

The RAW Filter: An Improvement to the Robert–Asselin Filter in Semi-Implicit Integrations

PAUL D. WILLIAMS

Department of Meteorology, University of Reading, Reading, United Kingdom

(Manuscript received 20 August 2010, in final form 29 November 2010)

ABSTRACT

Errors caused by discrete time stepping may be an important component of total model error in contemporary atmospheric and oceanic simulations. To reduce time-stepping errors in leapfrog integrations, the Robert–Asselin–Williams (RAW) filter was proposed by the author as a simple improvement to the widely used Robert–Asselin (RA) filter. The present paper examines the behavior of the RAW filter in semi-implicit integrations. First, in a linear theoretical analysis, the stability and accuracy are interrogated by deriving analytic expressions for the amplitude errors and phase errors. Then, power-series expansions are used to interpret the leading-order errors for small time steps and hence to identify optimal values of the filter parameters. Finally, the RAW filter is tested in a realistic nonlinear setting, by applying it to semi-implicit integrations of the elastic pendulum equations. The results suggest that replacing the RA filter with the RAW filter could reduce time-stepping errors in semi-implicit integrations.

1. Introduction

Contemporary atmospheric and oceanic numerical simulations are typically unconverged as the time step is reduced. For example, Williamson and Olson (2003) have found that an aqua-planet atmospheric general circulation model is sensitive to the time step of the parameterized physics. In particular, different time steps produce different zonally averaged precipitation patterns, such that the common double intertropical convergence zone (ITCZ) problem appears to be fixed by a simple reduction of the time step. Furthermore, Pfeffer et al. (1992) and Zhao and Zhong (2009) have found that different time-stepping schemes in atmospheric general circulation models produce substantially different climates. Therefore, errors caused by discrete time stepping may be an important component of total model error, as suggested by Teixeira et al. (2007).

In the atmospheric and oceanic sciences, the leapfrog time-stepping scheme remains in wider use than the more accurate alternatives, such as the Adams–Bashforth schemes (e.g., Durran 1991) and the Runge–Kutta schemes

(e.g., Kar 2006). The leapfrog scheme exhibits a well-known spurious computational mode (e.g., Lilly 1965; Young 1968; Mesinger and Arakawa 1976; Haltiner and Williams 1980; Durran 1999; Kalnay 2003), which may grow unphysically in nonlinear integrations. Unfortunately, many proposed strategies for controlling the computational mode, including the filtered leapfrog–trapezoidal scheme (Dietrich and Wormeck 1985) and the weighted filtered leapfrog–trapezoidal scheme (Roache and Dietrich 1988), incur a reduction in formal accuracy.

The computational mode of the leapfrog scheme may be controlled without loss of second-order accuracy, by alternating between leapfrog steps and other second-order steps, as in the leapfrog–trapezoidal scheme (Kurihara 1965) and the leapfrog–Adams–Bashforth scheme (Magazenkov 1980). Traditionally, however, the preferred strategy has been to apply a Robert–Asselin filter (Robert 1966; Asselin 1972) after each leapfrog step. Although in wide use (e.g., Schlesinger et al. 1983; Tandon 1987; Ford 1994; Caya and Laprise 1999; Griffies et al. 2001; Bartello 2002; Cordero and Staniforth 2004; Fraedrich et al. 2005; Hartogh et al. 2005; Williams et al. 2009), the Robert–Asselin filter is associated with some counterintuitive behavior (e.g., Déqué and Cariolle 1986; Robert and Lépine 1997) and a reduction to first-order accuracy.

Corresponding author address: Paul D. Williams, Department of Meteorology, University of Reading, Earley Gate, Reading RG6 6BB, United Kingdom.
E-mail: p.d.williams@reading.ac.uk

Williams (2009) proposed a simple modification to the Robert–Asselin filter, designed to preserve the second-order accuracy in leapfrog integrations. To understand the modification, suppose that the filter is to operate on three time-consecutive values of a dependent variable, denoted x_{n-1} , x_n , and x_{n+1} . Typically, x_{n+1} will have been calculated from x_{n-1} and x_n using a leapfrog step, and x_{n-1} and x_n may already have been filtered after previous leapfrog steps. To shorten what follows, let the unit of displacement be $\nu(x_{n-1} - 2x_n + x_{n+1})/2$, where ν is a dimensionless filter parameter that is assumed to be small and positive. With this definition, the original Robert–Asselin filter displaces x_n through one unit but does not displace x_{n+1} . In contrast, the modified filter proposed by Williams (2009) simultaneously displaces x_n through α units and x_{n+1} through $\alpha - 1$ units, where α is a second dimensionless filter parameter satisfying $0 \leq \alpha \leq 1$. Hereafter, for brevity and following Amezcua et al. (2011), the original filter will be labeled the Robert–Asselin (RA) filter and the modified filter will be labeled the Robert–Asselin–Williams (RAW) filter.

For the special case of $\alpha = 1$, the RAW filter is identical to the RA filter. For the special case of $\alpha = 1/2$, the RAW filter perturbs x_n and x_{n+1} through equal and opposite displacements, conserving the three-time-level mean state, $(x_{n-1} + x_n + x_{n+1})/3$. When the latter special case replaces the former in leapfrog integrations, the phase accuracy remains second order but the amplitude accuracy increases from first to third order (Williams 2009). Therefore, the RAW-filtered leapfrog scheme with $\alpha = 1/2$ exactly matches the accuracy of the Kurihara (1965) and Magazenkov (1980) alternating schemes (Durran 1999), but benefits from being nonalternating, less cumbersome, more intuitive, and simpler to implement. The upgrade from RA filtering to RAW filtering in an existing computer code is trivial (Fig. 1). The RAW filter is currently being tested in various general circulation models of the atmosphere and ocean, and the results will be reported in due course.

Williams (2009) studied the RAW filter in explicit leapfrog integrations. However, practical applications of the leapfrog scheme in meteorological and oceanographic models are usually accompanied by the semi-implicit modification (Robert 1969; Kwizak and Robert 1971; Staniforth 1997). The semi-implicit discretization method uses the explicit leapfrog scheme for the advection terms and the implicit trapezoidal scheme for the pressure gradient and divergence terms. In numerical analysis, this technique falls into the category of implicit–explicit (IMEX) methods. The implicitly treated terms are crucial for high-frequency gravity wave oscillations, which are often of secondary importance and, which, if treated explicitly, would violate the Courant–Friedrichs–Lewy

```

! Compute tendency at this time step
x_tend = [...]

! Perform leapfrog step
x_next = x_last + 2*delta_t*x_tend

! Compute filter displacement
d = nu*(x_last - 2*x_this + x_next)/2

! Apply filter
x_this = x_this + d*alpha
x_next = x_next + d*(alpha-1)

```

FIG. 1. The essential time-stepping lines of a typical leapfrog computer code, imagined to appear inside a time-stepping loop. The unshaded code performs an explicit leapfrog step and applies the RA filter. The upgrade from RA filtering to RAW filtering is achieved via the trivial insertion of the shaded code.

(CFL) stability condition unless impractically short time steps were taken.

Semi-implicit numerical schemes are used widely in practical applications, because they suffer from neither the computational expense of fully implicit schemes nor the demanding time step requirement of fully explicit schemes (although the latter requirement is partly alleviated with time-split methods; e.g., Skamarock and Klemp 1992). Therefore, strategies for improving semi-implicit integrations are highly sought after by modelers. Note, however, that the third-order Adams–Bashforth scheme is known to be an unsuitable substitute for the leapfrog scheme in semi-implicit integrations (Durran 1991).

Motivated by the above introductory discussion, the goal of the present paper is to study the RAW filter in semi-implicit integrations. First, a linear theoretical analysis of the RAW-filtered semi-implicit scheme is presented (section 2). The stability and accuracy are interrogated by deriving analytic expressions for the amplitude errors and phase errors (section 2a). Power-series expansions are used to interpret the leading-order errors for small time steps and hence to identify optimal values of the filter parameters (section 2b). Finally, the RAW filter is tested in a realistic nonlinear setting, by applying it to semi-implicit integrations of the elastic pendulum equations (section 3). The paper concludes with a summary and discussion (section 4).

