Chapter 1

ON THE USE OF HYBRID VERTICAL COORDINATES IN OCEAN CIRCULATION MODELING

Rainer Bleck
Los Alamos National Laboratory, Los Alamos, New Mexico, USA
bleck@lanl.gov

Abstract

The rationale for building hybrid-coordinate ocean circulation models is discussed in the context of various approaches presently taken to improve mathematical and physical aspects of ocean models. Design choices made in formulating the vertical grid generator, the core component of hybrid models, are laid out. A new experimental approach toward minimizing numerical errors and inconsistencies during simultaneous advection of temperature, salinity and density is presented.

1. Introduction

Motion systems interacting in the ocean range in size from centimeters to planetary, a spread of 9 orders of magnitude. Ocean circulation modeling therefore is a textbook example of a multiscale problem. Since even our fastest computers cannot spatially resolve the earth's surface by more than a few thousand grid points in each direction, only about one-third of those 9 orders of magnitude can be explicitly resolved in state-of-the-art ocean models, meaning that the other two-thirds are relegated to the "subgrid" scale. Hence, a perpetual challenge to ocean modeling is the need to parameterize the interaction of physical processes across the threshold between resolved and unresolved scales of motion. This need will be with us until computers become $(10^6)^3 = 10^{18}$ times faster than they are today – i.e., practically forever.

The ability of today's ocean models to correctly simulate essential aspects of the global circulation has led to a proliferation of applications where the models are used outside their "design range" and therefore generate results that do not always live up to the expectations of the various user communities. Ocean models clearly need to be developed further to satisfy those communities.

There are essentially three ways in which ocean models can be improved. One can

1 increase grid resolution;

2 improve model physics;

3 improve model numerics.

Increasing grid resolution (item 1), in theory, allows a numerically obtained solution to approach that of the underlying differential equation. However, given the huge spectral range of unresolved processes in the ocean, truncation errors continue to cast their shadow even over what we call "high-resolution" models.

Efforts in model physics improvement (item 2) in most cases boil down to improvements in the parameterization of spatially unresolved processes. This is a never-ending process: parameterization schemes tend to be sensitive to where the transition between resolved and unresolved scales occurs and therefore must evolve in lockstep with refinements in mesh size.

Improved numerics (item 3) is a non-exclusive alternative to higher grid resolution; both approaches lower the truncation error in the finite difference equations that constitute the ocean model.

Model numerics can be improved in several ways. One approach is to switch to higher-order finite difference approximation ("order" here refers to the power of the spatial or temporal mesh size in truncation error expressions; the higher the power, the faster the error goes to zero with increased grid resolution). Another approach is to revive techniques developed in pre-computer days when a problem became solvable only by transforming the equations into a coordinate system that exploited some symmetry or conservation laws inherent in the underlying physics. The geophysical fluid dynamics community is doing some of this already by formulating its equations in a coordinate system whose z axis points in the direction of gravity rather than, say, the center of our galaxy. But the concept of manipulating the equations to make them easier to solve or improve the accuracy of the solutions can be extended much further.

The recent proliferation of unconventional vertical coordinates in geophysical modeling should be viewed as one particular attempt to improve ocean model numerics along the lines of item 3 above. Repeating the phrase just used, the new coordinates currently being experimented with seek to exploit "some symmetry or conservation laws inherent in the underlying physics" with the goal of improving the accuracy of the numerical solution. Foremost among those conservation laws is the one stating that adiabatic motion follows surfaces of constant entropy or its proxy, potential density. Thus, if one uses potential density as vertical coordinate, adiabatic flow that is 3-dimensional in Cartesian space is rendered 2-dimensional in potential density space. This makes it particularly easy to satisfy adiabatic constraints while modeling lateral transport of tracers, including temperature and salinity.

The above advantage actually holds for resolved as well as a wide range of *un*resolved scales of motion. Lateral stirring processes in the ocean are known to mix properties predominantly along isentropic surfaces. Hence, a model based on potential density (which in this particular case must be locally referenced) should be able to simulate the effects of subgridscale stirring, typically parameterized as an eddy diffusion process, more accurately than a model in which lateral isentropic stirring must be projected onto the Cartesian x, y, z axes.

