INTERNATIONAL JOURNAL FOR NUMERICAL METHODS IN FLUIDS
Int. J. Numer. Meth. Fluids 2005; **49**:301–329

Reducing numerical diffusion in interfacial gravity wave simulations

O. B. Fringer[1,*,†], S. W. Armfield[2,‡] and R. L. Street[1,§]

[1]*Environmental Fluid Mechanics Laboratory, Stanford University, Stanford CA 94305-4020, U.S.A.*
[2]*Aerospace, Mechanical, and Mechatronic Engineering, The University of Sydney, Sydney, NSW 2006, Australia*

SUMMARY

We demonstrate how the background potential energy is an excellent measure of the effective numerical diffusion or antidiffusion of an advection scheme by applying several advection schemes to a standing interfacial gravity wave. All existing advection schemes do not maintain the background potential energy because they are either diffusive, antidiffusive, or oscillatory. By taking advantage of the compressive nature of some schemes, which causes a decrease in the background potential energy, and the diffusive nature of others, which causes an increase in the background potential energy, we develop two background potential energy preserving advection schemes that are well-suited to study interfacial gravity waves at a density interface between two miscible fluids in closed domains such as lakes. The schemes employ total variation diminishing limiters and universal limiters in which the limiter is a function of both the upwind and local gradients as well as the background potential energy. The effectiveness of the schemes is validated by computing a sloshing interfacial gravity wave with a nonstaggered-grid Boussinesq solver, in which QUICK is employed for momentum and the pressure correction method is used, which is second-order accurate in time. For scalar advection, the present background potential energy preserving schemes are employed and compared to other TVD and non-TVD schemes, and we demonstrate that the schemes can control the change in the background potential energy due to numerical effects. Copyright © 2005 John Wiley & Sons, Ltd.

KEY WORDS: interfacial gravity wave; miscible; advection; universal limiter; TVD; Boussinesq

---

[*] Correspondence to: O. B. Fringer, Department of Civil and Environmental Engineering, Stanford University, Stanford CA 94305-4020, U.S.A.
[†] E-mail: fringer@stanford.edu
[‡] E-mail: armfield@aeromech.usyd.edu.au
[§] E-mail: street@stanford.edu

Contract/grant sponsor: NSF; contract/grant numbers: ITR-0113111; ATM-0073395
Contract/grant sponsor: University of the Western Cape
Contract/grant sponsor: ONR; contract/grant number: N0014-99-1-0413

Received 16 September 2002
Revised 9 November 2004
Accepted 12 April 2005

1. INTRODUCTION

A host of literature exists on the study of surface gravity waves or gravity waves at an interface separating immiscible fluids. Free-surface simulations that do not involve breaking map the domain onto one that follows the surface with the use of a ζ-coordinate system, such as the large-eddy simulation of nonlinear free-surface waves of Hodges and Street [1]. These simulations do not allow overturning because of the monotonicity requirement of the curvilinear grid. Breaking interfacial gravity waves are more appropriately simulated with methods such as the marker and cell technique developed by Harlow and Welch [2], or the interface tracking method of Puckett *et al.* [3]. Chen *et al.* [4] implement the volume of fluid (VOF) method to simulate breaking free-surface gravity waves, and also simulate interfacial gravity waves, but the interface is immiscible and infinitesimal in thickness. Other methods, such as the particle level set method [5], yield excellent results for surface or interfacial gravity wave problems, but they too do not allow for a finite-thickness miscible interface.

Gravity waves at an interface between miscible fluids can be simulated with any hydrodynamic equation solver that computes the evolution of the density field. However, if the advection scheme that is used to compute the evolution of the density field is to capture the interfacial gravity wave dynamics realistically, it cannot generate excessive diffusion at the interface, nor can it impose antidiffusion and generate unphysical overshoots. The problem with advection schemes, however, is that they can be either diffusive or antidiffusive in the presence of sharp fronts.

The numerical diffusion or antidiffusion of the density field resulting from the numerical method can be quantified with the evolution of the background potential energy. According to Winters *et al.* [6], the potential energy can be split into its available and background components with

$$E_{\mathrm{p}} = E_{\mathrm{a}} + E_{\mathrm{b}} \tag{1}$$

The available potential energy, E_{a}, represents the potential energy of a system that is available to be converted into motion. The background potential energy, E_{b}, represents the potential energy of a system in its background state, that is, in a state in which no more available potential energy exists that can be converted into kinetic energy. As an example, consider the sloshing of a standing interfacial gravity wave in a laboratory tank that contains a layer of fresh water (blue) over a layer of heavier salty water (red), as shown in Figure 1(a). As the wave oscillates, some of its energy is lost to viscosity, while at the same time the interface between the two layers thickens as the salt diffuses into the fresh water. Eventually, the wave comes to rest due to viscosity, and the thickness of the diffuse layer is greater than it was when it started, as shown in Figure 1(b). If the same situation is repeated, but this time with a higher diffusivity of salt in water, then the wave comes to rest in the same manner, but in this case the interface is thicker due to the higher diffusivity, as shown in Figure 1(c). Now consider the case in which the wave oscillates in the *absence* of salt diffusivity. In this case, the wave comes to rest due to viscosity and the thickness of the interface does not change, as shown in Figure 1(d).

The potential energy of the density fields shown in Figures 1(b)–1(d) is given by

$$E_{\mathrm{p}} = g \int_V \rho z \, \mathrm{d}V \tag{2}$$

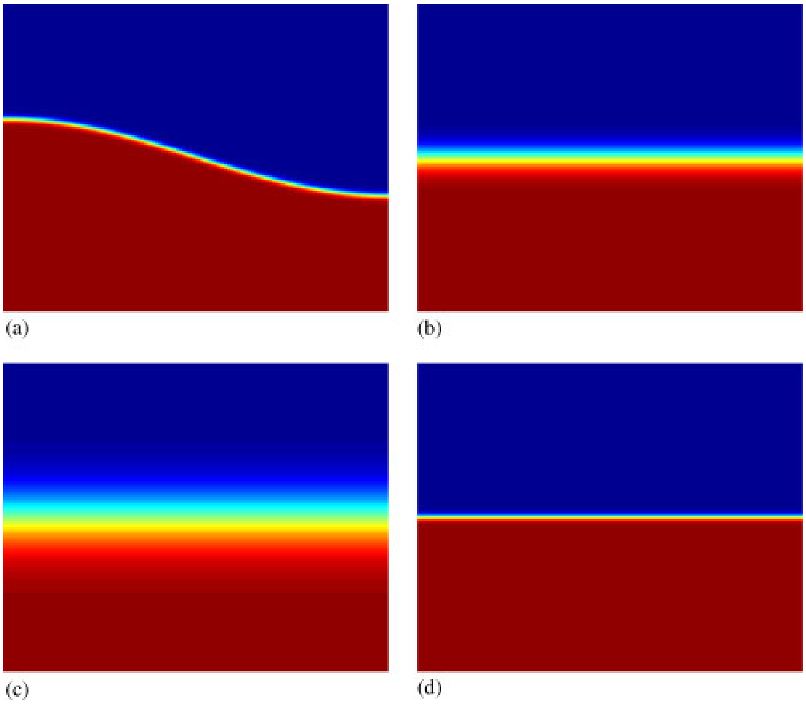

Figure 1. Depiction of a standing interfacial gravity wave when it begins (a), and after it has come to rest due to viscosity (b), (c), and (d). Because the diffusivity is much higher in case (c), the interface thickens substantially more than it does in case (b). Case (d) results when the diffusivity is zero.

where g is the gravitational body force per unit mass, ρ is the density of the fluid, and z is the height above some datum. Because the density fields in Figures 1(b)–1(d) do not have any horizontal variation, there is no available potential energy that can be converted into kinetic energy, so E_a vanishes and $E_\mathrm{p}=E_\mathrm{b}$. Since $\rho=\rho(z)$, the potential energy reduces to

$$E_\mathrm{p}=E_\mathrm{b}=gLW\int_{-d}^{0}\rho(z)z\,\mathrm{d}z \tag{3}$$

where L, W, and d are the length, width, and depth of the domain shown in Figure 1. By writing the centre of mass of each density field as

$$z_\mathrm{c}=\frac{LW}{M}\int_{-d}^{0}\rho(z)z\,\mathrm{d}z \tag{4}$$

where the mass is given by

$$M=LW\int_{-d}^{0}\rho(z)\,\mathrm{d}z \tag{5}$$

we can write Equation (3) as

$$E_b = Mgz_c \tag{6}$$

This is the background potential energy of the fluids at rest. The density fields with the thicker interfaces have a higher centre of gravity, and therefore they have a larger background potential energy. The density field with the lowest background potential energy is that shown in Figure 1(d). This density field was created in the absence of salt diffusivity, and hence the density field in Figure 1(d) represents the background density field associated with the density field shown in Figure 1(a). At any point in time during the history of oscillation of the interfacial gravity wave shown in Figure 1(a), its background potential energy distribution will be given by that shown in Figure 1(d) in the absence of salt diffusivity.

From this we can define the background density field as the density field that results if at some instant in time the thermal and salt diffusivities in a flow field vanished but the viscosity remained finite, and the flow was allowed to come to rest in some statically stable state. Winters *et al.* [6] define this background potential energy as

$$E_b = g \int_V \rho z^*(\mathbf{x}, t)\, dV \tag{7}$$

where $z^*(\mathbf{x}, t)$ is the height of a fluid parcel with density ρ in its background state, as shown, for example, in Figure 1(d). In Figures 1(b) and 1(c), the background potential energy is larger as a result of diffusion at the interface. Diffusion causes the interface to thicken, and hence causes the background potential energy to increase. In a closed domain and in the absence of heat or mass sources, if there is no thermal or salt diffusivity then the background potential energy must remain constant. In the absence of physical diffusion, if an advection scheme that is applied to the transport equation for density alters the background potential energy, then it is imposing unphysical numerical diffusion or antidiffusion on the density field. An increase in the background potential energy implies numerical diffusion, while a decrease in the background potential energy implies either numerical antidiffusion or the development of nonmonotonicity.

Unlike mass conservation, it is impossible for a numerical scheme to maintain the background potential energy of a flowfield from one time step to the next unless specific countermeasures are employed. Numerical methods that are not monotonic cause a decrease in the background potential energy, while methods that are monotonic either cause an increase in the background potential energy if they must diffuse the interface in order to maintain monotonicity, or cause a decrease in the background potential energy if they are overcompressive. As an example, in their simulations of internal gravity waves in lakes, Laval *et al.* [7] use the ULTIMATE-QUICKEST scheme [8] because they found it to be the most well-behaved for field-scale environmental flow simulations where sharp grid-scale density fronts are commonplace. But due to the presence of grid-scale density fronts, ULTIMATE-QUICKEST is forced to employ numerical diffusion to maintain monotonicity and thus causes an increase in the background potential energy. Because they use ELCOM [9], which employs operator splitting and treats the advection and diffusion steps separately, Laval *et al.* quantify the global diffusion incurred by the advection scheme by computing the background potential energy before and after the advection step. Using this unphysical change as a base (since the change should be identically zero), they employ a filter that effectively imposes antidiffusion on the density field to remove the effects of the numerical diffusion on the interface and bring the background potential energy back down to its correct value. The filter effectively adds a step to the advection process that removes the diffusive character of the ULTIMATE-QUICKEST scheme.