2. Theoretical analysis

a. Derivation of amplitude errors and phase errors

Semi-implicit time-stepping schemes are traditionally examined by applying them to the split-frequency oscillation equation for the complex variable $x(t)$,

$$\frac{dx}{dt} = i\omega_{\text{low}}x + i\omega_{\text{high}}x, \quad (1)$$

where $i = \sqrt{-1}$ and ω_{low} and ω_{high} are given (real) angular frequencies (e.g., Durran 1991, 1999). The semi-implicit numerical approximation to (1) uses the explicit leapfrog scheme to discretize the ω_{low} term and the implicit trapezoidal (Crank–Nicolson) scheme to discretize the ω_{high} term. With time step Δt , and using the RAW filter to control the computational mode, the semi-implicit numerical scheme is

$$\frac{x(t + \Delta t) - \bar{x}(t - \Delta t)}{2\Delta t} = i\omega_{\text{low}}\bar{x}(t) + i\omega_{\text{high}}\left[\frac{x(t + \Delta t) + \bar{x}(t - \Delta t)}{2}\right], \quad (2)$$

$$\bar{x}(t) = \bar{x}(t) + \frac{\nu\alpha}{2}[\bar{x}(t - \Delta t) - 2\bar{x}(t) + x(t + \Delta t)], \quad \text{and} \quad (3)$$

$$\bar{x}(t + \Delta t) = x(t + \Delta t) + \frac{\nu(\alpha - 1)}{2}[\bar{x}(t - \Delta t) - 2\bar{x}(t) + x(t + \Delta t)]. \quad (4)$$

In this three-stage method, (2) applies a semi-implicit time step to calculate $x(t + \Delta t)$ from $\bar{x}(t - \Delta t)$ and $\bar{x}(t)$, and (3) and (4) apply the RAW filter with parameters ν and α . In (2)–(4), x denotes an unfiltered provisional value, \bar{x} denotes a singly filtered provisional value, and $\bar{\bar{x}}$ denotes the doubly filtered definitive value.

To proceed with the analysis, the complex amplification factor A is defined by

$$A = \frac{x(t + \Delta t)}{x(t)} = \frac{\bar{x}(t + \Delta t)}{\bar{x}(t)} = \frac{\bar{\bar{x}}(t + \Delta t)}{\bar{\bar{x}}(t)}. \quad (5)$$

Rewriting (2)–(4) with function evaluations at time t only, using (5), yields a homogeneous matrix equation for the vector $[x(t), \bar{x}(t), \bar{\bar{x}}(t)]$. For nontrivial solutions, the determinant of the 3×3 matrix of coefficients must vanish, yielding

$$aA^2 + bA + c = 0, \quad (6)$$

where

$$a = 1 - i\omega_{\text{high}}\Delta t, \quad (7)$$

$$b = -\nu + \nu(1 - \alpha)i\omega_{\text{high}}\Delta t + [\nu(1 - \alpha) - 2]i\omega_{\text{low}}\Delta t, \quad (8)$$

$$c = \nu - 1 - (1 - \nu\alpha)i\omega_{\text{high}}\Delta t + \nu\alpha i\omega_{\text{low}}\Delta t. \quad (9)$$

For the special case of $\nu = 0$, which corresponds to the unfiltered semi-implicit scheme, (6)–(9) reduce to the quadratic equation derived by Kwizak and Robert (1971), as expected. For the alternative special case of $\omega_{\text{high}} = 0$, which corresponds to the RAW-filtered explicit leapfrog scheme, (6)–(9) reduce to the quadratic equation derived by Williams (2009), as expected. For the general case, the quadratic equation yields two roots, $A_{\pm}(\nu, \alpha, i\omega_{\text{low}}\Delta t, i\omega_{\text{high}}\Delta t)$, where A_+ and A_- are the complex amplification factors for the physical and computational modes, respectively. Explicit formulas for A_+ and A_- are omitted here, for brevity. For comparison, the complex amplification factor for the exact solution is found from (1) to be

$$A_{\text{exact}}(i\omega_{\text{low}}\Delta t, i\omega_{\text{high}}\Delta t) = \exp[i(\omega_{\text{low}} + \omega_{\text{high}})\Delta t]. \quad (10)$$

A hypothetical, error-free numerical scheme would have $A_+ = A_{\text{exact}}$ and would lack computational modes. This perfect scheme would neither amplify nor dissipate the unforced and undamped oscillations modeled by (1), because $|A_{\text{exact}}| = 1$, and it would advance the phase of the oscillations at a rate equal to the sum of the two frequencies, because $\arg(A_{\text{exact}}) = (\omega_{\text{low}} + \omega_{\text{high}})\Delta t$. However, like all numerical schemes, the RAW-filtered semi-implicit scheme has $A_+ \neq A_{\text{exact}}$. Its amplitude and phase errors are characterized by the magnitude and argument, respectively, of A_+ relative to those of A_{exact} .

To assess the amplitude errors, an explicit formula for A_+ is obtained by solving the quadratic equation in (6)–(9). The variation of $|A_+| - |A_{\text{exact}}|$ with $\omega_{\text{low}}\Delta t$ and $\omega_{\text{high}}\Delta t$ will be discussed in detail in the following paragraphs, and is plotted in Fig. 2 for a typical value of ν and three values of α . Positive values of the plotted quantity indicate numerical amplification of the physical mode, and negative values indicate numerical damping. Note that the plotted quantity is typically at least an order of magnitude smaller for $\alpha \approx 1/2$ than for $\alpha \approx 1$, across the range of $\omega_{\text{low}}\Delta t$ and $\omega_{\text{high}}\Delta t$ shown.

The leading-order behavior on the fully explicit axis ($\omega_{\text{high}}\Delta t = 0$) was studied by Williams (2009). For $\alpha = 1$, corresponding to the original RA filter, strong (quadratic) damping yields stability but poor accuracy. In contrast, for $\alpha = 1/2$, weak (quartic) amplification yields good accuracy but weak instability. Neither of these outcomes is especially desirable, although physical damping in dissipative systems may control any weak numerical instability. Worse still, for $\alpha = 0$ (not shown), strong (quadratic) amplification yields both instability and poor accuracy. However, for $\alpha \gtrsim 1/2$, competition between quadratic damping and quartic amplification yields both stability and good accuracy, over a tunable finite range of frequencies that widens with increasing α . Consequently, for fully explicit simulations of unforced and

