straightforward and we refer the interested reader to the existing literature (e.g. see references mentioned by Griffies et al. [9] and Song and Hou [23]). In contrast, coordinates that are function of dependent variables (e.g., density) have traditionally been implemented in a layered formulation. Here, the layer densities are restored to target values by means of entrainment and detrainment schemes, which can be worked out in two ways. One technique is local and consists in displacing water masses between adjacent layers until target values are reached. This technique is the common approach [12,2,15] and implicitly assumes a piecewise constant representation within each layer. There are two drawbacks to this technique. (1) Because entrainement and detrainement are upwind-weighted first-order accurate schemes, numerical diffusion is large. (2) Once target densities are reached, it is typical to only remap one of the variables (either temperature or salinity) and, given the layer density, diagnose the variable that was not remapped [15]. This, however, results in the non-conservation of the variable that is not remapped, which is unacceptable. Building on this piecewise layered representation, one might want to extend this approach by reconstructing a vertical profile. Here, the objective is to determine the new grid such that the integral of density over the new layers matches target values. This scenario is illustrated in Fig. 2. It turns out that there is no unique solution defining the set of layers. To avoid this problem, we choose to use a continuous representation of density (i.e., not layered), which allows for unambiguous interpolation. We can then specify coordinates via interface target densities. We should emphasize that target densities (whether interface or layer) define the grid but do not have any physical meaning. In that respect, using target interface densities is no more arbitrary than using target layer densities.

In a regridding–remapping context, building continuous isopycnal coordinates presents two major difficulties. (1) Density depends on salinity and temperature via the equation of state. When these are remapped, each layer of the new grid inherits a new density, which alters the very density profile upon which the new grid was based. The problem is thus coupled and

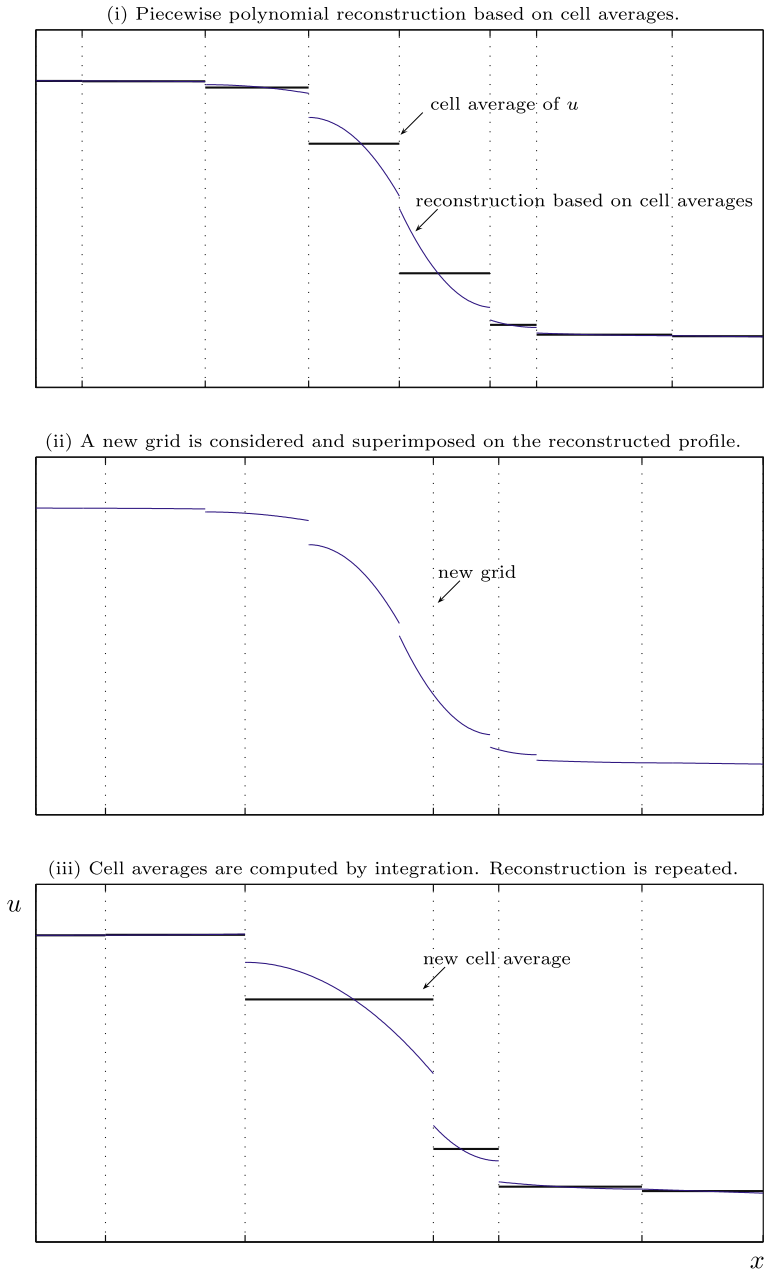

Fig. 1. A regridding–remapping algorithm occurs in three steps. The schematics illustrate a generic situation where the grid is defined by the x coordinates (horizontal axis) and depicted by the dotted lines and where u is the variable that is being remapped (vertical axis). (i) Piecewise polynomial profiles are reconstructed based on cell averages on a given grid. (ii) A new grid is considered and superimposed on the reconstructed profile. (iii) Analytic integration of the reconstructed profile over the cells of the new grid gives the cell averages for this new grid. The reconstruction step is then repeated. This illustration depicts the general case of reconstructions on nonuniform grids featuring discontinuities across cell interfaces.

achieving convergence is essential to ensure the stability of the model. (2) For practical reasons, the adiabatic character of the flow must be satisfied as accurately as possible [10]. Any spurious mixing resulting from regridding–remapping inaccuracies must be minimized. These two difficulties are addressed in this paper. In contrast, layered isopycnal models are built in such a way that they are intrinsically adiabatic.

2.1. Continuous isopycnal coordinates

A few notational conventions simplify the exposition of this work.

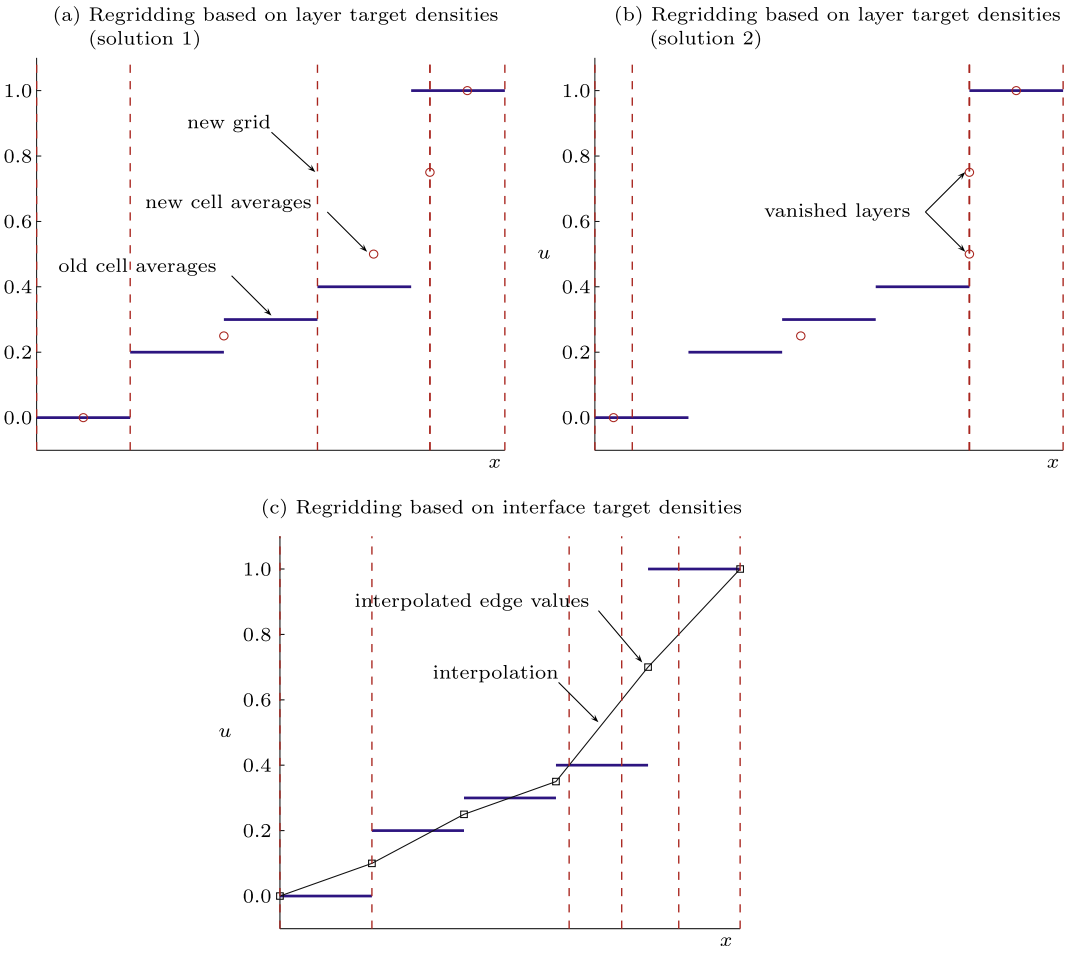

Fig. 2. Illustration of the difference in the grids obtained when using target layer values versus target interface values. In all panels, the initial grid is uniform in space and the initial cell averages are represented by thick (blue) lines. The objective is to determine the new grid (represented by the vertical dashed lines) based on target values. In panels (a) and (b), the new grid is determined such that the integral of density over the new layers matches the target values {0.0, 0.25, 0.50, 0.75, 1.0}. Depending on what direction of integration is chosen (from left to right as in (a) or from right to left as in (b)), the grid for which the new cell averages are the target values may be different. Notice that some layers have vanished. On the other hand, (c) depicts the case where a global reconstruction based on the cell averages is computed and the grid is determined by finding the location of the target interface values {0.0, 0.2, 0.4, 0.6, 0.8, 1.0}. This third method always yields a unique grid. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

Notation 1. $\mathcal{G}_N$ refers to a one-dimensional, nonuniform grid made up of N cells of widths $h_k, k = 1, \ldots, N$, and coordinates $x_k, k = 1, \ldots, N+1$, such that $h_k = x_{k+1} - x_k$.

Notation 2. We work in a finite-volume sense and every variable u is defined in terms of its cell values $\bar{u}_k$ on $\mathcal{G}_N$.

Notation 3. Within each cell of width $h_k = x_{k+1} - x_k$, use will be made of a local coordinate $\xi \in [0, 1]$ such that

$$x = x_k + (x_{k+1} - x_k)\xi = x_k + h_k\xi, \tag{1}$$

where x is the globally continuous coordinate.

The terms *conservative reconstruction* and *continuous reconstruction* will be used throughout the text. A reconstruction is defined as a piecewise polynomial representation of the data known at the cell level. A reconstruction is said to be conservative when the average of the polynomial over each cell is equal to the cell value. The conservative reconstruction over cell k is noted $R_k(x)$ and satisfies

$$\frac{1}{h_k}\int_{x_k}^{x_{k+1}} R_k(x)\,\mathrm{d}x = \int_0^1 R_k(\xi)\,\mathrm{d}\xi = \bar{u}_k. \tag{2}$$

A reconstruction is said to be continuous when it is continuous over the entire domain. The derivatives of continuous reconstructions are not required to be continuous. Note that building a profile that satisfies monotonicity, continuity and conservation is generally not possible. We therefore choose the monotonic reconstruction to satisfy conservation *or* continuity, but not both.

The regridding–remapping algorithm can be summarized as follows. We are given a grid $\mathcal{G}_N$ and cell averages $\bar{u}_k$ representing some scalar u. A function $f(u)$ and $k+1$ target values t_k are given. Cellwise values of f are noted $\bar{f}_k = f(\bar{u}_k)$. The objective is to determine the new grid $\widetilde{\mathcal{G}}_N$ such that $f_{\text{int}}(\tilde{x}_k) = t_k$, where $\tilde{x}_k$ are the coordinates defining the grid $\widetilde{\mathcal{G}}_N$ and f_{int} is a reconstruction that depends on $\bar{f}_k$ and is used to interpolate the grid. The scalar u is then remapped from $\mathcal{G}_N$ onto $\widetilde{\mathcal{G}}_N$ and a new reconstruction may be determined in order to obtain another grid. This procedure is iterated until convergence to a given tolerance. While remapping schemes must be based on conservative reconstructions, regridding schemes may rely on either conservative or continuous reconstructions.

In the context of an isopycnal coordinate ocean model, f is the density and u can be salinity or temperature. Density is the variable used for the regridding while salinity and temperature are remapped. For the sake of clarity, we assume that the function f depends only on one variable. We also assume that cell values $\bar{f}_k = f(\bar{u}_k)$ are monotonically increasing on the initial grid $\mathcal{G}_N$. This is a reasonable assumption for stably stratified flows (density increases with depth) and is required to uniquely determine the grid based on target interface values t_k. Note that the reconstruction f_{int} depends on cellwise values $\bar{f}_k = f(\bar{u}_k)$. An alternative would be to first compute the reconstruction u_{int} for the independent variable u and, then, define $f_{\text{int}} = f(u_{\text{int}})$. This method, however, does not guarantee that f_{int} is monotonically increasing when a nonlinear equation of state is used. For the sake of simplicity and perhaps at the cost of a small loss of accuracy, we choose the first approach.

2.2. Regridding schemes

As previously mentioned, regridding can be based either on conservative or continuous reconstructions. By allowing continuous, but not necessarily conservative, reconstructions, we are able to expand the suite of schemes presented by White and Adcroft [24]. Both types of reconstruction are acceptable since neither local conservation nor global continuity are required for the regridding step of the regridding–remapping algorithm. We require these profiles to be globally monotonic. Each cellwise polynomial must be monotonic; this property is termed local monotonicity. When combined with the requirement that all discontinuities of the edge values (if any) are monotonic, global monotonicity follows. Global monotonicity must be satisfied to provide unique locations of target interface values. Table 1 can be consulted for naming conventions of remapping and regridding schemes.

2.2.1. Conservative reconstructions

All remapping schemes presented by White and Adcroft [24], such as variations of the piecewise parabolic method (PPM) and the new piecewise quartic method (PQM) are monotonic and conservative and can therefore also be used for the regridding. Because these schemes are constrained to be conservative with respect to the cell values, they generally are not continuous. Regarding edge-value and edge-slope estimates, the same notation as that used by White and Adcroft [24] will be

Table 1
Naming convention for remapping and regridding schemes used in the paper. Continuity refers to continuity across cell interfaces. Conservation refers to cellwise conservation. Note that only P1M and P3M are new while all other schemes as well as discussions on explicit and implicit edge-value and edge-slope estimates were presented by White and Adcroft [24].