In this paper we employ several commonly used advection schemes to compute the standing interfacial gravity wave problem, and discuss the benefits and drawbacks of each method with regard to maintaining the background potential energy of the flow. Then, we develop a similar technique to that of Laval *et al.* [7], in which we use the background potential energy as a measure of the numerical diffusion. However, instead of using a filter, we employ a limiter that is a function of the background potential energy as well as the local gradients. In this manner, both monotonicity and the background potential energy can be maintained automatically without employing a filter.

2. GOVERNING EQUATIONS AND NUMERICAL METHOD

2.1. Discretization

We solve the two-dimensional Euler equations with the Boussinesq approximation,

$$\frac{\partial u_i}{\partial t} + \frac{\partial}{\partial x}(u_i u_j) = -\frac{1}{\rho_0}\frac{\partial p}{\partial x_i} + \frac{g}{\rho_0}(\rho - \rho_\mathrm{r})\delta_{i3} \tag{8}$$

subject to the continuity constraint

$$\frac{\partial u_i}{\partial x_i} = 0 \tag{9}$$

where $j = 1, 3$, u_i are the velocity components in the $i = 1, 3$ directions and the Einstein summation convention is assumed, g is the gravitational body force per unit mass, and ρ_0 is the reference density. The transport equation for density in a thermally and salt stratified environment is derived by assuming an equation of state of the form $\rho = \rho(s, T, p)$, where s is the salinity and T is the temperature, and using mass transport to obtain an evolution equation for salinity and the heat equation to obtain an equation for temperature. In this work we neglect the effects of pressure on the density and assume that the fluid is salt stratified. Using the transport equation of salt of the form

$$\frac{\partial s}{\partial t} + \frac{\partial}{\partial x_j}(s u_j) = 0 \tag{10}$$

and employing a linear equation of state for density such that

$$\frac{\rho - \rho_0}{\rho_0} = \beta(s - s_0) \tag{11}$$

where $\beta = 7 \times 10^{-4}\,\mathrm{psu}^{-1}$ is the saline expansion coefficient of water and s_0 is some constant reference salinity, the transport equation for density becomes

$$\frac{\partial \rho}{\partial t} + \frac{\partial}{\partial x_j}(\rho u_j) = 0 \tag{12}$$

It is assumed that the pressure field represents a departure from some arbitrary hydrostatic reference state. If p_{T} represents the total pressure, then it is related to the pressure in Equation (8) via

$$p_{\mathrm{T}}(x,y,z,t)=p(x,y,z,t)+p_{\mathrm{r}}(z) \tag{13}$$

where $p_{\mathrm{r}}(z)$ is the reference pressure field and is related to the reference density field $\rho_{\mathrm{r}}(z)$ by

$$\frac{\partial p_{\mathrm{r}}(z)}{\partial z}=-\rho_{\mathrm{r}}(z)g \tag{14}$$

This substitution is useful for computations of stratified flows in which the solution is started from rest and the reference pressure field is taken as the initial hydrostatic pressure field.

The numerical discretization of the momentum equations is similar to that carried out by Zang *et al.* [10], except the pressure correction method [11] is employed to obtain second-order accuracy in time. Without it, the approximate projection method used by Zang *et al.* is first-order accurate in time. Advection of momentum is computed with the QUICK scheme of Leonard [12] and discretized temporally with the second-order Adams–Bashforth scheme, and the pressure-Poisson equation is solved with the multigrid method with a normalized residual of 10^{-8}.

2.2. *Computation of the background potential energy*

Given a two-dimensional domain with a discrete density distribution given by $\rho_{i,k}$ and a cell volume distribution $\delta V_{i,k}$, the mass of the domain is given by

$$m=\sum_{i,k=1}^{N_i,N_k}\rho_{i,k}\delta V_{i,k} \tag{15}$$

where N_i and N_k are the total number of grid points in the i and k directions, respectively. If the height of the cell centres corresponding to each cell is given by $z_{i,k}$, then the total potential energy of the domain is given by a discretization of (2) to yield

$$E_{\mathrm{p}}=g\sum_{i,k=1}^{N_i,N_k}\rho_{i,k}z_{i,k}\delta V_{i,k} \tag{16}$$

Likewise, the background potential energy is given by the discrete form of (7), i.e.

$$E_{\mathrm{b}}=g\sum_{n=1}^{N_k\times N_k}\rho_{\mathrm{n}}^{*}z_{\mathrm{n}}^{*}\delta V_{\mathrm{n}}^{*} \tag{17}$$

where ρ_{n}^{*} is the sorted density distribution that is obtained by sorting the two-dimensional density field $\rho_{i,k}$ in decreasing order, and $\delta V_{\mathrm{n}}^{*}$ is the volume of the cell with density ρ_{n}^{*}. The sorted height z_{n}^{*} is computed with

$$z_{\mathrm{n+1}}^{*}=z_{\mathrm{n}}^{*}+\frac{\delta V_{\mathrm{n+1}}^{*}}{A(z_{\mathrm{n}}^{*})} \tag{18}$$

where the planform area $A(z)$ of the domain is known *a priori* and $n \in \{2, \dots, N_{ik}\}$, where $N_{ik} = N_i \times N_k$ is the total number of cells in the domain. The initial height z_1^* is given by

$$z_1^* = z_0 + \frac{\delta V_1^*}{2A(z_0)} \tag{19}$$

where z_0 corresponds to the bottom of the domain.

Computation of the background potential energy is expensive because it requires a sorting algorithm to obtain the sorted density field ρ_n^*. Tseng and Ferziger [13] have developed a probability density function approach that reduces the expense of computing the background potential energy by sorting the density field into N_b bins (which form the pdf) instead of N_{ik} bins. Their approach requires $\mathcal{O}(N_{ik}N_b)$ instead of $\mathcal{O}(N_{ik}^2)$ calculations to compute the background potential energy. We employ the quicksort algorithm to sort the density field, which requires $\mathcal{O}(\log_2 N_{ik})$ operations, and thus computation of the background potential energy requires $\mathcal{O}(N_{ik} \log_2 N_{ik})$ operations. From a computational point of view, the two methods are equivalent when the number of bins for the pdf approach is given by $N_b = \log_2 N_{ik}$.

3. ADVECTION SCHEMES

3.1. Existing schemes

Integrating the transport equation (12) in time and over a two-dimensional control volume (in x–z) gives the explicit update in conservative form as

$$\rho_{i,k}^{n+1} = \rho_{i,k}^{n} + C_w\rho_w - C_e\rho_e + C_s\rho_s - C_n\rho_n \tag{20}$$

where $\rho_{i,k}$ is a volume-averaged quantity and the subscripts denote East $(i+1/2)$, West $(i-1/2)$, North $(k+1/2)$, or South $(k-1/2)$ 'effective' face values that result from averaging the fluxes in time and over each face [14], and $C_w = u_w\Delta t/\Delta x$ and $C_s = v_s\Delta t/\Delta y$ are the Courant numbers defined at the faces. Using this discretization, conservation of mass is automatically guaranteed by computing unique flux-face values $\rho_{e(i)}$ and then using $\rho_{w(i)} = \rho_{e(i-1)}$ (likewise for ρ_n and ρ_s). If u and v are split into their positive and negative components such that

$$u^{\pm} = \tfrac{1}{2}(u \pm |u|) \tag{21}$$

$$v^{\pm} = \tfrac{1}{2}(v \pm |v|) \tag{22}$$

then (20) becomes

$$\begin{aligned}\rho_{i,k}^{n+1} = \rho_{i,k}^{n} &+ C_w^+\rho_w^+ + C_w^-\rho_w^- - C_e^+\rho_e^+ - C_e^-\rho_e^- \\ &+ C_n^+\rho_n^+ + C_n^-\rho_n^- - C_s^+\rho_s^+ - C_s^-\rho_s^-\end{aligned} \tag{23}$$

The objective of any finite-volume Eulerian advection scheme is to interpolate the volume-averaged values of $\rho_{i,k}$ to obtain the effective face values ρ_e. In this paper we will employ one-dimensional schemes to interpolate the volume-averaged quantities in each direction and employ the update in one step to obtain $\rho_{i,k}^{n+1}$. While this method has been shown to be unstable for schemes other than first-order upwinding unless operator splitting is employed and each coordinate direction is updated sequentially [14], the present formulation and the applications of References [10, 15] have been stable because of the low Courant numbers employed, and because the instability is a long wavelength instability that does not have time to grow in highly unsteady flows. We will employ explicit one-dimensional interpolation techniques that are guaranteed to be monotonic in one dimension but are not guaranteed to be monotonic in two or more dimensions. Again, because of the small Courant number we employ, we have found that one-dimensional interpolations remain monotonic in two dimensions except in extreme cases in which the velocity field is highly deformational. Monotonicity preservation in multidimensions can be guaranteed for arbitrary Courant numbers with the use of the methodology outlined by Leonard *et al.* [14].

The simplest one-dimensional approach is first-order upwinding, for which, assuming $C>0$ and $\rho=\rho^+$, yields

$$\rho_e = \rho_i \tag{24}$$

This method is highly diffusive [16] and only first-order accurate in space and time, but is simple to employ and guarantees monotonocity, even for the multidimensional case [14]. Second-order spatial accuracy can be achieved by employing central-differencing in space and forward-differencing in time (FTCS) with

$$\rho_e = \tfrac{1}{2}(\rho_i + \rho_{i+1}) \tag{25}$$

This method is unstable for pure advection unless diffusion is added to ensure that the grid Peclet number satisfies $Pe_\Delta = u\Delta x/\kappa \leqslant 2/C$ [17], where κ is the diffusion coefficient. We restrict ourselves to schemes that remain stable in the nondiffusive limit. The QUICK scheme of Leonard [12] is a third-order spatially accurate alternative to FTCS because, while it is theoretically unstable in the nondiffusive limit, the upper bound on the stability limit of $Pe_\Delta \leqslant 2/C$ results from a low wavenumber instability that is likely not excited in practice [18]. While QUICK was designed for steady flows, it has been applied to highly unsteady large-eddy simulations of turbulent flows and found to behave very well when coupled with the second-order Adams–Bashforth scheme [10]. In QUICK, an upwind-weighted curvature term is used to cancel the third-order truncation error present in the central-differencing scheme. Assuming $C > 0$, the face values using QUICK are given by

$$\rho_e = \tfrac{1}{2}(\rho_i + \rho_{i+1}) - \tfrac{1}{8}(\rho_{i-1} - 2\rho_i + \rho_{i+1}) \tag{26}$$

For unsteady flows, Leonard [12] presents the QUICKEST scheme, which is theoretically stable in the nondiffusive limit because it adds additional terms to stabilize the QUICK scheme. The face values using QUICKEST when $C>0$ are given by

$$\rho_e = \tfrac{1}{2}[(\rho_i + \rho_{i+1}) - C(\rho_{i+1} - \rho_i)] - \tfrac{1}{6}(1 - C^2)(\rho_{i-1} - 2\rho_i + \rho_{i+1}) \tag{27}$$

To guarantee monotonicity for steady flows, Leonard [19] proposed the SHARP scheme, which also behaves quite well for highly unsteady flows when used with the second-order Adams–Bashforth scheme [10], despite very subtle, but negligible, overshoot behaviour in the presence of sharp fronts. The ULTIMATE-QUICKEST scheme of Leonard [8] is the unsteady version of SHARP, which guarantees monotonicity in highly unsteady convective flow. This scheme employs the QUICKEST scheme with a universal limiter that limits the magnitude of the higher-order terms (in a very small fraction of the flow) to maintain monotonicity of the advected quantity.