The principal design element of isopycnic (potential density) coordinate models, in relation to conventional Cartesian coordinate models, is that depth (alias layer thickness) and potential density trade places as dependent and independent variables. This switch does not affect the number of prognostic equations, nor does it alter the familiar mix of wave modes and processes by which information is transmitted in the ocean. This is to say that both model types solve the same physical problem. However, since the two models are based on different sets of differential equations, their numerical properties should be expected to be very different as well.

Wind-forced process models framed in isopycnic coordinates have been in use since the 1960s (e.g., Welander, 1966; Holland and Lin, 1975a,b; Bleck and Boudra, 1981). With the addition of thermohaline forcing (Bleck et al., 1992; Oberhuber, 1993; Hu, 1997; Sun and Bleck, 2001; Cheng et al., 2004), these models have become sufficiently comprehensive to be useful in studying the oceanic general circulation.

The focus in this article is on one particular extension of the isopycnic coordinate concept that addresses certain shortcomings of potential density as vertical coordinate. These shortcomings are

- coordinate surfaces intersecting the sea surface (the outcropping problem);
- lack of vertical resolution in unstratified water columns.

Note that these are two sides of the same coin. In a global model, most density surfaces required to span the top-to-bottom density range in the subtropics are out of place in the relatively unstratified high-latitude waters and hence must exit the domain somewhere between the subtropics and the subpolar latitude bands.

The hybridized depth-isopycnic coordinate (Bleck, 2002) developed for the isopycnic model MICOM (Bleck et al., 1992) alleviates both shortcomings just mentioned. It does so by not allowing coordinate layers to outcrop, but rather forcing layers that are isopycnic in character at low latitudes to turn into fixed-depth layers at high latitudes. Stated differently, a coordinate layer associated with a chosen "target" isopycnal adheres to the depth of that isopycnal as long as the latter exists in a given water column. Near the outcrop latitude of the target isopycnal, the coordinate layer turns horizontal and becomes a fixed-depth layer. (The term fixed-depth here covers both constant-depth and bottom-following layer configurations.)

The approach just outlined creates fixed-depth coordinate layers at the same rate at which the model loses isopycnic layers between the equator and the poles. The new fixed-depth layers provide vertical resolution in unstratified high-latitude regions and on coastal shelves, allowing the hybrid coordinate model to simulate turbulent mixing and buoyant convection in a manner similar to z and σ coordinate models.

The term "hybrid vertical coordinate" has more than one meaning. Often it refers to configurations (e.g., Bleck 1978) where the model domain is divided into a stack of subdomains, each filled with coordinate surfaces defined independently of those in other subdomains. A simple example is the combination, popular in weather prediction models, of terrain-following coordinate surfaces in the lower atmosphere and isobaric surfaces in the upper atmosphere.

Our present hybrid scheme, which dates back to Bleck and Boudra (1981), does not rely on rigidly defined subdomains. Instead, it permits temporal and lateral transitions between different coordinate types – terrain-following, constant-depth, isopycnal – based on local conditions such as layer thickness and vertical density contrast. The scheme has much in common with the ALE (Arbitrary Lagrangian-Eulerian) technique of Hirt et al. (1974) but adds one important element to that scheme, namely, a mechanism for keeping coordinate layers aligned with, or for nudging them toward, their designated target isopycnals wherever possible. The original ALE scheme only concerns itself with maintaining a finite separation between adjacent coordinate surfaces. While the flexibility of coordinate placement in ALE-type schemes is disconcerting to some users because grid point location in physical space cannot be expressed in terms of a simple analytic formula, the flexibility inherent in the ALE concept allows it to maximize the size of the isopycnic subdomain in the model. This is a major advantage.

2. The Grid Generator

At the core of ALE-type hybrid models is a utility called the vertical grid generator. Most efforts to produce a robust hybrid model are directed at refining and "shockproofing" this unique utility. A thorough discussion of the grid generator used in HYCOM is therefore in order. For the convenience of the reader, some details already presented in Bleck (2002) will be repeated here, but emphasis will be on recent improvements.

Models like HYCOM belong to the so-called *layer* model class where, as mentioned earlier in the context of isopycnic-coordinate models, the depth z of coordinate surfaces is treated as a dependent variable.[1] Having lost z as independent variable, layer models need a new independent variable capable of representing the 3rd (vertical) model dimension. This variable is traditionally called "s".