Name	Characteristics	Continuous	Conservative	Accuracy
PCM	Piecewise constant reconstruction	No	Yes	h
PLM	Piecewise linear reconstruction	No	Yes	h^2
P1M	Piecewise linear reconstruction	Yes	No	h^2
ih_4	4th-order accurate implicit edge values			
PPM	Piecewise parabolic reconstruction	No	Yes	h^3
h_4	4th-order accurate explicit edge values			
PPM	Piecewise parabolic reconstruction	No	Yes	h^3
ih_4	4th-order accurate implicit edge values			
P3M	Piecewise cubic reconstruction	Yes	No	h^4
ih_4ih_3	4th-order accurate implicit edge values			
	3rd-order accurate implicit edge slopes			
PQM	Piecewise quartic reconstruction	No	Yes	h^4
ih_4ih_3	4th-order accurate implicit edge values 3rd-order accurate implicit edge slopes			
PQM	Piecewise quartic reconstruction	No	Yes	h^5
ih_6ih_5	6th-order accurate implicit edge values 5th-order accurate implicit edge slopes			

used here. Hence, h_n and ih_n respectively refer to explicit and implicit nth-order accurate estimates. Explicit estimates h_n are computed by fitting a polynomial in a finite-volume sense through the data within n contiguous cells and evaluating the polynomial at the location of the edge. Implicit estimates are based on compact schemes and require the solution of a tri-diagonal system. For a given order of accuracy, implicit estimates are more accurate than explicit estimates. For a given scheme, when both the edge values and the edge slopes are used, their order of accuracy is mentioned following the scheme name (e.g., PQM ih_4ih_3 means that ih_4 edge values and ih_3 edge slopes are used).

2.2.2. Continuous reconstructions

A continuous linear reconstruction based on the cell values (referred to as P1M, not to be confused with the discontinuous PLM) requires the determination of the single edge value at the interfaces between cells. A continuous cubic reconstruction (referred to as P3M) requires the estimation of the edge slopes in addition to the edge values. The slopes do not need to be uniquely defined at each edge: the slopes may differ when the edge is approached from the left and right. These two schemes are second- and fourth-order accurate, respectively. To be monotonic, the continuous linear reconstruction simply requires

Table 2
Orders of accuracy of a selection of regridding–remapping schemes considered in the paper. An order of accuracy h^n means that the scheme can exactly retrieve a global polynomial profile of degree $n-1$ based on the cell values. In other words, it means that the new grid obtained via regridding will be exact *and* the remapping between the old and new grids will be exact too. The schemes are the following: P1M ih_4 (continuous piecewise linear reconstruction with ih_4 edge-value estimates), PLM (classical conservative piecewise linear reconstruction), PPM h_4 (classical conservative piecewise parabolic reconstruction, as presented by Colella and Woodward [5]), PPM ih_4 (conservative piecewise parabolic reconstruction with ih_4 edge-value estimates), P3M ih_4ih_3 (continuous piecewise cubic reconstruction with ih_4 edge-value estimates and ih_3 edge-slope estimates), PQM ih_4ih_3 (conservative piecewise quartic reconstruction with ih_4 edge-value estimates and ih_3 edge-slope estimates), PQM ih_6ih_5 (conservative piecewise quartic reconstruction with ih_6 edge-value estimates and ih_5 edge-slope estimates).

Regridding	Remapping					
	PCM	PLM	PPM h_4	PPM ih_4	PQM ih_4ih_3	PQM ih_6ih_5
P1M ih_4	h	h^2	h^2	h^2	h^2	h^2
PLM	h	h^2	h^2	h^2	h^2	h^2
PPM ih_4	h	h^2	h^3	h^3	h^3	h^3
P3M ih_4ih_3	h	h^2	h^3	h^3	h^4	h^4
PQM ih_6ih_5	h	h^2	h^3	h^3	h^4	h^5

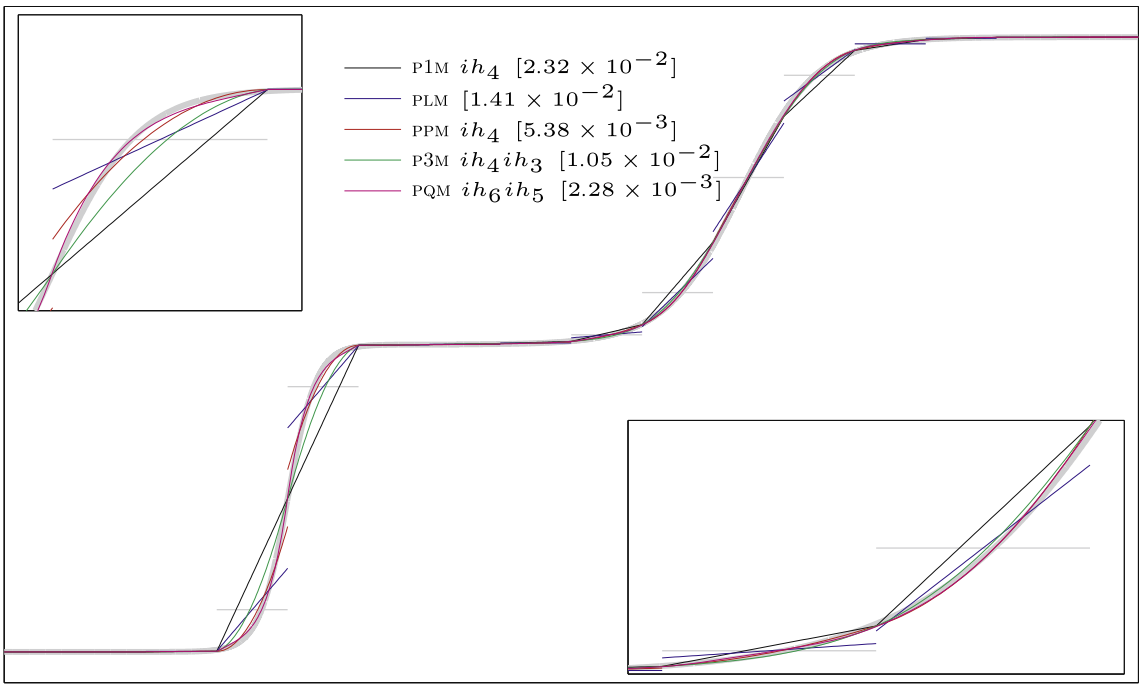

Fig. 3. Comparison of continuous (P1M ih_4 and P3M ih_4ih_3) and conservative (PLM, PPM ih_4 and PQM ih_6ih_5) reconstruction schemes. The exact profile is depicted by a thick, light gray line. The reconstruction schemes are based on the analytical cell values represented by horizontal light gray lines. There are 16 uniformly-distributed cells. PLM and PPM are discontinuous while P1M and P3M are continuous (by construction). Because PQM allows the representation of sharper curvatures, it remains continuous while being conservative. The L_2-norm of reconstruction errors are indicated in the legend. Overall, discontinuous schemes perform much better than continuous schemes of the same order because they can represent sharp gradients, with PQM being the best. The profile was so chosen to include two sharp elbows (on the left) and two smoother elbows (on the right), separated by a plateau in the interior.

the edge values to lie between neighboring cell averages. The variety of high-order edge-value estimates can be used here, together with the technique to bound the edge values to ensure monotonicity [24]. In contrast, the continuous cubic reconstruction needs to be properly limited, as described in detail in the appendix.

2.3. A hierarchy of regridding–remapping schemes

Given the regridding schemes introduced above and the remapping schemes presented by White and Adcroft [24], many choices are available for the regridding–remapping algorithm. Table 2 summarizes the selection of 30 regridding–remapping schemes that are considered in this paper, together with their order of accuracy. An order of accuracy h^n means that the scheme can exactly retrieve a global polynomial of degree $n-1$ based on the cell values. In other words, for this polynomial of degree $n-1$, the new grid obtained during the regridding step (via interpolation) will be exact *and* the remapping between the old and new grids will be exact too. A comparison between a selection of continuous and conservative reconstruction schemes is provided in Fig. 3. The test profile in Fig. 3 is defined on [0, 1] as follows:

$$\begin{cases} \frac{1}{4}[1+\tanh(40(x-0.25))] & \text{if } x \leqslant 0.4, \\ \frac{1}{2}+\frac{1}{4}[1+\tanh(15(x-0.65))] & \text{otherwise.} \end{cases} \tag{3}$$

Conservative schemes outperform continuous ones of the same order because they can represent sharp gradients more accurately by allowing discontinuities.

3. Extrapolation schemes at the boundaries

To preserve monotonicity of cell values, the reconstruction within boundary cells is necessarily limited to piecewise constants. In other words, boundary cell values are treated as extrema. A boundary cell is a cell whose one of the edges coincides with a boundary (see Fig. 4). As will be shown in test cases, there is a lot to gain from relaxing the monotonicity constraint within boundary cells. This relaxation can be done by carefully extrapolating the boundary edge values (and edge slopes when needed) to reconstruct the profile.

In the following, it is assumed that the reconstruction within all interior cells is known and boundary cell reconstructions need to be determined. The following procedures apply to the left boundary cell only (the cell with the boundary to the immediate left). Extension to the right boundary cell is straightforward. The indexes 0 and 1 refer to the boundary cell and the cell next to it, respectively. Cell widths are noted h_0 and h_1. Cell values are noted $\bar{u}_0$ and $\bar{u}_1$. The reconstructions are noted $R_0(\xi)$ and $R_1(\xi)$, where use is made of the local coordinate. $R_0(\xi)$ must be determined and $R_1(\xi)$ is known. A comparison of several extrapolation schemes is illustrated in Fig. 4.

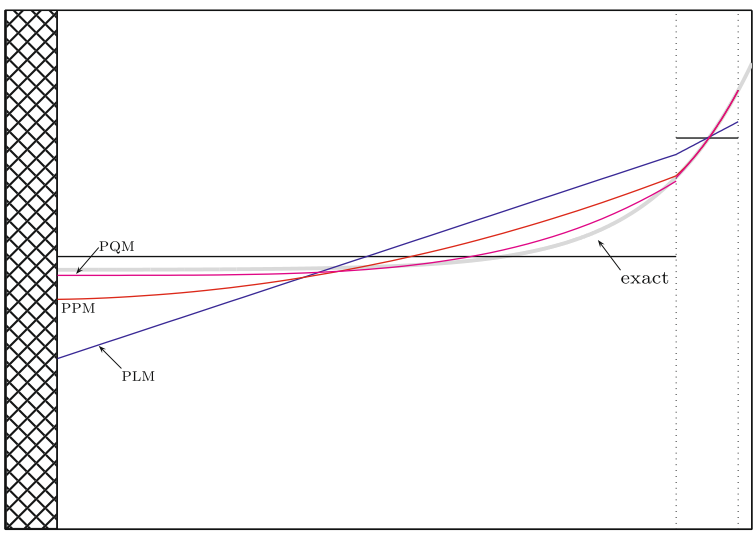

Fig. 4. Comparison of PLM, PPM and PQM extrapolation schemes within the left boundary cell. The thick, light gray line represents the exact solution. Cell averages are depicted by black, horizontal lines. The grid is represented by vertical dotted lines. The crosshatched area on the left indicates the location of the boundary. PQM performs the best.

3.1. Extrapolation for the P1M scheme

Two edge values are needed for P1M. The right edge value u_R is known and the left one must be extrapolated. The one-sided slope, expressed in the local coordinate system of the boundary cell, is first computed:

$$\sigma = 2\frac{\bar{u}_1 - \bar{u}_0}{h_0 + h_1} h_0.$$

The P1M reconstruction is then given by

$$R_0(\xi) = u_R + \sigma(\xi - 1).$$

3.2. Extrapolation for the PLM scheme

Only the slope must be determined. Enforcing local conservation fixes the second degree of freedom. The slope is computed based on the right edge value u_R and the cell average $\bar{u}_0$. Expressed in the local coordinate system of the boundary cell, we have

$$\sigma = 2(u_R - \bar{u}_0).$$

The PLM reconstruction is then given by

$$R_0(\xi) = \bar{u}_0 + \sigma\left(\xi - \frac{1}{2}\right).$$

3.3. Extrapolation for the PPM scheme

To determine a parabola within the boundary cell, we need to determine three degrees of freedom. The right edge value u_R and slope u'_R are known from the adjacent cell. The constraint of local conservation closes the system. Given these three parameters, the left edge value is given by

$$u_L = 3\bar{u}_0 + \frac{1}{2}u'_R - 2u_R.$$

Given the parameters u_L, u_R and $\bar{u}_0$, the resulting parabola is limited using the standard procedure developed by Colella and Woodward [5].

3.4. Extrapolation for the P3M scheme

Higher-order extrapolations, such as cubic, are trickier because the use of directionally-biased data often leads to large-amplitude oscillations. As a consequence, a cubic based on local conservation as well as on the right edge value, slope and curvature yields inaccurate estimates for the left edge value and slope and a very inaccurate reconstruction overall. A method of extrapolation less subject to oscillation is needed, which rational functions can fulfill. The rational functions presented by Xiao et al. [27] are monotonic by construction and possess three degrees of freedom. The general form is

$$L(\xi) = \frac{a + 2b\xi + \beta b\xi^2}{(1 + \beta\xi)^2}, \tag{4}$$

where a, b and β are parameters to be determined. This determination can be done by enforcing local conservation and the right edge value and slope (which are known). Doing so, we obtain

$$\beta = \frac{2(u_R - \bar{u})}{u_R^{prime}} - 1,$$
$$b = u_R(\beta + 1) - \bar{u},$$
$$a = \bar{u}(\beta + 1) - b.$$

Once the rational function is known, the left edge value and slope are given by

$$u_L = a$$

and

$$u'_L = 2(b - a\beta)$$

respectively. Given both edge values and both edge slopes, the cubic is completely determined and limited following the procedure described in the appendix.