Monotonicity is also guaranteed with the use of total variation diminishing, or so-called shock-capturing schemes [20–22], which are five-point second-order accurate schemes that maintain the total variation of the advected quantity. From Hirsch [23], a scheme is total variation diminishing if the total variation of a solution to a nondiffusive advection equation does not increase in time, such that $\mathrm{TV}(\rho^{n+1}) \leqslant \mathrm{TV}(\rho^n)$, where the total variation of ρ is given by

$$\mathrm{TV}(\rho)=\sum_i |\rho_{i+1}-\rho_i| \tag{28}$$

For the TVD schemes used in the present paper, the flux-face values are computed with the upwind values plus the addition of a higher-order term with

$$\rho_\mathrm{e}^+ = \rho_i + \tfrac{1}{2}\Psi(r_i^+)(1-C_\mathrm{e}^+)(\rho_{i+1}-\rho_i) \tag{29}$$

$$\rho_\mathrm{e}^- = \rho_{i+1} - \tfrac{1}{2}\Psi(r_i^-)(1+C_\mathrm{e}^-)(\rho_{i+1}-\rho_i) \tag{30}$$

where r represents the ratio of the upwind gradient to the local gradient such that

$$r_i^+ = \frac{\rho_i-\rho_{i-1}}{\rho_{i+1}-\rho_i} \tag{31}$$

$$r_i^- = \frac{\rho_{i+2}-\rho_{i+1}}{\rho_{i+1}-\rho_i} \tag{32}$$

The limiting function Ψ defines the particular scheme that is used, a host of which have been developed in the literature. We employ the following limiters that are presented in Hirsch [23]:

$$\Psi(r)=\begin{cases} 0 & \text{Upwind} \\ 1 & \text{Lax–Wendroff} \\ \max[0,\min(2r,1),\min(r,2)] & \text{Superbee} \\ (r+|r|)/(1+r) & \text{Van Leer} \\ \min(r,1) & \text{MINMOD} \\ \max\left[0,\min\left(2,2r,\dfrac{1+r}{2}\right)\right] & \text{MUSCL} \end{cases} \tag{33}$$

Therefore, if $\Psi(r)=0$, the scheme reduces to the upwind scheme, which is first-order accurate in space and time, and if $\Psi(r)=1$, it is the Lax–Wendroff scheme, which is second-order accurate in space and time. All other schemes are second-order accurate in both space and time as long as $\Psi(1)=1$. The ULTIMATE-QUICKEST scheme of Leonard [8] can also be expressed in this way using

$$\Psi(r)=\max\left[0,\min\left(\frac{1}{2}(1+r)+\frac{1}{6}(1-r)(1-2|C|),\frac{2}{1-|C|},\frac{2r}{|C|}\right)\right] \tag{34}$$

where the Courant number C is evaluated at the same flux face as ρ_e. One advantage of the ULTIMATE-QUICKEST scheme is that it is not as restrictive as the TVD schemes, since the upper bound of $\Psi(r)$ is boundless with the introduction of the Courant number in the denominator, while it is limited to a maximum value of 2 for the TVD schemes. Two compressive schemes that take advantage of the less restrictive boundary of this so-called universal limiter are the second-order accurate Super-C scheme, for which

$$\Psi(r)=\begin{cases}\min\left(\frac{2r}{|C|},1\right) & 0\leqslant r\leqslant 1\\ \min\left(r,\frac{2}{1-|C|}\right) & r>1\\ 0 & \text{otherwise}\end{cases} \tag{35}$$

and the maximally compressive first-order accurate Hyper-C scheme, for which

$$\Psi(r)=\begin{cases}\min\left(\frac{2r}{|C|},\frac{2}{1-|C|}\right) & r>0\\ 0 & \text{otherwise}\end{cases} \tag{36}$$

These schemes are overly compressive, yet the latter proves to be very useful as a scheme to reverse the effects of numerical diffusion, as discussed in the next section. An excellent comparison between the TVD and ULTIMATE schemes can be found in the work of Leonard [8], where their behaviour is quantified with respect to one-dimensional advection. In Section 4 we demonstrate the behaviour of these schemes when applied to a sloshing interfacial gravity wave and discuss the effects the schemes have on the background potential energy.

3.2. A background potential energy preserving formulation

As we will show in the next section, in the absence of salt diffusion, no existing advection schemes maintain the background potential energy. Schemes that do not preserve monotonicity cause a decrease in the background potential energy because overshoots in the density profile are equivalent to numerical antidiffusion. Monotonicity-preserving schemes, on the other hand, cause an increase in the background potential energy if they are even mildly diffusive, which is a requirement of such schemes if they are not among the ‘compressive’ of the monotonic schemes. The compressive schemes always cause a decrease in the background potential energy because steepening of density fronts is also effectively numerical antidiffusion, even if this effect is monotonic. Despite these adverse effects, we will show that it is possible to maintain the background potential energy by using a combination of diffusive and antidiffusive schemes, the choice of the particular scheme depending on the value of the background potential energy. In the absence of scalar diffusion, if the background potential energy is larger than its initial value, then this indicates the presence of numerical diffusion, whereas if the background potential energy is less than its initial value, then this indicates the presence of numerical antidiffusion. Because it is always desirable to maintain monotonicity from a physical perspective, regardless of whether or not nonmonotonicity causes a decrease in the background potential energy, we require that the schemes we choose to combine to create the background potential energy produce negligible overshoots. This is to guarantee that the change in the background potential energy is resulting from either numerical diffusion or antidiffusion and not overshoots in the density profile.

We employ two forms of the background potential energy formulation. For the first, we use the Hyper-C universal limiter as the compressive scheme, and for the diffusive scheme, we use the Van Leer limiter, so that this background potential energy preserving scheme becomes a limited scheme in which the limiter is a function of the background potential energy E_b and the upwind gradient ratio r. If $E_{\mathrm{b}0}$ is the initial background potential energy, then this limiter is given by

$$\Psi(r,E_\mathrm{b})=\begin{cases}(r+|r|)/(1+r) & E_\mathrm{b}\leqslant E_{\mathrm{b}0}\\ \min\left(\dfrac{2r}{|C|},\dfrac{2}{1-|C|}\right) & r>0 \text{ and } E_\mathrm{b}>E_{\mathrm{b}0}\\ 0 & \text{otherwise}\end{cases}\tag{37}$$

This scheme is referred to as BPEP1 in what follows.

As an alternative to this scheme, because the Hyper-C scheme can be overly compressive, we employ a second scheme, which we refer to as BPEP2, that employs the Superbee scheme in place of the Hyper-C scheme, so that its limiter is given by

$$\Psi(r,E_\mathrm{b})=\begin{cases}(r+|r|)/(1+r) & E_\mathrm{b}\leqslant E_{\mathrm{b}0}\\ \max[0,\min(2r,1),\min(r,2)] & r>0 \text{ and } E_\mathrm{b}>E_{\mathrm{b}0}\\ 0 & \text{otherwise}\end{cases}\tag{38}$$

4. IMPLEMENTATION OF ADVECTION SCHEMES

4.1. Computation set-up

We apply the schemes discussed in the previous section to a standing interfacial gravity wave in order to evaluate their capability to maintain the background potential energy. The test case is a finite-amplitude deep-water standing wave in an inviscid fluid, and the initial condition for the interface ζ to second order in the steepness ka is given by Thorpe [24]

$$k\zeta(x)=ka\left[\left(1-\frac{(ka)^2}{64}\right)\cos kx-\frac{(ka)^2}{8}\cos 3kx\right]\tag{39}$$

where the second-order dispersion relation that describes the frequency of the wave ω as a function of the wavenumber $k=2\pi/L$, nondimensional interface thickness $k\delta$, and steepness

ka, is given by

$$\omega^2 = \omega_{k\delta}^2 \left[1 - \frac{(ka)^2}{8} \right] \tag{40}$$

and $\omega_{k\delta}$ is the zero-steepness, finite-interface thickness frequency. The effect of a finite interface is to reduce the frequency, which, to first order, can be approximated, following Phillips [25], by

$$\omega_{k\delta}^2 = \frac{g'k}{2} \left(1 + \frac{k\delta}{4} \right)^{-1} \tag{41}$$

In Equations (39)–(41), a is the maximum wave amplitude, δ is the interface thickness, and $g' = g\Delta\rho/\rho_0$ is the reduced gravity, where $\Delta\rho$ is the density difference between the layers.

The evolution of the two-dimensional standing wave is computed in a 1.0 m × 1.0 m tank with an initial steepness of $ka = 0.1$ and nondimensional interface thicknesses of $k\delta = 0.01\pi$ and $k\delta = 0.05\pi$, as shown in Figure 2. The 80 × 80 grid is set up so that 4 cells resolve the $k\delta = 0.01\pi$ interface while 10 cells resolve the $k\delta = 0.05\pi$ interface. The grid spacing in the vertical is a minimum at the interface and is stretched from mid-depth with a maximum vertical stretching factor of 1.1, where the vertical stretching factor is defined by

$$\beta_k = \frac{z_{k+1} - z_k}{z_k - z_{k-1}} \tag{42}$$

The density difference between the layers is $\Delta\rho/\rho_0 = 0.03$, and the initial density distribution is given by

$$\rho(x,z) = -\frac{\Delta\rho}{2} \tanh \left[\frac{2\tanh^{-1}\alpha}{k\delta} (kz - k\zeta + kd/2) \right] \tag{43}$$

where $\alpha = 0.99$ and $k\zeta$ is defined in (39).

