Available online at www.sciencedirect.com
ScienceDirect
Ocean Modelling 20 (2008) 270–292
Ocean Modelling
www.elsevier.com/locate/ocemod

A pretty good sponge: Dealing with open boundaries in limited-area ocean models

J.W. Lavelle [a,*], W.C. Thacker [b]

[a] *Pacific Marine Environmental Laboratory, 7600 Sand Point Way NE, Seattle, WA 98115, USA*
[b] *Atlantic Oceanographic and Meteorological Laboratory, 4301 Rickenbacker Causeway, Miami, FL 33149, USA*

Received 19 July 2007; received in revised form 28 September 2007; accepted 2 October 2007
Available online 22 October 2007

Abstract

The problem of computing within a limited domain surrounded by open boundaries is discussed within the context of the shallow-water wave equations by comparing three different treatments, all of which surround the domain by absorbing zones intended to prevent reflections of outgoing waves. The first, which has attracted a lot of attention for use in electromagnetic and aeroacoustic applications, is intended to prevent *all* reflections. However, it has not yet been developed to handle the second important requirement of open boundaries, namely the ability to pass information about external conditions into the domain of interest. The other two treatments, which absorb differences from a specified external solution, allow information to pass through the open boundary in both directions. One, based on the flow relaxation scheme of [Martinsen, E.A., Engedahl, H., 1987. Implementation and testing of a lateral boundary scheme as an open-boundary condition in a barotropic ocean model. Coastal Eng. 11, 603–627] and termed here the "simple sponge," relaxes all fields toward their external counterparts. The other, a simplification and generalization of the perfectly matched layer, referred to here as the "pretty good sponge," avoids absorbing the component of momentum parallel to the open boundary. Comparisons for a case that is dominated by outgoing waves shows the pretty good sponge to perform essentially as well as the perfectly matched layer and better than the simple sponge. In comparisons for a geostrophically balanced eddy passing through open boundaries, the pretty good sponge out-performed the simple sponge when the only external information available was about the advecting flow, but when information about the nature of the eddy in the sponge zones was also available, the simple sponge performed better. For the case of an equatorial soliton passing through the boundary and no information provided about its nature outside the open domain, again the pretty good sponge performed better. Proving useful for situations governed by nonlinear equations forced by external conditions and being easy to implement, the pretty good sponge should be considered for use with existing limited-area ocean models.
Published by Elsevier Ltd.

Keywords: Perfectly matched layer; Open boundary; Absorbing layer; Sponge layer; External forcing; Flow relaxation; Regional model

* Corresponding author. Tel.: +1 206 526 6182.
E-mail address: J.William.Lavelle@noaa.gov (J.W. Lavelle).

1463-5003/$ - see front matter Published by Elsevier Ltd.
doi:10.1016/j.ocemod.2007.10.002

1. Introduction

The issue of how to treat the boundary when computing within limited domains has received a great deal of attention for a wide variety of problems. Information from outside the computational domain essential to the solution must be allowed to enter, and what should leave the domain must be allowed to depart. While there are rare situations where resources allow the size of the domain to be expanded far beyond the region of interest to minimize adverse impacts of misspecified boundary conditions, practicality usually dictates that the computational domain be not much larger than the region of interest, raising concern regarding how the open boundaries should be treated.

Techniques for treating open boundary generally fall into two categories. The first encompasses the radiation and characteristic boundary conditions, which should prevent outgoing waves from reflecting from the artificial boundary back into the computational domain (Orlanski, 1976; Flather, 1976; Engquist and Majda, 1977; Camerlengo and O'Brien, 1980; Higdon, 1994; Lie, 2001; Marchesiello et al., 2001; McDonald, 2002; Givoli and Neta, 2003; van Joolen et al., 2004; Blayo and Debreu, 2005). For these, dispersive waves approaching the boundary from a non-normal direction can be a problem (Raymond and Kuo, 1984), as can the issue of distinguishing between waves and slow modes (Nycander and Döoš, 2003). The second category encompasses techniques involving the use of absorbing layers, which allow disturbances to pass out of the region of interest into a limited exterior region where they can be dissipated before reflecting from the computational boundaries (Davies, 1976; Philander et al., 1987; Chassignet et al., 2003). Absorbing layers are generally easier to use, and they allow a simple way for passing information about the solution external to the region of interest (Martinsen and Engedahl, 1987; Bodony, 2006). Tsynkov (1998) reviews the early literature, but much work has been done in the intervening decade. In particular, the absorbing-layer technique referred to as the perfectly matched layer has been the focus of much attention.

Berenger (1994) introduced the first perfectly matched layer technique within the context of electromagnetic waves, splitting one variable into contributions associated with directions parallel and perpendicular to the boundary and splitting the equation governing its temporal change into two equations, one for each contribution. While analysis indicates that the governing partial–differential equations are entirely free from reflections for waves from any direction,[1] finite-difference computations do have small reflections that can be controlled by the thickness and properties of the absorbing layer. Hu (1996) adapted the technique for aeroacoustics, splitting all of the variables and doubling the number of governing equations within absorbing layers. Again there were small reflections, and filtering was used to prevent numerical instability. Abarbanel and Gottlieb (1997) analyzed Berenger's perfectly-matched-layer equations and Hesthaven (1998) analyzed Hu's, both finding that the split set of equations is not strongly well-posed, thus the need for filtering. Within a context much closer to ocean modeling Darblade et al. (1997) presented a splitting technique for the shallow-water equations linearized about a state of no motion and demonstrated that waves passing into the absorbing layers should have no reflections, while Navon et al. (2004) examined a similar splitting technique for the shallow-water wave equations linearized about a steady, spatially uniform mean flow and found that filtering was needed to maintain stability. Although focusing on elastodynamics Bécache et al. (2003) showed for general hyperbolic problems that the equations for perfectly matched layers were unstable unless components in the direction normal to the open boundary of group velocity and wave vector have the same sign. Their analysis for the linearized Euler equations, which resemble shallow-water wave equations with a mean flow, suggest that advection causes perfectly matched layers to become unstable. Hu (2001) reaches a similar conclusion. The perfectly matched layer can be formulated by introducing auxiliary variables rather than by splitting the original variables (Turkel and Yefet, 1998; Abarbanel et al., 1999; Hu, 2001). While this alternative formulation remains only weakly well-posed and unstable, adding terms that cause auxiliary variables to decay can control the problem without seriously degrading the layer's anti-reflective property.

In a recent modeling study Lavelle (2006) needed to prevent internal gravity waves generated by the interaction of the barotropic tide with a seamount from reflecting from the boundaries of the limited

[1] There should be no reflection as waves pass into the absorbing layer, but they might reflect at the outer boundary after passing through the layer. The width of the layer and the strength of the absorption parameters determine the amplitudes of the reflected wave that re-enter the domain.

computational domain. That problem clearly illustrates the two requirements that must be addressed via the treatment of the computational boundaries: (1) the need to prevent reflections of outgoing waves and (2) the need to specify the external conditions that force the solution. Within relatively thin sponge regions he nudged the solution toward a background solution for the tidal flow in the vicinity of the boundaries, successfully suppressing unwanted reflections and achieving negligible adverse impact in the vicinity of the seamount. His technique might be regarded as a stripped-down version of the un-split perfectly matched layer of Hu (in press), one requiring no auxiliary variables. It also resembles the flow relaxation scheme of Martinsen and Engedahl (1987), especially in its specification of an external solution for the tides. The purpose of this paper is to discuss his approach in a bit more depth, comparing it to those of Hu (in press) and Martinsen and Engedahl (1987) within the simpler context provided by the shallow-water wave equations.

The shallow-water wave equations, which can be regarded as proxies for the equations describing three-dimensional stratified geophysical flows,[2] were chosen as the context for the comparisons. These equations present the opportunity to focus on two aspects that are important for oceanographic problems, neither of which have received attention by the perfectly-matched-layer community. The first is the need to specify external information that is crucial for the solution within the limited-area open domain. In the motivating example, that information is the tidal flow through the boundaries; if that were not known, there would be no scattering and no waves that must depart the domain. Perfectly-matched-layer techniques have not yet been developed to address this issue; even though we expect no major obstacle to such development, that task is beyond the scope of this comparison. The second difference presented by the shallow-water wave equations is the fact that only a part of an initial disturbance radiates away. Due to the earth's rotation, steady currents can balance the slope of the sea surface via the Coriolis force, with the radiating waves being the mechanism for achieving geostrophic balance. For many geophysical applications the geostrophically balanced part of the flow is of greatest interest and the adjustment waves are regarded as noise. At open boundaries, the adjustment waves should be allowed to pass out, but special care may be needed with the balanced part of the flow.

Comparisons are made for several different situations. Section 2 considers the case of an initially motionless mound of water that collapses into outgoing waves. Even though this case has been chosen to emphasize the problem of allowing waves to exit, a residual geostrophically balanced flow should remain. For this case all three formulations of absorbing layers are compared: a perfectly-matched-layer formulation following the approach of Hu (in press), a "simple" sponge similar to that of Martinsen and Engedahl (1987), and the "pretty good sponge" like that of Lavelle (2006). The pretty good sponge can be seen to be a minimalist approach to the perfectly matched layer, retaining its peculiar feature of *not* absorbing the component of velocity parallel to the boundary. The simple sponge, on the other hand, absorbs all variables equally. Both the pretty good sponge and the simple sponge are much easier to implement than the perfectly matched layer and less plagued with issues of instability and ill-posedness, but they make no attempt to guarantee that all outgoing waves pass into the absorption zone without reflecting.

Section 3 shifts the focus to the second difficulty presented by open boundaries, namely the specification of external influences that must be incorporated into the computations. Two cases are considered, both involving a geostrophically balanced anti-cyclonic eddy advected by a larger-scale flow. While it would have been possible to describe these cases using linear equations, like those of the first example but with additional terms for the mean flow, the nonlinear shallow-water wave equations were chosen, because they are considerably more general. Rather than having the advecting flow be specified within the linearized equations, it must be provided as external information. This information is provided only within the absorbing zones, where departures from the advecting flow are absorbed. As the perfectly-matched-layer technique has not yet been developed either for specifying external flows or for general nonlinear equations, it is not compared. The first case to be considered is that of the eddy passing out of the open domain. Both the simple sponge and the pretty good sponge allow the eddy to exit and the flow to return to the background advecting flow when differences from the advecting flow are absorbed, with the pretty good sponge performing somewhat better than the simple sponge. However, both methods resulted in an unwanted loss of amplitude during the eddy's exit. As this

[2] Numerical models like HYCOM, Chassignet et al. (2003) approximate the ocean as a stack of layers governed by the shallow-water wave equations.

attenuation can be attributed to an incorrect specification of the external situation, one which ignores the presence of the eddy in the absorption zone, the second case explores the value of better external information. In analogy with nested-grid simulations, where the solution from one grid is transferred to another, the external solution within the absorption zones is taken from a reference solution employing a much larger computational domain. Armed with this information about what is really transpiring inside the absorbing region, this case examines an eddy *entering*, crossing, and exiting the limited domain. As expected, results benefit greatly from the improved information, this time with the simple sponge giving better results.

As the advecting eddy could have been treated using linear equations, Section 4 examines the case of an equatorial soliton, for which nonlinearity is essential. The external solution is taken to be that of no flow and no surface elevation with no attempt made to incorporate the presence of the soliton within the absorbing zones.[3] Because the simple sponge absorbs the component of velocity parallel to the boundary, its performance is noticeably worse than that of the pretty good sponge. Its loss of amplitude leads to a slower propagation speed and slightly delayed arrival at the open boundary. Neither passed through the boundary perfectly, but both did remarkably well considering that the information provided in the absorbing layer included no knowledge of their presence.