The oscillation-free nature and accuracy provided by rational functions come at a price. A rational function such as Eq. (4) has a singularity at $\xi_p = -1/\beta$, the location of which matters. Of course, the singularity must not lie in [0, 1]. However, this constraint alone does not prevent erratic behavior. Let us assume that the data set is monotonically increasing ($\bar{u}_1 > \bar{u}_0$). In that case, enforcing local conservation constrains the boundary edge-value estimate to be smaller than $\bar{u}_0$. If $\xi_p < 0$, there is

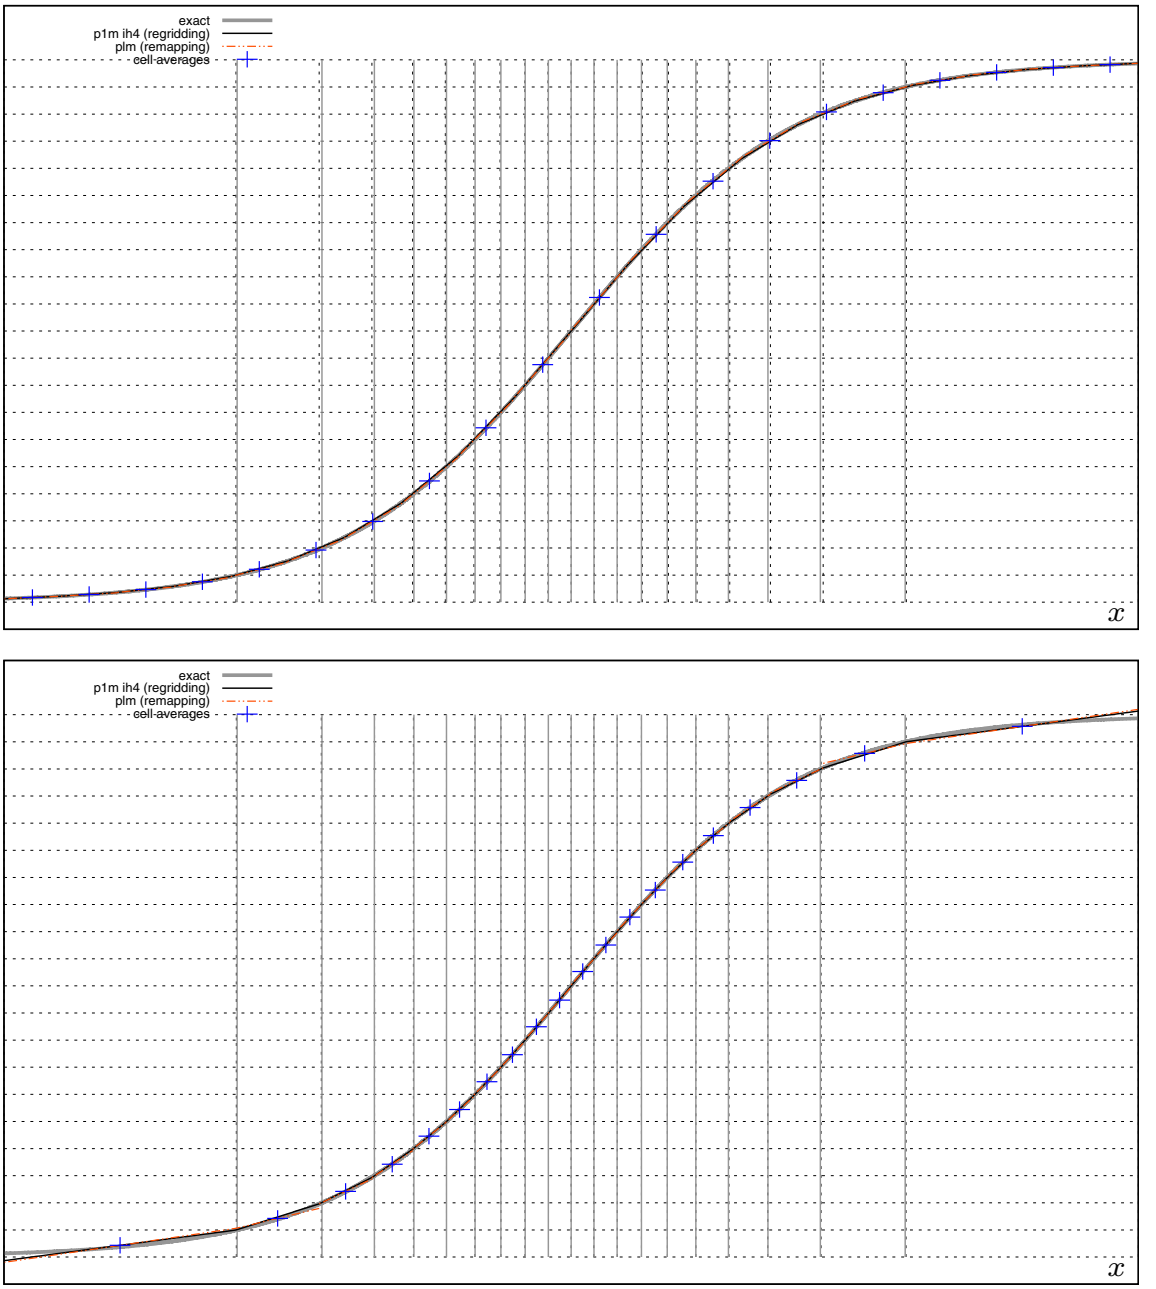

Fig. 5. Illustration of the first full iteration for the regridding–remapping algorithm used for continuous isopycnal coordinates. P1M ih_4 and PLM are used for regridding and remapping, respectively. The exact profile is depicted by a thick gray line. The cell averages are indicated by "+" symbols, which are located at the cell centers of the previous grid. In the top panel, cell averages are known on a uniform grid (this is the initial condition of the problem). The P1M ih_4 scheme is used to find the location of interfacial target values, which are represented by the horizontal dashed lines. The corresponding grid is represented by vertical dashed lines. The dashed lines depict the new grid onto which the remapping occurs based on the PLM reconstruction. The exact grid is drawn with gray vertical lines. When the vertical dotted and gray lines coincide, it means that the approximate grid is close to the exact grid. The bottom panel shows the situation after remapping has occurred; the new cell averages have been computed. A new P1M ih_4 reconstruction is determined based on the new cell averages in order to determine the next grid, which is depicted by the dashed line. PLM-based remapping then occurs to compute the new cell averages. Iterations continue until convergence (to a tolerance of 10^{-6}) of the grid is attained.

no guarantee that the value of the rational function at the boundary will be bounded at all. This behavior can be likened to that of the function $-1/(x+\epsilon)$ where $\epsilon \ll 1$. The desired behavior would be obtained by requiring $\xi_p > 1$. In that case, it is guaranteed that the rational function is asymptotic within the cell and, therefore, bounded from below. An example function is $-1/(x-1-\epsilon)$. By examining Eq. (4), the singularity is located at:

$$\xi_p = \frac{u'_R}{u'_R - 2(u_R - \bar{u}_0)}.$$

Since we have assumed $u'_R > 0$, to ensure that $\xi_p > 1$ requires that $u_R > \bar{u}_0$ and

$$u'_R > 2(u_R - \bar{u}_0). \tag{5}$$

Note that when $u'_R = 2(u_R - \bar{u}_0), \beta = 0$ and the rational function reverts to PLM. When the rational function does not behave properly, i.e., when Eq. (5) is not obeyed, the boundary edge value and slope are estimated using PPM extrapolation, as described above.

3.5. Extrapolation for the PQM scheme

Extrapolation for PQM works in a way very similar to the cubic extrapolation scheme. A quartic has five degrees of freedom. We use local conservation and the right edge value and slope, as calculated from the quartic in the adjacent cell. The boundary edge value and slope are determined using the rational function, Eq. (4). When the latter is not appropriate because of the location of the singularity, left edge value and slope are estimated using PPM. The quartic is then limited following the procedure detailed by White and Adcroft [24].

4. Convergence and error analysis

Convergence and error analyses of remapping schemes alone have been investigated by White and Adcroft [24]. In this section, we briefly explore the convergence properties of regridding–remapping schemes for use in the context of continuous isopycnal coordinates. Idealized one-dimensional test cases are considered.

A full regridding–remapping iteration is presented in Fig. 5 for the profile defined by

$$u(x) = \frac{1}{2}[1 + \tanh(5(x - 0.5))] \tag{6}$$

Table 3
Performance evaluation of 30 regridding–remapping schemes used to determine the location of a grid based on the profile given by Eq. (6) and on 21 uniformly-distributed target values. Each scheme is assessed with the help of four numbers. Each of these numbers is explained and defined in Section 4. The first one is the number of iterations required to attain grid convergence, which happens when the deviation between successive grids, as defined by Eq. (7), is small enough. The second one is a measure of the error on the last grid between the exact profile and the reconstruction used for remapping. It is defined by Eq. (8). The third one, as defined by Eq. (9), is a measure of the error between the last approximate and exact grids. The fourth, and last, one is a measure of the error on the last grid between the approximate cell values and exact cell values. It is defined by Eq. (10). See text for explanations as to the likely cause for the loss of accuracy (see the grid error) observed for PQM ih_6ih_5–PQM ih_6ih_5 compared with PQM ih_6ih_5 (regridding)–PQM ih_4ih_3 (remapping).

Regridding	Remapping					
	PCM	PLM	PPM h_4	PPM ih_4	PQM ih_4ih_3	PQM ih_6ih_5
P1M ih_4	38	7	8	8	8	7
	1.72(−2)	4.80(−3)	1.34(−3)	1.37(−3)	5.22(−4)	2.04(−4)
	6.91(−3)	3.19(−4)	1.53(−4)	1.40(−4)	1.49(−4)	1.54(−4)
	5.81(−3)	3.30(−4)	5.16(−5)	2.59(−5)	1.14(−5)	1.08(−6)
PLM	>50	6	8	8	7	8
	1.87(−2)	4.49(−3)	2.82(−3)	2.93(−3)	5.40(−4)	1.82(−4)
	8.33(−3)	2.36(−3)	1.68(−3)	1.67(−3)	2.04(−3)	2.03(−3)
	7.94(−3)	4.21(−4)	2.98(−4)	3.02(−4)	1.48(−5)	6.40(−6)
PPM ih_4	>50	4	3	3	3	3
	4.48(−2)	4.88(−3)	1.35(−3)	1.39(−3)	5.31(−4)	2.06(−4)
	2.90(−2)	2.03(−4)	5.01(−5)	3.29(−5)	2.90(−5)	2.60(−5)
	3.04(−2)	3.41(−4)	5.32(−5)	2.61(−5)	1.16(−5)	7.69(−7)
P3M ih_4ih_3	>50	7	4	4	3	4
	1.90(−2)	4.82(−3)	1.43(−3)	1.48(−3)	5.32(−4)	2.02(−4)
	1.18(−2)	3.94(−4)	1.24(−4)	1.07(−4)	1.59(−4)	1.50(−4)
	9.79(−3)	3.46(−4)	5.74(−5)	3.26(−5)	1.16(−5)	9.17(−7)
PQM ih_6ih_5	>50	9	6	4	3	3
	4.57(−2)	5.39(−3)	1.60(−3)	1.40(−3)	5.31(−4)	2.08(−4)
	7.53(−3)	2.74(−3)	2.74(−4)	3.15(−5)	2.07(−5)	1.12(−4)
	4.44(−3)	3.57(−4)	7.41(−5)	2.67(−5)	1.16(−5)	8.56(−7)

and using P1M ih_4 for the regridding step and PLM for the remapping step. Note that, in this particular case, we are simply assuming $f(u) = u$, which eases the understanding of the algorithm. A detailed explanation of the iterative procedure is given in the caption of Fig. 5.

In Table 3, several error measures are used to evaluate the performance of 30 regridding–remapping schemes in seeking the location of the 21 target values $\{0.0, 0.05, \ldots, 1.0\}$ for the profile defined by Eq. (6). Note that when a given target value is out of range, its location is set to be equal to that of the overtaken boundary. Each cell of Table 3 contains four numbers, which are explained hereafter.

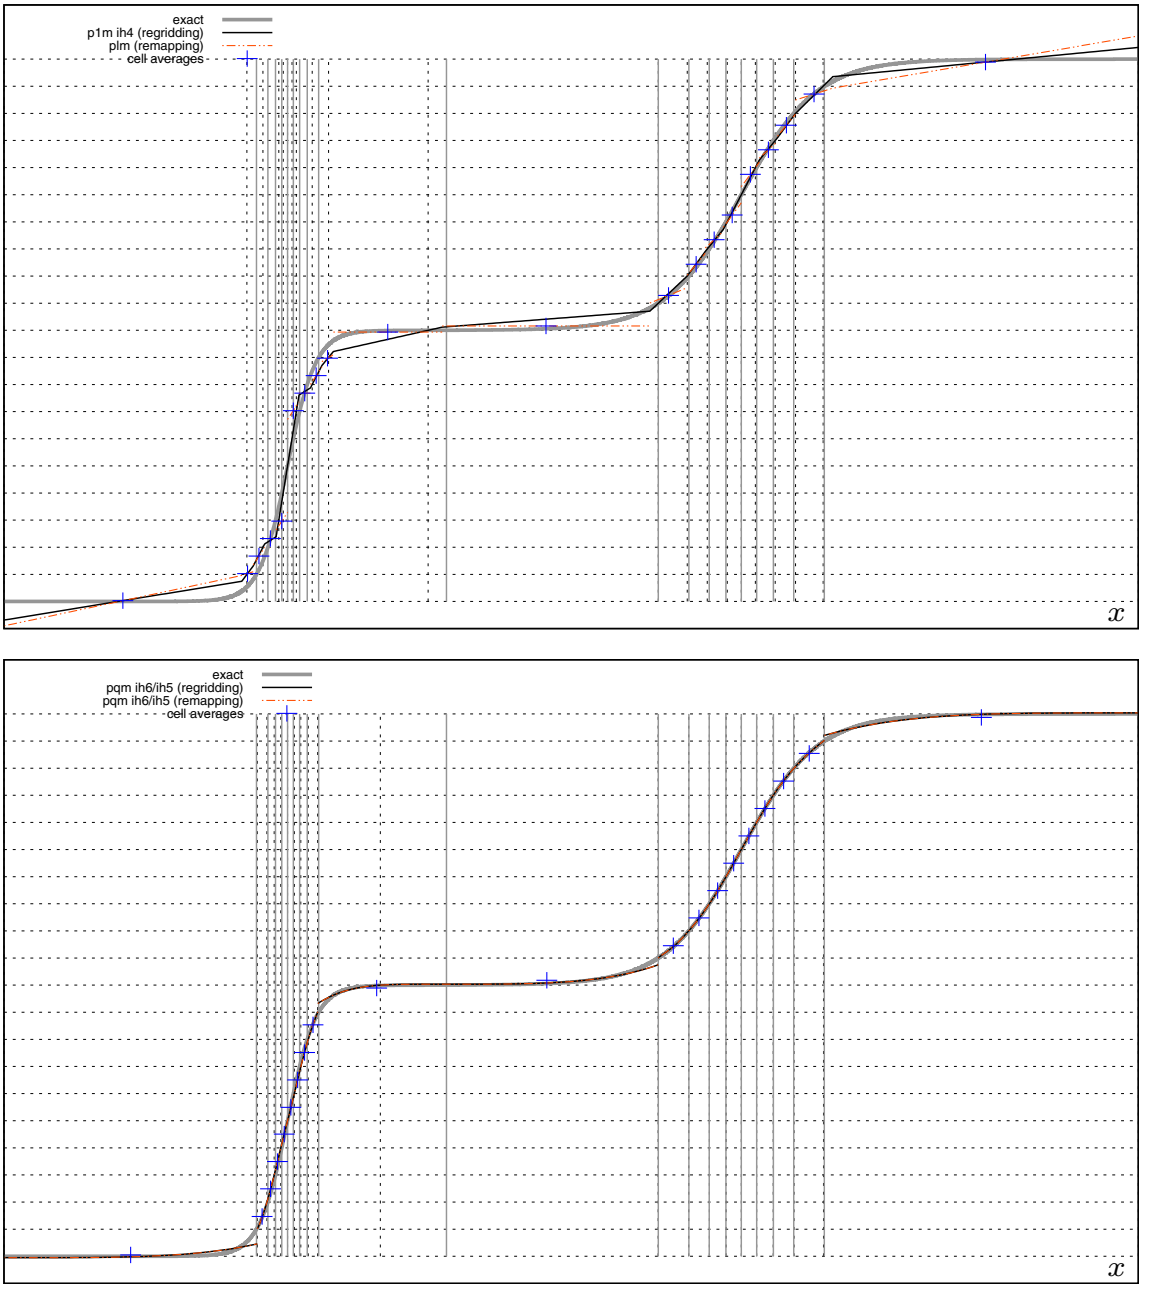

Fig. 6. Comparison between low-order (top) and high-order (bottom) regridding–remapping schemes after the first iteration for a profile containing sharp features (same profile as that of Fig. 3). In both cases, the reconstructions are based on the cell averages symbolized by "+". These averages were obtained after the first iteration. Notice the inability of low-order schemes to represent sharp features and to properly extrapolate within boundary cells. These shortcomings lead to slower convergence and low accuracy caused by rapid degradation of the solution.