With the number of unknowns increased by one (namely, layer interface depth), the model needs one additional diagnostic equation. The logical choice is an equation linking s to the other model variables. In purely isopycnic coordinate models, s is equated with potential density. In ALE-type hybrid ocean models, s becomes a "target" potential density value assigned to each coordinate layer. The grid generator's task is to move the coordinate layer toward the depth where the target density is observed to occur. Once a layer is aligned with its target isopycnal, it becomes a material layer whose subsequent evolution is no longer governed by the grid generator until interior mixing or surface thermohaline forcing cause the density to drift away from its target value.

If the target density lies outside the density range in the water column, the layer set in motion by the grid generator will encounter the sea surface or sea floor before finding its target density, depending on whether the target density lies below or above the density range found in the column. Instead of rendering layers whose target density does not exist dynamically invisible by deflating them, the grid generator is designed to impose a minimum layer thickness. This constraint also affects

[1]MICOM and HYCOM actually express layer depth in terms of pressure p, the weight per unit area of the water column. Replacing z by p simplifies matters if the vertical model dimension is expressed in terms of a thermodynamic variable. This also means that the *Boussinesq* approximation is not needed in these models.

layers which during their vertical migration impinge on layers already converted to fixed-depth layers.

In summary, grid generation in HYCOM is a two-step process. In step 1, interfaces are set in motion to restore the target density in individual layers. In step 2, this migration is checked for compatibility with an imposed minimum thickness constraint. Step 2 overrides step 1.

The details of the minimum thickness constraint are essentially the model designer's choice. For example, during HYCOM development the decision was made to impose nonzero-thickness constraints only at the surface but allow layers at the bottom to become massless. This is done for a good reason. Steeply inclined coordinate layers, like those following the bathymetry, are prone to errors in the horizontal pressure force calculation. In a layer containing no mass, such errors are dynamically inconsequential.

Complications arising from the grid generator's ad-hoc strategy for vertical grid point placement are minimal. This can be demonstrated as follows (Bleck, 1978). Recall that material vertical motion in hydrostatic models is inferred from mass continuity — specifically, from the vertically integrated horizontal mass flux divergence. The material vertical motion so diagnosed is then decomposed into motion of the coordinate surface and motion relative to the coordinate surface:

$$\begin{pmatrix}\text{vertical motion}\\ \underline{\textit{of}}\\ s\text{ surface}\end{pmatrix} + \begin{pmatrix}\text{vertical motion}\\ \underline{\textit{through}}\\ s\text{ surface}\end{pmatrix} = \begin{pmatrix}\text{vertically integrated}\\ \text{horizontal mass flux}\\ \text{divergence}\end{pmatrix} \tag{1.1}$$

After diagnosing the right-hand side of (1.1) at a given time step, the hydrostatic model needs one additional condition to distribute this quantity among the two terms on the left. In a material coordinate system, for example, the second term on the left is zero by definition; hence, the vertically integrated mass flux divergence yields the rate at which a coordinate surface moves up or down in space. The other extreme is the spatially fixed grid where, by definition, the first term on the left is zero; the vertically integrated mass flux divergence in that case yields the vertical velocity.

The point of this discussion is that compromise solutions between the two extremes just mentioned can easily be accomodated in a model on a grid-point-by-grid-point and time-step-by-time-step basis. Whatever vertical motion the grid generator prescribes for a given grid point during a given model time step is simply used in (1.1) in conjunction with the vertically integrated mass flux divergence to compute the appropriate