2. Absorbing layers for linear shallow-water wave equations

A mound of water collapsing into a train of outgoing waves provides an example for examining how well waves are allowed to pass out of the region of interest without reflection. As with all absorbing-layer approaches to open boundary conditions, the computational domain is extended beyond the domain of interest to include an absorbing zone within which the governing equations are modified by the addition of dissipative terms. The computational boundary becomes the outer boundary of the absorbing layer, and the interface between the domain of interest and the absorbing layer marks a transition in the nature of the equations that are to be solved. Fig. 1 shows a small domain of interest that extends ±510 km in both coordinate directions from the origin surrounded by an absorbing layer that is about 130 km wide, which is embedded within a larger computational domain to be used to generate reference solutions uncontaminated by reflections.

The evolution of an initially motionless Gaussian-shaped mound of water is described by the linear shallow-water wave equations:

$$\frac{\partial \eta}{\partial t} + \frac{\partial (Hu)}{\partial x} + \frac{\partial (Hv)}{\partial y} = 0, \tag{1}$$

$$\frac{\partial u}{\partial t} + g\frac{\partial \eta}{\partial x} - fv = 0, \tag{2}$$

$$\frac{\partial v}{\partial t} + g\frac{\partial \eta}{\partial y} + fu = 0, \tag{3}$$

where η is the surface elevation, u and v are components of the velocity in the x- and y-directions, respectively, t is time, H is the basin depth, g is the acceleration of gravity, and f is the Coriolis parameter accounting for the influence of the earth's rotation on the flow. If there were no rotation, all traces of the initial mound should ultimately disappear, leaving the surface flat $\eta = 0$ and motionless $u = v = 0$ within the region of interest. With $f \neq 0$, the radiated waves can be considered to be part of a geostrophic adjustment process, which leaves a residual circulation with the slope of the surface balanced by the Coriolis force. The boundary treatment should allow what should exit to exit and what should remain to remain.

Good results can be anticipated for the collapsing mound when using perfectly matched layers, as they are designed to prevent reflections of outgoing waves. Hu (in press) derives the perfectly matched layer from a complex coordinate transformation, which can be summarized by the replacements:

[3] Because of the nonlinearity of the equations, again no comparison is possible for the perfectly-matched-layer technique.

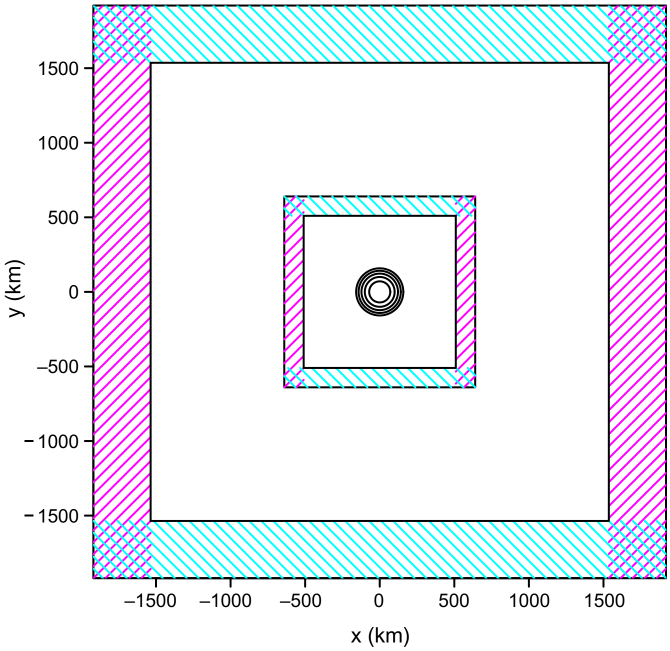

Fig. 1. Solutions within a square 1020 km on a side, which is extended with an absorbing zone 130 km wide, can be compared with solutions for the larger 3072 km square and its surrounding absorbing zone. For both the perfectly matched layer and the pretty good sponge, hatching with 45° slope indicates where σ_x attenuates u and η, while that with the −45° slope indicates where σ_y attenuates v and η. For the simple sponge, u, v, and η are all attenuated by σ wherever there is hatching of either slope. Circles at the center are contours of initial surface elevation for the Gaussian mound of water of the first example; the elevation is 1 m at the center and decreases radially by a factor of 0.01 with each contour. The larger domain provides reference solutions during the interval of time before waves reflecting from its computational boundary return to the smaller domain of interest.

$$\frac{\partial}{\partial x} \rightarrow \frac{\partial}{\partial x} \Big/ \left(1 + \sigma_x \frac{\partial}{\partial t}^{-1}\right), \tag{4}$$

$$\frac{\partial}{\partial y} \rightarrow \frac{\partial}{\partial y} \Big/ \left(1 + \sigma_y \frac{\partial}{\partial t}^{-1}\right), \tag{5}$$

where multiplication can clear the denominators and where $\frac{\partial}{\partial t}^{-1}$ is the temporal anti-derivative operator. For absorbing layers crossing the x-axis, (4) should be used; for those crossing the y-axis, (5); and for the corners, both should be used; see Fig. 2. The absorption coefficients σ_x and σ_y are functions of x and y, respectively; they grow from zero within the region of interest, absorbing waves more strongly the deeper they penetrate into the layer. The thickness of the layer and the strength of the absorption affect the chances that the wave might reflect at the outer boundary.

When applied to shallow-water wave equations linearized about the no-flow state, this prescription for a perfectly matched layer yields:

$$\frac{\partial \eta}{\partial t} + \frac{\partial (Hu)}{\partial x} + \frac{\partial (Hv)}{\partial y} = -(\sigma_x + \sigma_y)\eta - \sigma_x \sigma_y q_\eta - \sigma_y \frac{\partial (Hq_u)}{\partial x} - \sigma_x \frac{\partial (Hq_v)}{\partial y}, \tag{6}$$

$$\frac{\partial u}{\partial t} + g\frac{\partial \eta}{\partial x} - fv = -\sigma_x u + \sigma_x f q_v, \tag{7}$$

$$\frac{\partial v}{\partial t} + g\frac{\partial \eta}{\partial y} + fu = -\sigma_y v - \sigma_y f q_u. \tag{8}$$

The terms on the left-hand side of the equations correspond to the original shallow-water wave equations, while those on the right are the modifications for the perfectly matched layers. The temporal anti-derivative operators in (4) and (5) necessitate new fields q_η, q_u, and q_v, which are governed by additional equations:

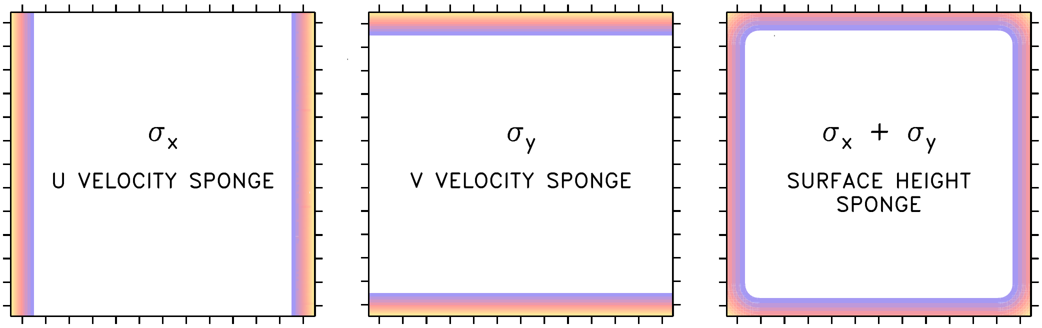

Fig. 2. Contours of absorption coefficients for the perfectly matched layers, which are the same as those for the pretty good sponge. There is no absorption of the component of velocity parallel to the boundaries, and the rate at which the normal component is absorbed increases with distance into the absorbing zone. Surface height is absorbed at all open boundaries, and the sum $\sigma_x + \sigma_y$ has contribution from both terms only in the corners.

$$\frac{\partial q_\eta}{\partial t} = \eta - \lambda q_\eta, \tag{9}$$

$$\frac{\partial q_u}{\partial t} = u - \lambda q_u, \tag{10}$$

$$\frac{\partial q_v}{\partial t} = v - \lambda q_v, \tag{11}$$

where, going beyond Hu's prescription, the positive parameter λ has been introduced to control instability.

It is instructive to compare the perfectly matched layer to a simpler formulation like that of Martinsen and Engedahl (1987):

$$\frac{\partial \eta}{\partial t} + \frac{\partial (Hu)}{\partial x} + \frac{\partial (Hv)}{\partial y} = -\sigma(\eta - \eta_e), \tag{12}$$

$$\frac{\partial u}{\partial t} + g\frac{\partial \eta}{\partial x} - fv = -\sigma(u - u_e), \tag{13}$$

$$\frac{\partial v}{\partial t} + g\frac{\partial \eta}{\partial y} + fu = -\sigma(v - v_e). \tag{14}$$

The absorption coefficient σ of the simple sponge acts on all variables, whereas by distinguishing between σ_x and σ_y the perfectly matched layer prevents absorption of v at x boundaries and absorption of u at y boundaries. Second, the simple sponge introduces no new fields. Martinsen and Engedahl (1987) incorporated tidal forcing via an external solution (η_e, u_e, v_e), deviations from which are absorbed in the sponge region; for an external solution with all variables zero, the damping is to zero like with the perfectly matched layer. The absorption coefficient σ varies with x or y as appropriate, increasing with distance into the absorbing layers. While the simple sponge introduces a stable complex frequency via the transformation $\partial/\partial t \to \partial/\partial t + \sigma$, care may be needed numerically to guarantee that the size of the time step τ is consistent with a Courant condition $\sigma\tau \lesssim 1$.

What happens if the auxiliary fields of the perfectly matched layer were all set to zero? Eqs. (9)–(11) are not needed, and what is left of (6)–(8), when generalized to allow for an external solution, describe the pretty good sponge introduced by Lavelle (2006):

$$\frac{\partial \eta}{\partial t} + \frac{\partial (Hu)}{\partial x} + \frac{\partial (Hv)}{\partial y} = -(\sigma_x + \sigma_y)(\eta - \eta_e), \tag{15}$$

$$\frac{\partial u}{\partial t} + g\frac{\partial \eta}{\partial x} - fv = -\sigma_x(u - u_e), \tag{16}$$

$$\frac{\partial v}{\partial t} + g\frac{\partial \eta}{\partial y} + fu = -\sigma_y(v - v_e). \tag{17}$$

Just as for the perfectly matched layer, only u and η are absorbed in the x-direction and v and η in the y-direction,[4] with σ_x varying with x and σ_y varying with y.

For all three methods boundary conditions must be specified at the outer limit of the absorbing zone. One possibility is to use radiation conditions to terminate the sponge layer (Petropoulos, 1998), but if there has been sufficient absorption, that should not be necessary. With strong absorption the solution is forced toward the external solution, and as the increase over the last grid cell can be regarded as achieving an infinite absorption rate, the external solution can provide values at the boundary. An easily implemented alternative is to require that derivatives normal to the boundary of differences from the external solution should vanish. That is the approach taken in the examples below.

The values chosen for the absorption rates σ_x and σ_y can depend on the details of the discrete formulation of the equations, as numerical instability might limit their values. For this case involving the collapsing Gaussian mound, as well as for those to follow, the maximum absorption rate was set by $\sigma_{\max}\tau = 0.9$, where τ is the time step used for the temporal integration. Within the layers absorption increases quadratically with distance d from the inner edge of the absorbing zone (Fig. 2); e.g., $\sigma_x = \sigma_{\max}(d/D)^2$, where D is the width of the absorbing zone where $\sigma_x \neq 0$. For all cases considered here, the widths of the absorbing zones were taken to be approximately[5] 10% of the span of the computational domain in the corresponding direction. Values for σ_x and σ_y were always the same for the pretty good sponge as for the perfectly matched layer. For the simple sponge, σ was taken to be the same as the sum $\sigma_x + \sigma_y$ in all sponge zones, so that its absorption of η would be the same as for the pretty good sponge. When using perfectly matched layers, a value must be assigned to λ that will control the instability attributed to the introduction of auxiliary fields; trial and error led to the choice of $\lambda\tau = 1.66$.