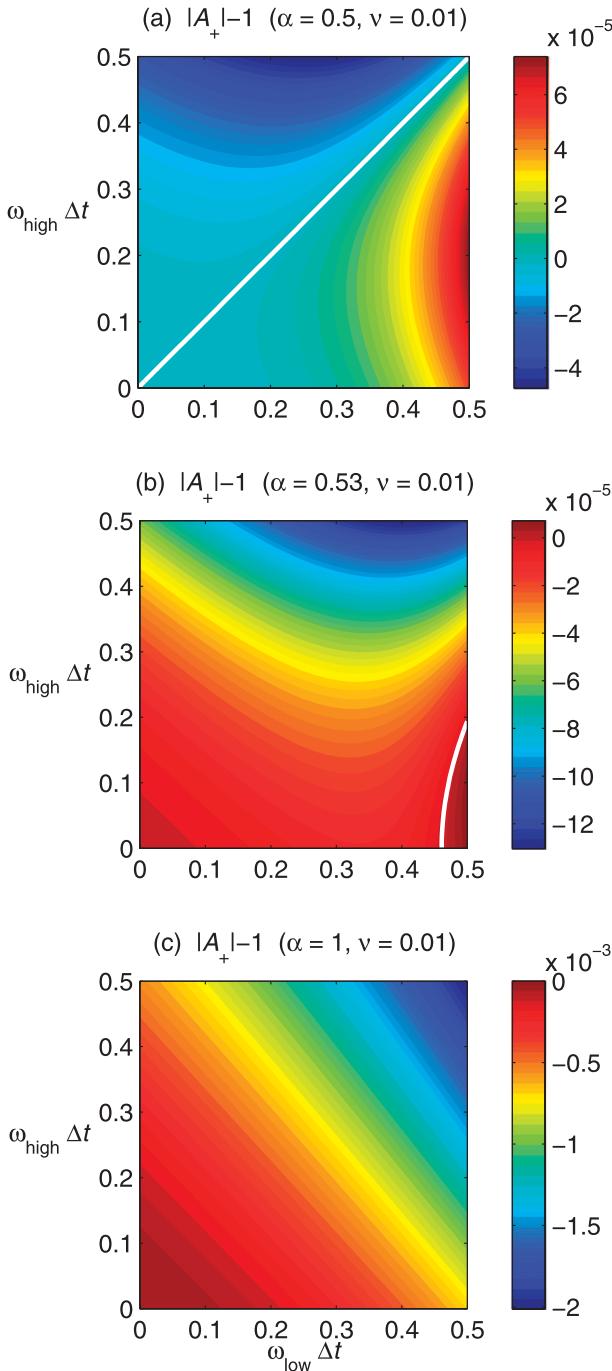


FIG. 2. Magnitude of the amplification factor for the physical mode of the RAW-filtered semi-implicit numerical scheme ($|A_+|$) relative to that for the exact solution ($|A_{\text{exact}}| = 1$). The scheme is applied to the split-frequency complex oscillation equation in (1), with the ω_{low} term treated explicitly and the ω_{high} term treated implicitly. The values of the RAW filter parameters are (a) $\alpha = \frac{1}{2}$, (b) $\alpha = 0.53$, and (c) $\alpha = 1$, each with $\nu = 0.01$. White contours denote the curves of neutral stability ($|A_+| - 1 = 0$). Note the use of different color maps: the plotted quantity is typically at least an order of magnitude smaller in (a) and (b) than in (c).

undamped oscillations, $\alpha \gtrsim \frac{1}{2}$ was recommended by Williams (2009).

The behavior on the fully implicit axis ($\omega_{\text{low}}\Delta t = 0$) is qualitatively different. For all values of $\alpha \gtrsim \frac{1}{2}$, damping of various strengths is obtained at all points on the axis, and amplification is never obtained. Therefore, arbitrarily long time steps may be taken without risking numerical instability of the implicitly treated mode. (Poor accuracy for the implicitly treated mode with long time steps is expected and uninteresting.) Unconditional linear stability is the desirable characteristic of implicit schemes, and it is preserved by the RAW filter.

The behavior in the semi-implicit plane ($\omega_{\text{high}}\Delta t > 0$, $\omega_{\text{low}}\Delta t > 0$) is arguably of greater practical interest than the behavior on either of the axes. For all values of α , the $\omega_{\text{low}}\Delta t$ coordinate of the neutral stability curve increases as $\omega_{\text{high}}\Delta t$ increases, and so the finite stable range of explicitly treated frequencies widens. For example, referring to Fig. 2b, the fully explicit case of $(\omega_{\text{low}}\Delta t, \omega_{\text{high}}\Delta t) = (0.5, 0)$ is unstable but the semi-implicit case of $(\omega_{\text{low}}\Delta t, \omega_{\text{high}}\Delta t) = (0.5, 0.5)$ is stable. Therefore, the inclusion of the implicitly treated mode appears to stabilize the numerical scheme.

The behavior in the top-left triangle of the semi-implicit plane ($\omega_{\text{high}}\Delta t \geq \omega_{\text{low}}\Delta t$) is of special practical interest. For all values of $\alpha \gtrsim \frac{1}{2}$, damping of various strengths is obtained at all points in the triangle, and amplification is never obtained. Therefore, if the frequency of the implicitly treated mode equals or exceeds the frequency of the explicitly treated mode, arbitrarily long time steps may be taken without risking numerical instability. The strength of the damping depends sensitively upon α , with $\alpha = \frac{1}{2}$ generally being much less dissipative than $\alpha = 1$. Indeed, for the case $\omega_{\text{high}}\Delta t = \omega_{\text{low}}\Delta t$, the amplitude of the oscillations is preserved exactly with $\alpha = \frac{1}{2}$.

To assess the phase errors, the variation of $\arg(A_+) - \arg(A_{\text{exact}})$ with $\omega_{\text{low}}\Delta t$ and $\omega_{\text{high}}\Delta t$ is plotted in Fig. 3. Positive values of the plotted quantity indicate numerical phase advancement of the physical mode, and negative values indicate numerical phase recession. The structure and magnitude of the plotted quantity appear indistinguishable for all values of α , indicating that phase errors, unlike amplitude errors, are relatively insensitive to α . The curve of zero phase error appears to be approximately $\omega_{\text{high}}\Delta t = 0.5\omega_{\text{low}}\Delta t$. The region between the curve of zero phase error and the fully implicit axis exhibits numerical phase recession, and the region between the curve of zero phase error and the fully explicit axis exhibits numerical phase advancement. The curves of zero amplitude error in Fig. 2 and the curves of zero phase error in Fig. 3 do not generally coincide.

