

Improving the Efficiency of a Digital Filtering Scheme for Diabatic Initialization

PETER LYNCH

Met Éireann, Dublin, Ireland

DOMINIQUE GIARD

CNRM/GMAP, Météo-France, Toulouse, France

VLADIMIR IVANOVICI

Romanian Institute of Meteorology and Hydrology, Bucharest, Romania

24 April 1996 and 2 December 1996

ABSTRACT

A method of diabatic initialization is described in which a nonrecursive digital filter is applied to both the backward (adiabatic) and forward (diabatic) steps. It has clear advantages over previously proposed schemes: the initialization requires significantly less computation time and the resulting changes in the analyzed fields are consistently smaller. Despite these smaller increments, the suppression of spurious high-frequency noise is at least as effective as for the earlier schemes.

1. Introduction

Digital filtering has been shown to be an effective means of initializing data for numerical weather prediction. In Lynch and Huang (1992, hereafter LH92) a simple filter was applied to a sequence of values centered on the initial time $t = 0$; these values were generated by two short adiabatic model integrations, one forward and one backward from $t = 0$. It was found that a total span of 6 h was required for effective elimination of high-frequency noise. Huang and Lynch (1993, hereafter HL93) showed, by means of an optimal filter, that equally effective noise control could be achieved with a span of only 3 h. They also described a means of incorporating diabatic effects: the backward adiabatic integration (of length 1.5 h) is followed by a diabatic forecast of twice the length, and the values generated by the diabatic integration are processed by the filter.

In Lynch and Huang (1994, hereafter LH94), recursive filters were applied to the initialization problem. Several schemes were investigated; in one, designated RAD—recursive adiabatic-diabatic—in LH94, the backward adiabatic integration was processed with a recursive filter, and the terminal output value was used to initiate a forward diabatic integration, to which re-

cursive filtering was again applied. The application of filtering to both the reverse and forward integrations is also possible using a nonrecursive filter. In this note we examine the advantages of this idea.

2. A simple conceptual model

Let us consider a function of time $x(t)$ governed by the oscillation equation

$$\frac{dx}{dt} - i\omega_0 x = 0 \quad (1)$$

with the initial value $x(0)$. The solution is $x(0) \exp(i\omega_0 t)$. Suppose this solution is filtered by convolution with an *impulse response* function $h(t)$:

$$\bar{x}(t) = h * x(t) = \int_{-\infty}^{\infty} h(\tau) x(t - \tau) d\tau.$$

Substituting the solution $x(0) \exp(i\omega_0 t)$ into the integral, it follows immediately that

$$\bar{x}(t) = H(\omega_0) x(t),$$

where the frequency response function $H(\omega)$ is the Fourier transform of $h(t)$. Now let us use a value of the filtered solution $y(0) = \bar{x}(0)$ as an initial value for a solution of the inhomogeneous equation

$$\frac{dy}{dt} - i\omega_0 y = f(t). \quad (2)$$

Corresponding author address: Dr. Peter Lynch, Met Éireann, Dublin 9, Ireland.
E-mail: plynch@irmet.ie




FIG. 1. Schematic illustration of original and modified filtering procedures. Thin lines indicate backward adiabatic integrations; thick lines indicate forward adiabatic integrations. (a) Scheme of Huang and Lynch (1993): adiabatic step is not filtered. (b) Modified scheme: both steps are filtered.

The solution is easily obtained by Laplace transformation or otherwise. For sinusoidal forcing, $f(t) = f_0 \times \exp(i\omega_1 t)$ ($\omega_0 \neq \omega_1$), it may be written

$$y(t) = y(0) \exp(i\omega_0 t) + f_0 \left[\frac{\exp(i\omega_0 t) - \exp(i\omega_1 t)}{i(\omega_0 - \omega_1)} \right]. \quad (3)$$

If this is filtered, the result is

$$\bar{y}(t) = [H(\omega_0)]^2 x(t) + f_0 \left[\frac{H(\omega_0) e^{i\omega_0 t} - H(\omega_1) e^{i\omega_1 t}}{i(\omega_0 - \omega_1)} \right], \quad (4)$$

whereas starting from the unfiltered value $x(0)$, as in HL93, it would be

$$\bar{y}(t) = [H(\omega_0)] x(t) + f_0 \left[\frac{H(\omega_0) e^{i\omega_0 t} - H(\omega_1) e^{i\omega_1 t}}{i(\omega_0 - \omega_1)} \right]. \quad (5)$$

Let us consider that (1) corresponds to the adiabatic model and (2) to the diabatic model, with physical forcing represented by $f(t)$. Suppose now that the adiabatic model is integrated and the solution $x(t)$ is filtered to give $\bar{x}(t)$. This is used to define initial data $y(0)$ for an integration of the diabatic model (2). The resulting solution $y(t)$ is again filtered, producing $\bar{y}(t)$. Then (4) shows that the adiabatic part of the solution is effectively filtered twice, while the diabatic part is filtered once. If the amplitude f_0 of the forcing is sufficiently small, a single filtering of the diabatic component should be adequate to reduce its high frequency components to an acceptable level.