To be able to evaluate how well the different absorbing layers prevent reflections, a solution without reflections is needed. This solution is generated using a grid that is much wider than the region of interest, and comparisons are made during an interval short enough to guarantee that no reflections from the boundaries of that larger grid have time to return. The region of interest is taken to be a square 1020 km on a side, and the reference solution is generated within a 3840 km square, as shown in Fig. 1. So if the first disturbance leaves the region of interest in 3 h, no reflections from the boundary of the larger grid can be expected for another 16 h.

The initial disturbance is a Gaussian-shaped mound of water: $\eta(x,y,t=0) = \eta_0 \exp(-(x^2+y^2)/R^2)$ with $\eta_0 = 1$ m, $R = 50$ km, and the coordinate origin at the center of the square domain. The water is initially at rest, not in geostrophic balance. With basin depth of $H = 100$ m, the shallow-water wave speed is $\sqrt{gH} \approx 113$ km/h and, with a Coriolis latitude of 45°, $f = 1.028 \times 10^{-4}\,\mathrm{s}^{-1}$ and the Rossby radius is $\sqrt{gH}/f \approx 300$ km. The adjustment to rotation consists of waves radiating outward and a small geostrophically balanced residual within the region of interest. The focus in this example is more on allowing the disturbance to depart than on the residual balanced circulation, the accuracy of which might be affected by the treatment of the open boundaries. So the external solution for the sponge region is taken to be the same homogeneous solution (all fields zero) as would be appropriate for the non-rotating case.

Fig. 3 illustrates the progression of the wave train as the mound collapses. The color-filled contours show results for the reference solution limited to the domain of interest at hourly intervals, starting 4 h after the initial conditions (lower-left panel) and ending at 9 h (upper-right). Overlaid are line contours at the same 2 cm spacing showing corresponding results for the perfectly-matched-layer solution. The red peak and trailing blue trough can be seen to exit the domain with minimal reflections, and the line contours can be seen to align well with changes in color. After the waves have exited, a residual geostrophically balanced mound of water remains, and far from its center the water elevation is essentially zero.

Fig. 4 shows another view of the same solution. Cross sections through the surface-elevation field along the x-axis are plotted for the same times shown in Fig. 3. For clarity, successive curves are displaced vertically, e.g., the amplitudes of the yellow curve for $t = 9$ h read 5 cm higher than they actually are, while the amplitudes of the red curve for $t = 4$ h are correctly indicated on the vertical axis. Note that the 1 m height of the initial mound leads to a first wave with peak amplitude of less than 8 cm entering the absorption layer

[4] Lavelle (2006) used only half the sum of σ_x and σ_y for absorbing η and obtained good results.

[5] The number and size of grid cells precluded a precise 10%.

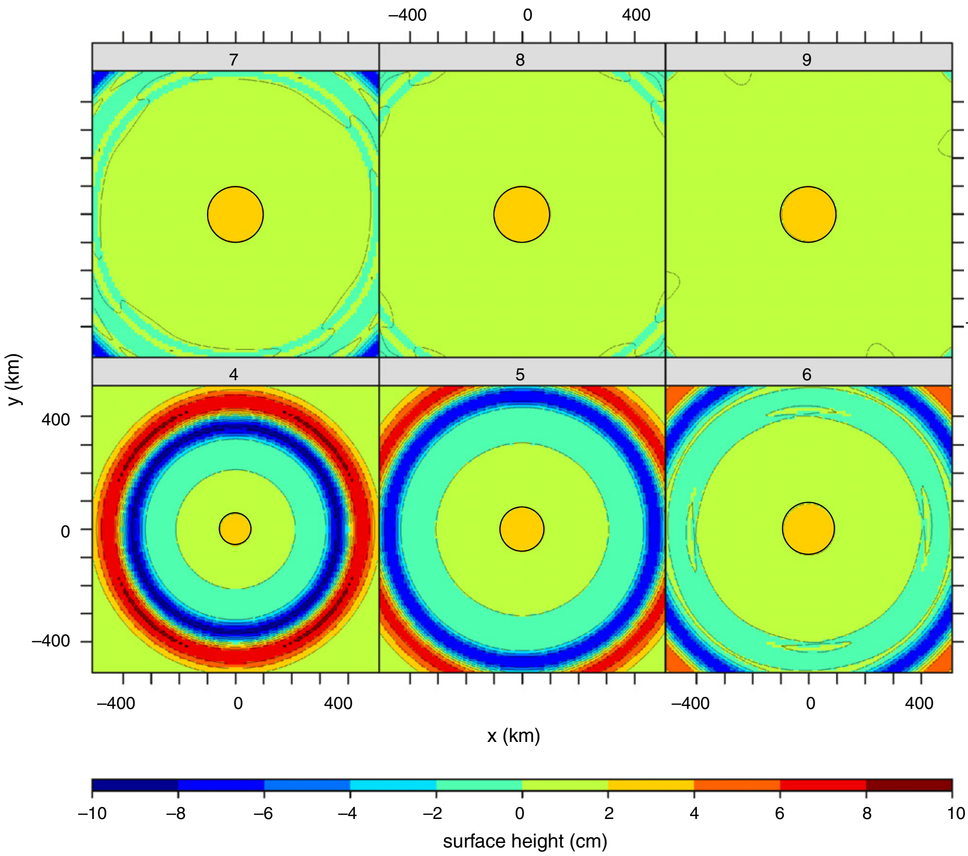

Fig. 3. Colors indicate the surface elevation (cm) for the reference solution for the collapsing Gaussian mound at times (h) indicated by the panel labels. Line contours indicate corresponding results for the solution with perfectly matched layers.

followed by a trough with slightly larger but negative amplitude. For each time, the perfectly-matched-layer solution has been superimposed over black curves representing the reference solution. Agreement within the domain of interest is sufficiently good that the superimposed curves generally obscure the reference solution. The largest differences are seen at the edges after the adjustment waves have left the domain. Damping of the residual geostrophically balanced circulation within the absorbing zones may contribute to these differences.

Fig. 5 shows results for the pretty good sponge and for the simple sponge, corresponding to those for the perfectly matched layer in Fig. 3. It is necessary to look quite closely at the upper panels to see differences between results using the pretty good sponge and those for the perfectly matched layer. However, results using the simple sponge, while not bad, are clearly worse than those using the other two techniques.

As the contour plots of the solutions look quite similar for all three open-boundary techniques, it is useful to focus on their errors, i.e., on their differences from the large-grid reference solution. Fig. 6 shows the maximum, minimum, and median errors of the surface elevation over the entire region of interest as a function of time for the three techniques. The peaks on the curves for the maximum errors occur when peaks of the wave train from the collapsing mound reach the corners of the domain, and the troughs on the curves for the minimum errors occur when troughs reach the corners. In both cases, extremes of the limited-area computations are less than those for the reference solution, indicating that damping within the absorbing zone has the side effect of also damping within the region of interest. In evaluating the relative performance of the three techniques, note that these statistics show the performance of computations using the simple sponge to have noticeably larger errors, while those using the pretty good sponge are very nearly the same as those when using a perfectly matched layer with the curves indicating maximum and median errors for the former almost obscuring those for the latter.

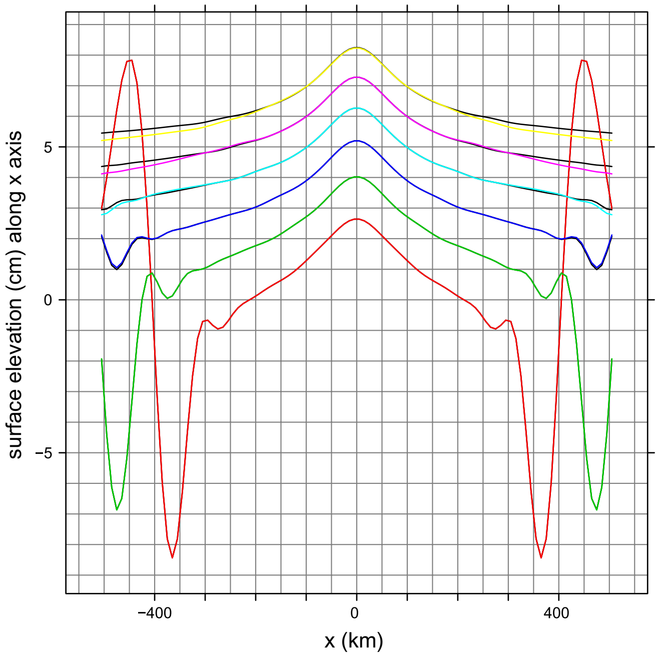

Fig. 4. Surface elevation along the x axis as computed using the perfectly matched layer are shown in color for the same sequence of times illustrated in Fig. 3 successively offset by 10 cm. Red indicates results at $t = 4$ h; green results at $t = 5$ h with 1 cm added; blue, $t = 6$ h with 2 cm added; ...; yellow, $t = 9$ h with 5 cm added. Also shown in black are the corresponding results for the reference solution, most of which are obscured by the colored curves.

3. Moving eddy

One issue that was not addressed by the example of the previous section was what to do when damping all fields to zero does not seem sensible. For example, if an eddy should propagate out of the region of interest, it is easy to imagine that reduction of its amplitude within an absorption layer would lead to an unwanted loss of amplitude for the part of the eddy still within the region of interest. To address such concerns, the absorption zone can damp the fields, not to zero, but to an external solution.

An eddy swept out of the open domain by an advecting flow provides a good example for examining this problem. Navon et al. (2004) and McDonald (2002) addressed this problem via the shallow-water equations linearized about a uniform background flow, the former using a variable-splitting version of the perfectly matched layer and the later using transparent boundary conditions without an absorbing layer. While we too could discuss this example within the same framework, we preferred to use the nonlinear shallow-water equations, as they provide a better proxy for a more general ocean model. Doing so raises the issue of how to maintain the steady flow that should advect the eddy. We do this by damping deviations from this flow in the sponge zones; there the external solution is characterized by a geostrophic surface tilt supporting a steady, uniform current.[6] If initial disturbances within the domain of interest are allowed to propagate away, then for a basin with a flat bottom what is left should be the advecting flow that is maintained by the external solution imposed within the absorbing layers. Thus, the external solution can be expected to get the information about the advecting flow into the computations. Departures from this external solution will be absorbed, so we can also expect that interior disturbances can exit. However, as this choice for the external solution provides no information about the nature of the eddy as it leaves the region of interest, we can expect the eddy's

[6] Note that the external information is provided everywhere in the sponge zones without diagnosing what should be coming in or going out. This can be contrasted to radiation methods' need to determine for which boundary points external information should be imposed.