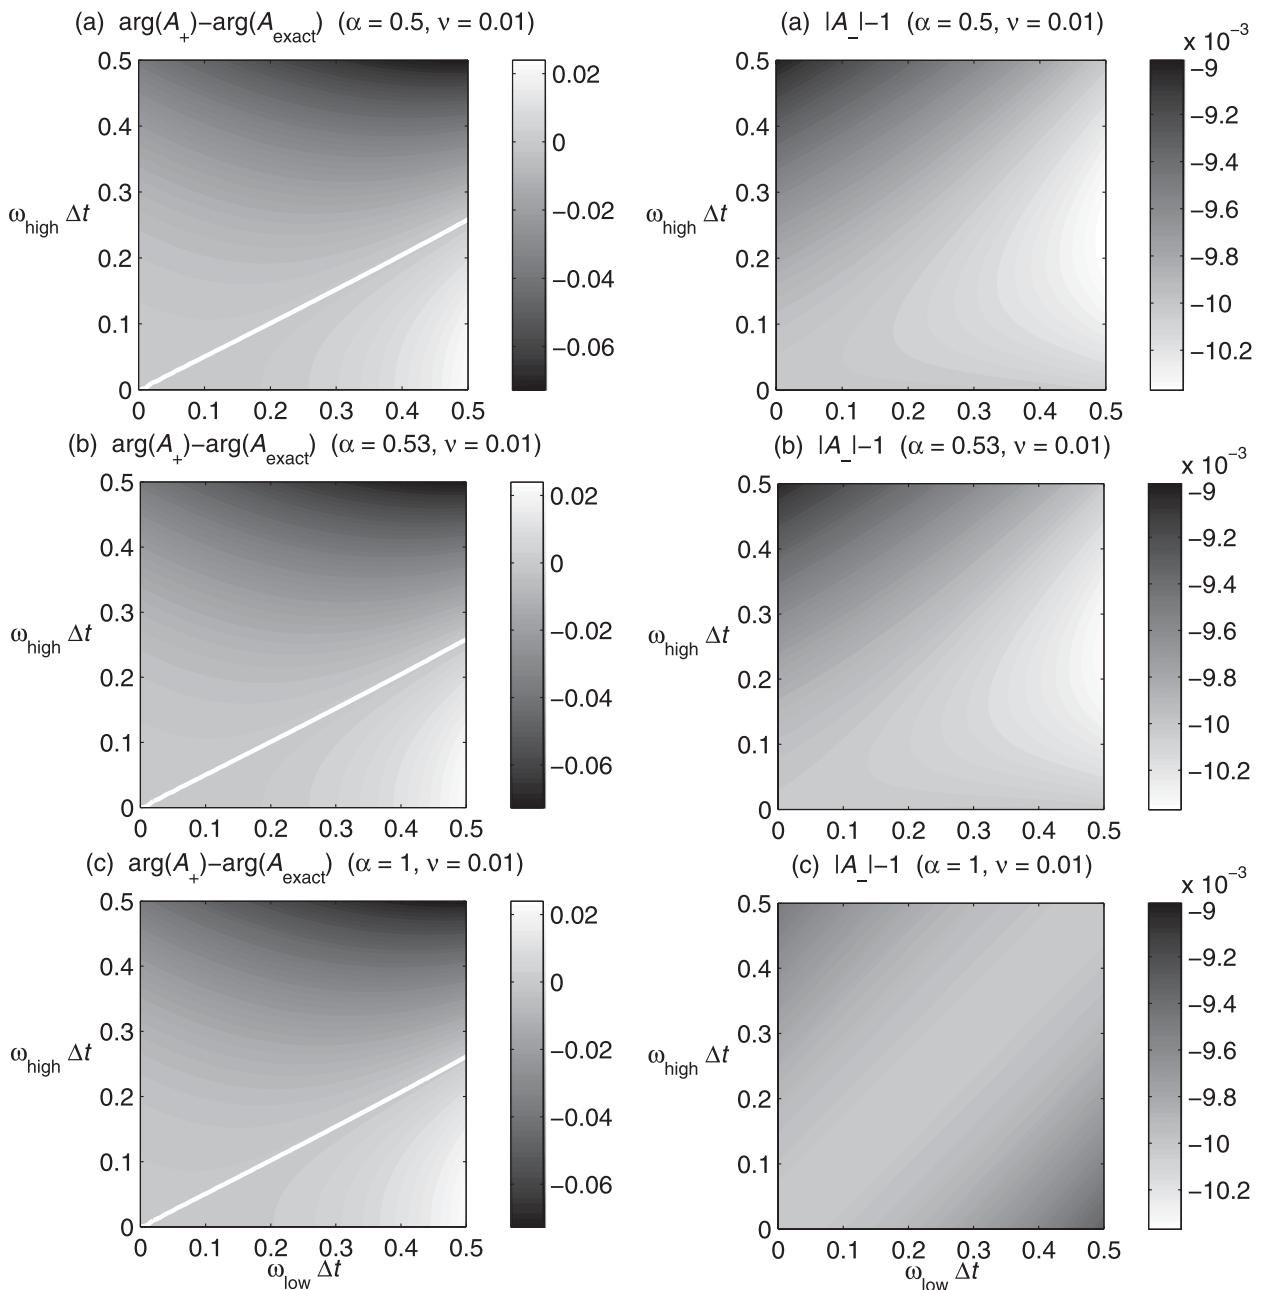


FIG. 3. Argument of the amplification factor for the physical mode of the RAW-filtered semi-implicit numerical scheme ($\arg A_+$) relative to that for the exact solution ($\arg A_{\text{exact}}$). The scheme is applied to the split-frequency complex oscillation equation in (1), with the ω_{low} term treated explicitly and the ω_{high} term treated implicitly. The values of the RAW filter parameters are (a) $\alpha = 1/2$, (b) $\alpha = 0.53$, and (c) $\alpha = 1$, each with $\nu = 0.01$. White contours denote the curves of zero phase error ($\arg A_+ - \arg A_{\text{exact}} = 0$).

Finally, to determine whether the computational mode is suppressed, the variation of $|A_-| - 1$ with $\omega_{\text{low}} \Delta t$ and $\omega_{\text{high}} \Delta t$ is plotted in Fig. 4. Positive values of the plotted quantity would indicate amplification of the computational

FIG. 4. Magnitude of the amplification factor for the computational mode of the RAW-filtered semi-implicit numerical scheme ($|A_-| - 1$) relative to unity. The scheme is applied to the split-frequency complex oscillation equation in (1), with the ω_{low} term treated explicitly and the ω_{high} term treated implicitly. The values of the RAW filter parameters are (a) $\alpha = 1/2$, (b) $\alpha = 0.53$, and (c) $\alpha = 1$, each with $\nu = 0.01$.

mode, and negative values would indicate damping. For all values of α , damping is obtained at all points in the plane, and amplification is never obtained. Therefore, the unconditional linear stability of the computational mode is preserved by the RAW filter, as desired.

b. Leading-order behavior for small time steps

As Δt decreases continuously to zero, transects of constant $\omega_{\text{high}}/\omega_{\text{low}}$ are traversed near the origin of the $(\omega_{\text{low}}\Delta t, \omega_{\text{high}}\Delta t)$ plane. To better understand the stability and accuracy of the physical mode in this limit, the leading-order behavior along these transects will now be derived. First, $\omega_{\text{high}} = r\omega_{\text{low}}$ is substituted into the expressions (7)–(9) for the coefficients of the quadratic equation in (6). Then, series-expansion solutions are sought for the magnitude and argument of $A_+(\nu, \alpha, r, i\omega_{\text{low}}\Delta t)$ in powers of $\omega_{\text{low}}\Delta t$, treating ν , α , and r as constants. The solutions are obtained using the symbolic manipulation capability of the Mathematica software (Wolfram Research 2008).

For the magnitude of A_+ , the power series obtained is

$$|A_+| = 1 + \frac{1}{4}\nu(1 - 2\alpha)(1 + r)^2(\omega_{\text{low}}\Delta t)^2 + \frac{1}{16}\nu[3 - 4\alpha - r(5 - 8\alpha)](1 + r)^3(\omega_{\text{low}}\Delta t)^4 + \mathcal{O}(\omega_{\text{low}}\Delta t)^6, \quad (11)$$

or, expressed equivalently as a power series in $\omega_{\text{high}}\Delta t$,

$$|A_+| = 1 + \frac{1}{4}\nu(1 - 2\alpha)\frac{(1 + r)^2}{r^2}(\omega_{\text{high}}\Delta t)^2 + \frac{1}{16}\nu[3 - 4\alpha - r(5 - 8\alpha)]\frac{(1 + r)^3}{r^4}(\omega_{\text{high}}\Delta t)^4 + \mathcal{O}(\omega_{\text{high}}\Delta t)^6. \quad (12)$$

The lowest-order terms in the power series are identical to $|A_{\text{exact}}| = 1$, as expected. In the coefficients of the quadratic and quartic terms, contributions $\mathcal{O}(\nu^2)$ and higher have been neglected, because $\nu \ll 1$ in practical applications. In this limit, the coefficients of the quadratic and quartic terms are each directly proportional to ν but depend nontrivially upon r and α . Therefore, the binary stability and the order of amplitude accuracy are each influenced only by r and α and not by ν (although the growth rate of any instability and the absolute amplitude error are each directly proportional to ν). The signs of the quadratic and quartic terms vary in the (r, α) plane as shown in Fig. 5. The lines $\alpha = 1/2$ and $\alpha = (3 - 5r)/(4 - 8r)$, along which the quadratic and quartic terms, respectively, are zero, intersect at the point $(r, \alpha) = (1, 1/2)$ and partition the plane into five regions with distinct stability and accuracy characteristics.

The stability characteristics will be considered first. In the region where the quadratic and quartic terms are each positive, the numerical scheme is unstable for the smallest time steps (at which the quadratic term dominates) and remains unstable for larger time steps (at

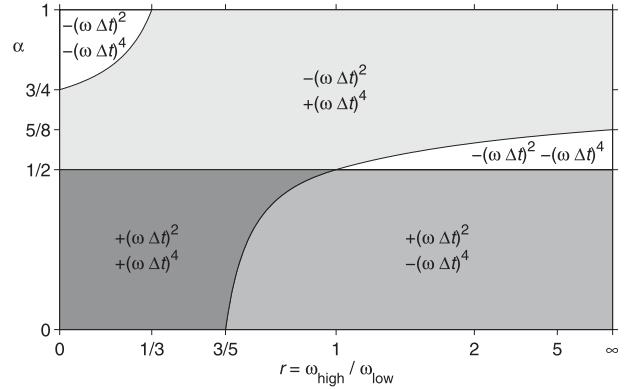


FIG. 5. Behavior in the (r, α) plane of the signs of the quadratic and quartic terms in the power series for $|A_+|$, given by (11) and (12) for $\nu \ll 1$. The lines $\alpha = 1/2$ and $\alpha = (3 - 5r)/(4 - 8r)$, along which the quadratic and quartic terms, respectively, are zero, partition the plane into five regions, which are shaded differently. No shading denotes a negative quadratic term and a negative quartic term, light shading denotes a negative quadratic term and a positive quartic term, medium shading denotes a positive quadratic term and a negative quartic term, and heavy shading denotes a positive quadratic term and a positive quartic term. The r axis is linear in r between $r = 0$ and $r = 1$ and linear in $1/r$ between $r = 1$ and $r = \infty$. The line $r = 0$ corresponds to fully explicit integrations and the line $r = \infty$ corresponds to fully implicit integrations. The line $\alpha = 1$ corresponds to the original RA filter.