The wave periods adjusted for amplitude and interface thickness using the dispersion relation (40) are $T = 9.28$ s for $k\delta = 0.01\pi$ and $T = 9.43$ s for $k\delta = 0.05\pi$. Since these wave periods

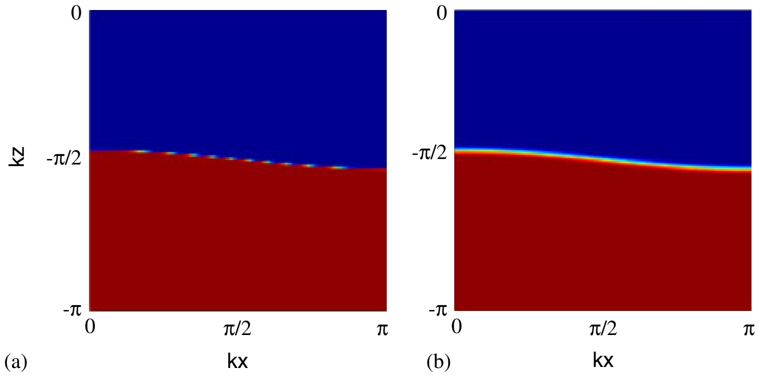

Figure 2. Initial conditions for the $ka = 0.1$ interfacial standing wave with an interface thickness of: (a) $k\delta = 0.01\pi$; and (b) $k\delta = 0.05\pi$ with a density difference of $\Delta\rho/\rho_0 = 0.03$ between the two layers.

are first-order approximations in $k\delta$ and second-order in ka, to obtain more accurate values for the periods, two oscillations of the internal seiche are first computed with a total of 800 time steps, corresponding to time steps of $\Delta t = 0.0232$ s for $k\delta = 0.01\pi$ and $\Delta t = 0.0236$ s for $k\delta = 0.05\pi$. This corresponds to a maximum wave Courant number of $C_{\mathrm{w}} = 0.04$ for each case, where the wave Courant number is defined by

$$C_{\mathrm{w}} = ka\frac{c\Delta t}{\Delta x} \tag{44}$$

and $c = \omega/k$ is the wave speed. From these simulations, the average periods are then computed to yield periods of $T = 9.74 \pm 0.06$ for $k\delta = 0.01\pi$ and $T = 9.82 \pm 0.05$ for $k\delta = 0.05\pi$, that corresponds to $\Delta t = 0.02435$ s and $\Delta t = 0.02455$ s, respectively.

4.2. Behaviour of advection schemes

Advection of the density field is computed with the advection schemes discussed in Section 3, namely (a) First-order upwind, (b) Lax–Wendroff, (c) QUICK, (d) QUICKEST, (e) SHARP, (f) ULTIMATE-QUICKEST, (g) MINMOD, (h) Van Leer, (i) MUSCL, (j) Superbee, (k) Super-C, (l) Hyper-C, (m) BPEP1 and (n) BPEP2. Figures 3 and 4 depict the surface plots for each scheme when applied to the $k\delta = 0.01\pi$ and $k\delta = 0.05\pi$ cases, and Figures 5 and 6 depict the vertical density profiles at $x = 0$ for each of the surface plots. The computations are compared to the exact solution, which is identical to the initial condition since the density field returns to its initial distribution after two periods of oscillation. Errors that arise in the computation of the density distribution change the frequency of oscillation of the interface and hence alter the final position of the interface after two computed oscillations. The errors in computing the frequency of oscillation, and hence the error in the final position of the interface, are shown in Figure 7. For the $k\delta = 0.01\pi$ case in Figure 7(a), the Lax–Wendroff and QUICK schemes induce the most pronounced underestimation of the period because the overshoots effectively reduce the interface thickness for this case, which translates to an increased frequency in Equation (41). For the $k\delta = 0.05\pi$ case, however, these schemes cause an overestimation of the wave period because they effectively add higher modes to the solution that propagate at slower speeds. Because it is the most diffusive and causes the most interface thickening, first-order upwind causes an overestimation in the wave period by roughly 1.5% for both interface cases, and because it is the most compressive, Hyper-C causes an underestimation in the wave period by roughly 1% for the thick interface and roughly 0.5% for the thin interface. While these errors are small, Figure 7 shows that there is a clear relationship between the character of the scheme and the resulting error in the computation of the wave period.

Errors in the computation of the wave period for the nonmonotonic schemes result from overshoots in the density field. These are most evident in Figures 3 and 5 for the Lax–Wendroff, QUICK, and QUICKEST schemes for the $k\delta = 0.01\pi$ interface, and are less drastic for the $k\delta = 0.05\pi$ interface, as shown in Figures 4 and 6. Overshoots generated by the SHARP scheme are negligible since they only appear at the beginning of the simulation for $k\delta = 0.01\pi$, and the SHARP scheme is monotonic for $k\delta = 0.05\pi$. For $k\delta = 0.01\pi$, once the interface thickens due to numerical diffusion, the SHARP scheme smoothes out these initial overshoots and the interface once again becomes monotonic. While the compressive schemes Superbee, Super-C and Hyper-C most accurately predict the shape of the $k\delta = 0.01\pi$ interface, they

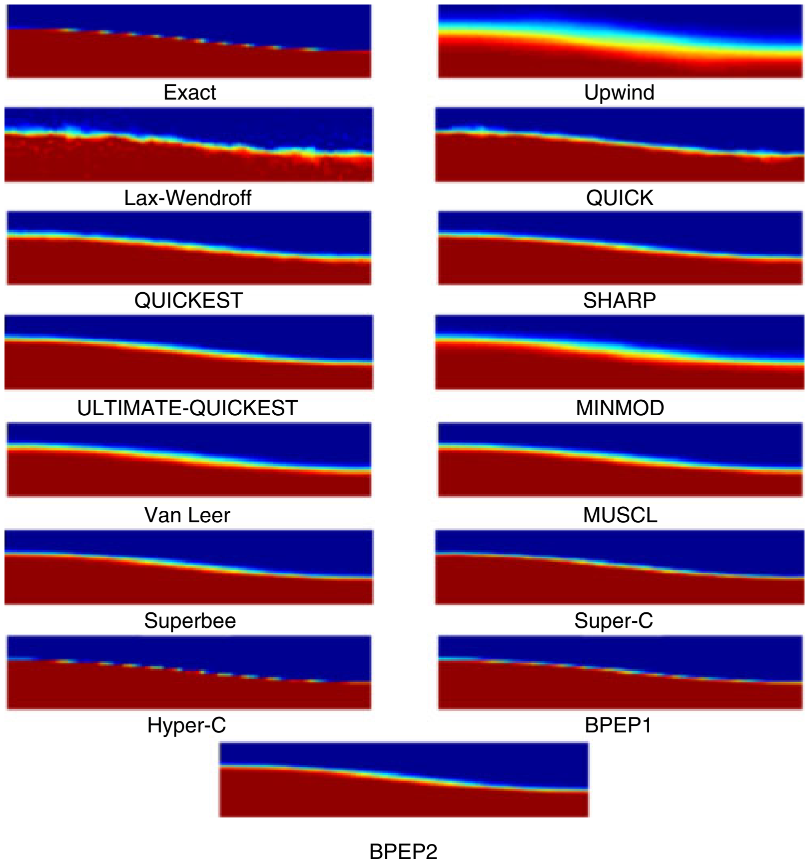
Exact
Upwind
Lax-Wendroff
QUICK
QUICKEST
SHARP
ULTIMATE-QUICKEST
MINMOD
Van Leer
MUSCL
Superbee
Super-C
Hyper-C
BPEP1
BPEP2

Figure 3. Density fields for the $k\delta = 0.01\pi$ interface thickness after two periods of oscillation for the advection schemes discussed in the text. The exact solution corresponds to the initial condition.

perform poorly for the $k\delta = 0.05\pi$ interface because they erroneously compress the density profile at the interface. In particular, the Hyper-C scheme produces multistep profiles in the density field, as can be seen in Figure 6. All of the other nonmonotonic schemes perform much better for the thicker interface case because interfacial diffusion is not necessary to maintain monotonicity. Interfacial diffusion, however, increases for the monotonic schemes in regions of the flowfield where the local Courant number is reduced. This can be seen most clearly in Figure 4, where the interface thickness is larger at the centre of the wave profile than it is at the edges. This is the region where the vertical and horizontal velocity is a minimum. This effect is pronounced for the TVD schemes such as Superbee because, while the Courant number is a minimum at the centre of the domain, it is a maximum at the edges, resulting in nonlinear steepening and interface thinning there.

From a qualitative point of view, the overly compressive schemes behave the best for the thin interface while the other monotonic schemes behave the best for the thick interface, and the nonmonotonic schemes are unacceptable for both cases. The BPEP1 scheme appears

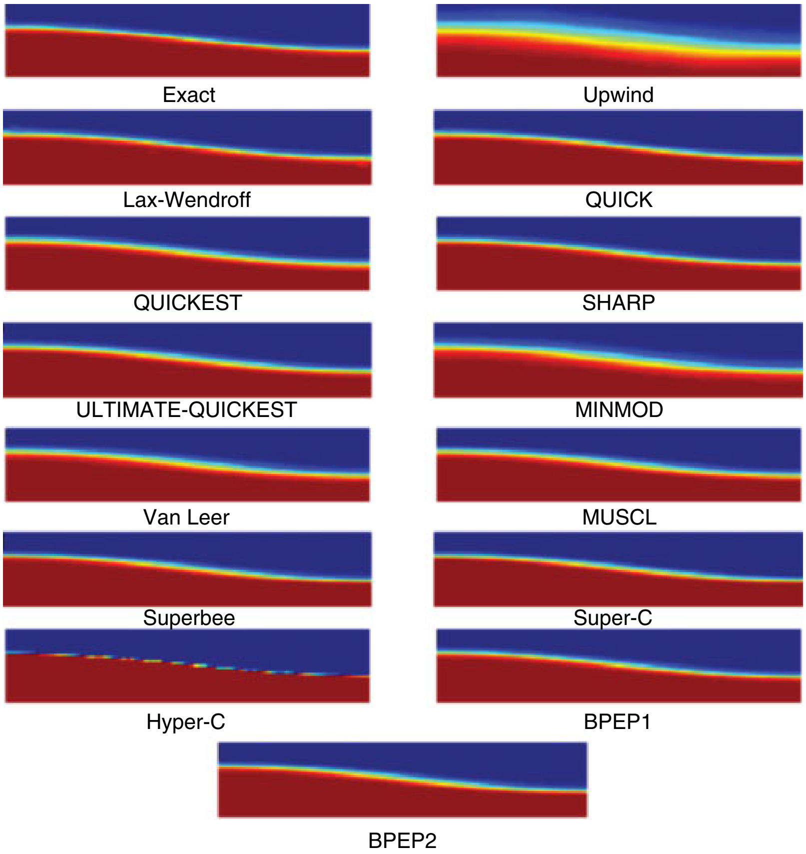

Exact; Upwind; Lax-Wendroff; QUICK; QUICKEST; SHARP; ULTIMATE-QUICKEST; MINMOD; Van Leer; MUSCL; Superbee; Super-C; Hyper-C; BPEP1; BPEP2

Figure 4. Density fields for the $k\delta = 0.05\pi$ interface thickness after two periods of oscillation for the advection schemes discussed in the text. The exact solution corresponds to the initial condition.