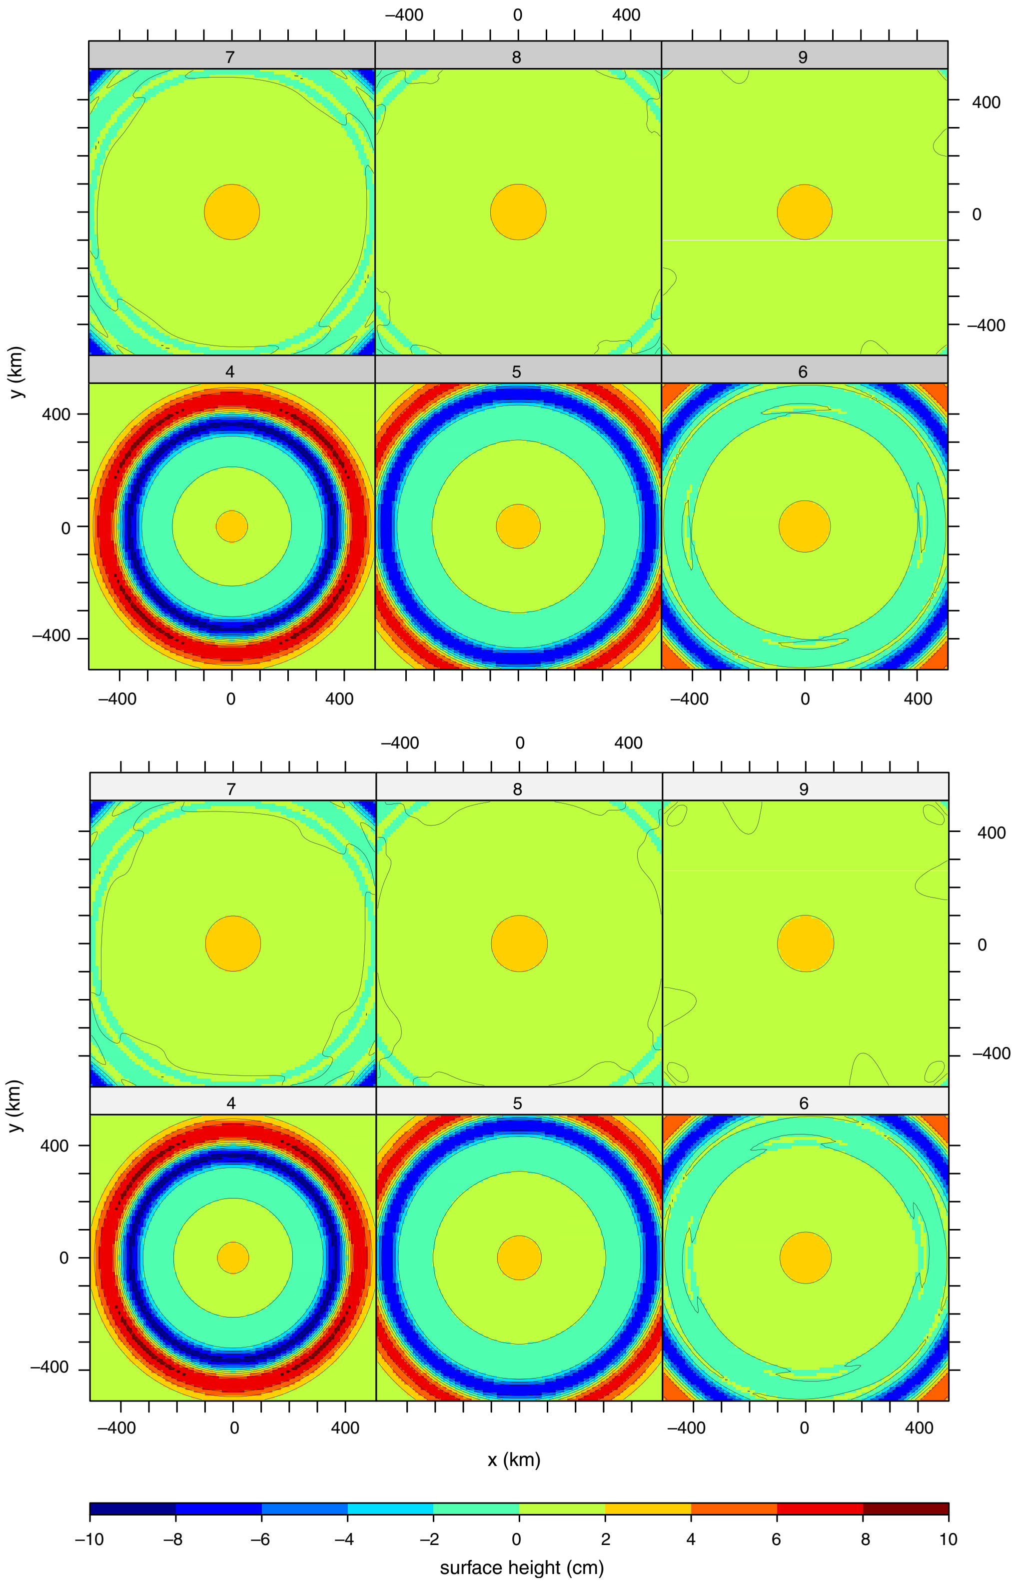

Fig. 5. Same as Fig. 3 but with line contours of the upper set of panels indicating results for the pretty good sponge and those for the lower set for the simple sponge.

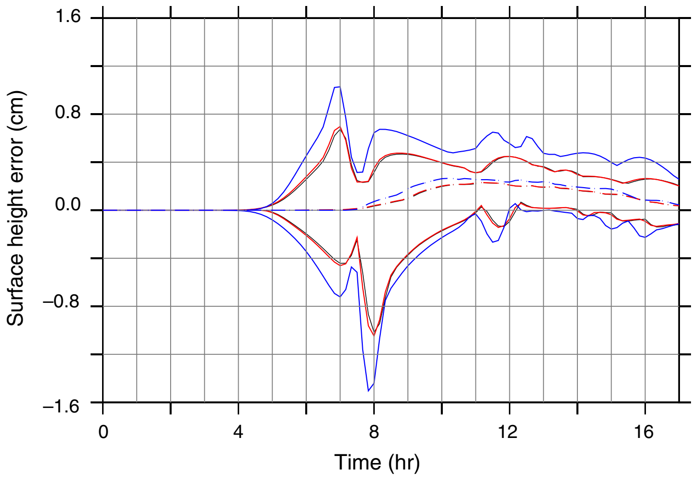

Fig. 6. Maximum (upper three curves), minimum (lower) and median (dashed) differences of surface elevation from the reference solution when using the perfectly matched layer (black), pretty good sponge (red), and the simple sponge (blue).

strength to be diminished as it passes through the sponge zone, and we can expect the weakened eddy's circulation to impact the solution within the region of interest.

The nonlinear shallow-water wave equations are:

$$\frac{\partial \eta}{\partial t} + \frac{\partial (H+\eta)u}{\partial x} + \frac{\partial (H+\eta)v}{\partial y} = 0, \tag{18}$$

$$\frac{\partial u}{\partial t} + u\frac{\partial u}{\partial x} + v\frac{\partial u}{\partial y} - fv + g\frac{\partial \eta}{\partial x} = 0, \tag{19}$$

$$\frac{\partial v}{\partial t} + u\frac{\partial v}{\partial x} + v\frac{\partial v}{\partial y} + fu + g\frac{\partial \eta}{\partial y} = 0. \tag{20}$$

In the sponge zones the zeros on the right-hand sides are replaced with the absorbing terms from the right-hand sides of (12)–(14) for the simple sponge and with those from (15)–(17) for the pretty good sponge. It is not clear how to obtain absorbing terms for the perfectly matched layer for the nonlinear equations. Even for the equations linearized around a state with uniform flow, problems with instabilities have required special treatment (Hu, 2001; Abarbanel et al., 1999).

Computations were made using the same large domain described in Section 2 to get a reference solution. The initial surface elevation was the same Gaussian used for the stationary mound of the previous examples added to a geostrophic tilt balancing a uniform current $u_e = v_e = 5$ m/s, while the initial velocity field was the sum of the external flow and a geostrophic flow balancing the slope of the initial mound. As this does not achieve nonlinear balance between the elevation and velocity fields, in the first hour adjustment waves propagate away, reducing the height of the rotating mound from 1.0 m to 0.9 m where it remains until the anticyclonic eddy has left the region of interest.[7] The color-filled contour plots of the reference solution for the surface elevation at 2-h intervals (Fig. 7) show the eddy, with steady amplitude and shape, as it approaches and passes out the northwest corner of the region of interest. Similar plots would show that both the simple sponge and the pretty good sponge produce solutions free of boundary reflections but with exiting eddies of diminished strength.

Figs. 8 and 9 show differences between the reference solution for the advected eddy and the solutions using, respectively, the simple sponge and the pretty good sponge. The positive errors indicate an eddy with diminished amplitude approaching the corner of the region of interest; about 50% more undesired attenuation can

[7] With the mean basin depth taken to be 500 m, adjustment waves quickly left the region of interest. Because sponge zones for the large domain controlled reflections, any adjustment waves re-entering the domain were negligible when compared to the eddy.

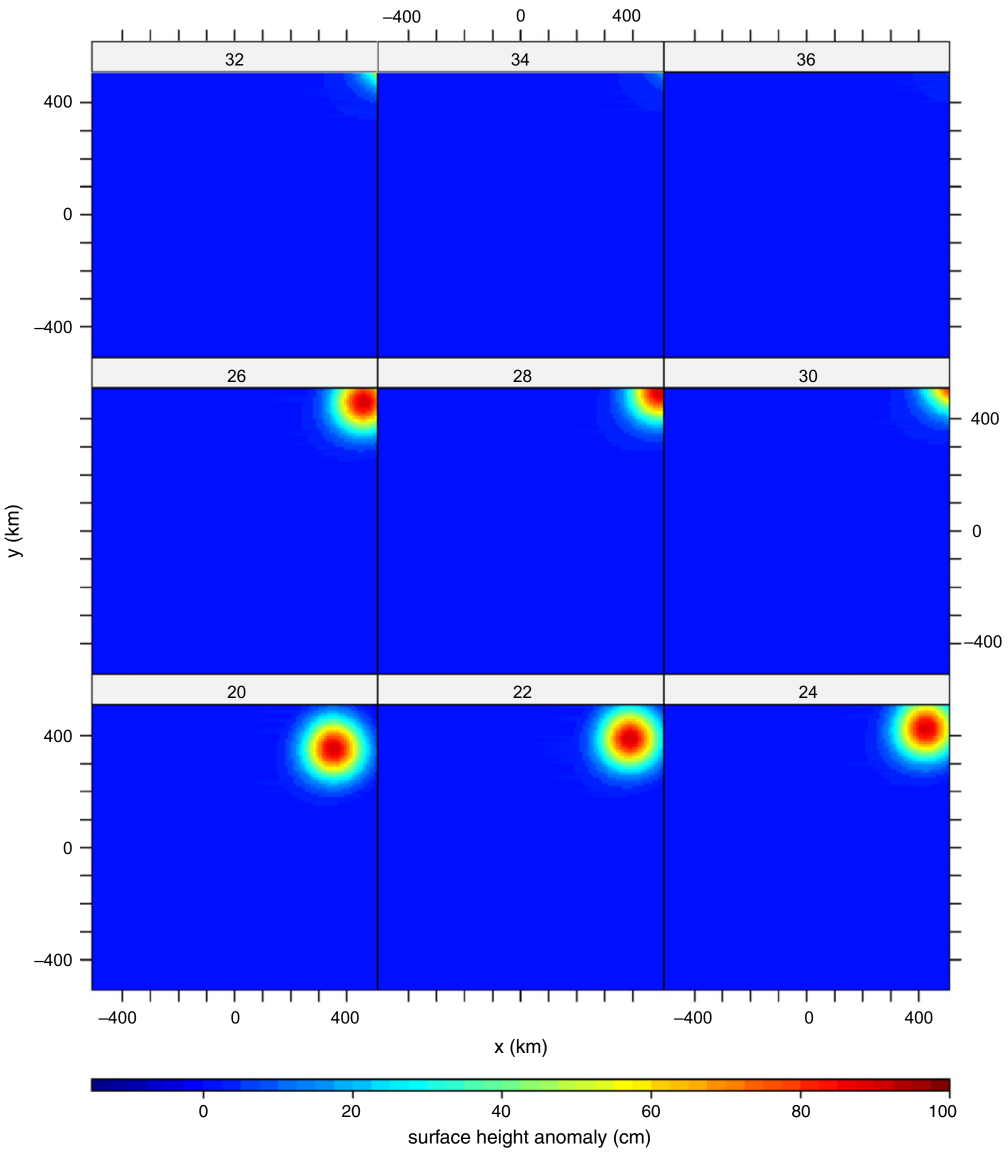

Fig. 7. Surface height anomaly (cm) for the reference solution at times (h) indicated by the panel labels. The geostrophic tilt associated with the uniform background current has been removed.

be seen for the simple sponge than for the pretty good sponge. Clearly, absorbing the eddy within the sponge zone is affecting the solution within the domain of interest. Nevertheless, even with no information about the nature of the eddy in the sponge zone, results are not too bad; reflected waves are not a problem, and after the eddy leaves the solution returns to the desired uniform flow.