which the quartic term dominates). In the region where the quadratic term is positive and the quartic term is negative, the scheme is unstable for the smallest time steps (at which the quadratic term dominates) but may be stabilized by the quartic term for a range of larger time steps (if the sextic term remains small). In the region where the quadratic term is negative and the quartic term is positive, the scheme is stable for the smallest time steps (at which the quadratic term dominates) but may be destabilized by the quartic term for a range of larger time steps (if the sextic term remains small). None of this stability behavior is especially desirable. In contrast, in the two regions where the quadratic and quartic terms are each negative, the scheme is stable for the smallest time steps (at which the quadratic term dominates) and remains stable for larger time steps (at which the quartic term dominates). Therefore, these two regions offer the best stability behavior.

The accuracy characteristics will be considered next. In most of the (r, α) plane, the leading-order amplitude errors for a single time step and for a given time interval vary with Δt quadratically and linearly, respectively, yielding first-order amplitude accuracy. The exception is the line $\alpha = 1/2$, along most of which the leading-order amplitude errors vary quartically and cubically, yielding third-order amplitude accuracy. The further exception is the point $(r, \alpha) = (1, 1/2)$, at which the leading-order

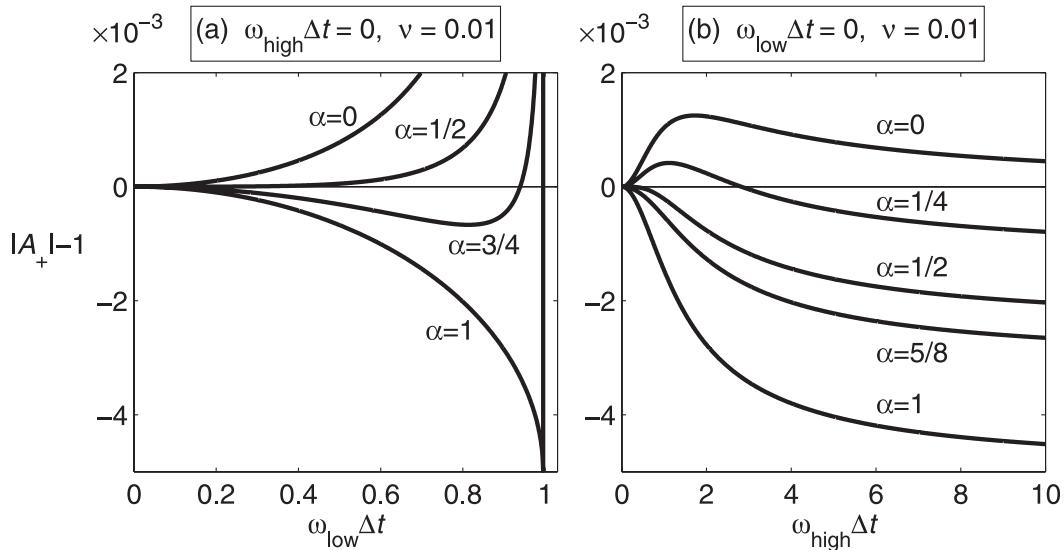


FIG. 6. Magnitude of the amplification factor for the physical mode ($|A_+|$) relative to that for the exact solution ($|A_{\text{exact}}| = 1$), for (a) fully explicit integrations and (b) fully implicit integrations. In the RAW filter, $\nu = 0.01$ and α takes various critical values suggested by inspection of Fig. 5 along (a) $r = 0$ and (b) $r = \infty$. The case $\alpha = 1$ corresponds to the original RA filter.

amplitude errors vary at worst sextically and quintically, yielding at least fifth-order amplitude accuracy. In fact, the amplitude errors are exactly zero at this point (Fig. 2a).

The above leading-order amplitude behavior was predicted from the truncated power-series expansions in (11) and (12). Nevertheless, it is a faithful representation of the unapproximated behavior, which is plotted in Fig. 6 using the full analytic solution to the quadratic equation in (6)–(9). Optimal values of α are determined by a tension between the needs for accuracy and stability. For fully explicit integrations ($\omega_{\text{high}}\Delta t = 0$) of unforced and undamped oscillations, the RAW filter with $\alpha = 3/4$ is perhaps the best compromise, because it offers better accuracy than the original RA filter (with $\alpha = 1$) and almost the same stability criterion. For fully implicit integrations ($\omega_{\text{low}}\Delta t = 0$) and typical semi-implicit integrations ($\omega_{\text{low}} \ll \omega_{\text{high}}$), the RAW filter with $\alpha = 1/2$ is perhaps the best compromise, because it offers better accuracy than the original RA filter (with $\alpha = 1$) and retains the unconditional stability. There is no case for which the optimal value is $\alpha = 1$.

Finally, for the argument of A_+ , the power series obtained is

$$\begin{aligned} \arg(A_+) &= (1+r)(\omega_{\text{low}}\Delta t) \\ &+ \frac{1}{6}(1-2r)(1+r)^2(\omega_{\text{low}}\Delta t)^3 + \mathcal{O}(\omega_{\text{low}}\Delta t)^5, \end{aligned} \tag{13}$$

or, expressed equivalently as a power series in $\omega_{\text{high}}\Delta t$,

$$\begin{aligned} \arg(A_+) &= \frac{(1+r)}{r}(\omega_{\text{high}}\Delta t) \\ &+ \frac{1}{6}\frac{(1-2r)(1+r)^2}{r^3}(\omega_{\text{high}}\Delta t)^3 + \mathcal{O}(\omega_{\text{high}}\Delta t)^5. \end{aligned} \tag{14}$$

The lowest-order terms in the power series are identical to $\arg(A_{\text{exact}}) = (\omega_{\text{low}} + \omega_{\text{high}})\Delta t$, as expected. In the coefficients of the cubic terms, contributions $\mathcal{O}(\nu)$ and higher have been neglected. In this limit, the coefficients of the cubic terms are dependent upon r but independent of ν and α , because all the dependency upon α enters through terms $\mathcal{O}(\nu)$ and higher. Therefore, the leading-order phase behavior is influenced only by r and not by ν or α . It follows that the phase behavior is insensitive to the RA and RAW filters. Numerical phase advancement is obtained for $r < 1/2$, and numerical phase recession for $r > 1/2$. For $r \neq 1/2$, the leading-order phase errors for a single time step and for a given time interval vary with Δt cubically and quadratically, respectively, yielding second-order phase accuracy. For $r = 1/2$, the leading-order phase errors vary quintically and quartically, yielding fourth-order phase accuracy.