to represent the interface most accurately for both cases because it takes advantage of the compressive behaviour of the Hyper-C scheme for the thin interface, and the Van Leer scheme for the thick interface. The BPEP2 scheme does not capture the thin interface case as well because it is only as compressive as the Superbee scheme. This is confirmed quantitatively by the results depicted in Table I, which presents several error metrics, which are computed as follows. Given that ρ^{2T} is the advected solution after two periods, and $\hat{\rho}$ represents the exact solution, or the solution at $t = 0$, then the L_1-error norm of the density field is given by

$$L_1 = \frac{\sum_{i=1}^{N_i} \sum_{k=1}^{N_k} |\rho_{i,k}^{2T} - \hat{\rho}_{i,k}| \delta A_{i,k}}{\sum_{i=1}^{N_i} \sum_{k=1}^{N_k} |\hat{\rho}_{i,k}| \delta A_{i,k}} \tag{45}$$

the L_2-error norm of the density field is given by

$$L_2^2 = \frac{\sum_{i=1}^{N_i} \sum_{k=1}^{N_k} (\rho_{i,k}^{2T} - \hat{\rho}_{i,k})^2 \delta A_{i,k}}{\sum_{i=1}^{N_i} \sum_{k=1}^{N_k} \hat{\rho}_{i,k}^2 \delta A_{i,k}} \tag{46}$$

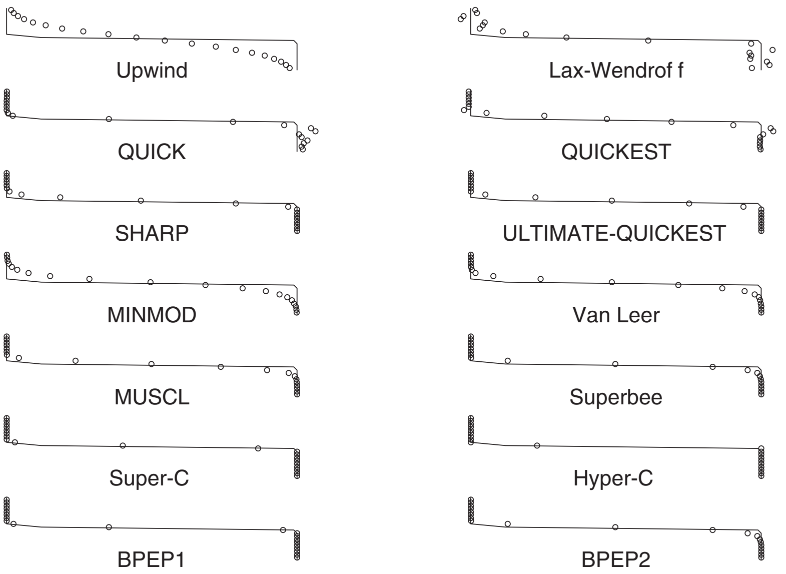

Figure 5. Density profiles at $x = 0$ for the $k\delta = 0.01\pi$ interface thickness after two periods of oscillation for the advection schemes discussed in the text. The initial or exact field is the solid line while the circles represent the advected profiles.

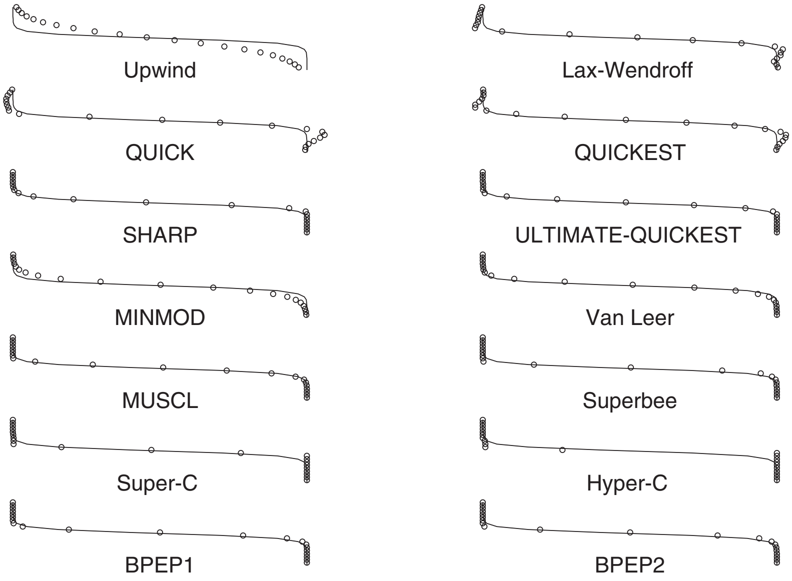

Figure 6. Density profiles at $x = 0$ for the $k\delta = 0.05\pi$ interface thickness after two periods of oscillation for the advection schemes discussed in the text. The initial or exact field is the solid line while the circles represent the advected profiles.

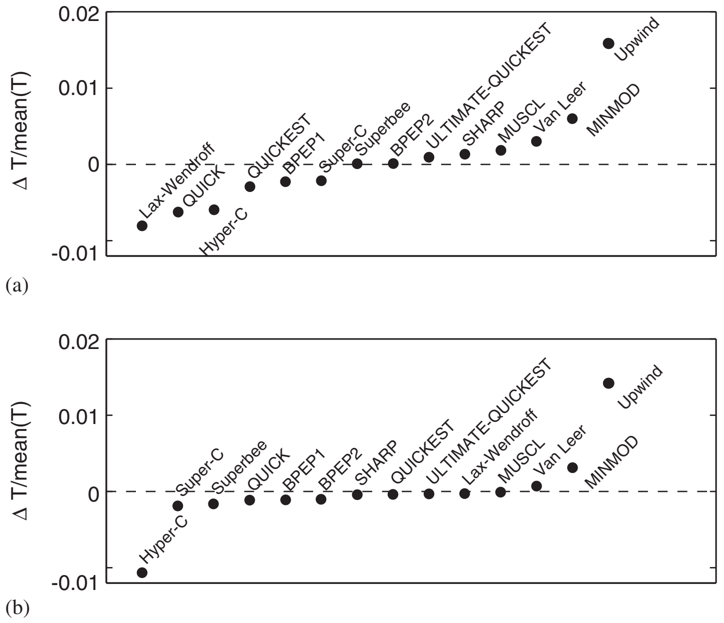

Figure 7. Normalized error in computing the wave period for the: (a) $k\delta = 0.01\pi$; and (b) $k\delta = 0.05\pi$ cases when using different advection schemes.

Table I. Per cent errors for different advection schemes after two interfacial wave oscillations for the $k\delta = 0.01\pi$ and $k\delta = 0.05\pi$ cases. The lowest values in the L_1 and L_2 columns are in bold.

	$k\delta = 0.01\pi$			$k\delta = 0.05\pi$		
Scheme	L_1	L_2	$L_{\max}$	L_1	L_2	$L_{\max}$
Lax–Wendroff	2.6	9.4	32.5	0.3	1.5	13.3
QUICK	1.5	6.9	14.6	0.5	2.2	15.4
QUICKEST	1.6	7.9	11.4	0.7	3.1	6.2
Upwind	4.8	16.4		3.9	12.6	
MINMOD	2.2	10.5		1.3	5.5	
Van Leer	1.6	8.9		0.7	3.7	
MUSCL	1.4	8.3		0.6	3.2	
ULTIMATE-QUICKEST	1.2	7.6		0.4	2.6	
Superbee	0.9	6.5		0.4	2.2	
Super-C	0.5	4.8	0.2	0.4	2.0	
Hyper-C	**0.1**	**1.4**	2.7	1.0	6.2	1.5
SHARP	1.1	6.8	1.2	**0.3**	**1.5**	
BPEP1	0.3	2.8	0.3	**0.3**	1.9	
BPEP2	0.9	6.5		0.4	2.3	

and the max-norm is given by

$$L_{\max}=\frac{\max(\rho^{2T})-\max(\hat{\rho})}{\max(\hat{\rho})} \tag{47}$$

where N_i and N_k are the number of grid points in the horizontal and vertical, respectively, and $\delta A_{i,k}$ is the area of grid cell i,k. The max-norm is a measure of the nonmonotonicity of the particular scheme and is identically zero for monotonic schemes.

The most surprising result in Table I is that the SHARP scheme produces the lowest L_1 and L_2 errors for both interface thicknesses if we do not include the overly compressive schemes Superbee, Super-C, and Hyper-C for the thin interface case, or the BPEP1 scheme. Although it was designed for steady flows, it is clear that SHARP is a superior scheme for the present simulations even though it does not guarantee monotonicity of the advected quantity. The overshoots it exhibits clearly have a minimal effect on the overall solution, and are substantially lower than those produced by the Lax–Wendroff, QUICK and QUICKEST schemes. Furthermore, it becomes monotonic for the thicker interface case. Overall, the BPEP1 scheme produces the lowest errors for both interface cases. Like the SHARP scheme, however, the BPEP1 scheme is not monotonic, nor are the Super-C and Hyper-C schemes, since they have positive max-norms in Table I. The $k\delta=0.01\pi$ interface case is being resolved by only four grid cells at the start of the simulation. The compressive schemes effectively reduce this resolution to one cell in some regions of the flow, and, as a result, the velocity field becomes highly oscillatory because of the QUICK scheme that is used to advect momentum as well as the central-differencing character of the underlying pressure correction method. Under such extreme conditions, the universal limiter does not guarantee monotonic behaviour of the density field for the Super-C and Hyper-C schemes. This is a result of the one-dimensional interpolation scheme to obtain the face values, which does not guarantee monotonicity, as outlined in Reference [14], especially for highly deformational velocity fields. This is shown in Figure 8, where oscillations exist in the velocity field for the $k\delta=0.01\pi$ case, and as

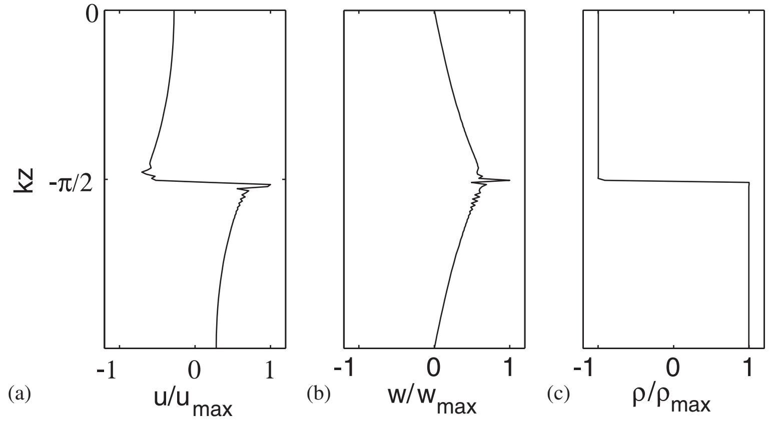

(a) (b) (c)

Figure 8. Oscillations in the: (a) u-velocity; (b) w-velocity fields that induce negligible nonmonotonic behaviour of the density field; and (c) when using the Hyper-C scheme for the $k\delta=0.01\pi$ interface. The profiles are shown at $kx=3\pi/4$ and $t=0.6T$.

a result, slight overshoots are introduced into the density field. This same character is also exhibited by the Hyper-C scheme for the $k\delta = 0.05\pi$ case because the scheme reduces the interface thickness due to its overcompressive behaviour. The BPEP1 scheme, however, is monotonic for the $k\delta = 0.05\pi$ case because the velocity fields are smoother, and the BPEP2 scheme is monotonic for both cases because the TVD schemes still maintain monotonicity.