In some circumstances information about the nature of the solution outside the region of interest is available from a larger-scale model. Here, the reference solution on the larger grid can provide such information.[8] Using it as the external solution should prevent the undesired attenuation of the interior solution, as only deviations from the sponge zone's advecting eddy would be absorbed.

[8] If the reference solutions were sampled at coarser resolution and interpolated, the computations would resemble those for grid nesting.

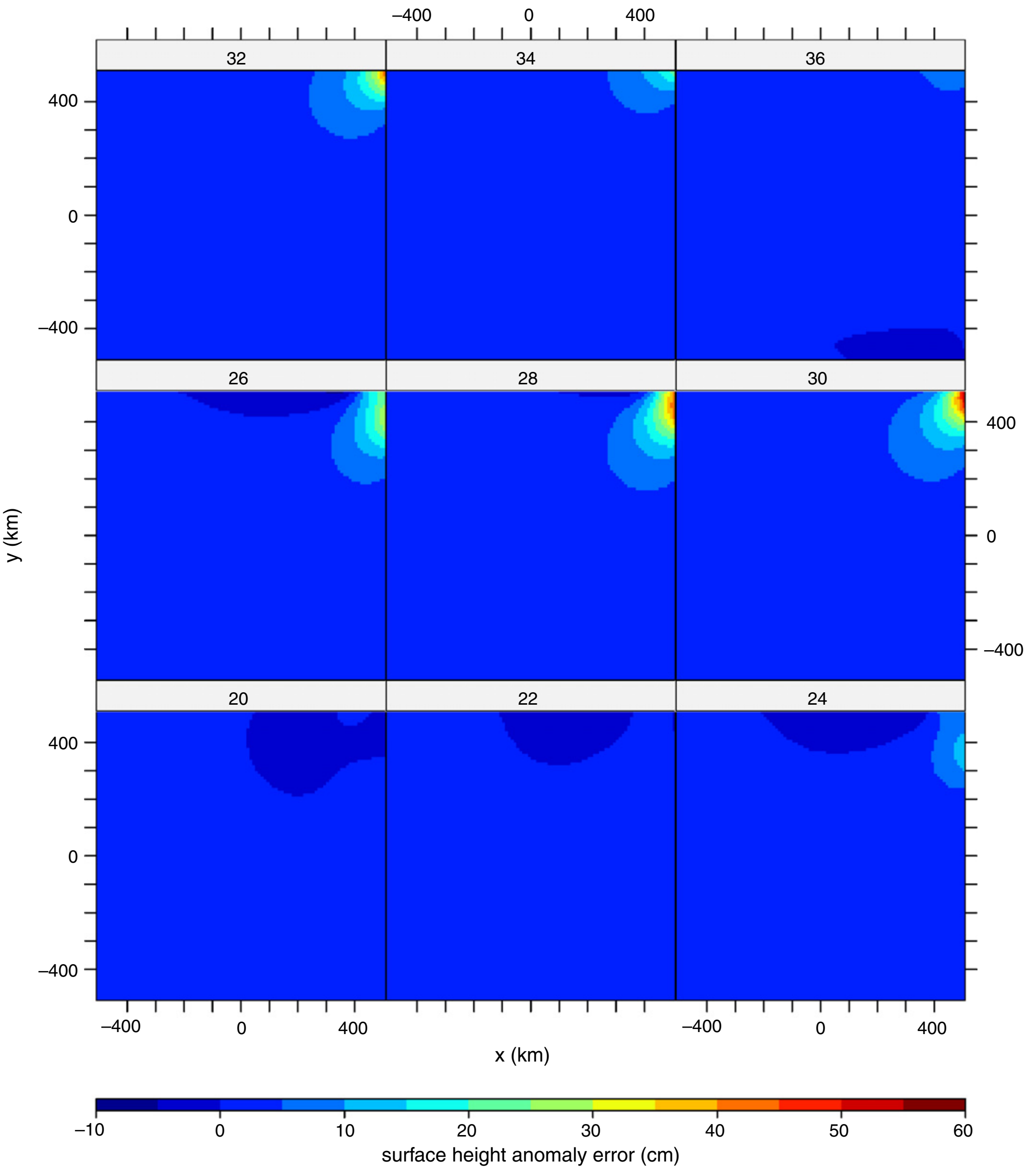

Fig. 8. Error (cm) of the surface height anomaly for the advected eddy when using the simple sponge.

Not only should improved information provide a better exit through the open boundary, information about an incoming eddy should allow it to enter. McDonald (2002) considers such a case, using boundary conditions to provide information about the incoming eddy. Here the parallel example, with absorbing layers, shows the value of a good external solution within the sponge zone in introducing an eddy into the interior region. After reflecting on the mechanism to get the eddy in, it is easier to understand the value of an external solution in the sponge region as the eddy exits.

Fig. 10 compares the solution obtained using the pretty good sponge with the reference solution from the extended grid for an eddy entering and then leaving the open domain. The computational set up is essentially the same as for the previous example, except that the center of the eddy is initially located 1358 km southwest of the center of the open domain. As the eddy approaches the region of interest, the external solution provides

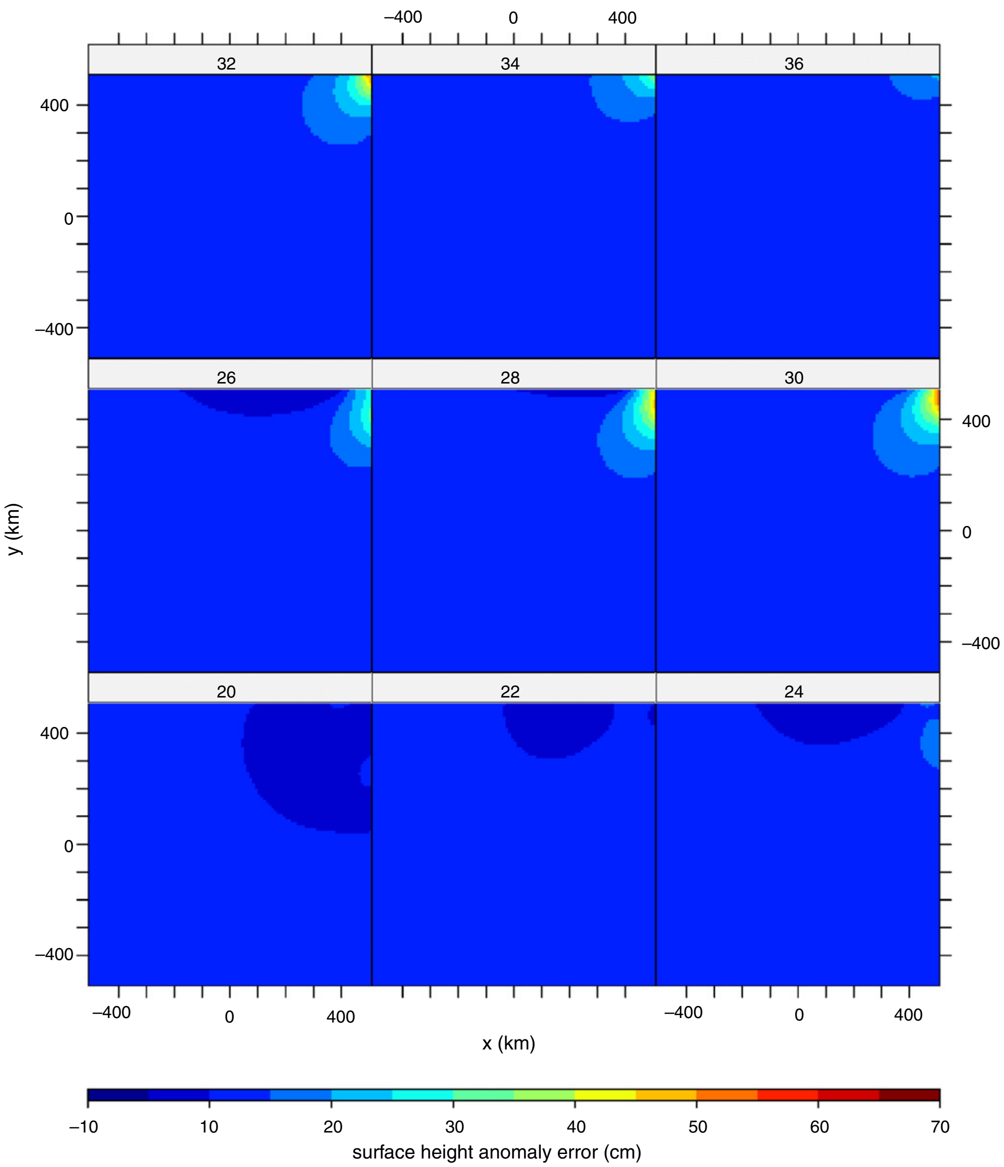

Fig. 9. Error (cm) of the surface height anomaly for the advected eddy when using the pretty good sponge.

the pretty good sponge information about its arrival, and later, after it progresses through the domain, about its departure. The lower-left panel shows the eddy entering, and subsequent panels show its progress at 8-h intervals. Colors indicate the reference solution, while the line contours at 10-cm intervals indicate the solution using the pretty good sponge. The line and color contours are nearly congruent, indicating that the pretty good sponge effectively passes information about the eddy as it enters and leaves as well as information about the advecting flow that carries it through the domain.

Fig. 11 gives a better view of the differences, showing them to be roughly three orders of magnitude smaller than the modeled feature. Largest errors can be seen along the southern boundary as the eddy enters; the surface elevation far from the eddy's center is too large, so errors are negative, indicating that the external solution within the sponge zone does not initiate a sufficiently strong outward flow. Eight hours later, after

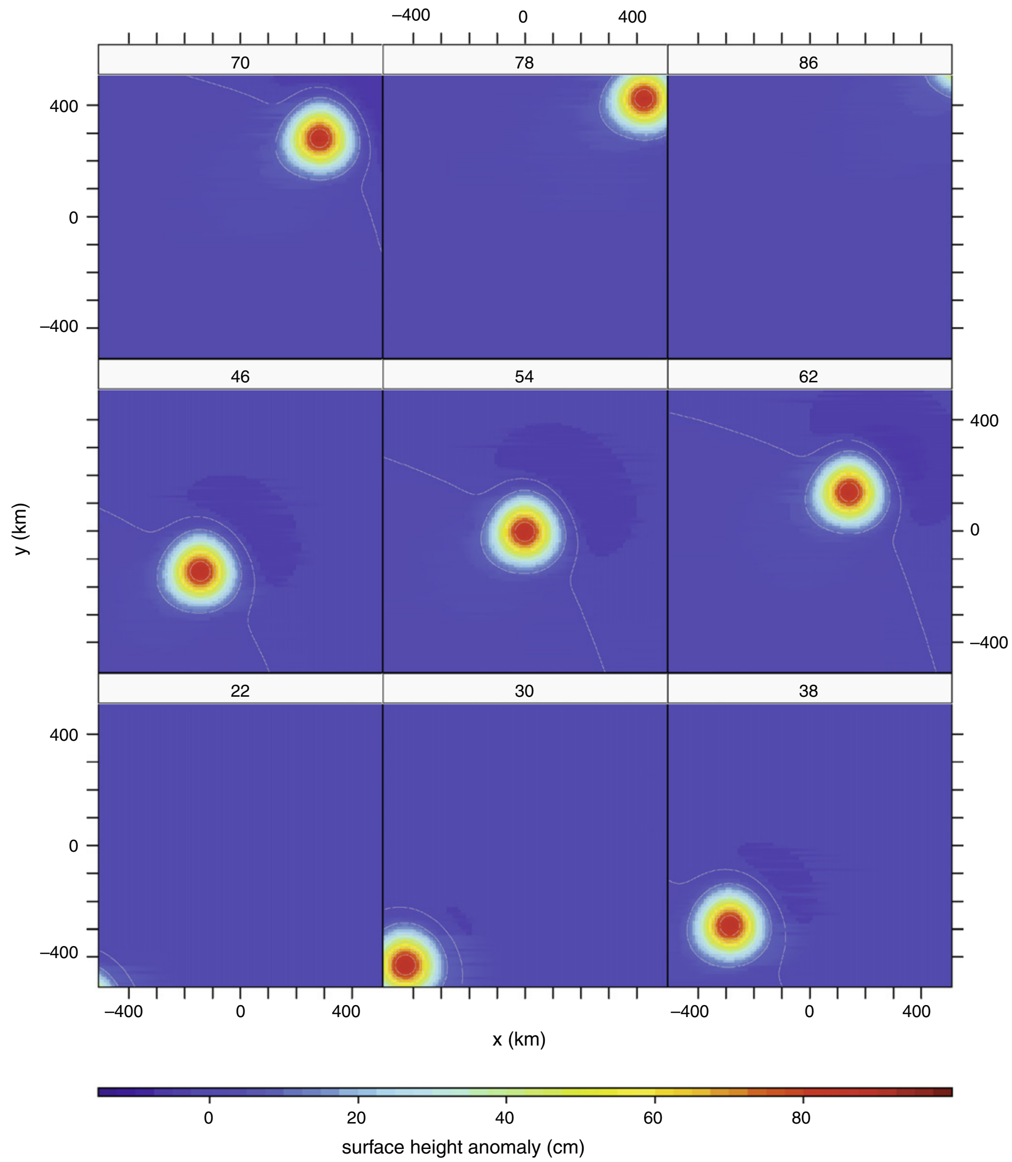

Fig. 10. Color indicates the reference solution for the entering/exiting eddy's surface height anomaly and line contours with 10 cm spacing indicate the corresponding results obtained using the pretty good sponge. Panel labels indicate the time (h) after initialization.