3. Numerical test integrations

The RAW-filtered semi-implicit scheme will now be tested in a realistic nonlinear setting, by applying it to numerical integrations of the ordinary differential equations for an elastic pendulum confined to a vertical

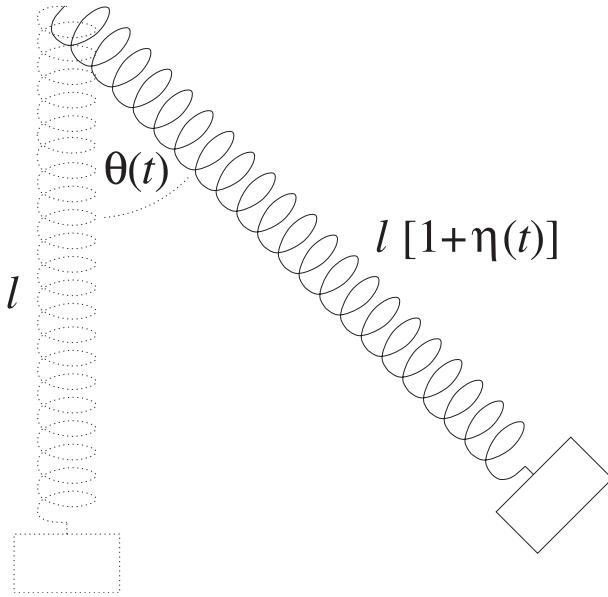


FIG. 7. Schematic diagram of the elastic pendulum, or swinging spring, showing the equilibrium position (dotted) and a non-equilibrium position defined by the values of η and θ (solid).

plane. Unlike the corresponding rigid pendulum, which exhibits only a rotational mode about the point of suspension, the elastic pendulum also exhibits a relatively fast vibrational mode, to which the relatively slow rotational mode is nonlinearly coupled. The elastic pendulum is known informally as the swinging spring, and its dynamics have been investigated extensively (e.g., Lynch 2002), motivated partly by an analogy with interactions in rotating fluids between slow balanced modes and fast inertia-gravity modes.

The elastic pendulum is shown schematically in Fig. 7. In the configuration to be studied here, a massless spring of unstretched length $l_0 = 1$ m and force constant $k = 100$ N m⁻¹ is loaded with a point mass of $m = 0.1$ kg in a gravitational field of acceleration $g = 10$ m s⁻². The spring may be stretched along its axis but not bent. The equilibrium length of the loaded spring is $l = l_0 + mg/k = 1.01$ m and the linear angular frequencies of the slow and fast modes are $\omega_{\text{low}} = \sqrt{g/l} \approx 3.1$ rad s⁻¹ and $\omega_{\text{high}} = \sqrt{k/m} \approx 31.6$ rad s⁻¹, respectively.

The dynamics are studied using cylindrical polar coordinates, with the origin at the point of suspension and the zero of azimuth aligned with the gravitational field. Writing the radial coordinate of the point mass as $l[1 + \eta(t)]$ and the azimuthal coordinate as $\theta(t)$, the state of the system is completely specified by the values of $\eta(t)$ and $\theta(t)$ and the values of their time derivatives, $\nu_\eta(t)$ and $\nu_\theta(t)$, respectively. The equilibrium position is $\eta = 0$ and $\theta = 0$. In the simulations to be studied here, the point mass is released from rest with $\eta = 0.01$ and $\theta = 1.0$

rad at time $t = 0$, and its evolution is tracked from $t = 0$ to $t = 10$ s.

The nonlinear equations of motion for $[\eta(t), \nu_\eta(t), \theta(t), \nu_\theta(t)]$ are

$$\dot{\eta} = \underline{\nu_\eta}, \tag{15}$$

$$\dot{\nu}_\eta = -\omega_{\text{low}}^2(1 - \cos\theta) - \underline{\omega_{\text{high}}^2 \eta} + (1 + \eta)\nu_\theta^2, \tag{16}$$

$$\dot{\theta} = \nu_\theta, \text{ and} \tag{17}$$

$$\dot{\nu}_\theta = \frac{-\omega_{\text{low}}^2 \sin\theta - 2\nu_\eta \nu_\theta}{1 + \eta}. \tag{18}$$

If (15)–(18) were linearized about the equilibrium position, then the slow and fast modes would decouple and $\theta(t)$ and $\eta(t)$ would each execute independent simple harmonic motion with angular frequencies ω_{low} and ω_{high} , respectively. The two underlined terms in (15) and (16) are crucial for the fast mode. Therefore, in the semi-implicit discretization used for the present integrations, the underlined terms are treated implicitly using the trapezoidal (Crank–Nicolson) scheme and all other terms are treated explicitly using the leapfrog scheme. Semi-implicit discretized equations for more complicated systems are generally solved by iteration, but here they can be rearranged to obtain explicit analytic expressions for the future state. Once the future state is thereby calculated, the RAW filter is applied to all four state variables.

The present numerical experiments use a time step of $\Delta t = 0.1$ s, giving $\omega_{\text{low}}\Delta t \approx 0.3$ and $\omega_{\text{high}}\Delta t \approx 3$. The experiments are found to be unstable with this time step when all the terms in (15)–(18) are treated explicitly, as expected because of the large value of $\omega_{\text{high}}\Delta t$. Two semi-implicit integrations will be studied, with $\alpha = 1$ and $\alpha = 1/2$, and each with $\nu = 0.2$. A single two-time-level forward step is used to initiate the three-time-level semi-implicit scheme. The two experiments will be compared to a reference numerical solution obtained using the semi-implicit scheme with $\Delta t = 10^{-6}$ s, $\nu = 0.2$, and $\alpha = 1/2$. The reference solution is found to be insensitive to all of these numerical choices, and so is taken to be a well-converged approximation to the exact solution for the purposes of the following comparison.

The results are shown in Fig. 8. Compared to the reference solution, the evolution of $\theta(t)$ is captured reasonably well in the simulation with $\alpha = 1/2$, but large amplitude and phase errors occur after a few periods of the slow oscillation in the simulation with $\alpha = 1$. In particular, the oscillations are damped by the RA filter (with $\alpha = 1$) but are undamped by the RAW filter (with $\alpha = 1/2$), consistent with Williams (2009). The evolution of $\eta(t)$ is wildly

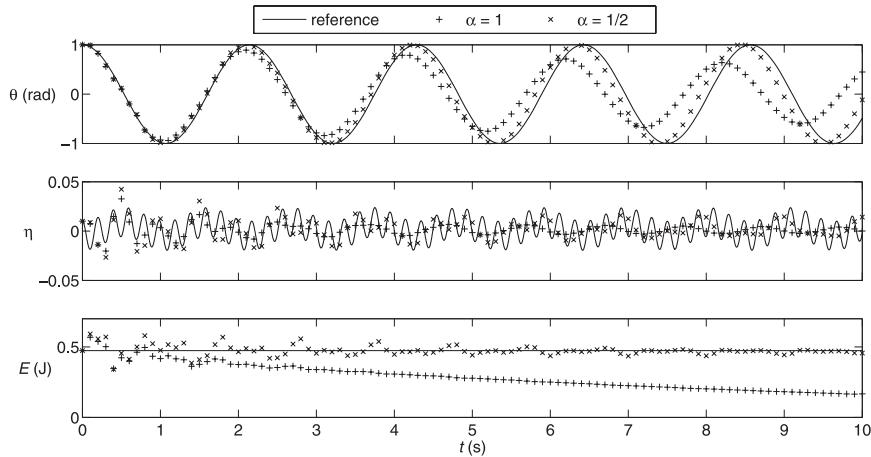


FIG. 8. Numerical integrations of the nonlinear elastic pendulum equations in (15)–(18), showing the angular and radial motion of the point mass and the total energy of the system. At time $t = 0$, the mass is released from rest with $\theta = 1.0$ rad and $\eta = 0.01$, and hence with $E \approx 0.47$ J. The RAW-filtered semi-implicit numerical scheme is used, with $\Delta t = 0.1$ s, $\nu = 0.2$, and (in two separate integrations) $\alpha = 1$ and $\alpha = 1/2$. The integrations are compared to a reference solution obtained with $\Delta t = 10^{-6}$ s.

inaccurate for both cases of α , because the time step is too long to resolve the fast oscillations, but the implicit discretization of the underlined terms in (15) and (16) keeps the integrations stable. Similar (but less pronounced) results are found for smaller values of ν .