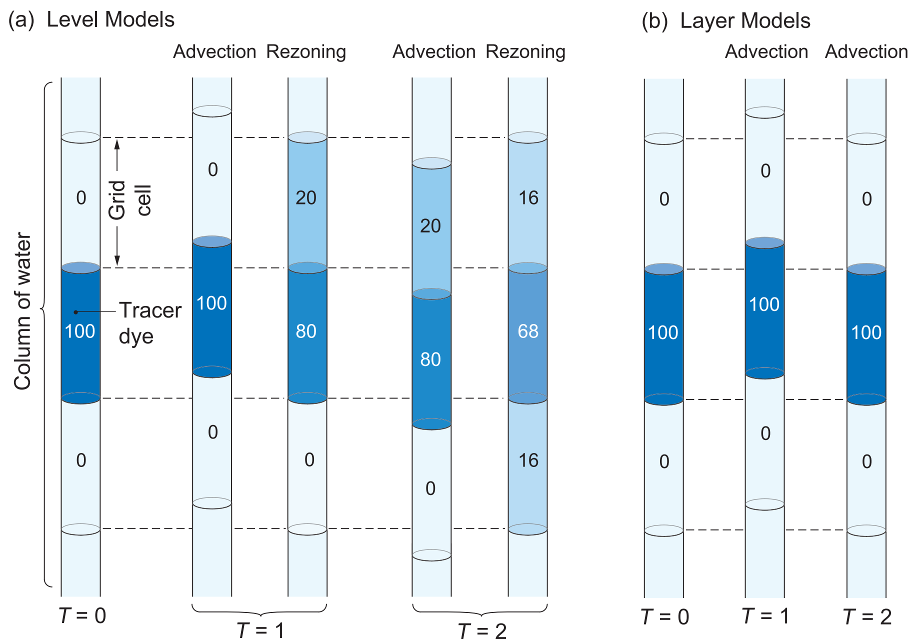

Figure 1.1. Schematic illustrating numerical dispersion in a water column subjected to gravity wave-induced oscillatory vertical motion. Shown is a stack of three grid cells. The initial state, T=0, is chosen to coincide with the wave trough, at which time the middle cell is assumed to be filled with a tracer of concentration 100. The approaching wave crest causes water to be advected upward by a distance chosen in this example to correspond to one-fifth of the vertical cell size (T=1). In level models (a), the clock is stopped momentarily to allow the tracer to be reapportioned ("rezoned") among the original grid cells. The next wave trough causes the water column to return to its original position (T=2). After renewed rezoning, tracer concentration in the middle cell has fallen to 68, the remainder having seeped into cells above and below. In layer models (b), the periodic rezoning steps are skipped, so tracer concentration remains unaffected by the wave motion.

generalized vertical velocity $ds/dt \equiv \dot{s}$. (By definition, $\dot{s}$ is the rate at which a fluid element moves up or down in s space. To avoid dimensional ambiguities created by the lack of a physical definition of s, the quantity actually diagnosed is the interlayer mass flux $\dot{s}\partial p/\partial s$ which is always in units of pressure per time.) The latter forms the basis for vertically advecting all prognostic variables in the model grid. The simplicity of the mechanism expressed by (1.1) is one of the factors that make ALE-type hybrid modeling attractive.

Many ideas have been put forth on how the minimum thickness constraint in ALE ocean models should be formulated. One option is to scale the vertical spacing of nonisopycnic coordinate surfaces by the depth of the turbulent surface mixed layer, thereby ensuring that coordinate surfaces exist throughout the mixed layer for evaluating turbulent exchange processes. This concept is attractive at first sight but has shortcomings if the mixed layer depth changes rapidly, as typically happens during transitions from surface cooling to warming. Not only do fluctuations in vertical grid spacing spawn fluctuations in the magnitude of the truncation error in the finite difference equations, but the large $\dot{s}$ values in (1.1) resulting from rapid coordinate movement are likely to unduly disperse water mass properties in the vertical. These effects are of particular concern if the mixed layer depth responds, as it does in nature, to the 24-hr cycle in solar radiation.

If suppression of excessive vertical migration of coordinate surfaces during the daily or annual heating-cooling cycle is deemed important, the optimal strategy is to “park” coordinate layers near the surface at those times (night or winter, respectively) when their target density does not exist. When surface warming makes the target density reappear, it will reappear at the sea surface. A coordinate layer lying in waiting near the surface can reattach itself to the target density with minimal vertical displacement. Reattachment will also take place sooner if intervening fixed-depth layers, associated with lighter target densities yet to appear, are kept as thin as possible.

False numerical dispersion of water mass properties caused by truncation errors in the finite-difference advection operators is a major concern in ocean modeling. The problem is particularly acute in the z direction where undulating vertical velocities associated with gravity wave trains can have a noticeable dispersive effect; in fact, elimination of vertical dispersion by gravity waves is often mentioned as one of the points in favor of using a material coordinate system. The dispersive effect of gravity waves is illustrated in Fig. 1.1.

While HYCOM’s coordinate surfaces are for the most part material and thus remain unaffected by the dispersion problem just mentioned, the enforcement of minimum layer thickness constraints in the upper part of the ocean does open the door to vertical dispersion. This is of particular concern if the vertical regridding process displaces coordinate surfaces over large distances, because a large first term in (1.1) is likely to produce an $\dot{s}$ term of similar magnitude. Dispersion in this case likely will be larger than in a fixed-grid model whose $\dot{s}$ is given by the right-hand side of (1.1) and thus is bounded by dynamic constraints.