- The first one is the number of iterations required to attain grid convergence to a tolerance of 10^{-6}. The deviation between successive grids is calculated as

$$D^{(m)} = \left[\frac{1}{N+1}\sum_{k=1}^{N+1}(x_k - \tilde{x}_k)^2\right]^{1/2}, \tag{7}$$

  where m is the iteration number, x_k are coordinates of the old grid and $\tilde{x}_k$ are coordinates of the new grid. Iterations stop when $D^{(m)}$ is smaller than a given tolerance.
- The second number of each cell in Table 3 is the L_2-norm of the error between the exact profile and the reconstruction used for remapping, computed on the last grid. It is defined as

$$\mathcal{E}_1 = \left[\sum_{k=1}^{N}\int_{x_k}^{x_{k+1}}(u(x) - R_k(x))^2\,\mathrm{d}x\right]^{1/2}, \tag{8}$$

  where $u(x)$ is the exact profile and $R_k(x)$ is the reconstruction on cell k. When the reconstruction used for remapping is exact, this reconstruction error vanishes. This property remains true even when the reconstruction used for regridding is inexact. In that case, the new grid will be inexact but since the grid-on-grid remapping is exact, cell averages and reconstructions remain exact.
- The third number of each cell in Table 3 is the grid error, which measures the error between the approximate grid obtained at the end of the iterative procedure and the exact grid based on the exact profile. The error measure is computed as follows:

$$\mathcal{E}_2 = \left[\frac{1}{N+1}\sum_{k=1}^{N+1}\left(x_k - x_k^{\text{exact}}\right)^2\right]^{1/2}. \tag{9}$$

  For the grid error to vanish, both the reconstruction used for regridding and the reconstruction used for remapping must be exact. It is therefore the most comprehensive error measure of all.
- The fourth, and last, number of each cell in Table 3 is a measure of the error between cell values obtained on the last grid and the exact cell values computed on that grid and based on the analytical profile. The error is calculated as follows:

$$\mathcal{E}_3 = \left[\sum_{k=1}^{N} h_k\left(\bar{u}_k - \bar{u}_k^{\text{exact}}\right)^2\right]^{1/2}, \tag{10}$$

  which vanishes when the reconstruction used for remapping is exact. In Table 3, the error measures are nondimensional.

The most reliable schemes are those that are able to quickly determine an accurate grid. These schemes have the smallest grid error, Eq. (9), and they converge to the final grid in few iterations. In addition, if the reconstruction error, Eq. (8), is small, any subsequent arbitrary regridding (i.e., not necessarily isopycnal) will yield accurate cell values through remapping. In view of these criteria, any scheme using PCM for remapping should be ruled out, if only for the lack of convergence. The

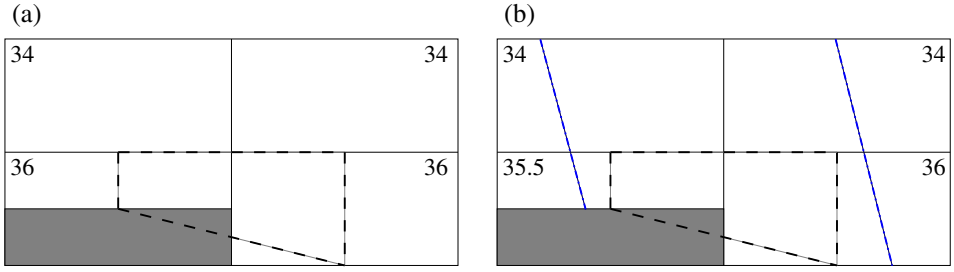

Fig. 7. Pressure gradient force issue and solution illustrated using a two-layer configuration with bottom topography occupying the lower half of the left bottom cell. Four tracer cells are shown. The number within each tracer cell is the salinity (constant value per cell) and we assume that density is equal to salinity. The pressure gradient force is computed at velocity points (staggered with respect to tracer points) by integration along the sides of the volume represented by the dashed line [1], which effectively assumes a piecewise linear representation of topography. The objective is to set up the initial conditions and modify the pressure gradient force calculation in order for the model to be motionless. (a) In the original layered approach, reconstruction is constant within each layer and density is constant along a layer. The right side integral exactly counterbalances the left side integral and the bottom side integral. The pressure gradient force is zero and there is no motion. However, in a regridding context, both layer thicknesses must be determined by seeking the location of the interface density 35. Let us assume PLM is used for regridding and remapping. If the original configuration is used, the location of interface density 35 will be higher than the current interior interface within the left column. In the right column, the interface will remain in place because it precisely corresponds to an interface density of 35 based on a linear reconstruction. If the left interior interface moves upward, the pressure force is no longer zero and motion occurs, which we want to avoid. (b) To avoid motion *and* maintain the current layer thicknesses, the salinity within the bottom left cell must be such that the piecewise linear reconstruction (shown with the dashed-dotted line) is the same within each column. In doing so, the regridding step will not displace the interior interface. For the pressure gradient force to be zero, side integrals must no longer assume constant reconstruction and must take into account the linear reconstruction.

schemes based on PQM remapping by far outperform their peers when PPM, P3M or PQM are used for regridding. Smaller grid errors are achieved by using various combinations of PPM and PQM schemes with the smallest being obtained via PQM ih_6ih_5 (regridding)–PQM ih_4ih_3 (remapping).

Estimating ih_4 and ih_6 edge values and ih_3 and ih_5 edge slopes require solving tridiagonal systems. As explained by White and Adcroft [24], for ih_4 edge values and ih_3 edge slopes, the entries of the systems are closed-form expressions. For ih_6 edge values and ih_5 edge slopes, however, the entries of the tridiagonal system are determined by solving 6×6 linear systems when the grid is nonuniform (for uniform grids, closed-form expressions are easy to obtain). The linear systems are composed of powers of coordinates, up to x^6, which tends to produce relatively inaccurate solutions when the grid contains thin cells and abrupt changes in resolution. This situation is most likely the cause for the loss of accuracy (see the grid error) observed for PQM ih_6ih_5–PQM ih_6ih_5 compared with PQM ih_6ih_5 (regridding)–PQM ih_4ih_3 (remapping) in Table 3.

Using higher-order schemes also becomes valuable when the profile contains sharp features that cause the resolution to be concentrated. Despite large variations in the grid resolution, higher-order schemes are still able to yield accurate reconstructions overall. This behavior is exemplified in Fig. 6 where a comparison between low-order (P1M ih_4-PLM) and high-order

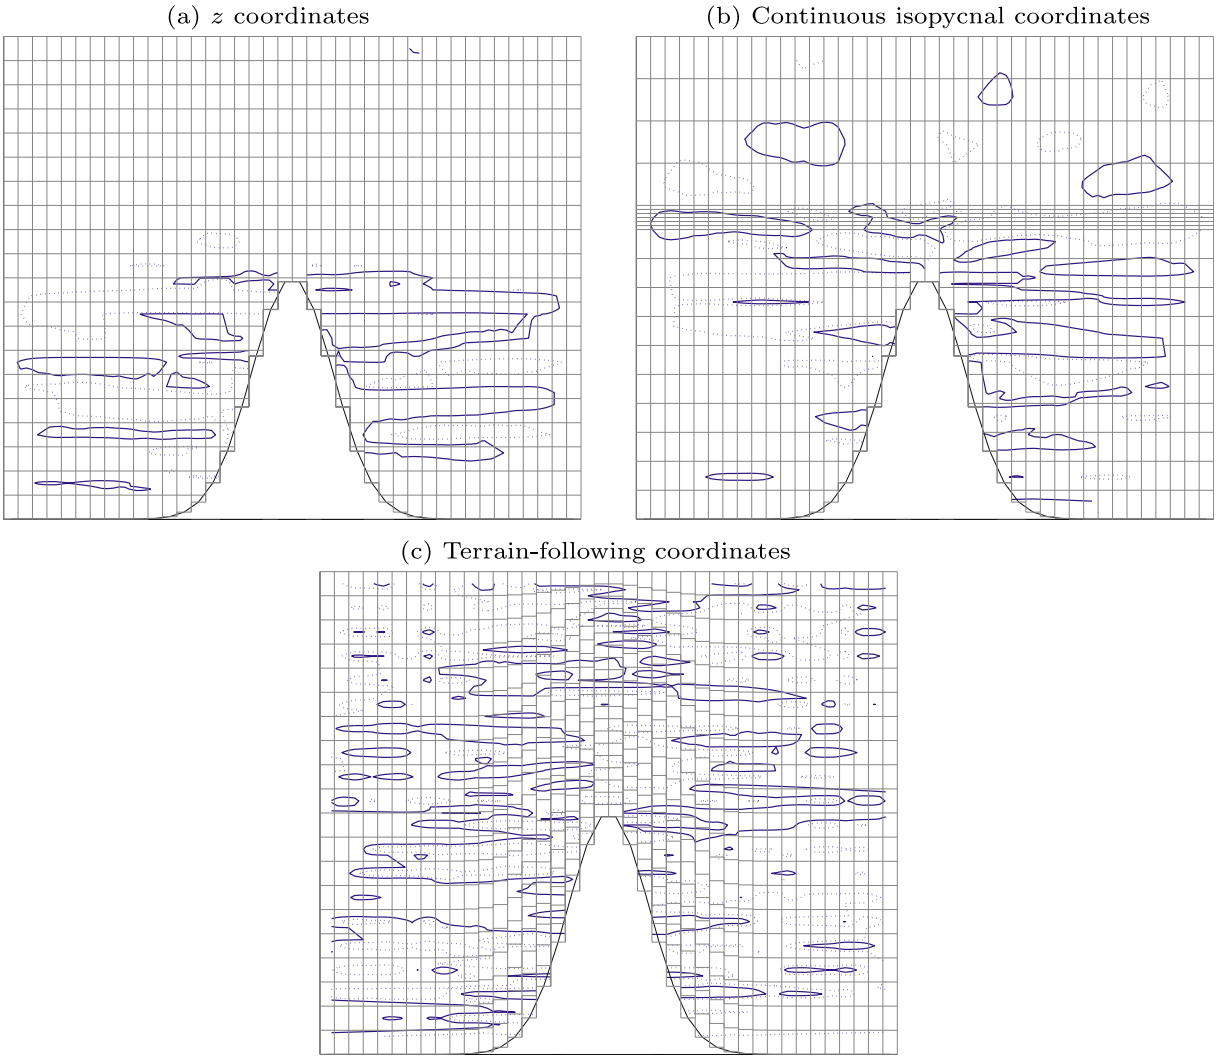

Fig. 8. Velocity contours representative of a motionless state using different coordinate systems. The Gaussian sea mount lies in a 4000-m deep, 200-km wide basin. Stratification is linear and a linear equation of state is used. There is no forcing and initial conditions are prescribed to ensure motionless evolution, as explained in Fig. 7. There is no viscosity, no diffusion and no bottom drag. There are 20 layers in the vertical and 40 cells across. The grid is depicted in light gray. The dynamical and thermodynamical time steps are both 900 s. Regridding–remapping is carried out at every time step using PPM ih_4 for remapping (and also regridding in (b)). Contours are drawn at velocities of 10^{-12} m s^{-1} (solid line) and -10^{-12} m s^{-1} (dotted line) at day 5 for z and continuous isopycnal coordinates. Contours of $\pm 10^{-13}$ m s^{-1} are used for terrain-following coordinates. The maximum absolute value is about 10^{-11} m s^{-1}, which is due to round-off errors originating in the pressure gradient force calculation. (a) z coordinates using partial cells. (b) Continuous isopycnal coordinates using a non-uniform target density distribution to illustrate the possibility of locally increasing vertical resolution. (c) Terrain-following coordinates. Velocity anomalies are smaller for the terrain-following coordinates because there is no vanishing layer.

(PQM ih_6ih_5–PQM ih_6ih_5) regridding–remapping schemes is presented after the first iteration. The low-order scheme fails to obtain an accurate grid, fails to capture sharp elbows and is inaccurate within boundary cells. The high-order scheme provides a remedy to these three flaws. In this example, the lack of resolution near boundaries is due to the isopycnal nature of the grid and may be problematic. If the grid is to be subsequently modified to include more cells near the boundaries – e.g., to refine the bottom boundary layer or the surface mixed layer – inaccurate boundary reconstructions would lead to inaccurate cell values and could impair the physical integrity of the solution. In that respect, we notice that Table 3 shows a decrease in the reconstruction error (second number in each cell) when, for a given regridding scheme, the remapping is improved. This error reduction is mostly due to an improved boundary representation since the smooth interior does not pose any particular problem.

In all experiments presented in the following section, only one regridding iteration is performed when building continuous isopycnal grids (as opposed to iterating until convergence to a small tolerance). In practical applications, iterating until convergence would be too computationally expensive because it would require several computations of reconstructed profiles for both the regridding and the remapping at each time step. We do care, however, about the ability of the algorithm to converge to a final state. In situations where a steady state is reached, it is important that successive regridding–remapping events do not alter the dynamics. This property of non-alteration is the motivation behind investigating the convergence of the regridding–remapping algorithm.

5. Test cases

In this section, two-dimensional test cases are considered where different vertical coordinate systems are compared and the effect of regridding–remapping schemes on continuous isopycnal coordinate configurations is investigated. We intentionally do not include any experiment using the full-fledged version of the ocean general circulation model. Such experiments will be the subject of a forthcoming paper focusing on the choice of the most adequate vertical grids for realistic, large-scale simulations, using the material presented in this paper.