The elastic pendulum under consideration is physically unforced and undamped. Therefore, the time-continuous equations, (15)–(18), exactly conserve the total energy,

$$E = \frac{1}{2}ml^2[v_\eta^2 + (1 + \eta)^2v_\theta^2] - mgl(1 + \eta)\cos\theta + \frac{1}{2}kl^2(\eta + mg/kl)^2 + mgl - \frac{1}{2}k(l - l_0)^2. \quad (19)$$

The first term in (19) is the kinetic energy, the second term is the gravitational potential energy, the third term is the elastic potential energy, and the two constant terms arbitrarily define the resting equilibrium energy to be zero. From the form of (19), energy accuracy is expected to be improved when amplitude accuracy is improved. The evolution of $E(t)$ in the numerical simulations is also shown in Fig. 8. Initially, $E \approx 0.47$ J. The reference solution conserves E to within 10^{-11} J over the integration period. In the two experimental simulations, the energy conservation is captured well with $\alpha = 1/2$, but most of the energy is lost to numerical dissipation by the end of the integration period with $\alpha = 1$.

Of particular interest are the magnitudes of the numerical errors after a given integration interval, and the rates at which they converge to zero as the time step is progressively refined. Such data, for the variable θ at

time $t = 10$ s, are displayed in Fig. 9. The errors for $\alpha = 0.4$ and $\alpha = 0.6$ are approximately equal in magnitude and opposite in sign, and the same for $\alpha = 0.3$ and $\alpha = 0.7$, and so on, but the errors for $\alpha = 1/2$ are typically orders of magnitude smaller. The leading-order errors scale quadratically with the time step for $\alpha = 1/2$, but

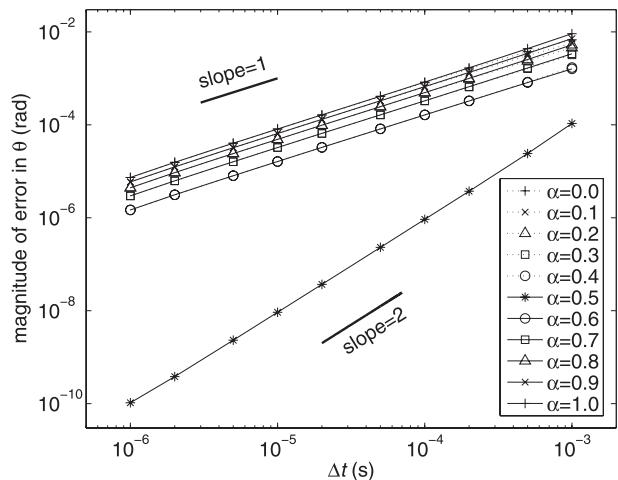


FIG. 9. Errors diagnosed from numerical integrations of the nonlinear elastic pendulum equations in (15)–(18). For each integration, at time $t = 0$ the mass is released from rest with $\theta = 1.0$ rad and $\eta = 0.01$, and the RAW-filtered semi-implicit numerical scheme with $\nu = 0.2$ is used. This figure is produced from 110 such integrations, one for each combination of 10 different values of Δt and 11 different values of α . The error in θ is defined to be the difference at time $t = 10$ s between the value obtained using time step Δt and a reference value obtained using time step 10^{-7} s. For guidance, straight lines of slopes 1 and 2 are also drawn.

linearly for all other $\alpha \neq 1/2$, consistent with the analysis of section 2b. The same error scalings are found for the other variables (η , ν_θ , ν_η , and E) and at other times.

4. Summary and discussion

Evidence suggests that discrete time stepping may be an important component of total model error in contemporary atmospheric and oceanic numerical simulations. Thus motivated, the Robert–Asselin–Williams (RAW) filter was proposed by Williams (2009) as a simple improvement to the Robert–Asselin (RA) filter, which is widely used to control the computational mode associated with the leapfrog time-stepping scheme. The present paper has applied the RAW filter to semi-implicit integrations, in which the explicit leapfrog scheme is used to discretize the low-frequency terms and the implicit trapezoidal (Crank–Nicolson) scheme is used to discretize the high-frequency terms.

Section 2 presented a linear theoretical analysis of the RAW-filtered semi-implicit scheme applied to the split-frequency complex oscillation equation. The stability and accuracy were interrogated in section 2a by deriving the complex amplification factors for the physical and computational modes, multiplication by which advances the numerical solution through one time step. The unconditional linear stability of the physical mode in the fully implicit limit was found to be preserved by the RAW filter. The unconditional linear stability of the computational mode was also found to be preserved by the RAW filter.

The stability and accuracy of the physical mode were further interrogated in section 2b by deriving power-series expansions for the amplitude errors and phase errors. The expansions were used to interpret the leading-order errors for small time steps and hence to identify optimal regions of the space spanned by the frequency ratio ($r = \omega_{\text{high}}/\omega_{\text{low}}$) and the filter parameters (α and ν) in the limit $\nu \ll 1$. For fully explicit integrations of unforced and undamped oscillations, the RAW filter with $\alpha = 3/4$ was found to be perhaps the best compromise between the needs for accuracy and stability, because it was found to offer better accuracy than the RA filter and almost the same stability criterion. For fully implicit integrations and typical semi-implicit integrations, the RAW filter with $\alpha = 1/2$ was found to be perhaps the best compromise, because it was found to offer better accuracy than the RA filter and to retain the unconditional stability.

Finally, section 3 tested the RAW-filtered semi-implicit scheme in a realistic nonlinear setting, by applying it to numerical integrations of the elastic pendulum equations. The implicit trapezoidal (Crank–Nicolson) scheme was

used to discretize the terms that are crucial for the fast vibrational mode, and the explicit leapfrog scheme was used to discretize the nonlinear terms and the terms that are crucial for the slow rotational mode. As expected, the simulated fast oscillations were found to be wildly inaccurate compared to a reference solution, because the time step was too long to resolve them. The simulated slow oscillations were found to be damped for $\alpha = 1$ (corresponding to the original RA filter) but to be undamped for $\alpha = 1/2$. Consequently, the total energy of the system, which is conserved exactly by the time-continuous equations, was found to be conserved for $\alpha = 1/2$ much better than for $\alpha = 1$.

The present paper has demonstrated that the RAW-filtered leapfrog scheme is an improvement to the RA-filtered leapfrog scheme. Of course, neither of these filtered leapfrog schemes is able to match the accuracy of third-order schemes (e.g., Runge–Kutta and Adams–Bashforth). What the RAW-filtered leapfrog scheme does offer, however, is suitability for semi-implicit integrations (unlike the third-order Adams–Bashforth scheme; Durran 1991) and arguably a higher numerical accuracy per unit of coding effort (because the implementation is trivial; Fig. 1).

The present paper gains practical verification from Amezcua et al. (2011), who have implemented the RAW filter in the Simplified Parameterizations, Primitive-Equation Dynamics (SPEEDY) atmospheric general circulation model (Molteni 2003). By default, SPEEDY uses a semi-implicit time-stepping scheme with the RA filter. The upgrade from RA filter (with $\nu = 0.1$) to RAW filter (with $\nu = 0.1$ and $\alpha = 0.53$) was reported to have required the insertion of a single new line of computer code, which did not noticeably affect the integration speed. In a series of tests, no significant changes were found to the model's monthly climatology, which meant that the parameterized physics required no retuning. In further tests, the RAW filter was found to significantly improve the skill of short- and medium-range weather forecasts. The improvements increased with lead time and were especially evident in the medium range. For example, for tropical surface pressure predictions, 5-day forecasts made using the RAW filter were found to have approximately the same skill as 4-day forecasts made using the RA filter.

In addition to improving atmosphere models, the RAW filter could also bring substantial benefits to ocean models, by correcting several known problems with the RA filter. For example, in an isopycnal-coordinate simulation of the steady-state global ocean circulation at 2° horizontal resolution, the RA filter has been found to generate numerical diapycnal volume fluxes that are as large as the physical diapycnal volume fluxes (P. J. van Leeuwen 2010, personal communication). It remains to

be seen whether the RAW filter could fix this behavior. Also, in leapfrog ocean models with time-varying vertical coordinates, the RA filter is known to violate tracer conservation. Leclair and Madec (2009) attempt to fix the nonconservation by reformulating the forcing and filtering terms. Possibly, an even more superior scheme will result from combining the reformulation of Leclair and Madec (2009) with the RAW filter. The exploration of these possibilities is left for future work.