In HYCOM’s original grid generator (see Bleck, 2002), the vertical “remapping” of prognostic variables following the “regridding” by the grid generator is formulated as a donor cell process: the amount of a variable X transferred from one grid cell to the next due to interface movement is computed under the assumption that property X is distributed uniformly within each grid cell. The donor cell scheme is known to be

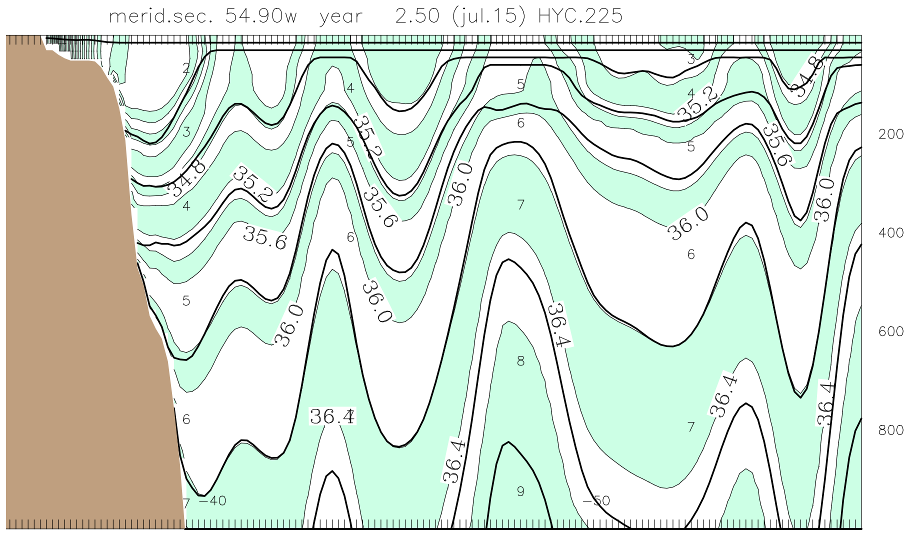

Figure 1.2. Sample vertical section through HYCOM solution extending south from Montevideo into the eddy-rich Brazil and Malvinas Current confluence region. Heavy lines: Layer interfaces. Shaded contours: potential density anomaly (σ_2, $\mathrm{kg\,m^{-3}}$). Tick marks along top and bottom indicate horizontal grid resolution [$0.225° \times \cos(lat.)$, approx. 15 km]. Vertical scale: 1000 m. Crowded isopycnals on continental shelf (upper left corner) are due to Rio de la Plata inflow.

quite diffusive (as illustrated, for example, in Fig. 1.1), and upper-ocean vertical dispersion in HYCOM therefore has been a persistent concern. In the presently supported HYCOM version, the donor cell scheme has been replaced by the more accurate and less diffusive Piecewise Linear Method (PLM).

An illustration of how hybrid coordinates work in practice is given in Fig. 1.2.

3. Mutually consistent T/S/ρ advection

HYCOM development is by no means complete. One problem still awaiting a satisfactory solution is created by the fact that sea water potential density ρ is a function of two independent and equally influential tracers, potential temperature T and salinity S. Advecting both T and S in an isopycnic layer is at best redundant and, at worst, a source of "coordinate drift" – inconsistencies between T, S-implied potential density and the prescribed coordinate value. Since HYCOM's coordinate layers for the most part are isopycnic, HYCOM has inherited this problem from its predecessor MICOM.

Early versions of MICOM (Bleck et al., 1992) addressed this issue by invoking a "coordinate maintenance" algorithm which was also part of Oberhuber's (1993) isopycnic model. The strategy adopted in subsequent MICOM versions was to alleviate both the redundancy and the consistency problem by advecting only one variable, S, and diagnose T from S and the coordinate value ρ. (This strategy clearly works only in isopycnic layers. MICOM's nonisopycnic slab mixed layer requires a second prognostic thermodynamic tracer aside from S.)

The pitfalls of diagnosing T from S and ρ are obvious and have prompted some MICOM users working on polar ocean circulation problems to switch from S to T as prognostic variable (e.g., Holland and Jenkins, 2001). Treating T as a diagnostic variable not subject to an explicitly enforced conservation law is also problematic in climate models used for predicting secular temperature changes in the ocean-atmosphere system. But the alternative, advecting T and diagnosing S everywhere in a global model, has its own drawbacks because of the strong correlation between ρ and T in the stratified low- to mid-latitude upper ocean which makes salinity a relatively poorly constrained diagnostic variable there.