The initialization scheme described in HL93 is illustrated schematically in Fig. 1a. The thin line represents the reverse adiabatic integration of duration T . This is not filtered; its terminal value, depicted by an open circle, is used to initiate a forward diabatic integration (the thick line) that is then subjected to filtering. This in-

tegration must be of duration $2T$ to ensure that the output of the symmetric filter is valid at $t = 0$. The modified scheme, to be analyzed here, is illustrated in Fig. 1b. The reverse adiabatic integration (thin curve) is filtered, yielding an output valid at $t = -T/2$. This value, depicted by the black spot, is used to initiate a forward diabatic integration (thick curve) of duration T , centered on the initial time $t = 0$.

Considering the adiabatic component of the solution, we see that the HL93 scheme involves a single application of a filter of span $2T$; the modified scheme involves two applications of a filter of span T . Comparison between the schemes can be made by considering the relationship between the magnitude of the response function, $|H_{2T}(\omega)|$, of the filter with the longer time span and the squared response function $|H_T(\omega)|^2$ of the filter with the shorter time span.

The diabatic initialization schemes apply filtering to trajectories that do not correspond, at time $t = 0$, to the original analysis. Let us suppose that both ω_0 and ω_1 are low frequencies, so that we may assume $H(\omega_0) \approx H(\omega_1) \approx 1$. Thus, the filtering does not substantially affect these components. The result of a backward adiabatic integration of length T followed by a forward diabatic integration of the same length may easily be deduced: integrate (1) backward from $x(0)$ to time $t = -T$ to get $x(-T) = x(0) \exp(-i\omega_0 T)$; use this value to initiate a forward integration of (2) from $t = -T$ to $t = 0$ to obtain

$$y(0) = x(0) + f_0 \left[\frac{\exp(i\omega_0 T) - \exp(i\omega_1 T)}{i(\omega_0 - \omega_1)} \right]. \quad (6)$$

For short time spans ($\omega_0 T$ and $\omega_1 T$ small), assuming the filter response is unity, we have

$$\bar{y}(0) \approx x(0) + f_0 T. \quad (7)$$

This differs from the original value $x(0)$ by the *diabatic discrepancy* $f_0 T$, which is proportional to the amplitude f_0 of the diabatic forcing and also to the time span T . Since the span of the new scheme is typically half that of the old scheme, the effect of this discrepancy is reduced. Of course, the problem of diabatic discrepancy disappears when using a *finalization* or *launch* technique as described in Fillion et al. (1995) and LH94, respectively; but then the filtered fields are not applicable at the initial time.

3. Some examples of filters

The filter used in HL93 was an optimal filter having pass-band and stop-band edges with periods $\tau_p = 15$ h and $\tau_s = 3$ h (recall that the period is given by $\tau = 2\pi/\omega = 2\pi\Delta t/\theta$). The total span was $2T = 3$ h and the time step was $\Delta t = 360$ s. As shown in Lynch (1997), the Dolph–Chebyshev filter with a stop-band edge $\tau_s = 3$ h and the same time span (3 h) is a simple optimal filter with virtually the same frequency response. We will

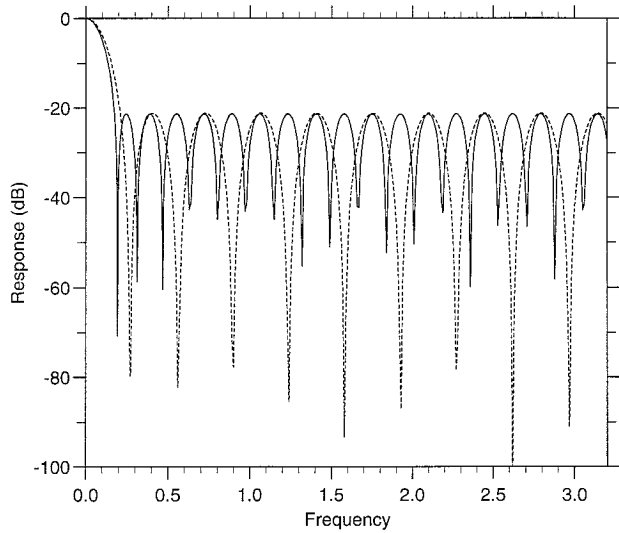


FIG. 2. Frequency response (dB) for Dolph filters with ripple ratio $r \approx 0.1$. Solid line: filter order $N = 2M + 1 = 37$; $\delta = 20 \log |H(\theta)|$ plotted. Dashed line: filter order $N = 2M + 1 = 19$; $\delta = 20 \log |H(\theta)|^2$ plotted.