the eddy has almost completed its entrance, the sign changes, indicating the far-field elevation along the southern boundary is now too small. After another eight hours, the sign changes again and the magnitudes decrease. Errors continue to decrease in magnitude as the eddy crosses the domain and appear to be a pattern attached to the advecting eddy.

When the same example was repeated using the simple sponge, results were similar but better, especially during and following the eddy's entrance into the open region, as can be seen in Fig. 12. It appears that providing information about the component of velocity parallel to the boundary helps to get the eddy into the open domain. This case suggests that the simple sponge might be better suited for nested-grid applications where high-quality external information is available. On the other hand, for applications where a minimal amount of external information exists, the previous cases suggest that the pretty good sponge might be better.

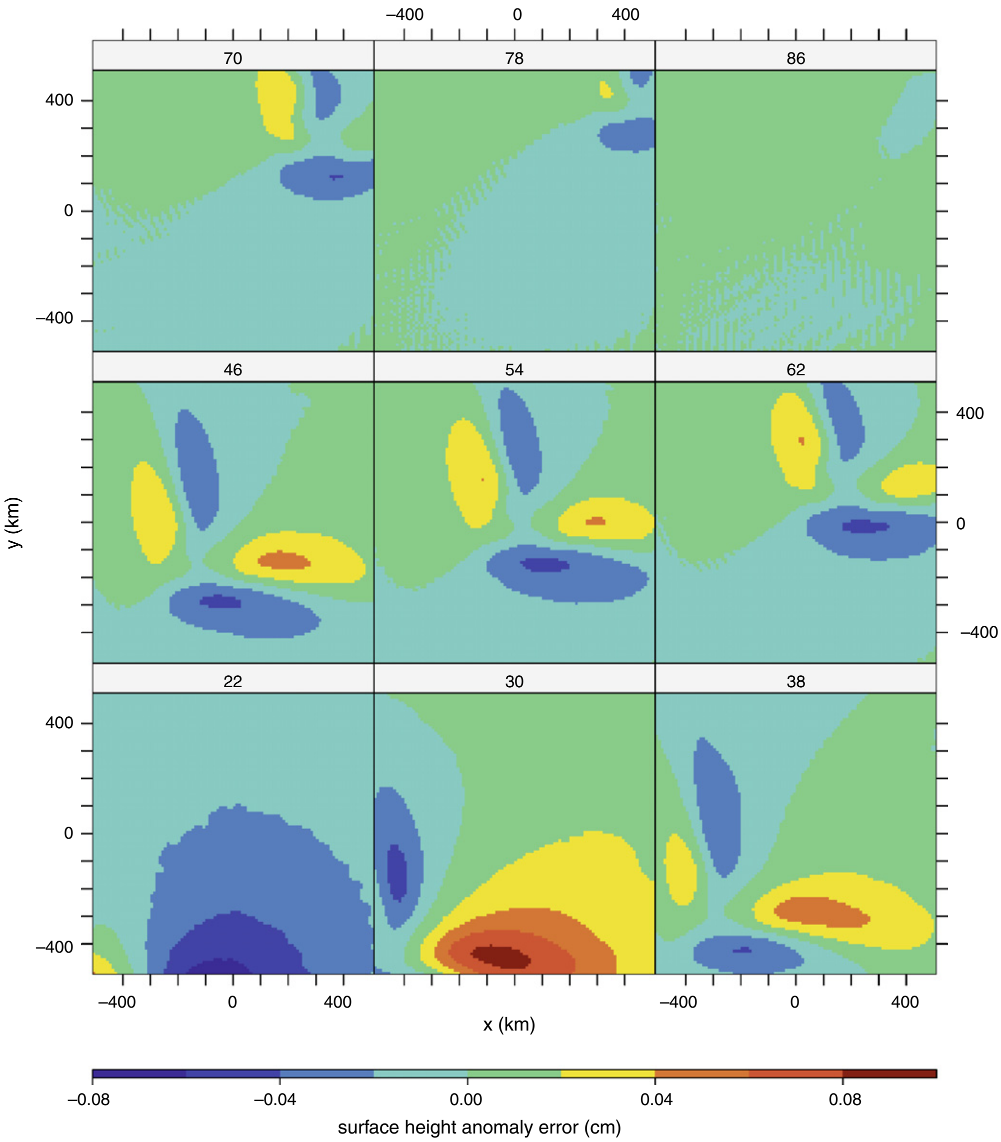

Fig. 11. Error (cm) of surface height anomaly for the pretty good sponge solution for the entering and exiting eddy.

4. Equatorial soliton

While the example of the advected eddy could have been treated with linear equations, an equatorial soliton provides an intrinsically nonlinear example. Haidvogel and Beckmann (1999) used it to examine the fidelity of numerical results in several of their ocean models, and Marchesiello et al. (2001) to test their open boundary conditions. As the soliton preserves its shape through a balance between nonlinearity and dispersion, it is interesting to see to what extent reducing its amplitude within the absorption zone destroys its shape. Just as with the propagating eddy, there is the issue of what external solution should be specified in the sponge zones. Damping differences from zero can be anticipated to reduce the soliton's amplitude as it passes through, possibly with negative consequences within the region of interest.

Using a perturbation approach Boyd (1980, 1985) sequentially provided the zeroth and first-order analytic expressions for a soliton on a beta plane. His lowest order solution is used here to initialize η, u, and v for a westward propagating equatorial soliton. It assumes the equivalent barotropic depth H to be uniform, it

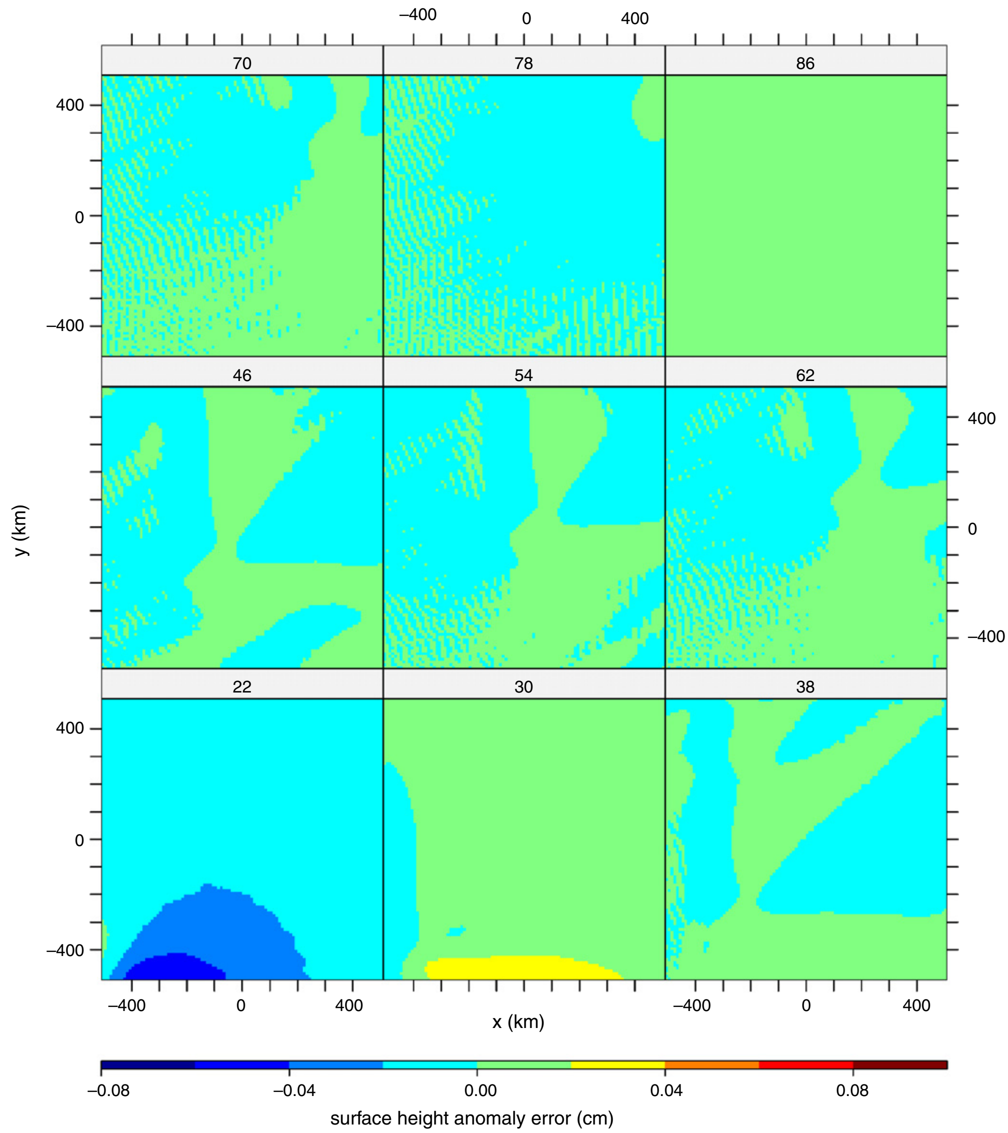

Fig. 12. Error (cm) of surface height anomaly for the simple sponge solution for the entering and exiting eddy.

depends on the earth's radius a and rate of rotation Ω through Lamb's non-dimensional parameter $E = 4\Omega^2 a^2 / gH$, and its amplitude is specified by the non-dimensional parameters B and $A = 0.772B^2$:

$$\eta = AH\text{sech}^2\left(\frac{B(x - x_0)}{L}\right)\frac{6y^2 + 3L^2}{4L^2}\exp\left(-\frac{y^2}{2L^2}\right), \tag{21}$$

$$u = A\sqrt{gH}\text{sech}^2\left(\frac{B(x - x_0)}{L}\right)\frac{6y^2 - 9L^2}{4L^2}\exp\left(-\frac{y^2}{2L^2}\right), \tag{22}$$

$$v = -2AB\sqrt{gH}\tanh\left(\frac{B(x - x_0)}{L}\right)\text{sech}^2\left(\frac{B(x - x_0)}{L}\right)\frac{2y}{L}\exp\left(-\frac{y^2}{2L^2}\right), \tag{23}$$

where the length scale $L = a/\sqrt[4]{E}$ is set by the radius of the earth, and where $y = 0$ corresponds to the equator.