These results show that accuracy comes at the cost of minor nonmonotonicity, since SHARP and BPEP1 are the most well-behaved, yet introduce overshoots in the density field for the $k\delta = 0.01\pi$ case. The second-order TVD schemes maintain monotonicity, but at the expense of diffusing the interface and hence increasing the L_1 and L_2 errors. As we will show in the next section, the diffusive character of the second-order TVD schemes requires the use of the Hyper-C scheme if the background potential energy is to be conserved for both the $k\delta = 0.01\pi$ and 0.05π cases.

4.3. *Effects on the background potential energy*

Using the technique described in Section 2.2, the background potential energy is computed at each time step and used to quantify the numerical diffusion associated with each of the advection schemes presented in Section 3. Figures 9 and 10 show the history of the background potential energy for the nonmonotonic schemes Lax–Wendroff, QUICK, QUICKEST and SHARP, and Figures 11 and 12 show the history of the background potential energy for the monotonic schemes. The Super-C, Hyper-C and BPEP1 schemes are included among the monotonic schemes because the overshoots they generate are minimal and do not significantly affect the background potential energy. The background potential energy is plotted as a departure from the initial background potential energy ΔE_b and normalized by the initial available potential energy E_{a0}.

Figures 9 and 10 show that the effect of the highly oscillatory schemes Lax–Wendroff, QUICK and QUICKEST is to cause a substantial decrease in the background potential energy. One can show that overshoots in the density profile lead to a decrease in the background potential energy by considering a quiescent two-layer fluid with a density distribution

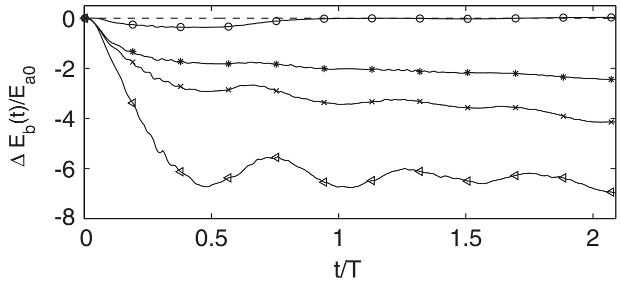

Figure 9. Background potential energy evolution of the internal seiche with initial interface thickness $k\delta = 0.01\pi$ normalized by the initial available potential energy E_{a0}, for the nonmonotonic schemes. Legend: Lax–Wendroff ◁, QUICK ×, QUICKEST ∗, SHARP ○.

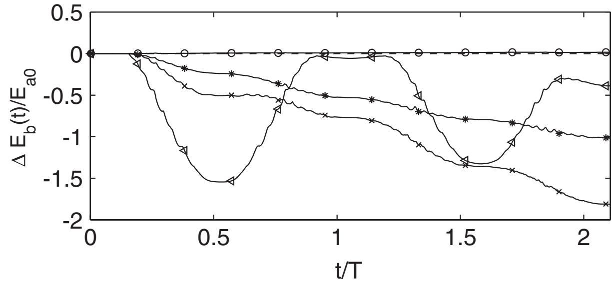

Figure 10. Background potential energy evolution of the internal seiche with initial interface thickness $k\delta = 0.05\pi$ normalized by the initial available potential energy E_{a0}, for the nonmonotonic schemes. Legend: Lax–Wendroff $\triangleleft$, QUICK $\times$, QUICKEST $*$, SHARP $\circ$.

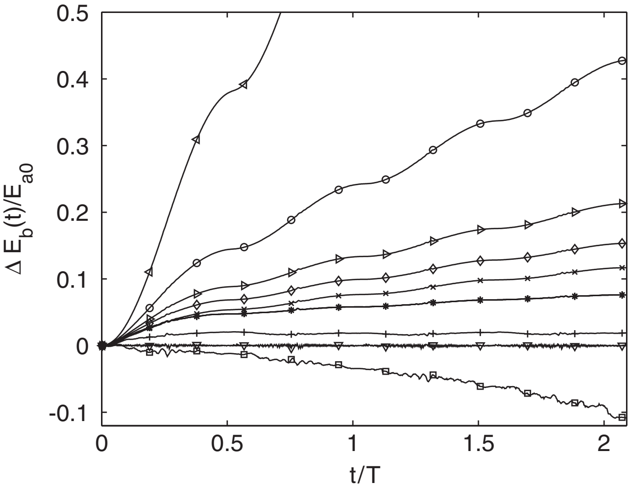

Figure 11. Background potential energy evolution of the internal seiche with initial interface thickness $k\delta = 0.01\pi$ normalized by the initial available potential energy E_{a0}, for the monotonic schemes. Legend: Upwind $\triangleleft$, MINMOD $\circ$, Van Leer $\triangleright$, MUSCL $\diamond$, ULTIMATE-QUICKEST $\times$, Superbee/BPEP2 $*$, Super-C $+$, Hyper-C $\Box$, BPEP1 $\bigtriangledown$.

given by

$$\rho(z) = \begin{cases} \dfrac{\Delta\rho}{2} & -d \leqslant z < -d/2 \\ -\dfrac{\Delta\rho}{2} & -d/2 \leqslant z < 0 \end{cases} \tag{48}$$

with $\Delta\rho$ being the density difference between the two layers and d being the total depth of the two layers. Given this distribution, the quiescent background potential energy E_{b0} is obtained using Equation (3) as

$$E_{b0} = -\tfrac{1}{8}Md \tag{49}$$

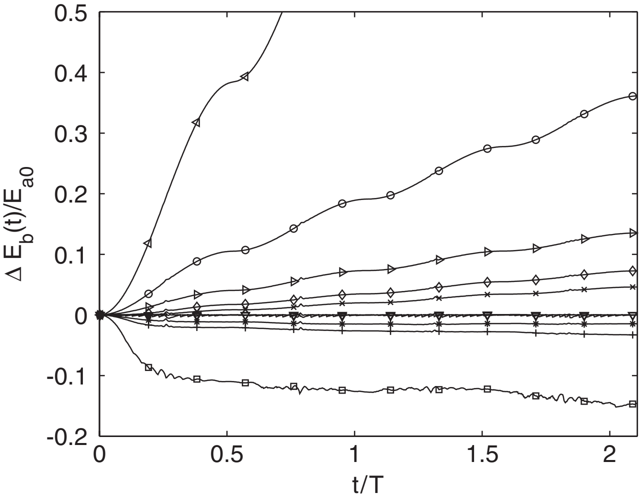

Figure 12. Background potential energy evolution of the internal seiche with initial interface thickness $k\delta = 0.05\pi$ normalized by the initial available potential energy E_{a0}, for the monotonic schemes. Legend: Upwind ◁, MINMOD ○, Van Leer ▷, MUSCL ◇, ULTIMATE-QUICKEST ×, Superbee ∗, Super-C +, Hyper-C □, BPEP1/BPEP2 ▽.

where M is the mass of the fluid, as in Equation (5). To understand the effects of overshoots on the background potential energy, we can approximate a numerical perturbation as affecting both the upper and lower layers uniformly by an amount $\delta\rho$, as shown in Figure 13. The resulting perturbed density field is given by

$$\rho(z) = \begin{cases} \dfrac{\Delta\rho}{2} + \delta\rho & -d \leqslant z < -d/2 \\ -\dfrac{\Delta\rho}{2} - \delta\rho & -d/2 \leqslant z < 0 \end{cases} \tag{50}$$

noting that this perturbation does not change the mass of the fluid M. The change in the background potential energy resulting from this perturbation is given by

$$\delta E_\mathrm{b} = -2|E_{\mathrm{b}0}|\frac{\delta\rho}{\Delta\rho} \tag{51}$$

which shows that overshoots in the density profile, represented by $+\delta\rho$, result in a decrease in the background potential energy. Diffusion of the profile, on the other hand, which is represented by $-\delta\rho$, results in a gain in the background potential energy.

The Lax–Wendroff scheme induces the largest overshoots, and hence it causes the most substantial change in the background potential energy for both the $k\delta = 0.01\pi$ and 0.05π cases. For the thicker interface case shown in Figure 10, however, the net change in the background potential energy for all three cases is substantially reduced as a result of the reduced overshoot. Oscillations in the background potential energy result from the periodic nature of the Courant number over the course of the simulation. The Lax–Wendroff scheme shows the most pronounced oscillations in the background potential energy because it is the

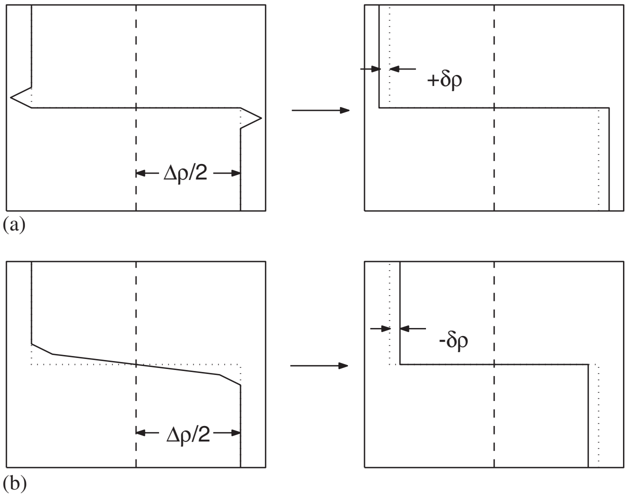

Figure 13. Approximate effect of a numerical perturbation to a quiescent density field, showing that overshoots can be approximated by increases in the density difference between the layers by an amount $+\delta\rho$ (a), while diffusion can be approximated by a decrease in the same amount, as shown in (b).

most dispersive of the schemes presented here. Changes in the background potential energy are the lowest at $t/T = 0$, 1/2, 1, 3/2 and 2, since these are the points at which the fluid velocity is a minimum and hence the Courant number is a minimum. As shown in Figure 14, the Lax–Wendroff scheme generates overshoots in the density field on the underside of the interface at $kx = 3\pi/4$ due to negative dispersion as it rises. This translates into an increase in the background potential energy between $t/T = 0$ and $t/T = 1/2$, but as the interface at $x = 3\pi/4$ falls between $t/T = 1/2$ and $t/T = 1$, positive dispersion removes the overshoots and causes an increase in the background potential energy. This occurs at $kx = \pi/4$ as well, but above the interface. From Figure 14, there is weak dispersion that induces overshoots above the interface at $kx = 3\pi/4$ as it falls, but these are much weaker because the upper half of the interface is smoothed as it rises between $t/T = 1/2$ and $t/T = 1$.