5.1. Ocean model description

The prototype ocean model we use is built on the layered isopycnal ocean model HIM (Hallberg Isopycnal Model) [14], which was enhanced with the regridding–remapping algorithm presented in this paper to provide generalized vertical coordinate capabilities. An Arakawa C grid is used to discretize the equations layerwise. The dynamics is split into barotropic and baroclinic modes using two different time steps [11] and both estimates of the free-surface elevation fields are reconciled using the approach advocated by Hallberg and Adcroft [13]. The pressure gradient force is analytically calculated following Adcroft et al. [1]. This algorithm assumes a constant representation of temperature and salinity within each layer, which had to be modified to account for higher-order polynomial reconstructions. As will be shown, this step is necessary to ensure consistency with respect to the initial conditions, which allows for the representation of a motionless state.

As implemented, the regridding–remapping toolbox and the original version of the model are black boxes to each other. If regridding is activated, a simple flag is switched on. Otherwise, the original layered isopycnal version of the model is used. When a new grid is to be defined, the model transfers the primitive model variables – layerwise velocity components, layer thicknesses, temperature, salinity, various tracers – to the regridding–remapping toolbox. The latter takes care of designing a

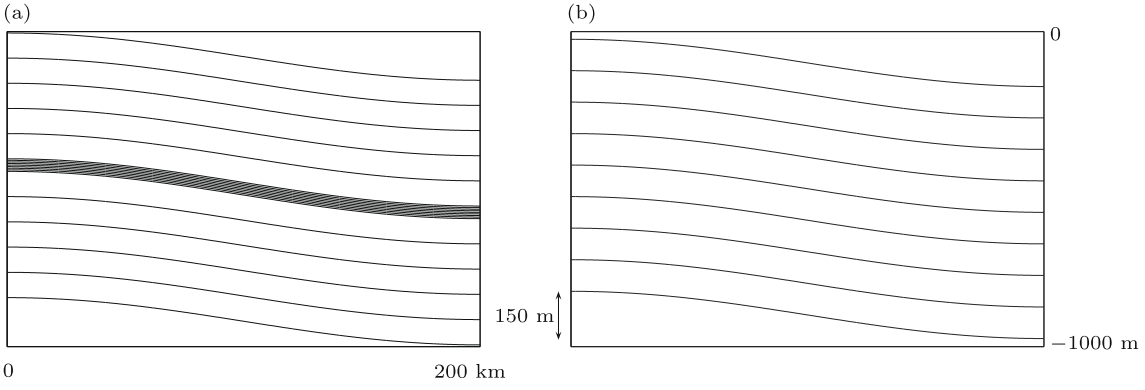

Fig. 9. Initial conditions used for the internal wave experiment. The basin is 1000 m deep and 200 km wide. There are 160 cells in the horizontal. In both cases, salinity is uniformly distributed across layers, independently from their thicknesses (the salinity increment from layer to layer is constant). Salinity varies from 34 at the top to 36 at the bottom. A linear equation of state is used and only depends on salinity. Initial interfacial displacements are prescribed according to a cosine function. The total displacement from left to right is 150 m. (a) There are 20 layers and a thin pycnocline comprises 8 of them. (b) Layer thicknesses are uniformly distributed.

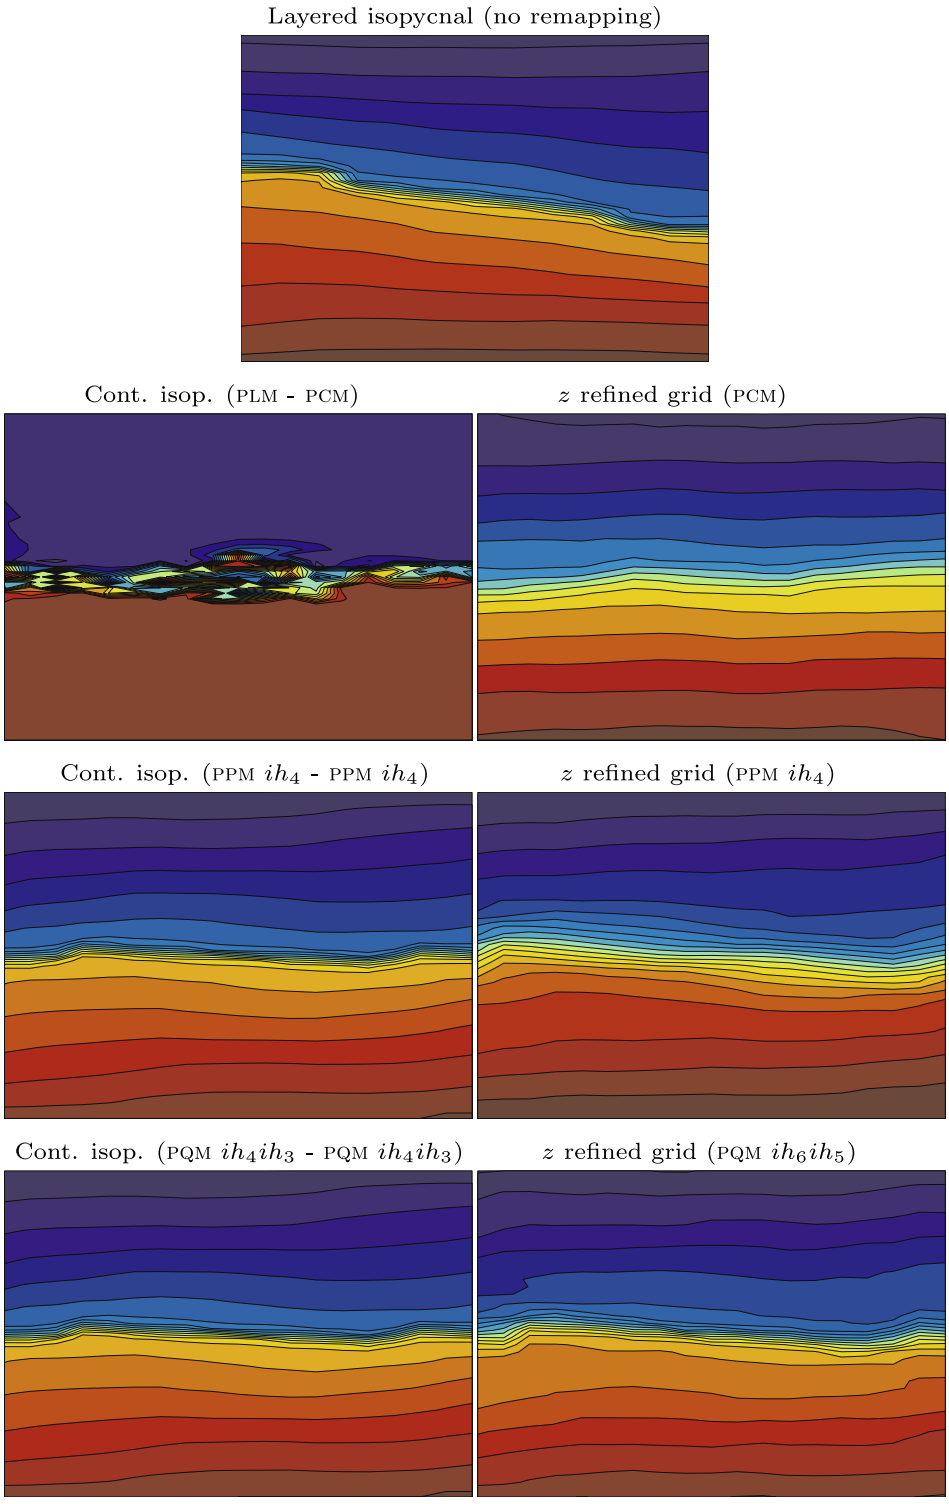

Fig. 10. Salinity contours (dark blue is 34, dark red is 36, increment is 0.2) at day 25, starting with the initial state represented in Fig. 9(a). Results from the layered isopycnal version of the model is shown in the top panel, in which case there is no remapping, no diabatic effects and the vertical structure of the pycnocline is well preserved. In this regard, this solution is the reference. Continuous isopycnals based on high-order regridding–remapping schemes are able to preserve the sharp pycnocline. Note that PLM–PCM is so diffusive that it erodes all stratification. When using z coordinates (with local refinement at mid-depth), only PQM is capable of decently resolving the sharp salinity gradients around the pycnocline. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

new grid and remapping all variables onto the new grid. Those variables are then transferred back to the model's dynamical core to continue the time integration. Regridding–remapping capabilities are called after one or several thermodynamical time steps. Between regridding events, the dynamics simply evolves according to the layered version of the model. When regridding–remapping is activated at every thermodynamical time step, the extra computational cost incurred varies between 5% and 10%, depending on which schemes are used.

5.2. Motionless state

In a model configuration without external forcing, where densities are constant layerwise and the pressure gradient force calculation assumes so, the model state remains motionless unless layer interfaces are inclined. This motionless state is represented in Fig. 7(a) for a very simple two-column, two-layer configuration. In a continuous isopycnal-coordinate framework where regridding and remapping are used, achieving this motionless state requires the initial state and the pressure gradient

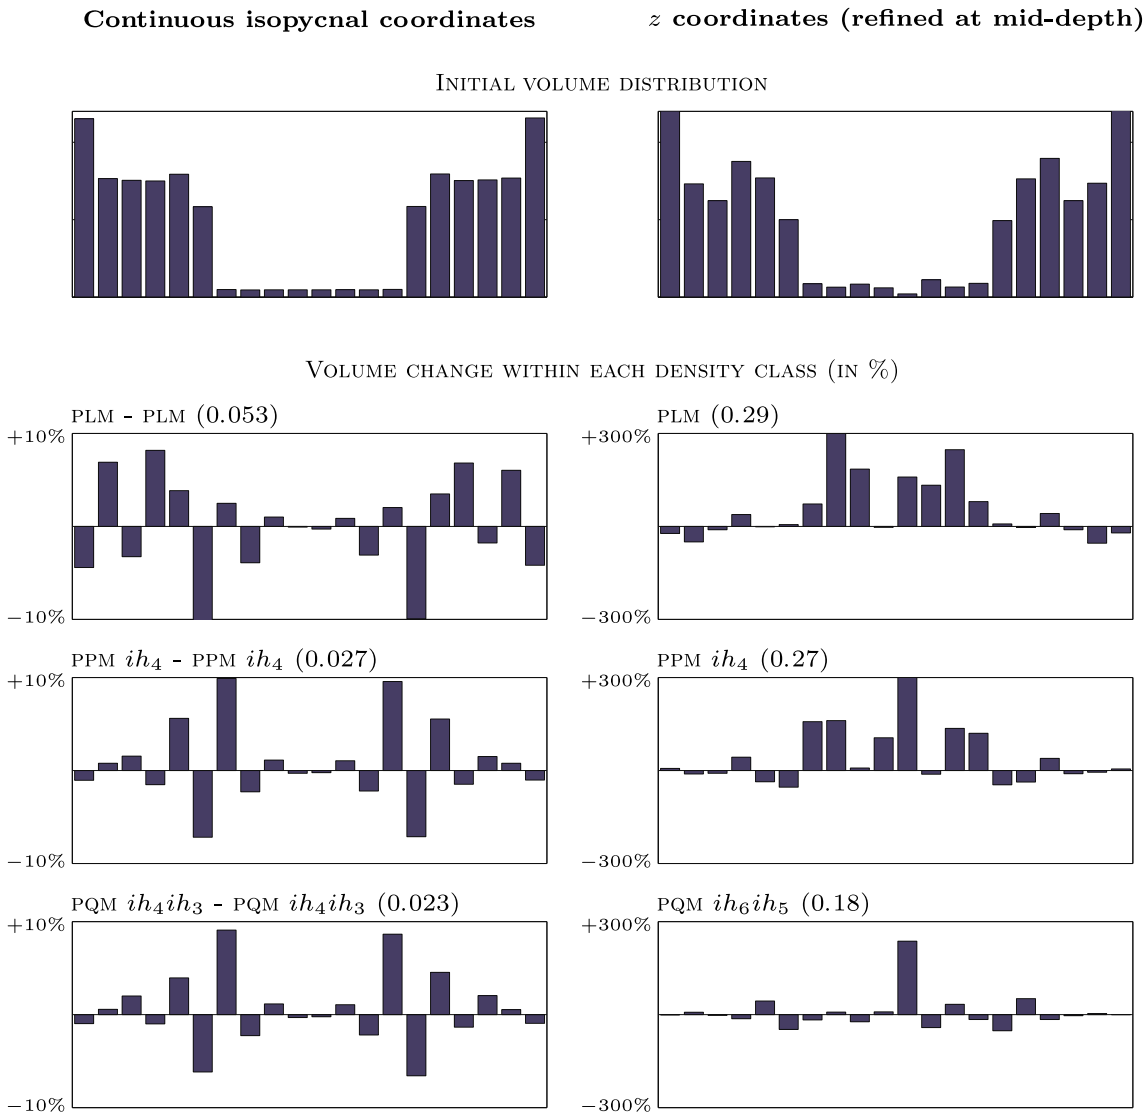

Fig. 11. Analysis of spurious diffusion introduced by the various regridding–remapping schemes used in the experiment described in Fig. 10. The initial volume distributions represent the initial volumes contained within each of the salinity classes uniformly spanning the range 34–36. Because of the thin pycnocline, most of the volume is contained within lower and higher salinity classes. Initial conditions are altered by an initial regridding–remapping step, which explains the differences between continuous isopycnals and *z*. The volumetric change is shown for these two coordinate systems and for different schemes. Note the scale difference (10% for continuous isopycnals and 300% for *z*). Volumetric changes for layered isopycnal coordinates are strictly zero, which serves as reference. The numbers between parentheses next to each scheme name represent the total amount of volume that has shifted between salinity classes compared with the initial state (total volume is 1). PQM only marginally decreases spurious diffusion with continuous isopycnals but the improvement is substantial with *z* coordinates.

force calculation to be consistent. In such a framework, the positions of layer interfaces depend on target densities and global reconstructions. If the layer densities lead to different global density reconstructions in different water columns, the locations of interface densities will be different as well. Interfaces will end up inclined and lead to spontaneous motion. To avoid this spurious behavior, layer densities must be initialized in a way that is consistent with the reconstruction used for the regridding. The densities must be chosen to generate the same global density profile within each column. Given this constraint, a motionless state will be achieved only when the pressure gradient force computation is modified to take into account the fact that across-layer reconstructions are no longer constant. These modifications are illustrated in Fig. 7(b).

The ability to preserve a motionless state in the regridding–remapping framework is important. It is a verification of consistency requirements, as explained above and illustrated in Fig. 7. It also assesses the robustness of remapping and regridding schemes since those should not trigger spontaneous motion. Fig. 8 shows that the regridding–remapping algorithm passes this sanity check when stratification is linear, a linear equation of state is used, initial conditions are consistently prescribed and the pressure gradient force calculation accounts for linear reconstructions of density *and* bottom topography.