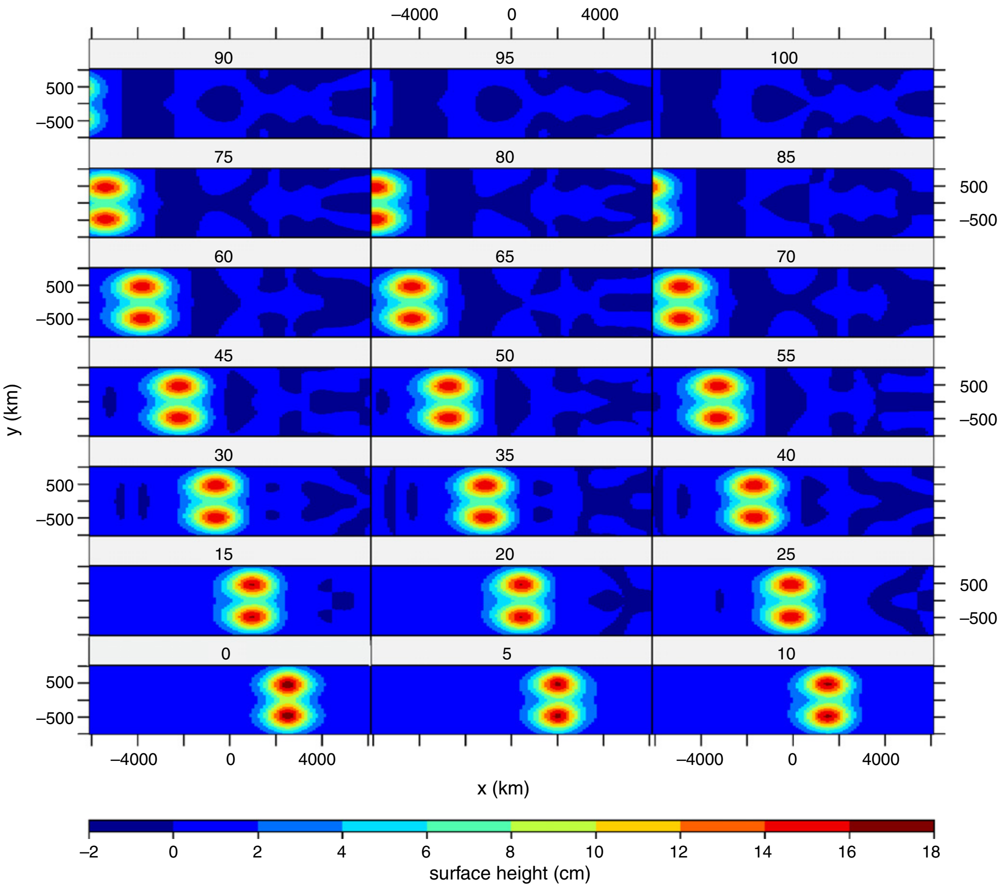

Fig. 13. Reference surface height (cm) for the equatorial soliton. Panel labels indicate the number of days following the initial time.

The domain of interest for the soliton computations spanned 12,240 km in longitude and 2040 km in latitude and was centered on the equator, with the initial center of the soliton at $x_0 = 2520$ km. The absorbing zones extended 1560 km in the $\pm x$ directions and 260 km in the $\pm y$ directions. The reference solution was computed in a 36,960 km by 4080 km domain surrounded by sponge zones with widths of 4560 km and 520 km, respectively. The equivalent barotropic depth H was set at 1 m. With $B = 0.394$, the initial soliton amplitude was 16.9 cm, giving a maximum westward current of 0.92 m s^{-1}. This configuration produced a soliton with non-negligible amplitude within the northern and southern sponge zones, mimicking the condition that Marchesiello et al. (2001) imposed on their test problem.

Fig. 13 shows the reference solution for the surface elevation, sampled every five days, as the soliton propagates to the west and out of the domain of interest. As the numerical soliton adjusts to its zero-order analytic initialization, it spawns a small wake that moves eastward out of the domain, leaving a slightly reduced amplitude that can be seen when examining results for day five.[9] As time progresses, the soliton can be seen to maintain its shape and speed. Low-amplitude variability unrelated to the soliton can also be seen in the reference solution.

Plots for the solutions using either the simple sponge or the pretty good sponge would look very much like those for the reference solutions, so they are not shown. Instead Fig. 14 presents the differences of these solution from the reference solution. The upper set of panels show errors for the pretty good sponge while the lower set show those for the simple sponge. The simple sponge solution's red hues near the northern and

[9] A similar adjustment was indicated by Haidvogel and Beckmann (1999) and by Marchesiello et al. (2001).

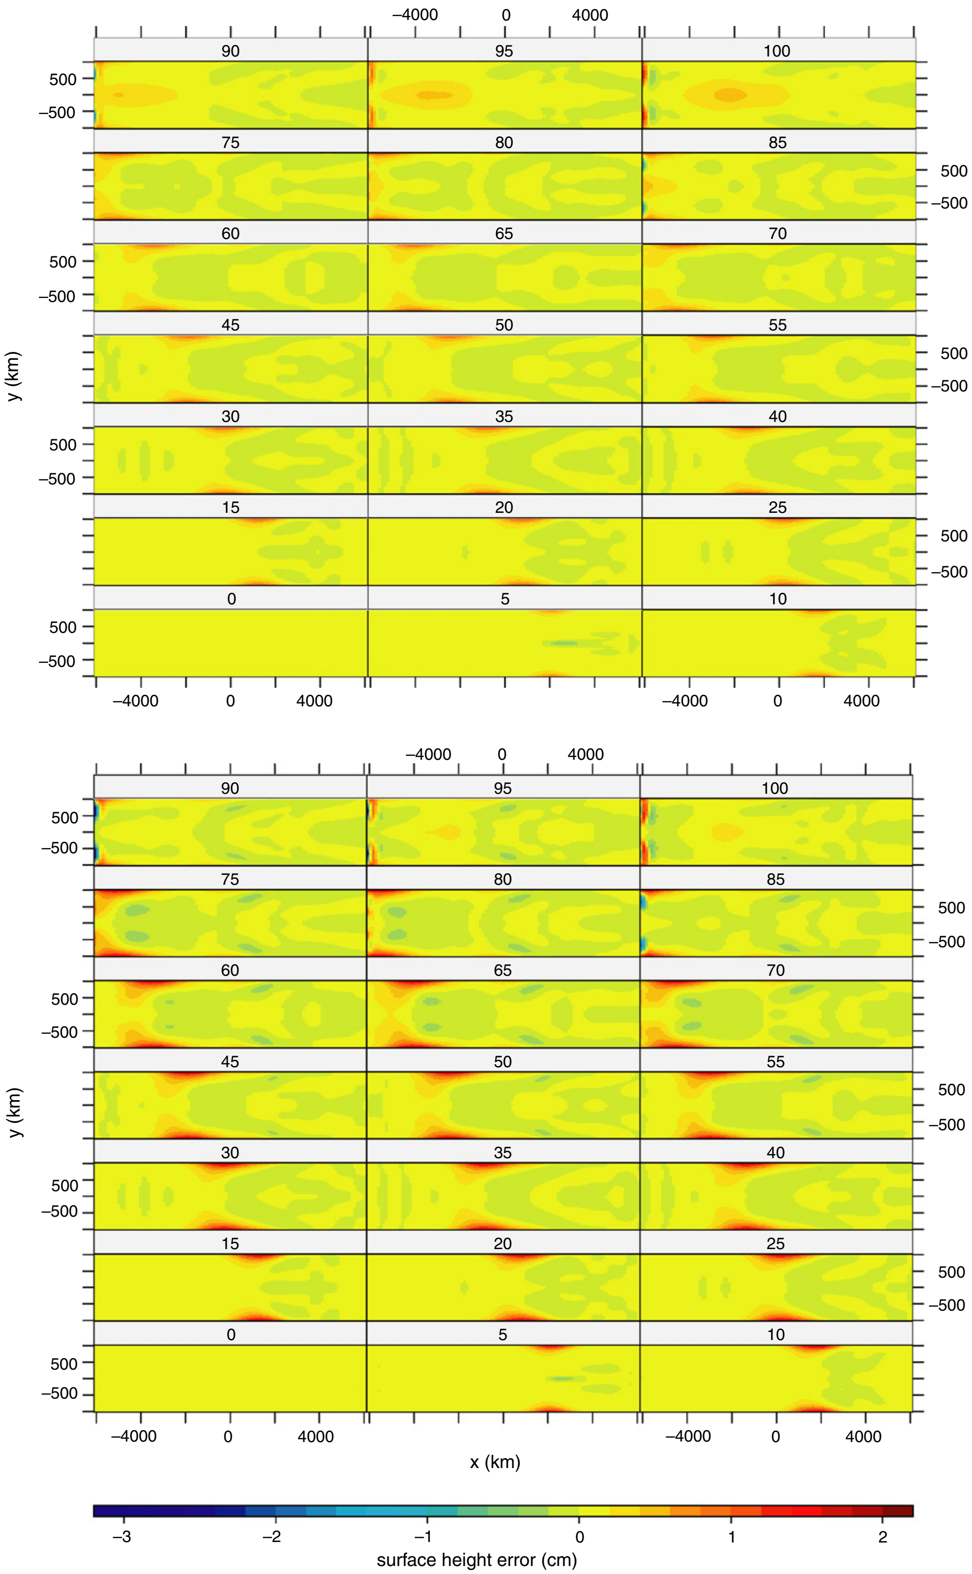

Fig. 14. Surface height errors (cm) for the pretty good sponge (upper panels) and for simple sponge (lower panels) solution for the equatorial soliton.

southern boundaries indicate that attenuation of the soliton's amplitude within the northern and southern absorbing zones reduce the soliton's amplitude within the domain of interest. Similar errors cannot be seen for the pretty good sponge. Blue hues at the western boundary as the soliton exits suggests that the speed of the soliton for the simple sponge solution is somewhat slower than for the reference solution. This could be due to the reduction in amplitude caused by the northern and southern sponge zones or caused by the western sponge zone. Similar but smaller blue patches appear in the panels for the pretty good sponge solution. The only situation where the errors for the pretty good sponge are noticeably larger than those for the simple sponge occurs after the soliton has exited; a patch of small positive error can be seen propagating from the western boundary along the equator in the panels for days 90–100.

Motivated by the extra cost associated with computing within the absorbing layers, an anonymous reviewer requested a comparison with at least one radiation condition. Because it is relatively easily implemented, we chose the popular Flather (1976) condition. Anticipating different accuracy would complicate a comparison of cost, we chose a configuration where computational costs would be similar so that accuracy could be compared; this was done by enforcing the radiation conditions at the outer boundaries of the (now absorption-free) sponge zones. Results using the Flather boundary conditions showed that the eastward moving waves from the initial adjustment reflected from the eastern boundary causing growing negative amplitudes, which became so large that it became impossible to simulate the soliton leaving the domain. While some other radiation/characteristic condition might provide stable and better results than sponge methods at comparable computational cost or comparable results at lower cost, it is worth noting that Marchesiello et al. (2001) had to supplement their radiation conditions with nudging and sponge layers to get acceptable solutions.

5. Conclusion

The comparisons of open-boundary techniques for simulating the collapse of an initially stationary mound of water show that using the pretty good sponge gives comparable results to those obtained when perfectly matched layers are used, whereas the results for the simple sponge are noticeably worse. For the perfectly-matched-layer technique the external solution does not enter the formalism explicitly; instead, absorbing outgoing waves until their amplitudes vanish implies an external solution for which all fields are zero. For the other two techniques the external solution was explicit: the waves should exit into a flat, motionless external domain. Due to the earth's rotation, the collapsing mound should ultimately reach geostrophic balance after adjustment waves have radiated away, but because the adjusted solution is not necessarily zero in the absorption zones, absorption can continue even though it is not needed.

A better examination of the role of the external information was provided by the example of an eddy moving through the open boundary. Unfortunately, as the state of the art of perfectly matched layers has not yet reached the point where it can be used for problems where external information must be specified, only the performance of the pretty good and simple sponges could be compared. When the external solution provided information only about the advecting flow but not about the nature of the part of the eddy outside the open domain, the pretty good sponge performed better than the simple sponge. However, when information was provided about the external part of the eddy, while results from both methods benefited substantially, the simple sponge gave the better solution. Together, these results suggest that the simple sponge might be preferred for situations where very good external information is available, such as for nested-grid applications, whereas the pretty good sponge might be preferred when less is known about the external situation.