Owing to its nonmonotonicity at the beginning of the simulation, the SHARP scheme initially causes a slight decrease in the background potential energy for the $k\delta = 0.01\pi$ case, as shown in Figure 9, but diffusion eventually causes a net increase in the background potential energy after two wave periods. Because it is monotonic for $k\delta = 0.05\pi$, however, the SHARP scheme causes a monotonic increase in the background potential energy for this case, as shown in Figure 10.

As shown in Figures 11 and 12, the background potential energy increases monotonically for first-order upwind and all of the TVD schemes except for Superbee. The Superbee limiter is compressive for the $k\delta = 0.05\pi$ case, as shown in Figure 12, and hence causes a decrease in the background potential energy. However, in order to guarantee monotonicity for the $k\delta = 0.01\pi$ case in Figure 11, interfacial diffusion causes a monotonic increase in the background potential energy. The only scheme that causes a decrease in the background potential energy for the $k\delta = 0.01\pi$ case in Figure 11 is Hyper-C, which lowers the background potential energy of any initial distribution with a finite interface thickness. As shown in Figure 12, however, in

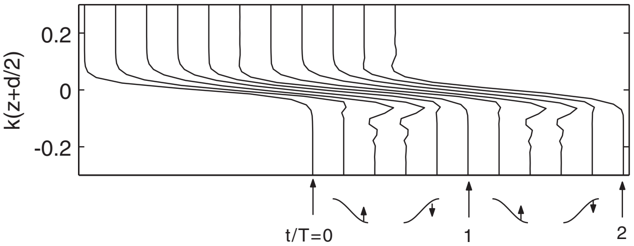

Figure 14. Density profiles over time at $kx = 3\pi/4$ for the $k\delta = 0.05\pi$ case, showing how the Lax–Wendroff scheme generates overshoots due to negative dispersive error as the profile at $kx = 3\pi/4$ rises, and effectively removes these overshoots with positive dispersion as the profile descends.

addition to the Superbee scheme, both the Super-C as well as the Hyper-C schemes cause a monotonic decrease in the background potential energy due to their compressive behaviour. For the thin interface case shown in Figure 11, the BPEP2 scheme does not maintain the background potential energy because it is only as compressive as the Superbee scheme, and hence the background potential energy for the BPEP2 scheme is identical to that for Superbee for the $k\delta = 0.01\pi$ case. Because Superbee can lower the background potential energy for $k\delta = 0.05\pi$, the BPEP2 scheme maintains the background potential energy for this case. The BPEP1 scheme, on the other hand, maintains the background potential energy for both cases in Figures 11 and 12 because it employs the Hyper-C scheme that is the only scheme capable of reducing the background potential energy for $k\delta = 0.01\pi$. Figure 12 demonstrates how the rate of increase in the background potential energy for the Van Leer scheme is roughly the same magnitude as the rate of decrease of the background potential energy for the Hyper-C scheme. This was the original motivation for using the Van Leer scheme as the diffusive limiter, despite a wealth of available choices. We have found through numerical experimentation that the Van Leer scheme yields the most favourable results.

We quantitatively determine the relative increase or decrease in the background potential energy by computing the total change in the background potential energy of each scheme after two periods with

$$\Delta E_{\mathrm{b}} = \frac{E_{\mathrm{b}f} - E_{\mathrm{b}0}}{E_{\mathrm{a}0}} \tag{52}$$

where $E_{\mathrm{b}0}$ is the initial background potential energy, $E_{\mathrm{b}f}$ is the background potential energy after two periods, and $E_{\mathrm{a}0}$ is the initial available potential energy, as defined in Equation (1). The results are presented in Table II. Positive values indicate a scheme that is diffusive on average, while negative values indicate schemes that either induce overshoots or are compressive. Because they generate overshoots for both interface cases, the values for the Lax–Wendroff, QUICK, and QUICKEST schemes are negative, while the Hyper-C scheme is negative for both cases because it is compressive, as are the Superbee and Super-C schemes for $k\delta = 0.05\pi$. Not surprisingly, the Super-C scheme induces the smallest change in the background potential energy after the BPEP1 scheme for the thin interface case. The BPEP2 scheme matches the value for Superbee when $k\delta = 0.01\pi$, since for this thin interface the BPEP2 scheme is effectively the Superbee scheme. While they are small, the values for the BPEP schemes are not identically zero because, as will be shown in the next section, they oscillate about the x-axis as they change back and forth between the compressive and Van Leer schemes.

Table II. Per cent error measures of the background potential energy for different advection schemes after two interfacial wave oscillations for the $k\delta_1 = 0.01\pi$ and $k\delta_2 = 0.05\pi$ cases.

Scheme	$k\delta = 0.01\pi$	$k\delta = 0.05\pi$
Lax–Wendroff	−696.0	−38.9
QUICK	−413.6	−181.1
QUICKEST	−243.8	−101.6
Upwind	166.5	167.0
MINMOD	42.8	36.2
Van Leer	21.3	13.5
MUSCL	15.3	7.3
ULTIMATE-QUICKEST	11.7	4.6
Superbee	7.6	−1.5
SHARP	3.3	1.6
Super-C	1.9	−3.3
Hyper-C	−10.8	−14.6
BPEP1	**0.1**	**0.0**
BPEP2	7.6	**0.0**

5. COST AND ACCURACY

As shown in Figure 15, the background potential energy oscillates about the x-axis, and hence the background potential energy is actually conserved in an average sense. The advantage of the BPEP2 scheme is that, because the Superbee scheme is not as compressive as the Hyper-C scheme, the background potential energy does not oscillate as drastically as it does for the BPEP1 scheme. As mentioned in Section 2.2, conservation of the background potential energy in this way incurs an added computational expense because the background potential energy is computed at every time step. The total computation time on an Intel Pentium 4 1.80 GHz cpu for the BPEP1 scheme is 192 s, of which 81% is spent on the pressure-Poisson equation, 13% on computation of the background potential energy, and the remaining 6% on advection and other computations, which indicates that the BPEP schemes incur an extra 15% computation time. It is possible to reduce the added expense by decreasing the frequency with which the background potential energy is computed. This is done by computing the background potential energy once every fixed number of time steps, rather than at every time step, and hence reducing the associated expense of computing the background potential energy in proportion to the number of skipped time steps. The drawback to not computing the background potential energy as frequently is that this effectively causes the BPEP schemes to revert to one of the limiters for a fixed number of time steps, regardless of the state of the background potential energy. As a result, the more infrequently it is computed, the more the background potential energy deviates from its initial value. This is depicted in Figure 16, which depicts the mean-square departure of the background potential energy from its initial value for $k\delta = 0.05\pi$ for both BPEP schemes. This error is given by

$$L^2_{E_\mathrm{b}} = \frac{1}{T}\int_0^T \left(\frac{E_\mathrm{b}(t) - E_{\mathrm{b}0}}{E_{\mathrm{a}0}}\right)^2 \mathrm{d}t \tag{53}$$

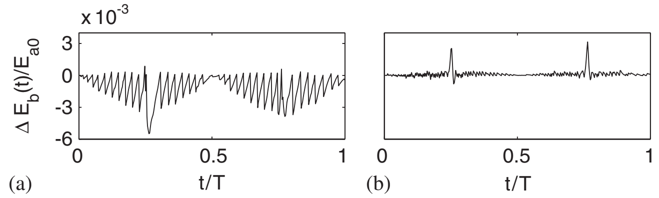

Figure 15. Zoomed in views of the background potential energy for the: (a) BPEP1; and (b) BPEP2 schemes for $k\delta = 0.05\pi$, showing how the background potential energy oscillates about the x-axis.

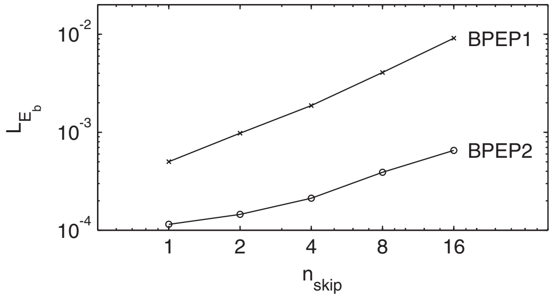

Figure 16. Mean-square departure of the background potential energy from its initial value as a function of the number of steps over which computation of the background potential energy is skipped. The wave Courant number for this case is $C_{\mathrm{w}} = 0.04$.

where T is the period of oscillation. As expected, the BPEP2 scheme yields mean-square departures of the background potential energy that are an order of magnitude less than the BPEP1 scheme. The added expense associated with computing the background potential energy can be reduced by a factor of 16, but this results in an increase in the mean-square departure of the background potential energy $L_{E_{\mathrm{b}}}$ by an order of magnitude when using the BPEP1 scheme. However, doing so for the BPEP2 scheme causes in increase in the mean-square departure by only a factor of 5.

When the background potential energy is computed for every time step, the mean-square departure in Equation (53) converges to first order in time at best, as shown in Figure 17, despite the second-order accuracy of the underlying solvers. In this figure, the time is non-dimensionalized and convergence is plotted as a function of the wave Courant number C_{w} from Equation (44). First-order convergence is to be expected, since the background potential energy is not a continuous function when the BPEP schemes are used. When the wave Courant number increases to roughly 0.1, the BPEP1 scheme maintains the background potential energy to less than 1% of its initial value, and the BPEP2 scheme does so to within 0.05%. Therefore, while the BPEP1 scheme is the most favourable overall scheme for the present simulations, if the interface thickness is not too small relative to the vertical grid spacing, it is best to use the BPEP2 scheme since it can maintain the background potential energy with a tighter tolerance.

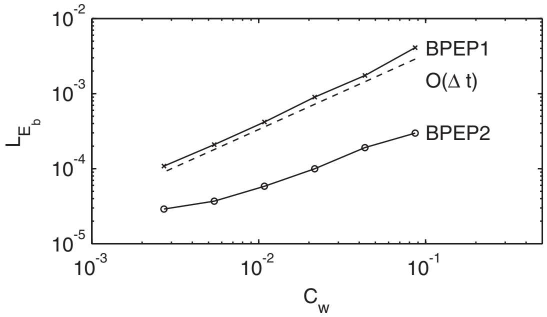

Figure 17. Temporal convergence of the background potential energy when the background potential energy is computed at every time step. The wave Courant number, C_w, varies by changing the time step and keeping all other parameters fixed.

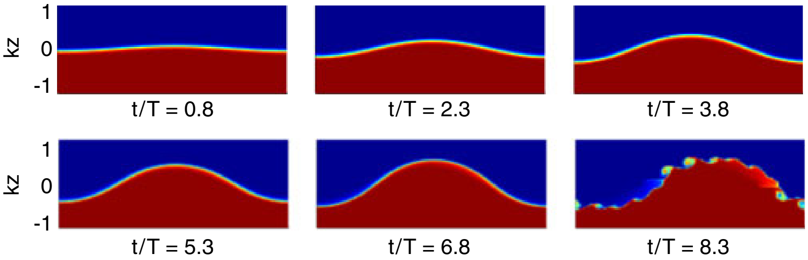

Figure 18. Evolution of a two-dimensional breaking interfacial wave with nondimensional interface thickness $k\delta = \pi/10$, with T being the wave period. The wave propagates from left to right.