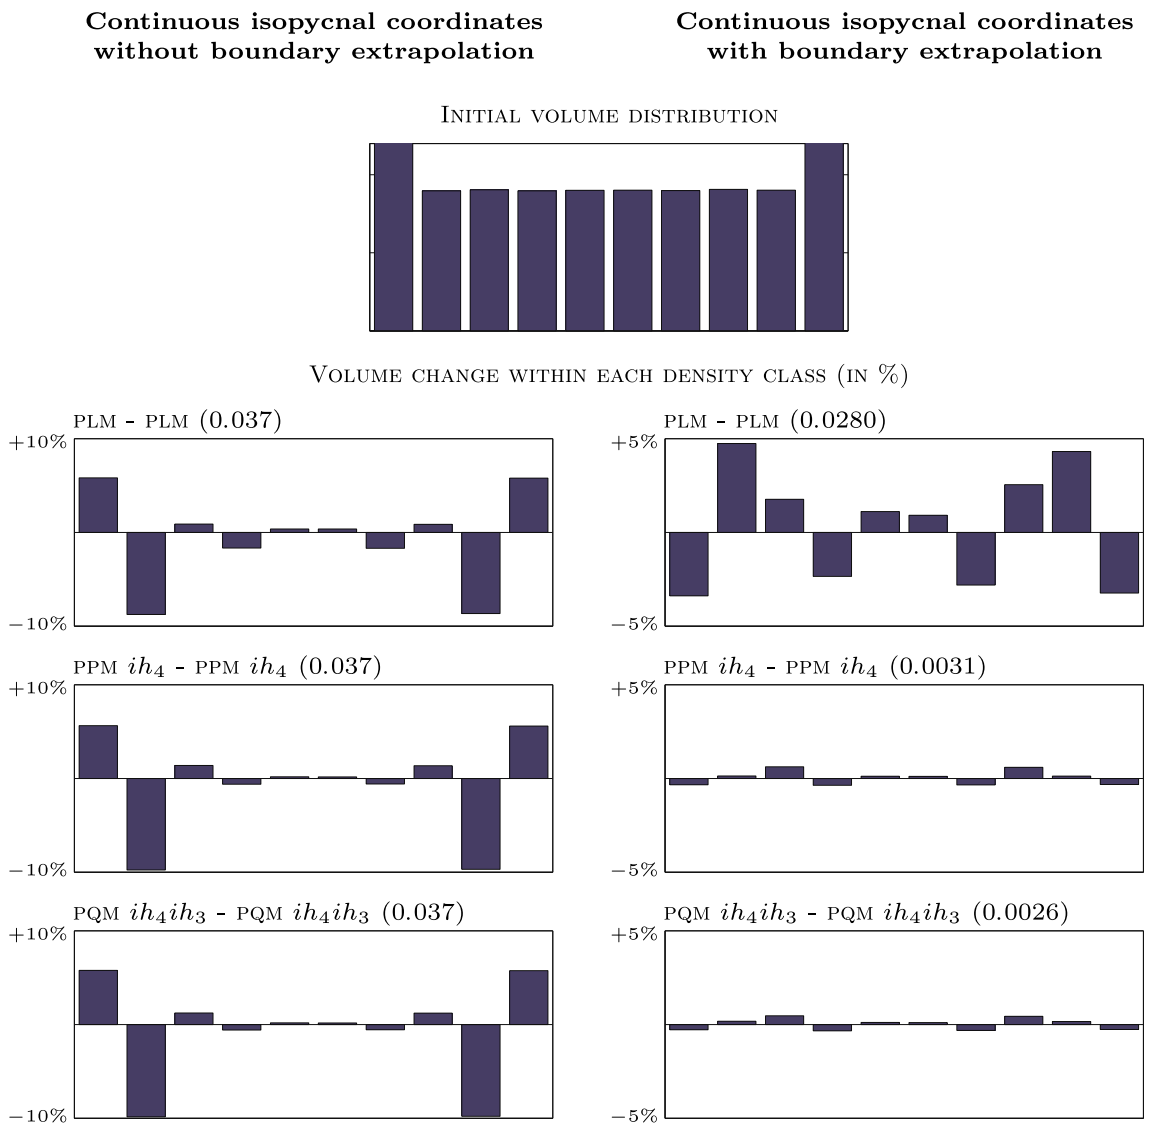

Fig. 12. The effect of using high-order reconstructions (i.e., extrapolating) within boundary cells on spurious diffusion is investigated using continuous isopycnal coordinates, starting with the initial conditions presented in Fig. 9(b). The initial volume distribution represents the initial volumes contained within each of the salinity classes uniformly spanning the range 34–36. Target densities have been chosen to yield thicker boundary cells. The volumetric change is shown for different remapping schemes when no boundary extrapolation is used (left) and when boundary extrapolation is switched on (right). Volumetric changes for layered isopycnal coordinates are strictly zero, which serves as reference. No matter which scheme is used, spurious diffusion caused by using PCM within boundary cells is dominant. Higher-order extrapolations substantially decrease the amount of spurious diffusion, leading to less than 0.3% of volume being displaced from initial salinity classes when using PQM.

Details on the numerical experiments are given in the figure's caption. As a consequence, terrain-following coordinates do not experience any pressure gradient errors when all properties are linear. However, any nonlinear departure will generate pressure gradient errors that are larger than the machine-precision threshold seen here.

To represent topography with continuous isopycnal and z coordinates, vanishing layers are used in such a way that the sum of all layer thicknesses is equal to the local ocean depth. In theory, the thickness of these vanishing layers should be zero. However, a finite minimum thickness is used in our model for two reasons. First, high-order edge-value and edge-slope estimates behave unstably when there is a large and abrupt change in layer thicknesses (of 6 or more orders of magnitude). In the following experiments, we set the minimum thickness to be 10^{-3} m. Vanished layers generated during the regridding are therefore inflated to comply with this threshold. This issue could be addressed by getting rid of vanished layers before regridding and remapping and adding them back. This procedure will be investigated in the future. Second, any inaccuracy in the pressure gradient force calculation – even if it is as low as machine-precision level, which is unavoidable – results in spontaneous motion. Very thin vanishing layers, say on the order of 10^{-10} m (as is used in HIM), are overly sensitive to this kind of spontaneous motion, even if it is negligible in magnitude. The effect is a large change in density (via advection of salt and temperature) within the vanishing layer and subsequent non-negligible motion. The latter is due to the regridding in response to this density change and to a nonzero pressure gradient force. As a consequence, for the experiment shown in Fig. 8, a minimum thickness of 10^{-1} m was necessary to keep the velocity magnitude as low as 10^{-12} m s^{-1}. When a minimum thickness of 10^{-3} m is used, the velocity magnitude increases to about 10^{-8} m s^{-1}, which is acceptable and recommended for realistic applications.

5.3. Internal waves

We now consider a 1000-m deep, 200-km wide, flat-bottomed rectangular basin to assess the accuracy of regridding–remapping schemes in the context of nonrotating, internal wave propagation. Two different initial states are used and are presented in Fig. 9. The first one (Fig. 9(a)) contains 20 layers and a thin pycnocline made up of 8 layers. With constant salinity increments from layer to layer, the vertical salinity structure features sharp gradients. We wish to assess the ability of the regridding–remapping schemes to resolve the sharp gradients with the least amount of spurious diffusion. The second one (Fig. 9(b)) is simpler and consists of 10 uniformly-distributed layers. It will be used to investigate the effect of boundary reconstructions on accuracy. In all cases, the horizontal grid contains 160 cells. This overly high resolution – 40 cells would be sufficient – is chosen as a way to limit the amount of diffusion caused by the PLM layerwise advection scheme. While this scheme is justified for layered isopycnal coordinates where along-layer gradients tend to be reduced, z-coordinate solutions may suffer a lot from spurious diffusion introduced by horizontal advection schemes. If the horizontal grid is too coarse, this spurious diffusion might dominate that from remapping, which would render our analyses irrelevant. The baroclinic and thermodynamical time steps are both 450 s. The barotropic time step is 5 s. The regridding–remapping algorithm is called every thermodynamical time step. There is neither momentum diffusion, nor tracer diffusion (vertical or horizontal). Density is set to be equal to salinity. Any spurious diffusion is a direct result of remapping, and to a much lesser extent, layerwise advection. All results are shown after 25 days.

Fig. 10 investigates the ability of continuous isopycnal and z regridding–remapping schemes to preserve the vertical salinity structure, when setting off the model with the thin pycnocline initial conditions. For the z-coordinate experiments,

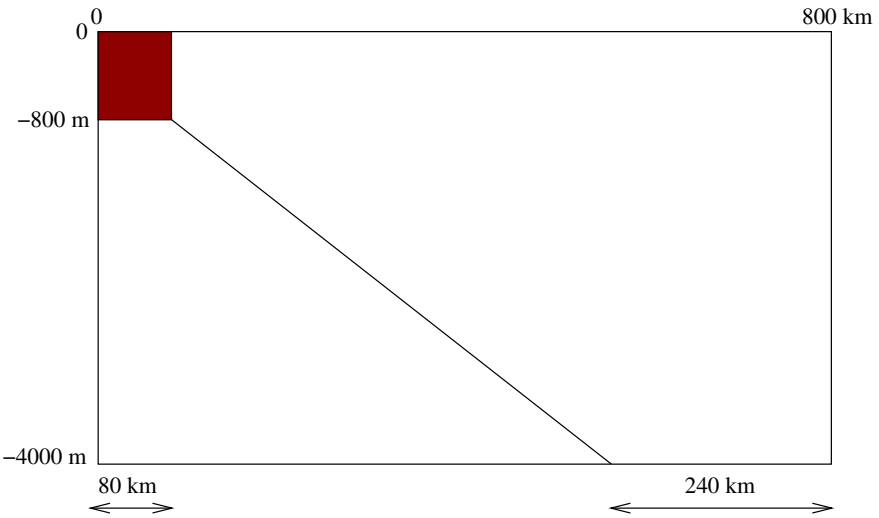

Fig. 13. Basin geometry and initial conditions for the dense overflow experiment. A linear stratification is considered where salinity varies from 34 at the surface to 35 at the bottom. Dense water (salinity of 35) is located in the shallow bay and flows downslope. In this experiment, density is equal to salinity.

the grid is refined at mid-depth. While PQM brings about only marginal improvement over PPM for continuous isopycnal coordinates, using PQM for z coordinates turns out to be necessary to obtain an acceptable solution in terms of the vertical salinity structure. The solution obtained with the layered isopycnal version of the model serves as a reference for the pycnocline thickness. Yet, we do not expect the solutions to look identical (there is a noticeable phase lag between the layered and continuous isopycnal solutions). Different vertical coordinates are known to have great influence on the solutions and

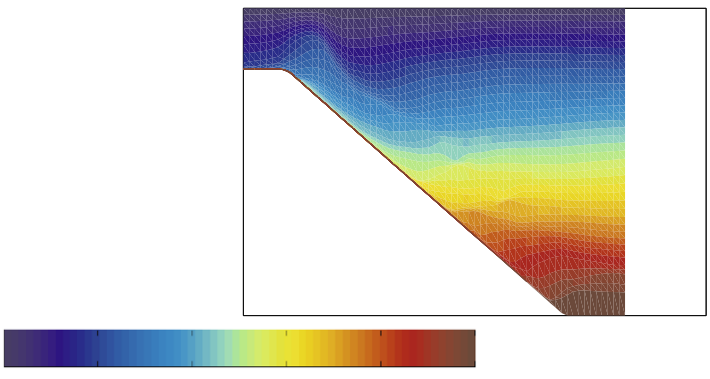

determining which coordinate system is the best is out of the scope of this paper. We use the layered isopycnal version of the model as reference for its inherent adiabatic nature and ability in preserving the vertical salinity structure.

An assessment of the accuracy of these schemes is proposed in Fig. 11 where the volume change of a range of salinity classes is used as a proxy for measuring spurious diffusion. By construction, the layered isopycnal version is adiabatic and since no remapping is involved, each salinity class strictly conserves its volume. The closer a regridding–remapping scheme approaches this behavior, the more accurate it is. The results of Fig. 11 confirm that the improvement gained by using PQM

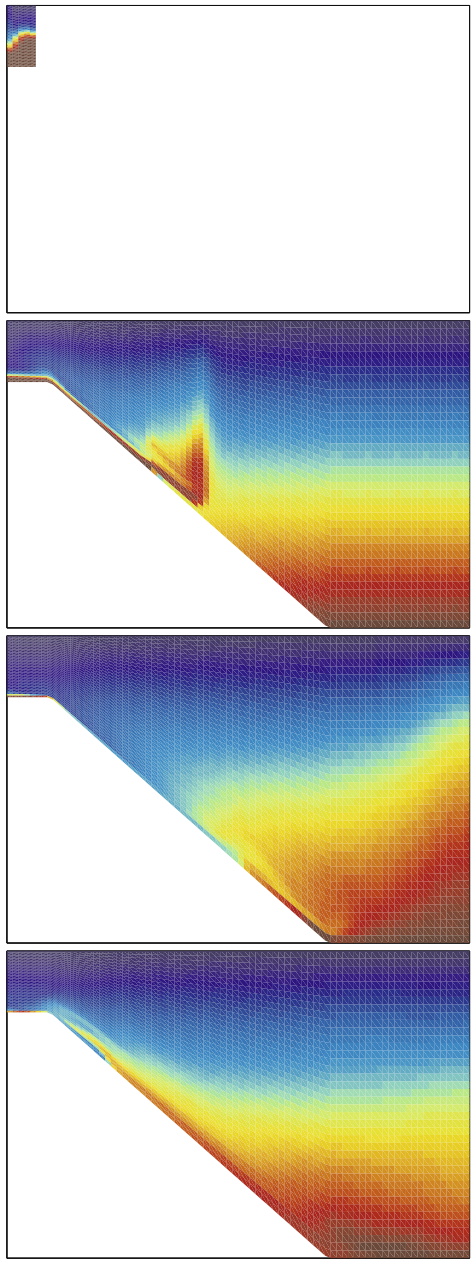

over PPM is not significant for continuous isopycnal coordinates, while being substantial for z coordinates. The fact that accuracy saturation is reached with PPM for continuous isopycnal coordinates may not be surprising. Continuous isopycnal coordinates try to mimic the behavior of impermeable surfaces in a different way. When accurate high-order regridding–remapping schemes are used, the layers that are obtained via regridding may be close enough to that obtained by letting the model evolve as a layered model from the last time regridding–remapping was performed. In this respect, regridding

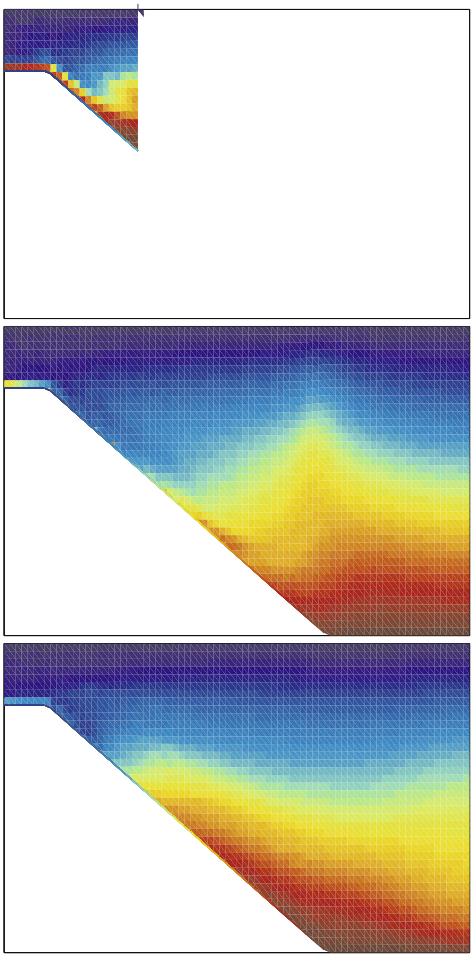

barely modifies the grid and the remapping step has negligible effects. The spurious diffusion that we experience beyond the saturation level attained with PPM may well be a dominance of spurious diffusion caused by the PLM layer advection scheme. Using higher-order advection schemes will shed light on this issue, which is future work.