Dispensing of a conservation equation for S may also be more detrimental to dynamic stability than dispensing of one for T because of the somewhat stronger control exerted by the atmosphere on the oceanic T field. This is to say that spurious salinity transients are harder to control in a model (in the absence of artificial restoring boundary conditions, that is) because of the lack of a natural restoring process on salinity akin to thermal relaxation. Given that salinity is more likely to act dynamically as a "loose cannon", one can argue that, globally speaking, S conservation is more important than T conservation in situations where a choice must be made between the two.

Since there is no guarantee that ρ is spatially uniform in any given HYCOM coordinate layer, HYCOM must everywhere carry two prognostic thermodynamic tracers. The strategy adopted in the production version is to treat both T and S as prognostic variables and delegate the coordinate maintenance task to the grid generator. Unfortunately, this choice is not optimal in all respects.

The real ocean has a tendency toward "density compensation", meaning that T, S fields evolve in a manner which minimizes the dynamic effects of T, S contrasts on the buoyancy field. This is to say that salinity fronts are often accompanied by compensating temperature fronts. The main problem with advecting T, S in a numerical model (any model, not just HYCOM) is that numerical shortcomings of the transport algorithm can and will destroy the spatial coherence of T, S fronts. In HYCOM, this will lead to localized ρ anomalies which the grid generator, in an attempt to restore target density, will convert into undulations in the layer thickness field. This adjustment in turn causes additional vertical dispersion.

It is for this reason that experiments continue in which T, S are replaced as prognostic variables in HYCOM by pairs of thermodynamic tracers which are more likely to maintain the coherence of T, S fronts during advection. In MICOM's isopycnic interior, this is presently achieved by advecting only one variable, S or T, and diagnosing the other one knowing ρ. The most straightforward generalization of this concept to the HYCOM case is to advect, as is done in MICOM's slab mixed layer, either the pair ρ, S or the pair ρ, T. An alternative approach, which addresses the pitfalls of diagnosing either T from S or S from T, is discussed below.

Spiciness

Because the density of near-freezing sea water is mainly a function of salinity, diagnosing T from ρ and S is an ill-posed problem in polar oceans – in the sense that small changes in S or ρ can bring about large changes in the diagnosed value of T. If advecting ρ in HYCOM is deemed important for maintaining the spatial coherence of T, S fronts, then the ideal second variable to be advected should be one whose isolines are everywhere orthogonal to isopycnals in T, S space. Such a variable exists and has become known as spiciness (Flament, 2002).

Orthogonality means that we need to construct a function χ satisfying

$$\begin{pmatrix} \partial\chi/\partial S \\ \partial\chi/\partial T \end{pmatrix} \cdot \begin{pmatrix} \partial\rho/\partial S \\ \partial\rho/\partial T \end{pmatrix} = 0.$$

Since

$$\begin{pmatrix} -\partial\rho/\partial T \\ \partial\rho/\partial S \end{pmatrix} \cdot \begin{pmatrix} \partial\rho/\partial S \\ \partial\rho/\partial T \end{pmatrix} = 0,$$

χ and ρ are connected through

$$\partial\chi/\partial S = -\partial\rho/\partial T \qquad\qquad \partial\chi/\partial T = \partial\rho/\partial S.$$

For dimensional consistency, T and S must appear in these expressions in nondimensional form. With this in mind, we can construct a perfect differential of χ,

$$d\,\chi = \left(-\frac{\partial\rho}{\partial(T/T_0)}\right) d(S/S_0) + \left(\frac{\partial\rho}{\partial(S/S_0)}\right) d(T/T_0), \qquad (1.2)$$