6. APPLICATION TO A BREAKING PROGRESSIVE INTERFACIAL GRAVITY WAVE

As an example of an application of the BPEP schemes, we apply the BPEP2 scheme to the generation and breaking of progressive interfacial gravity waves, details of which can be found in the work of Fringer and Street [15]. In that work, periodic interfacial gravity waves are generated by applying a forcing function to the horizontal momentum equation until the waves reach their critical breaking steepness. An example of a breaking interfacial wave with interface thickness $k\delta = \pi/10$ is shown in Figure 18. In this simulation, the domain is similar to that shown in Figure 2, except the width is 0.2 m and the depth is 0.3 m, and the total number of grid cells is 256^2. Prior to application of the forcing function, the initial density field is given by Equation (43) with $ka = 0.0$ and $\Delta\rho/\rho_0 = 0.03$, and the velocity field is quiescent. In Figure 18, the density field is plotted in a frame that moves with the wave, but in the simulation the wave propagates through the periodic domain roughly eight times before breaking begins. The growth rate is kept to a minimum in order to prevent transient oscillations from contaminating the solution. During wave growth, the salt diffusivity is set to zero and the BPEP2 scheme is employed to reduce the effects of numerical diffusion on the interface to keep the nondimensional interface thickness as close to its initial value as possible as well as to prevent the background potential energy from increasing during wave growth. Upon wave breaking, the salt diffusivity is set to its molecular value and the SHARP scheme of Leonard [19] is then employed. Maintenance of the background potential energy during wave growth is essential for these simulations because the breaking criterion is dependent on a realistic value of the potential energy at the time of breaking.

7. CONCLUSIONS

Advection schemes do not in general maintain the background potential energy of a standing interfacial gravity wave because they are either compressive, diffusive, or oscillatory. Compressive schemes cause a reduction in the interface thickness because of steepening of the density profile, and as a result lead to a decrease in the background potential energy. Diffusive schemes, on the other hand, thicken the interface and cause an increase in the background potential energy. Oscillatory schemes always lead to a decrease in the background potential energy because they are effectively antidiffusive. However, dispersive schemes can incur a decrease in the background potential energy, followed by an almost equal and opposite increase, if oscillations induced during one cycle of wave propagation are canceled out during another cycle of wave oscillation.

By combining the characteristics of diffusive and antidiffusive schemes, we have developed schemes that conserve the background potential energy to a specified tolerance. The first scheme, BPEP1, employs a combination of the Hyper-C universal limiter and the Van Leer TVD limiter to maintain the background potential energy to within 1% of its initial value for flows with wave Courant numbers less than 0.1. It maintains the background potential energy for a wide interface thickness range at the risk of generating a slight nonmonotonicity in the advected quantity. The other scheme, BPEP2, guarantees monotonicity of the advected quantity, but does not maintain the background potential energy for thin interfaces because it employs the Superbee limiter, which must employ numerical diffusion to maintain monotonicity of sharp fronts. If the interface thickness is large enough, the BPEP2 scheme is favourable, especially for large Courant numbers, since it can maintain the background potential energy by an order of magnitude less than the BPEP1 scheme.

Use of the BPEP schemes requires computation of the background potential energy and hence incurs roughly 15% extra computation time. This added expense can be reduced by lowering the frequency of computation of the background potential energy. For example, computing the background potential energy every 16 time steps reduces the added computation time associated with the BPEP schemes to less than 1%. However, this results in an increase in the mean-square departure of the background potential energy from its initial value by an order of magnitude for the BPEP1 scheme and a factor of 5 for the BPEP2 scheme.

The present scheme reduces the change in the background potential energy by adjusting the limiter in the formulation of the advection scheme. This is in contrast to the scheme of Laval *et al.* [7], in which the background potential energy is identically preserved with the use of a pycnocline filter that is independent of the advection scheme. That scheme entails an iterative calculation of the length scale required of the filter that removes the background potential energy added during the advection step and conserves mass to a specified tolerance (since the filter does not conserve mass). Laval *et al.* suggest that the expense associated with this iterative technique (which requires roughly 20 computations of E_b per time step) can be reduced by refreshing the sharpening filter less often, much in the same way we have outlined the method by which the cost of the BPEP schemes can be reduced by computing E_b less often. In principle, a BPEP formulation could be constructed that determines the appropriate limiter that identically conserves the background potential energy by blending two limiters, rather than switching between them [26]. For example, if Ψ_1 is a compressive limiter, while Ψ_2 is a diffusive limiter, then the 'blended' BPEP scheme would employ the limiter $\Psi = \gamma\Psi_1 + (1-\gamma)\Psi_2$ $(0 \leqslant \gamma \leqslant 1)$, where the value of γ that identically conserves E_b could be determined via a Newton iteration at each time step. This scheme is the subject of future work in which a background potential energy preserving scheme is being developed that preserves E_b by altering the limiter locally, rather than globally.

In general, it has been shown that the background potential energy is an excellent measure of the effective numerical diffusion or antidiffusion of an advection scheme. Physically speaking, it must remain constant in the absence of thermal or salt diffusivity or mass fluxes into and out of the domain. The second-order TVD schemes have been shown to either increase or decrease the background potential energy, depending on the particular nature of the limiter involved. Likewise, schemes that employ the universal limiter can also cause an increase or decrease in the background potential energy, yet they do not guarantee monotonicity for very sharp fronts, especially when the velocity field is highly oscillatory. However, the presence of slight nonmonotonicity has been shown to produce good behaviour, especially for the BPEP1 scheme. Moreover, despite its slight nonmonotonicity, the SHARP scheme has been shown to produce the best results overall, second to the BPEP1 scheme. It causes an increase in the background potential energy by only 1.6% for the thin interface case.

ACKNOWLEDGEMENTS

The authors would like to thank B. P. Leonard for his extensive and insightful comments, as well as two anonymous reviewers, whose comments also substantially improved the quality of this paper. OBF is grateful for his support as a DOE Computational Science Graduate Fellow, his support by NSF Grant ITR-0113111 (Program Director: Dr E. Itsweire), and his support as a South Africa Teaching Fellow at the University of the Western Cape. SWA acknowledges support under NSF Grant ATM-0073395 (Program Director: R. Rogers). RLS appreciates his support under the NSF grants and by ONR Grant N0014-99-1-0413 (Scientific Officer: S. Murray).

REFERENCES

1. Hodges BR, Street RL. On simulation of turbulent nonlinear free-surface flows. *Journal of Computational Physics* 1999; **151**:425–457.
2. Harlow FH, Welch JE. Numerical calculation of time-dependent viscous incompressible flow of a fluid with free surface. *Physics of Fluids* 1965; **8**:2182–2189.
3. Puckett EG, Almgren AS, Bell JB, Marcus DL, Rider WJ. A high-order projection method for tracking fluid interfaces in variable density incompressible flows. *Journal of Computational Physics* 1997; **130**:269–282.
4. Chen G, Kharif C, Zaleski S, Li J. Two-dimensional Navier–Stokes simulation of breaking waves. *Physics of Fluids* 1999; **11**:121–133.
5. Enright D, Fedkiw R, Ferziger J, Mitchell I. A hybrid particle level set method for improved interface capturing. *Journal of Computational Physics* 2002; **183**:83–116.
6. Winters KB, Lombard PN, Riley JJ, D'Asaro EA. Available potential energy and mixing in density-stratified flows. *Journal of Fluid Mechanics* 1995; **289**:115–128.
7. Laval B, Hodges BR, Imberger J. Reducing numerical diffusion effects with pycnocline filter. *Journal of Hydraulic Engineering* (ASCE) 2003; **129**:215–224.
8. Leonard BP. The ULTIMATE conservative difference scheme applied to unsteady one-dimensional advection. *Computer Methods in Applied Mechanics* 1991; **88**:17–74.
9. Hodges BR. Numerical techniques in CWR-ELCOM. *Technical Report No. WP 1422-BH*, Centre for Water Research, 2000.

10. Zang Y, Street RL, Koseff JR. A non-staggered grid, fractional step method for time-dependent incompressible Navier–Stokes equations in curvilinear coordinates. *Journal of Computational Physics* 1994; **114**:18–33.
11. Armfield SW, Street RL. Fractional step methods for the Navier–Stokes equations on non-staggered grids. *ANZIAM Journal* 2000; **42**(E):C134–C156.
12. Leonard BP. A stable and accurate convective modelling procedure based on quadratic upstream interpolation. *Computer Methods in Applied Mechanics* 1979; **19**:59–98.
13. Tseng YH, Ferziger JH. Mixing and available potential energy in stratified flows. *Physics of Fluids* 2001; **13**:1281–1293.
14. Leonard BP, Lock AP, MacVean MK. Conservative explicit unrestricted-time-step multidimensional constancy-preserving advection schemes. *Monthly Weather Review* 1996; **124**:2588–2606.
15. Fringer OB, Street RL. The dynamics of breaking progressive interfacial waves. *Journal of Fluid Mechanics* 2003; **494**:319–353.
16. Gross ES, Koseff JR, Monismith SG. Evaluation of advective schemes for estuarine salinity simulations. *Journal of Hydraulic Engineering* 1999; **125**:32–46.
17. Leonard BP. Note on the von Neumann stability of the explicit FTCS convective diffusion equation. *Applied Mathematical Modelling* 1980; **4**:401–402.
18. Paolucci S, Chenoweth DR. Stability of the finite difference transport equation. *Journal of Computational Physics* 1982; **47**:489–496.
19. Leonard BP. Simple high-accuracy resolution program for convective modelling of discontinuities. *International Journal for Numerical Methods in Fluids* 1988; **8**:1291–1318.
20. Argrow BM. Computational analysis of dense gas shock tube flow. *Shock Waves* 1996; **6**:241–248.
21. Roe PL. Generalized formulation of TVD Lax–Wendroff schemes. *Technical Report*, NASA Langley Research Center, 1984. *ICASE Report 84-53*. NASA CR-172478.
22. Tamamidis P, Assanis DN. Evaluation of various high-order-accuracy schemes with and without flux limiters. *International Journal for Numerical Methods in Fluids* 1993; **16**:931–948.
23. Hirsch C. Numerical computation of internal and external flows. *Computational Methods for Inviscid and Viscous Flows*, vol. 2. Wiley: New York, 1990.
24. Thorpe SA. On standing internal gravity waves of finite amplitude. *Journal of Fluid Mechanics* 1968; **32**:489–528.
25. Phillips OM. *The Dynamics of the Upper Ocean*. Cambridge University Press: Cambridge, 1977.
26. Leonard BP. Personal communication, 2004.