Finally, Fig. 12 investigates the effect of using high-order reconstructions (i.e., extrapolating) within boundary cells on spurious diffusion using continuous ispoycnal coordinates, starting with the initial conditions presented in Fig. 9(b). When a simple PCM reconstruction is used within boundary cells, the global error does not decrease no matter which regridding–remapping scheme is used. In that case, spurious diffusion near boundaries completely dominates any spurious diffusion that might occur in the interior. It turns out that employing high-order reconstruction schemes within boundary cells is

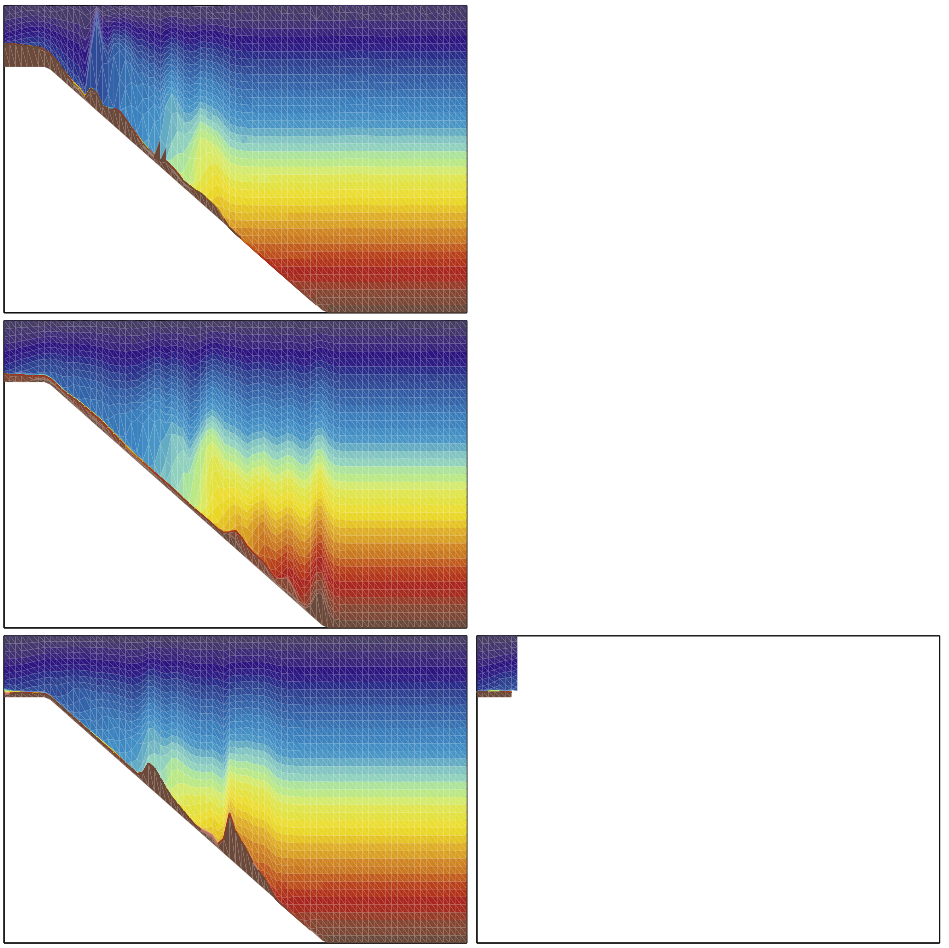

critical for not wasting the overall potential of high-order regridding–remapping schemes. This experiment exemplifies one more time the superiority of PPM and PQM over PLM, as illustrated by the very small amount of volume change (less than 0.3% overall) in salinity classes and hence, the quasi-adiabatic nature of the solution. As a reference, running the layered isopycnal version with a coefficient of diapycnal diffusion of 2×10^{-5} m^2 s^{-1} yields the same global amount of spurious mixing. Since the initial vertical displacement in our experiment (about 100 m) is about one order of magnitude larger than what is typically experienced within the thermocline, we expect the amount of spurious diffusion caused by regridding–remapping to be even smaller in more realistic, larger-scale numerical simulations and to approach values akin to molecular processes.

5.4. Dense overflow

Resolving dense overflows has become a classic benchmark for ocean models and, in particular, to study the influence of vertical coordinates. In this paper, we do not aim at drawing any conclusion as to which vertical grid is the best to represent overflows. Past studies have largely contributed to the conclusion that isopycnal coordinates are ideal for this kind of problem [26,22,17]. We consider a nonrotating 4000-m deep, 800-km wide basin with a 800-m deep shallow bay containing dense water (Fig. 13). A linear stratification is prescribed for which salinity varies from 34 at the surface to 35 at the bottom. The equation of state simply assumes that density is equal to salinity. Water in the bay has a salinity of 35. The model is run for 10 days with a baroclinic and thermodynamical time step of 900 s. The barotropic time step is 10 s. Horizontal and vertical momentum diffusion are set to 10^4 m^2 s^{-1} and 10^{-4} m^2 s^{-1}, respectively. There is no tracer diffusion. The horizontal grid resolution is 10 km (80 cells) and 40 layers span the vertical space.

Figs. 14 and 15 present four snapshots (days 1, 2, 5 and 10) of the solutions obtained with the layered isopycnal, continuous isopycnal, z and sigma versions of the model. The layered isopycnal version is considered the reference solution and the continuous isopycnal version yields a similar solution. By contrast, the z and sigma versions are not-surprisingly very different and lag both isopycnal versions. The extremely large amount of numerical entrainment in the z-simulation is explicable by the relatively coarse vertical and horizontal resolution compared with the bottom boundary layer thickness [26]. As shown in Fig. 16, our z-coordinate results are extremely similar to that obtained using MITgcm [19,20], which is a state-of-the-art z-coordinate, C-grid ocean general circulation model. The MITgcm simulation is hydrostatic and uses the seventh-order horizontal advection scheme OS7 presented by Daru and Tenaud [6], which is less diffusive than the PLM layerwise advection scheme used in our model. Differences in simulation outputs are due to, and may not be limited to, different parameterizations, different topography representations and different advection schemes. Though similar overall, these differences emphasize the importance of using a single framework when studying the effect of vertical coordinates. In Fig. 17, the solution obtained with the continuous isopycnal version is presented at day 2 when using six different regridding–remapping schemes. Using PQM for both the regridding and the remapping is necessary to fully capture the dense plume along the slope. It also turns out that continuous reconstruction schemes (i.e., P1M – not shown – and P3M) are very inappropriate for this kind of flow featuring sharp density gradients.

6. Conclusions

We have developed a consistent and efficient high-order regridding–remapping algorithm for use in generalized coordinate ocean models. The regridding–remapping algorithm conserves volume (or mass in a non-Boussinesq version), momentum, salt, temperature and any other passive tracer to machine precision. It is a necessary ingredient for long-term ocean climate modeling. Our regridding–remapping framework has been used in a series of idealized one-dimensional numerical experiments as well as two-dimensional internal wave and overflow test cases. In all cases, PQM schemes provide the most accurate solutions for regridding and remapping. Our model is capable of replicating z-, sigma- and isopycnal-coordinate results. To our knowledge, we have presented the first ocean model capable of representing continuous isopycnal coordinates.

Particular emphasis has been put on the design of the continuous isopycnal framework, which uses the traditional layered isopycnal paradigm in the underlying representation of the system between regridding–remapping events. Using continuous isopycnals allows a more elegant and accurate construction of truly hybrid vertical coordinates, which open the door to many other choices of vertical coordinates. This study will be the subject of a forthcoming paper. However, moving away from a layered framework has consequences, one of which being that along-layer density gradients no longer vanish. The Montgomery potential approach for computing pressure gradient forces is no longer valid and the finite volume form of the pressure force calculation must be used with the modifications presented in this paper to account for across-layer density variations. In contrast with layered isopycnals, layerwise advection now plays a more influential role as it can alter the layer-averaged densities. PLM may no longer be an acceptable choice for layerwise advection in a hybrid vertical coordinate model. In a hybrid coordinate framework, diffusion tensors need to be rotated to minimize spurious diapycnal mixing. We believe, however, that the benefits of the more general character of our approach far outweigh these additional costs. It should also be noted that a number of physical processes may be more easily added to the model when the vertical grid is not constrained to layered isopycnals. These processes include geothermal heating, double diffusion and interior heat sources and sinks.

In developing this regridding–remapping algorithm, high-order reconstructions within boundary cells turned out to be crucial for obtaining sensible results and for reducing spurious diffusion near boundaries. Advection schemes may borrow

the ideas put forward in this paper regarding boundary extrapolations to increase their accuracies near boundaries *and* extrema. We also note that, in our model, vertical advection is implicitly embedded in the remapping step and directly benefits from high-order schemes.

The issue of which vertical coordinate system is best for large-scale ocean simulations remains outstanding [9]. Though intercomparison exercises have taken place in the past [4,25] as an attempt to address this issue, they generally involve different models, which, apart from the vertical grid, differ in many other aspects (Fig. 16). This limitation raises the question as to whether differences in model solutions are only caused by different vertical grids and stresses the need for a single framework in which to evaluate the impact of coordinate choice. Only a few models are able to use z and σ coordinates within the same framework and none of them include an isopycnal representation. The hybrid coordinate paradigm presented in this paper, though still incomplete for full-fledged studies of the large-scale ocean circulation, enables the comparison of many different vertical coordinate systems within a single framework.

Acknowledgments

The authors would like to thank Maxim Nikurashin for running the MITgcm simulation and providing us with the results. We also thank Sonya Legg and, especially, Steve Griffies for valuable comments on the final draft. An anonymous reviewer provided useful suggestions to improve the presentation of this paper. Laurent White is supported by the ECCO2 project (http://ecco2.org/), "Estimating the Circulation and Climate of the Ocean, Phase II: High Resolution Global-Ocean and Sea-Ice Data Synthesis", NASA award number NNG06GC28G. Laurent White is an honorary postdoctoral researcher with the Belgian National Fund for Scientific Research (FNRS).

Appendix A. The continuous cubic interpolation scheme (P3M)

A cubic has four degrees of freedom. Both edge values and slopes are used to determine them. Because a cubic is potentially fourth-order accurate, we require the edge-value and edge-slope estimates to be fourth-order accurate as well. As a sanity check, a global cubic profile must be exactly retrieved via piecewise cubic interpolation. A cubic can be written locally as

$$C(\xi) = a_0 + a_1\xi + a_2\xi^2 + a_3\xi^3. \tag{11}$$

Given the left and right edge values u_L and u_R, respectively, and the left and right edge slopes u'_L and u'_R, respectively, the four coefficients in Eq. (11) are given by:

$$\begin{aligned} a_0 &= u_L, \\ a_1 &= u'_L, \\ a_2 &= 3(u_R - u_L) - u'_R - 2u'_L, \\ a_3 &= u'_R + u'_L + 2(u_L - u_R), \end{aligned} \tag{12}$$

where

$$\begin{aligned} u'_L &= \frac{\partial C}{\partial \xi}\Big|_{\xi=0} = \frac{\partial C}{\partial x}\Big|_{x=x_k} h_k, \\ u'_R &= \frac{\partial C}{\partial \xi}\Big|_{\xi=1} = \frac{\partial C}{\partial x}\Big|_{x=x_{k+1}} h_k. \end{aligned}$$

The above relationships are easily derived from Eq. (1).

In the following, use will be made of the following slope definitions. Note that the variable u is generic here. Given a cell of width h_C and left and right neighboring cells of widths h_L and h_R, respectively, the limited PLM slope σ is defined as

$$\sigma = \begin{cases} \operatorname{sign}(\sigma_C)\min(|\sigma_L|, |\sigma_R|, |\sigma_C|) & \text{if } \sigma_L\sigma_R > 0 \\ 0 & \text{otherwise}, \end{cases} \tag{13}$$

where σ_L and σ_R are the left and right one-sided slopes, respectively, and σ_C is the centered slope. The sign function is equal to 1 for positive arguments, -1 for negative arguments and 0 otherwise. The one-sided and centered slopes are defined as

$$\begin{aligned} \sigma_L &= 2\frac{\bar{u}_C - \bar{u}_L}{h_L + h_C} \times \frac{h_L + h_C}{h_C} = 2\frac{\bar{u}_C - \bar{u}_L}{h_C}, \\ \sigma_R &= 2\frac{\bar{u}_R - \bar{u}_C}{h_C + h_R} \times \frac{h_C + h_R}{h_C} = 2\frac{\bar{u}_R - \bar{u}_C}{h_C}, \\ \sigma_C &= 2\frac{\bar{u}_R - \bar{u}_L}{h_L + 2h_C + h_R}, \end{aligned} \tag{14}$$

where $\bar{u}_L, \bar{u}_C$ and $\bar{u}_R$ are the cell values associated with the left, center and right cells, respectively. Note that the slopes defined by Eq. (14) are the traditional van Leer limited PLM slopes (e.g., [18]), written for nonuniform grids.

Once unique edge-value estimates have been computed, they are bounded by neighboring cell values when needed. The ensuing edge values will not be modified afterwards. Edge-slope estimates are then computed and modified, if necessary, to ensure consistency with the limited PLM slope σ, as defined by Eq. (13). These steps are very similar to those involved with the PQM limiter [24]. For example, if the PLM slope is nonnegative, both edge slopes should be nonnegative as well. If one of them is inconsistent, it is set to zero.