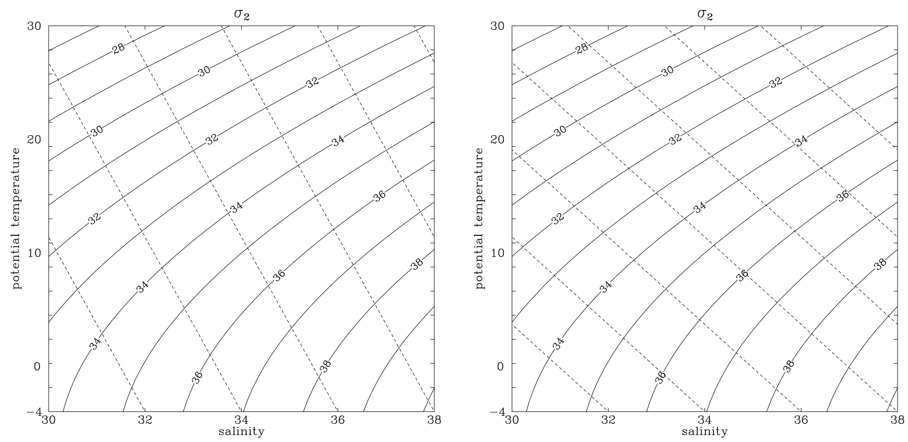

Figure 1.3. T, S diagrams showing ispopycnals referenced to 2000 m (solid) and two renditions of linearized spiciness χ' (dashed). Left: λ=-0.13 psu/°C; right: λ=-0.26 psu/°C.

which upon integration yields the sought-after spiciness function χ. A multiplicative or additive constant can be incorporated into the definition of χ at will.

The appearance of T_0, S_0 in (1.2) implies that orthogonality of ρ and χ is not "universal" but rather a matter of scaling. To visualize this dependence, consider two sets of orthogonal lines plotted in a T, S diagram, one set representing ρ and one representing χ. These lines lose their orthogonality as soon as the diagram is stretched in one or the other direction. The stretching operation is equivalent to changing T_0 and/or S_0.

The need in MICOM to recover T diagnostically from known values of ρ and S, and to do this in a noniterative fashion (iteratively obtained solutions tend to be unreliable in ill-posed problems), requires that the equation of state be approximated by a polynomial of at most 4^{th} degree in T. The approximation traditionally used is of 3^{rd} order in T and 1^{st} order in S (Brydon et al., 1999):

$$\rho(S,T) = c_1 + c_2 T + c_3 S + c_4 T^2 + c_5 ST + c_6 T^3 + c_7 ST^2. \qquad (1.3)$$

Assuming that T, S are already nondimensionalized, it is easy to derive from this polynomial the differential expression (1.2) and integrate the latter to obtain a polynomial expression for spiciness:

$$\chi(S,T) = -(c_2 S + 2c_4 ST + \frac{1}{2}c_5 S^2 + 3c_6 ST^2 + 2c_7 S^2 T^2) + \\ +(c_3 T + \frac{1}{2}c_5 T^2 + \frac{1}{3}c_7 T^3). \qquad (1.4)$$