The equatorial soliton provided an intrinsically nonlinear case. Even with no external information about the soliton in the absorbing layers, the results were quite good, with those for the pretty good sponge being consistently better than those for the simple sponge. Much of its superiority can be attributed to *not* absorbing the component of momentum parallel to the open boundary. The simple sponge, which *does* absorb this component, has noticeable errors near the northern and southern boundaries as the soliton propagates toward the west.

The examples used here to examine the performance of the pretty good sponge had all been used previously to test other methods for handling open boundaries. The collapsing mound and the advected eddy were used by McDonald (2002) to test his radiation method and by Navon et al. (2004) to test their perfectly-matched-layer approach, while the equatorial soliton was used by Marchesiello et al. (2001) to test their radiation boundary condition approach. In every case the results obtained here using the pretty good sponge and simple boundary conditions agree favorably with the results shown in these papers.

Why is the performance of the pretty good sponge so good? The intent here has been only to explore its ability with a few examples, not to present a mathematical analysis of its properties, but a few comments in that direction may be appropriate. Recall that the prescription for the pretty good sponge was taken from that for the perfectly matched layer, which is based on the complex wave number transformation (4) and (5). The idea behind this transformation is to introduce spatial absorption while preserving the dispersion equation for waves of all frequencies. The analysis of perfectly matched layers in terms of complex wave numbers and real frequencies, which is intended to guarantee the absence of reflections, is at odds with the stability analyses, which deal with complex frequencies and real wave numbers; and the fact that instability has been an issue for perfectly matched layers indicates that deeper understanding is needed. Empirically, the instability can be controlled by introducing decay terms to control the growth the auxiliary variables introduced by the transformation. On the other hand, eliminating the auxiliary terms entirely leads to the pretty good sponge. The distinction between absorption at x and y boundaries is all that remains of the complex wave number transformation. For the shallow-water wave equations linearized about a state of no motion, this distinction manifests as a different treatment for u and v: the component of velocity parallel to the boundary is *not* absorbed. Why this is effective still deserves further study, but what is clear is that its origin is in the dispersion equation.

Because the complex wave number transformation makes explicit reference to the coordinate axes, both components of velocity are absorbed in the corner zones, albeit at different rates. Moreover, the largest errors for the collapsing mound were near these corners. This suggests that the perfectly matched layer might benefit from a coordinate-free formalism. For example, (4) and (5) might be replaced by a single expression suitable for a general curvilinear boundary with the normal direction at any point within the sponge zone being toward the closest point on the boundary.

While the success of the pretty good sponge for the collapsing mound example might be attributed to its origin in the perfectly matched layer, as its form was obtained from the linear equations used in that example, there is still the question of why it should work so well for the nonlinear examples. For those involving the advecting eddy, a plausible explanation might be that the solution within the sponge zone closely resembles the solution for equations linearized about the specified external solution, so that the dispersion equation for the disturbances that are damped is qualitatively like those of the linear example. However, this explanation is not convincing for the example of the equatorial soliton, where nonlinearity balances dispersion to maintain the shape and speed of the wave. No guidance is available from the community developing perfectly matched layers, as they have not yet addressed nonlinear equations nor the use of external information.

The examples with advected eddies clearly show that the external solution specified in the absorbing layers plays two roles. The first is to maintain the flow needed to advect the eddies and the second is to introduce information about there being an eddy within the advecting flow. The first is mandated by the use of nonlinear equations, which require boundary information to sustain the advection. If these examples had been addressed using linearized equations as did Navon et al. (2004) and McDonald (2002), the advection would be built in and boundary conditions would be focused on the second role. Specifying the advecting flow of this example is like specifying climatic conditions as the external solution in more general situations, bringing in the minimum information needed to have a meaningful solution. For the motivating problem of topographic scattering of tides, the tidal flow at the boundary sets the context, but the problem of allowing the waves to exit remains.

It should be emphasized that the external information is provided everywhere within all sponge zones. To see why this should work consider the simpler case of the advection equation, which captures the essence of wave propagation along characteristics. As the advecting flow brings the field ξ into the domain, the external solution ξ_e determines how much is brought in; on the other hand, when the flow is out of the open domain, ξ_e has no way to influence the interior. As the shallow-water wave equations are hyperbolic, they can be analyzed in terms of characteristic variables, each obeying an advection-like equation. Consequently, even though the characteristics might enter or exit the domain at any angle, along any given characteristic the associated variable should get its value when entering from the external solution and should be allowed to exit even though an external solution has been specified. The principal worry is that the specified external solution should supply the needed information and not inject misinformation. When the information is guaranteed to be excellent, like in the example with the entering/exiting eddy, it is not surprising that the simple sponge did better, as it exploited all of the information. What is less clear is why the pretty good sponge injected less misinformation when the external solution provided no information about the eddy.

This ability of absorbing-layer methods to use the provided external information directly can be contrasted with radiation methods' need to diagnose sites of inflow where external information should be used. That need originates from the necessity of providing boundary conditions locally. In contrast, sponges allow the external information to be introduced through the governing equations. This obviates the problems of mis-diagnosing inflow/outflow and of selecting which waves pass out most effectively. However, this advantage comes at the computational price associated with extending the domain beyond the primary region of interest. On the other hand, the cost of implementing absorbing-layer methods is quite low, whereas considerably more care is often needed to make radiation methods compatible with the ocean model's numerical scheme.

The conclusion to be drawn from these examples is that the pretty good sponge can be used to handle open boundaries in rather general circumstances. Not only does it allow outgoing disturbances to leave the domain, but it also provides a way to incorporate external information, even for nonlinear flows. As it is easy to implement, it can be incorporated into existing limited-area ocean models with very little effort, allowing its utility for a wide range of problems to be gauged.

Acknowledgements

This work was supported by the National Oceanic and Atmospheric Administration's (NOAA) Vents Program, by the Pacific Marine Environmental Laboratory (PMEL), and by the Atlantic Oceanographic and Meteorological Laboratory (AOML). Contribution 3095 from NOAA/Pacific Marine Environmental Laboratory.

References

Abarbanel, S., Gottlieb, D., 1997. A mathematical analysis of the PML method. J. Comput. Phys. 134, 357–363.

Abarbanel, S., Gottlieb, D., Hesthaven, J.S., 1999. Well-posed perfectly matched layers for advective acoustics. J. Comput. Phys. 154, 266–283.

Bécache, E., Fauqueux, S., Joly, P., 2003. Stability of perfectly matched layers, group velocities and anisotropic waves. J. Comput. Phys. 188, 399–433.

Berenger, J.-P., 1994. A perfectly matched layer for the absorption of electromagnetic waves. J. Comput. Phys. 114, 185–200.

Blayo, E., Debreu, L., 2005. Revisiting open boundary conditions from the point of view of characteristic variables. Ocean Modell. 9, 231–252.

Bodony, D.J., 2006. Analysis of sponge zones for computational fluid mechanics. J. Comput. Phys. 212, 681–702.

Boyd, J.P., 1980. Equatorial solitary waves. Part I: Rossby solitons. J. Phys. Oceanogr. 10, 1699–1717.

Boyd, J.P., 1985. Equatorial solitary waves. Part 3: Westward-traveling modons. J. Phys. Oceanogr. 15, 46–54.

Camerlengo, A.L., O'Brien, J.J., 1980. Open boundary conditions in rotating fluids. J. Comput. Phys. 35, 12–35.

Chassignet, E.P., Smith, L.T., Halliwell Jr., G.R., Bleck, R., 2003. North Atlantic simulations with the hybrid coordinate ocean model (HYCOM): impact of the vertical coordinate choice and resolution, reference density, and thermobaricity. J. Phys. Oceanogr. 33, 2504–2526.

Darblade, G., Baraille, R., Le Roux, A.-Y., Carton, X., Pinchon, D., 1997. Conditions limites non réflechissantes pour un modèle de Saint-Venant bidimensionnel barotrope linéarisé. C. R. Acad. Sci. Paris 324, 485–490.

Davies, H.C., 1976. A lateral boundary formulation for multi-level prediction models. Quart. J. Roy. Meteor. Soc. 102, 405–418.

Engquist, B., Majda, A., 1977. Absorbing boundary conditions for the numerical simulation of waves. Math. Comp. 31, 629–651.

Flather, B.A., 1976. A tidal model of the northwest European continental shelf. Mem. Soc. R. Sci. Liège 6 (10), 141–164.

Givoli, D., Neta, B., 2003. High-order nonreflecting boundary conditions for the dispersive shallow water equations. J. Comput. Appl. Math. 158, 49–60.

Haidvogel, D.B., Beckmann, A., 1999. Numerical Ocean Circulation Modeling. Imperial College Press.

Hesthaven, J.S., 1998. On the analysis and construction of perfectly matched layers for the linearized Euler equations. J. Comput. Phys. 142, 129–147.

Higdon, R.L., 1994. Radiation boundary conditions for dispersive waves. SIAM J. Numer. Anal. 31, 64–100.

Hu, F.Q., 1996. On absorbing boundary conditions for linearized Euler equations by a perfectly matched layer. J. Comput. Phys. 129, 201–219.

Hu, F.Q., 2001. A stable, perfectly matched layer for linearized Euler equations in unsplit physical variables. J. Comput. Phys. 173, 455–480.

Hu, F.Q., in press. Development of PML absorbing boundary conditions for computational aeroacoustics: a progress review. Computers and Fluids. doi:org/10.1016/j.compfluid.2007.02.012.

Lavelle, J.W., 2006. Flow, hydrography, turbulent mixing, and dissipation at Fieberling Guyot examined with a primitive equation model. J. Geophys. Res. 111, C07014. doi:10.1029/2005JC003224.

Lie, I., 2001. Well-posed transparent boundary conditions for the shallow water equations. Appl. Numer. Math. 38, 445–474.

Marchesiello, P., McWilliams, J.C., Shchepetkin, A., 2001. Open boundary conditions for long-term integration of regional oceanic models. Ocean Modell. 3, 1–20.

Martinsen, E.A., Engedahl, H., 1987. Implementation and testing of a lateral boundary scheme as an open boundary condition in a barotropic ocean model. Coastal Eng. 11, 603–627.

McDonald, A., 2002. A step toward transparent boundary conditions for meteorological models. Mon. Wea. Rev. 130, 140–151.

Navon, I.M., Neta, B., Hussaini, M.Y., 2004. A perfectly matched layer approach to the linearized shallow water equations models. Mon. Wea. Rev. 132, 1369–1378.

Nycander, J., Döös, K., 2003. Open boundary conditions for barotropic waves. J. Geophys. Res. 108 (C5). doi:10.1029/2002JC001529.

Orlanski, I., 1976. A simple boundary condition for unbounded hyperbolic flows. J. Comput. Phys. 21, 251–269.

Petropoulos, P.G., 1998. On the termination of the perfectly matched layer with local absorbing boundary condition. J. Comput. Phys. 143, 665–673.

Philander, S.G.H., Hurlin, W.J., Seigel, A.D., 1987. Simulation of the seasonal cycle of the tropical Pacific Ocean. J. Phys. Oceanogr. 17, 1986–2002.

Raymond, W.H., Kuo, H.L., 1984. A radiation boundary condition for multi-dimensional flows. Quart. J. Roy. Meteor. Soc. 110, 535–551.

Tsynkov, S.V., 1998. Numerical solution of problems on unbounded domains. A review. Appl. Numer. Math. 27, 465–532.

Turkel, E., Yefet, A., 1998. Absorbing PML boundary layers for wave-like equations. Appl. Numer. Math. 27, 533–557.

van Joolen, V., Neta, B., Givoli, D., 2004. High-order nonreflecting boundary conditions. Comput. Math. Appl. 48, 1167–1180.