At this point, edge values are bounded and edge slopes are consistent. Yet, this does not guarantee monotonicity. The existence of a local extremum is equivalent to the existence of an inconsistent inflexion point, namely an inflexion point where the slope is inconsistent with the PLM slope. When that happens, the edge slopes must be modified in order to move the inflexion point away from the interior and onto one of the edges.

The second derivative $C^{(2)}(\xi)$ of Eq. (11) is given by

$$C^{(2)}(\xi) = b_0 + b_1\xi, \tag{15}$$

where $b_0 = 2a_2$ and $b_1 = 6a_3$, as defined by Eq. (12). If $b_1 = 0$, the second derivative of $C(\xi)$ is single-signed: $C(\xi)$ is parabolic and monotonic. The algorithm stops here. If $b_1 \neq 0$, there is an inflexion point located at $\xi_{ip} = -b_0/b_1$. If $\xi_{ip} \notin (0,1)$, the curvature in $[0,1]$ is single-signed and the cubic is monotonic. The algorithm stops here. If $\xi_{ip} \in (0,1)$, the slope of the cubic at $\xi = \xi_{ip}$, $C'(\xi_{ip})$, must be computed. If $C'(\xi_{ip})$ has the same sign as the PLM slope, the cubic is monotonic and the algorithm stops here. If the sign of $C'(\xi_{ip})$ is opposite to the sign of the PLM slope, the cubic is not monotonic. These four cases are illustrated in Fig. 18. In the last case, the inflexion point is moved away from the interior and toward one of the edges, according to the following rule:

$$\begin{aligned} &|\sigma_L| \leqslant |\sigma_R| \Rightarrow \text{move inflexion point onto left edge},\\ &|\sigma_R| < |\sigma_L| \Rightarrow \text{move inflexion point on right edge}, \end{aligned} \tag{16}$$

where the slopes are defined by Eq. (14). White and Adcroft [24] provide a justification for using the above rule. Shifting the inflexion point must be done by adjusting the edge slopes. Both cases are now explained.

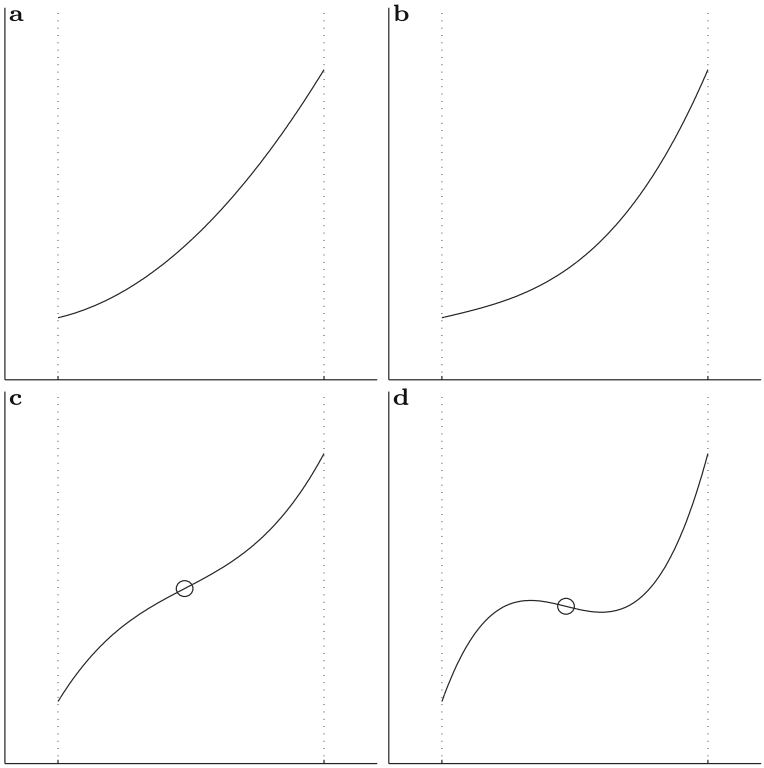

Fig. 18. In each panel, the cubic has consistent edge slopes. Yet, this edge-slope consistency does not guarantee monotonicity. The presence of a local extremum – and the breakdown of monotonicity – depends on whether the slope at the inflexion point (marked by an empty circle) is consistent (here positive). (a) The cubic has no inflexion point: it is degenerated into a parabola. (b) The cubic has an inflexion point lying outside the cell. Thus, the curvature is single-signed within the cell and the cubic is monotonic. (c) The inflexion point lies within the cell but the slope at the location of the inflexion point is consistent. The cubic is monotonic. (d) The slope at the location of the inflexion point is inconsistent. The cubic is nonmonotonic and needs further limiting.

A.1. Shifting the inflexion point onto the left edge

Requiring the inflexion point of the cubic to be located on the left edge means that we must enforce Eq. (15) to vanish at $\xi = 0$. Hence, we must have $a_2 = 0$, or

$$3(u_R - u_L) - u'_R - 2u'_L = 0. \tag{17}$$

Because both edge slopes may be adjusted, Eq. (17) is overdetermined and only one slope can be modified. Solving Eq. (17), tentative adjusted slopes (identified by an asterisk) are given by

$$u'^{*}_L = \frac{3}{2}(u_R - u_L) - \frac{1}{2}u'_R,$$
$$u'^{*}_R = 3(u_R - u_L) - 2u'_L.$$

Three situations may occur. (1) u'^{*}_L is consistent, (2) u'^{*}_R is consistent and (3) both u'^{*}_L and u'^{*}_R are inconsistent. The case where both tentative slopes are consistent is included in either (1) or (2).

When case (1) occurs, we keep the new left edge slope and compute the right one according to Eq. (17). The new slopes are defined as follows:

$$u'_L = u'^{*}_L,$$
$$u'_R = 3(u_R - u_L) - 2u'^{*}_L.$$

When case (2) occurs, we keep the new right edge slope and compute the left one according to Eq. (17). The new slopes are defined as follows:

$$u'_L = \frac{3}{2}(u_R - u_L) - \frac{1}{2}u'^{*}_R,$$
$$u'_R = u'^{*}_R.$$

Finally, when case (3) occurs, the left edge slope is set equal to zero (see justification by White and Adcroft [24]) and the right edge slope is coputed according to Eq. (17). The new slopes are

$$u'_L = 0,$$
$$u'_R = 3(u_R - u_L).$$

An illustration of this limiter is provided in Fig. 19.

A.2. Shifting the inflexion point onto the right edge

Requiring the inflexion point of the cubic to be located on the right edge means that we must enforce Eq. (15) to vanish at $\xi = 1$. Hence, we must have $a_2 + 3a_3 = 0$, or

$$3(u_R - u_L) + 2u'_R + u'_L = 0. \tag{18}$$

Because both edge slopes may be adjusted, Eq. (18) is overdetermined and only one slope can be modified. Solving Eq. (18), tentative adjusted slopes (identified by an asterisk) are given by

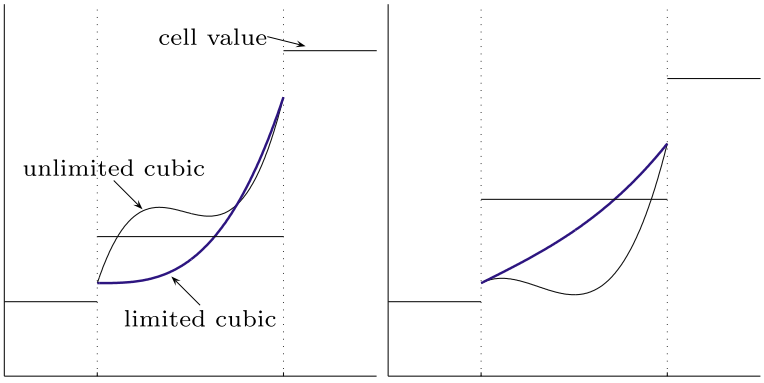

Fig. 19. P3M limiter in action when the initial cubic has consistent edge slopes (both edge slopes are positive and the PLM slope is positive) but has an inconsistent inflexion point. In both panels, the inflexion point is shifted onto the left edge.

$$u_L'^* = 3(u_R - u_L) - 2u_R',$$

$$u_R'^* = \frac{3}{2}(u_R - u_L) - \frac{1}{2}u_L'.$$

Three situations may occur. (1) $u_L'^*$ is consistent, (2) $u_R'^*$ is consistent and (3) both $u_L'^*$ and $u_R'^*$ are inconsistent. The case where both tentative slopes are consistent is included in either (1) or (2).

When case (1) occurs, we keep the new left edge slope and compute the right one according to Eq. (18). The new slopes are defined as follows:

$$u_L' = u_L'^*,$$

$$u_R' = \frac{3}{2}(u_R - u_L) - \frac{1}{2}u_L'^*.$$

When case (2) occurs, we keep the new right edge slope and compute the left one according to Eq. (17). The new slopes are defined as follows:

$$u_L' = 3(u_R - u_L) - 2u_R'^*,$$

$$u_R' = u_R'^*.$$

Finally, when case (3) occurs, the right edge slope is set equal to zero (see justification by White and Adcroft [24]) and the left edge slope is computed according to Eq. (18). The new slopes are

$$u_L' = 3(u_R - u_L),$$

$$u_R' = 0.$$

References

[1] A. Adcroft, R. Hallberg, M. Harrison, A finite volume discretization of the pressure gradient force using analytic integration, Ocean Model. 22 (2008) 106–113.
[2] R. Bleck, An oceanic general circulation model framed in hybrid isopycnic-Cartesian coordinates, Ocean Model. 4 (1) (2002) 55–88.
[3] R. Bleck, D.B. Boudra, Initial testing of a numerical ocean circulation model using a hybrid (quasi-isopycnic) vertical coordinate, J. Phys. Oceanogr. 11 (1981) 755–770.
[4] E.P. Chassignet, H. Arango, D. Dietrich, T. Ezer, M. Ghil, D.B. Haidvogel, C.C. Ma, A. Mehra, A.M. Paiva, Z. Sirkes, DAMEE-NAB: the base experiments, Dynam. Atmos. Oceans 32 (2000) 155–183.
[5] P. Colella, P.R. Woodward, The piecewise parabolic method (ppm) for gas-dynamical simulations, J. Comput. Phys. 54 (1984) 174–201.
[6] V. Daru, C. Tenaud, High order one-step monotonicity-preserving schemes for unsteady compressible flow calculations, J. Comput. Phys. 193 (2004) 563–594.
[7] J. Donea, A. Huerta, J.-P. Ponthot, A. Rodríguez-Ferran, Arbitrary Lagrangian–Eulerian methods, in: E. Stein, R. de Borst, T.J.R. Hughes (Eds.), Encyclopedia of Computational Mechanics, John Wiley and Sons, 2004 (Chapter 14).
[8] R. Gerdes, A primitive equation ocean circulation model using a general vertical coordinate transformation 1. Description and testing of the model, J. Geophys. Res. 98 (C8) (1993) 14683–14701.
[9] S.M. Griffies, C. Böning, F.O. Bryan, E.P. Chassignet, R. Gerdes, H. Hasumi, A. Hirst, A.-M. Treguier, D. Webb, Developments in ocean climate modelling, Ocean Model. 2 (2000) 123–192.
[10] S.M. Griffies, R.C. Pacanowski, R.W. Hallberg, Spurious diapycnal mixing associated with advection in a z-coordinate ocean model, Mon. Weather Rev. 128 (2000) 538–564.
[11] R. Hallberg, Stable split time stepping scheme for large-scale ocean modeling, J. Comput. Phys. 135 (1997) 54–65.
[12] R. Hallberg, Time integration of diapycnal diffusion and Richardson number-dependent mixing in isopycnal coordinate ocean models, Mon. Weather Rev. 128 (2000) 1402–1419.
[13] R. Hallberg, A. Adcroft, Reconciling estimates of the free surface height in Lagrangian vertical coordinate ocean models with mode-split time stepping, Ocean Model 29 (1) (2009) 15–26.
[14] R. Hallberg, A. Gnanadesikan, The role of eddies in determining the structure and response of the wind-driven southern hemisphere overturning: results from the modeling eddies in the southern ocean (MESO) project, J. Phys. Oceanogr. 36 (2006) 2232–2252.
[15] G.R. Halliwell, Evaluation of vertical coordinate and vertical mixing algorithms in the HYbrid-Coordinate Ocean Model (HYCOM), Ocean Model 7 (2004) 285–322.
[16] A. Kasahara, Various vertical coordinate systems used for numerical weather predictions, Mon. Weather Rev. 102 (1974) 509–522.
[17] S. Legg, R.W. Hallberg, J.B. Girton, Comparison of entrainment in overflows simulated by z-coordinate, isopycnal and non-hydrostatic models, Ocean Model. 11 (2006) 69–97.
[18] R.J. Leveque, Finite Volume Methods for Hyperbolic Problems, Cambridge University Press, 2002.
[19] J. Marshall, A. Adcroft, C. Hill, L. Perelman, C. Heisey, A finite-volume, incompressible Navier Stokes model for studies of the ocean on parallel computers, J. Geophys. Res. 102 (1997) 5753–5766.
[20] J. Marshall, C. Hill, L. Perelman, A. Adcroft, Hydrostatic, quasi-hydrostatic, and non-hydrostatic ocean modeling, J. Geophys. Res. 102 (1997) 5733–5752.
[21] J.M. Oberhuber, Simulation of the atlantic circulation with a coupled sea ice-mixed layer-isopycnal general circulation model. Part I: model description, J. Phys. Oceanogr. 23 (1993) 808–829.
[22] M.P. Papadakis, E.P. Chassignet, R.W. Hallberg, Numerical simulations of the Mediterranean sea outflow: impact of the entrainment parameterization in an isopycnic coordinate ocean model, Ocean Model. 5 (2003) 325–356.
[23] Y.T. Song, T.Y. Hou, Parametric vertical coordinate formulation for multiscale, Boussinesq, and non-Boussinesq ocean modelling, Ocean Model. 11 (2006) 298–332.
[24] L. White, A. Adcroft, A high-order finite volume remapping scheme for nonuniform grids: the piecewise quartic method (PQM), J. Comput. Phys. 227 (2008) 7394–7422.

[25] J. Willebrand, B. Barnier, C. Boning, C. Dieterich, P.D. Killworth, C. Le Provost, Y. Jia, J.-M. Molines, A.L. New, Circulation characteristics in three eddy-permitting models of the North Atlantic, Prog. Oceanogr. 48 (2001) 123–161.
[26] M. Winton, R. Hallberg, A. Gnanadesikan, Simulation of density-driven frictional downslope flow in *z*-coordinate ocean models, J. Phys. Oceanogr. 28 (1998) 2163–2174.
[27] F. Xiao, T. Yabe, X. Peng, H. Kobayashi, Conservative and oscillation-less atmospheric transport schemes based on rational functions, J. Geophys. Res. 107 (2002) 4609–4619.