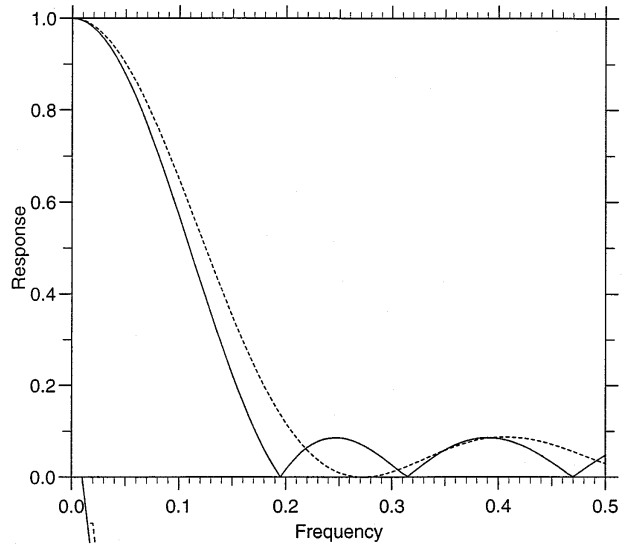


FIG. 3. Amplitude response for Dolph filters as in Fig. 2 above. Solid line: filter order $N = 2M + 1 = 37$; $|H(\theta)|$ plotted. Dashed line: filter order $N = 2M + 1 = 19$; $|H(\theta)|^2$ plotted.

therefore use simple Dolph–Chebyshev filters for the comparison below. We choose a time step $\Delta t = 300$ s so that $M = T/\Delta t = 18$ is an even number. Let us recall that the frequency response of a Dolph–Chebyshev filter with θ_s as stop-band edge frequency is

$$H(\theta) = rT_{2M} \left(\frac{\cos \frac{\theta}{2}}{\cos \frac{\theta_s}{2}} \right), \quad (8)$$

where $T_{2M}(x)$ is a Chebyshev polynomial of order $2M$ and the total time span is $2T = 2M\Delta t$. The ripple ratio r measures the maximum response in the stop band ($|\theta| > \theta_s$) and is determined by requiring $H(0) = 1$:

$$\frac{1}{r} = T_{2M} \left(\frac{1}{\cos \frac{\theta_s}{2}} \right) = \cosh \left[2M \cosh^{-1} \left(\frac{1}{\cos \frac{\theta_s}{2}} \right) \right] \quad (9)$$

(see Lynch 1997).

In Fig. 2, two response functions are shown. The ordinate is the filter attenuation, $\delta = 20 \log |H(\theta)|$ (dB). The solid curve is the response of a Dolph–Chebyshev filter (denote it H_{2T}) of total span $2T = 3$ h; the stop-band edge is $\tau_s = 3$ h and $M = 18$. The dashed curve is the squared response of a Dolph–Chebyshev filter (H_T^2) of total span $T = 1.5$ h; thus, $M = 9$. The stop-band edge $\tau_s = 2.5$ h has been chosen by experiment so that the total attenuation in the stop band was about the same in both cases: -21.3 dB for the filter with the longer span and -21.2 dB for the squared response of the filter with the shorter time span. A more systematic

means of choosing τ_s will be described below. The amplitudes of the frequency responses of the two filters are shown in Fig. 3 with the θ axis expanded for clarity: the responses are very similar, but the pass band of the squared short filter is somewhat broader.

If a semi-Lagrangian advection scheme is used, relatively long time steps are possible. We consider a comparison between filters H_{2T} and H_T^2 having the same stop bands and total spans as above, but with a time step $\Delta t = 900$ s. Thus, $M = 6$ for H_{2T} and $M = 3$ for H_T^2 . The frequency responses are shown in Fig. 4. Once again,