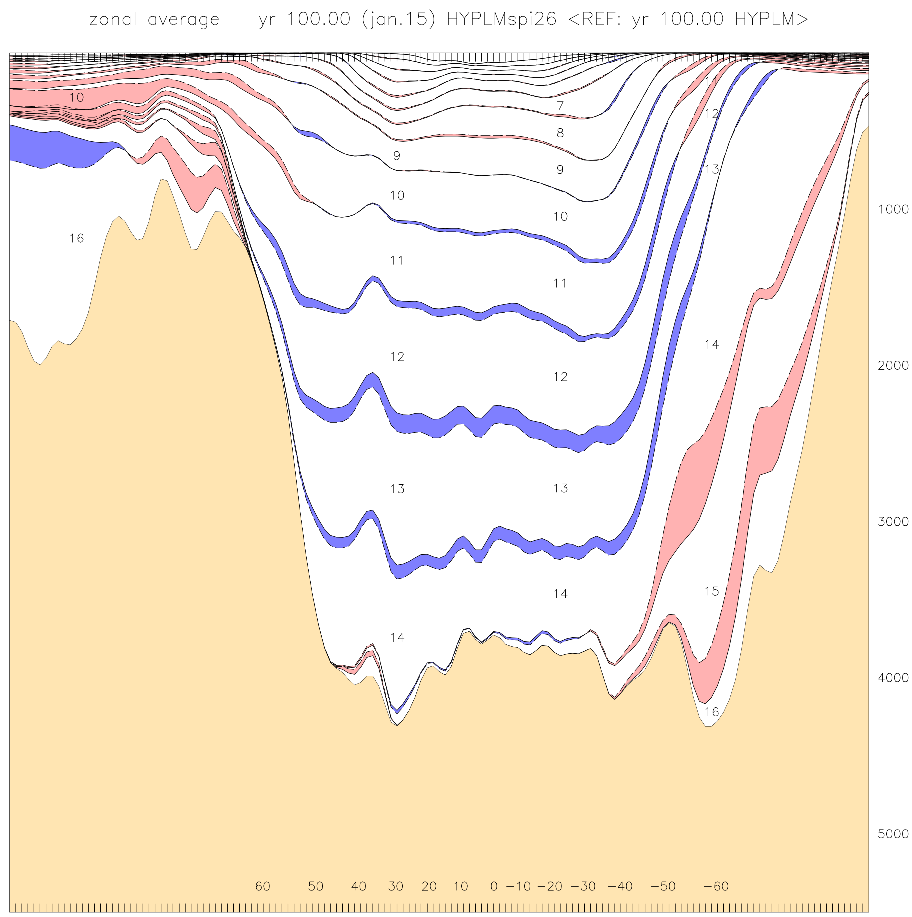

Figure 1.4. Vertical section through the zonally averaged density field after a 100-year, global, coarse mesh integration of HYCOM forced by monthly climatology [mesh size $2° \times \cos(lat.)$]. Colors highlight differences in isopycnal layer depth resulting from using two different λ values, -0.26 psu/°C and 0. Blue/red: interfaces in λ=-0.26 psu/°C run are at shallower/greater depth, respectively, than interfaces in $\lambda = 0$ run.

Recall that the goal of the present exercise is to advect buoyancy-related properties in terms of the pair ρ, χ instead of T, S. While computing ρ, χ from T, S at the beginning of each advection step is trivial, the inverse, i.e., recovering T, S from the advected ρ, χ fields by jointly solving (1.3) and (1.4) in noniterative fashion seems impossible. One way to overcome this obstacle is to use a linear approximation of (1.4) as the second advected variable. Let us write this variable as

$$\chi' = S + \lambda T \tag{1.5}$$

where λ is a free parameter which should be chosen to mimic the orthogonality of ρ and χ across the T, S range encountered in the world ocean. Solving the coupled system (1.3),(1.5) for T, S, given ρ and χ', is no more complicated than diagnosing T from (1.3), given ρ and S.

The dashed lines in Fig. 1.3 show two choices of χ' which roughly optimize ρ–χ' orthogonality at high and low temperatures. The corresponding λ values are -0.13 psu/°C and -0.26 psu/°C, respectively. (Note that χ' is expressed here in salinity units.)

The advantages of replacing ρ, S advection by ρ, χ' advection are difficult to quantify, mainly because the complexities of the wind- and thermohaline-forced ocean circulation make it hard to identify numerical solutions that are clearly impacted by errors resulting from advecting S and treating T as a diagnostic variable.

Fig. 1.4 illustrates the extent to which two HYCOM solutions, one based on ρ, S advection and one on ρ, χ' advection using λ = -0.26 psu/°C, diverge during a global 100-year coarse-mesh global simulation forced by monthly climatology. In this figure, the gap between isopycnal layer interfaces in the two simulations is colored red or blue depending on whether the interfaces in the χ'-based solution are at a greater or shallower depth, respectively, than the corresponding interfaces in the reference solution based on S advection ($\lambda = 0$). Fig. 1.4 shows that the use of spiciness leads to a very slight density increase at low-to mid-latitudes (blue coloration), while density is seen to decrease to a somewhat larger degree at high latitudes (red coloration).

Differences between a model run based on λ = -0.13 psu/°C and the S-advecting reference run are roughly half as large as those shown in Fig. 1.4.

It is virtually impossible to judge whether the ρ, χ'-based solution represents an improvement over the ρ, S-based one. We are able to state, however, that the switch from S to χ' advection induces much smaller changes (at least in this particular experiment) than what is typically seen in 100-year experiments when surface forcing fields or aspects of model physics are changed. This is a positive result, because it indicates that the increased robustness of the algorithm for diagnosing T from the two prognostic mass field variables is not achieved at the price of encountering new potentially harmful model sensitivities.

Perhaps the best evidence that advecting χ' instead of S leads to improved model performance is seen in a multi-century integration of a coupled ocean-atmosphere model consisting of HYCOM and the GISS atmospheric general circulation model. An early experiment based on the traditional ρ, S-advecting version of HYCOM showed incidents of anomalous ice growth occurring roughly once per century (grey curve