Acknowledgments. The author is funded through a University Research Fellowship from the Royal Society (UF080256). The paper was completed while the author was a Visiting Fellow at the Isaac Newton Institute for Mathematical Sciences, University of Cambridge, United Kingdom. The comments of two anonymous reviewers are gratefully acknowledged.

REFERENCES

- Amezcu, J., E. Kalnay, and P. D. Williams, 2011: The effects of the RAW filter on the climatology and forecast skill of the SPEEDY model. *Mon. Wea. Rev.*, **139**, 608–619.
- Asselin, R., 1972: Frequency filter for time integrations. *Mon. Wea. Rev.*, **100**, 487–490.
- Bartello, P., 2002: A comparison of time discretization schemes for two-timescale problems in geophysical fluid dynamics. *J. Comput. Phys.*, **179**, 268–285.
- Caya, D., and R. Laprise, 1999: A semi-implicit semi-Lagrangian regional climate model: The Canadian RCM. *Mon. Wea. Rev.*, **127**, 341–362.
- Cordero, E., and A. Staniforth, 2004: A problem with the Robert–Asselin time filter for three-time-level semi-implicit semi-Lagrangian discretizations. *Mon. Wea. Rev.*, **132**, 600–610.
- Déqué, M., and D. Cariolle, 1986: Some destabilizing properties of the Asselin time filter. *Mon. Wea. Rev.*, **114**, 880–884.
- Dietrich, D. E., and J. J. Wormeck, 1985: An optimized implicit scheme for compressible reactive gas flow. *Numer. Heat Transfer*, **8**, 335–348.
- Durran, D. R., 1991: The third-order Adams–Bashforth method: An attractive alternative to leapfrog time differencing. *Mon. Wea. Rev.*, **119**, 702–720.
- , 1999: *Numerical Methods for Wave Equations in Geophysical Fluid Dynamics*. Springer-Verlag, 482 pp.
- Ford, R., 1994: Gravity-wave radiation from vortex trains in rotating shallow-water. *J. Fluid Mech.*, **281**, 81–118.
- Fraedrich, K., H. Jansen, E. Kirk, U. Luksch, and F. Lunkeit, 2005: The Planet Simulator: Towards a user friendly model. *Meteor. Z.*, **14**, 299–304.
- Griffies, S. M., R. C. Pacanowski, M. Schmidt, and V. Balaji, 2001: Tracer conservation with an explicit free surface method for z-coordinate ocean models. *Mon. Wea. Rev.*, **129**, 1081–1098.
- Haltiner, G. J., and R. T. Williams, 1980: *Numerical Prediction and Dynamic Meteorology*. 2nd ed. Wiley, 496 pp.
- Hartogh, P., A. S. Medvedev, T. Kuroda, R. Saito, G. Villanueva, A. G. Feofilov, A. A. Kutepov, and U. Berger, 2005: Description and climatology of a new general circulation model of the Martian atmosphere. *J. Geophys. Res.*, **110**, E11008, doi:10.1029/2005JE002498.
- Kalnay, E., 2003: *Atmospheric Modeling, Data Assimilation and Predictability*. Cambridge University Press, 364 pp.
- Kar, S. K., 2006: A semi-implicit Runge–Kutta time-difference scheme for the two-dimensional shallow-water equations. *Mon. Wea. Rev.*, **134**, 2916–2926.
- Kurihara, Y., 1965: On the use of implicit and iterative methods for the time integration of the wave equation. *Mon. Wea. Rev.*, **93**, 33–46.
- Kwizak, M., and A. J. Robert, 1971: A semi-implicit scheme for grid point atmospheric models of the primitive equations. *Mon. Wea. Rev.*, **99**, 32–36.
- Leclair, M., and G. Madec, 2009: A conservative leapfrog time stepping method. *Ocean Modell.*, **30**, 88–94.
- Lilly, D. K., 1965: On the computational stability of numerical solutions of time-dependent non-linear geophysical fluid dynamics problems. *Mon. Wea. Rev.*, **93**, 11–26.
- Lynch, P., 2002: The swinging spring: A simple model of atmospheric balance. *Large-Scale Atmosphere–Ocean Dynamics II: Geometric Methods and Models*, J. Norbury and I. Roulstone, Eds., Cambridge University Press, 64–108.
- Magazenkov, L. N., 1980: Time integration schemes for fluid dynamics equations, effectively damping the high frequency components (in Russian). *Tr. Gl. Geofiz. Obs.*, **410**, 120–129.
- Mesinger, F., and A. Arakawa, 1976: *Numerical Methods Used in Atmospheric Models, Volume I*. Global Atmospheric Research Programme Publications Series, No. 17, GARP, 64 pp.
- Molteni, F., 2003: Atmospheric simulations using a GCM with simplified physical parametrizations. I: Model climatology and variability in multi-decadal experiments. *Climate Dyn.*, **20**, 175–191.
- Pfeffer, R. L., I. M. Navon, and X. L. Zou, 1992: A comparison of the impact of two time-differencing schemes on the NASA GLAS climate model. *Mon. Wea. Rev.*, **120**, 1381–1393.
- Roache, P. J., and D. E. Dietrich, 1988: Evaluation of the filtered Leapfrog–Trapezoidal time integration method. *Numer. Heat Transfer*, **14**, 149–164.
- Robert, A. J., 1966: The integration of a low order spectral form of the primitive meteorological equations. *J. Meteor. Soc. Japan*, **44**, 237–245.
- , 1969: The integration of a spectral model of the atmosphere by the implicit method. *Proc. WMO–IUGG Symp. on NWP*, Tokyo, Japan, Japan Meteorological Agency, 19–24.
- , and M. Lépine, 1997: An anomaly in the behaviour of a time filter used with the leapfrog scheme in atmospheric models. *Numerical Methods in Atmospheric and Oceanic Modelling: The André J. Robert Memorial Volume*, C. A. Lin, R. Laprise, and H. Ritchie, Eds., NRC Research Press, S3–S15.
- Schlesinger, R. E., L. W. Uccellini, and D. R. Johnson, 1983: The effects of the Asselin time filter on numerical solutions to the linearized shallow-water wave equations. *Mon. Wea. Rev.*, **111**, 455–467.
- Skamarock, W. C., and J. B. Klemp, 1992: The stability of time-split numerical methods for the hydrostatic and the nonhydrostatic elastic equations. *Mon. Wea. Rev.*, **120**, 2109–2127.
- Staniforth, A., 1997: André Robert (1929–1993): His pioneering contributions to numerical modelling. *Numerical Methods in Atmospheric and Oceanic Modelling: The André J. Robert Memorial Volume*, C. A. Lin, R. Laprise, and H. Ritchie, Eds., NRC Research Press, 25–54.
- Tandon, M. K., 1987: Robert’s recursive frequency filter: A re-examination. *Meteor. Atmos. Phys.*, **37**, 48–59.
- Teixeira, J., C. A. Reynolds, and K. Judd, 2007: Time step sensitivity of nonlinear atmospheric models: Numerical convergence, truncation error growth, and ensemble design. *J. Atmos. Sci.*, **64**, 175–189.

- Williams, P. D., 2009: A proposed modification to the Robert–Asselin time filter. *Mon. Wea. Rev.*, **137**, 2538–2546.
- , T. W. N. Haine, P. L. Read, S. R. Lewis, and Y. H. Yamazaki, 2009: QUAGMIRE v1.3: A quasi-geostrophic model for investigating rotating fluids experiments. *Geosci. Model Develop.*, **2**, 13–32.
- Williamson, D. L., and J. G. Olson, 2003: Dependence of aquaplanet simulations on time step. *Quart. J. Roy. Meteor. Soc.*, **129**, 2049–2064.
- Wolfram Research, 2008: Mathematica Version 7.0. Wolfram Research.
- Young, J. A., 1968: Comparative properties of some time differencing schemes for linear and nonlinear oscillations. *Mon. Wea. Rev.*, **96**, 357–364.
- Zhao, B., and Q. Zhong, 2009: The dynamical and climate tests of an atmospheric general circulation model using the second-order Adams–Bashforth method. *Acta Meteor. Sin.*, **23**, 738–749.