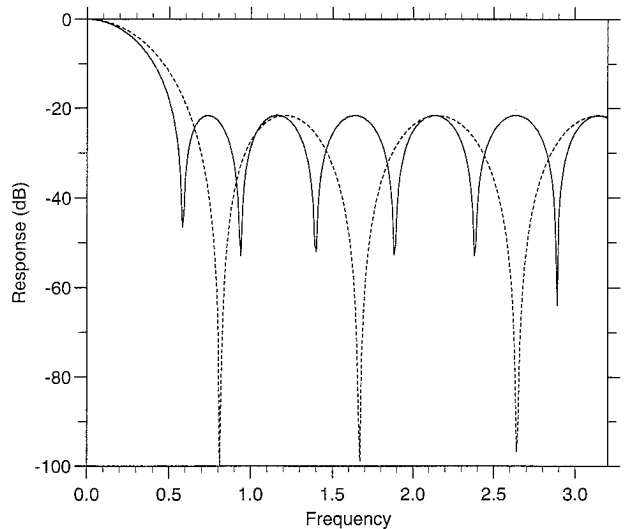


FIG. 4. Frequency response (dB) for Dolph filters with ripple ratio $r \approx 0.1$. Solid line: filter order $N = 2M + 1 = 13$; $\delta = 20 \log |H(\theta)|$ plotted. Dashed line: filter order $N = 2M + 1 = 7$; $\delta = 20 \log |H(\theta)|^2$ plotted.

the attenuation in the stop band is similar for both filters (-21.6 dB in each case) and the pass band of the squared short filter is slightly broader. The responses in the pass band for filters with the longer and shorter time steps were compared (by rescaling the θ axis) and were found to be very similar: thus, effective frequency discrimination is achievable with the longer time step.

4. Choosing the filter parameters

Methods of ensuring that singly and doubly applied filters have the same level of damping for high frequencies will now be described. Let unprimed and primed quantities refer to the filters with the longer and shorter time spans, respectively.

First, let us assume that the shorter span is exactly one-half of the longer one. We now determine the parameters for a filter with one-half of the span, having a ripple ratio $r' = \sqrt{r}$, so that its squared response will have the same high frequency attenuation as the longer filter. Its stop-band edge θ'_s is determined by inversion of (9) with appropriate parameter values:

$$\theta'_s = 2 \cos^{-1} \left[\cosh \left(\frac{1}{M} \cosh^{-1} \frac{1}{\sqrt{r}} \right) \right]^{-1}. \quad (10)$$

It was found in the above examples that θ'_s was somewhat larger than θ_s . It must be ascertained in each case whether θ'_s is small enough to provide the required frequency selectivity.

Numerical tests using (10) revealed that, in some cases, the widening of the stop band for the shorter filter was unacceptable. An alternative approach, which ensures that both filters have the same stop-band edge, will now be described. First we note that, if θ_s and r are specified, the required filter span can be deduced from (9): we have $T = 2M\Delta t$ with

$$2M = \frac{\cosh^{-1}(1/r)}{\cosh^{-1} \left(1 / \cos \frac{\theta_s}{2} \right)}.$$

Noting that, for small ϵ , $\cosh^{-1}(1/\cos \epsilon) \approx \epsilon$, as may be easily verified by standard Taylor expansions, and recalling that $\theta_s = 2\pi\Delta t/\tau_s$, we find that

$$T \approx \left(\frac{1}{\pi} \cosh^{-1} \frac{1}{r} \right) \tau_s, \quad (11)$$

provided $\theta_s \ll \pi$ or $2\Delta t \ll \tau_s$, which is always true in practice.

Now assume that the longer and shorter spans are $T = 2M\Delta t$ and $T' = 2M'\Delta t$, respectively, and that both filters have the same stop-band edge θ_s . The response functions are

$$H(\theta) = rT_{2M} \left(\frac{\cos \frac{\theta}{2}}{\cos \frac{\theta_s}{2}} \right),$$

where $r = 1/T_{2M}(1/\cos\theta_s/2)$ for the longer filter, and

$$H'(\theta) = r'T_{2M'} \left(\frac{\cos \frac{\theta}{2}}{\cos \frac{\theta_s}{2}} \right),$$

where $r' = 1/T_{2M'}(1/\cos\theta_s/2)$ for the shorter one. Since the filter with the shorter span is to be applied twice, we require $r' = \sqrt{r}$ to ensure that the stop-band attenuation is equivalent in both cases. Thus, from (11) we have

$$T' \approx \left(\frac{1}{\pi} \cosh^{-1} \frac{1}{\sqrt{r}} \right) \tau_s. \quad (12)$$

Taking typical values $\tau_s = 3$ h and $r = 0.05$ we find from (11) that the filter applied only once must have a span at least equal to 3.52 h, whereas (12) implies that the filter applied twice achieves the same high-frequency damping with a span of only 2.08 h. This leads to a substantial reduction in computation time.

5. Application to a forecast model

The new initialization scheme has been evaluated by application to the limited-area spectral model ALADIN (partially described in Bubnová et al. 1995). This model is run quasi-operationally at Météo-France and coupled with the global spectral NWP model ARPEGE. Horizontal representation of the variables is achieved by double Fourier series. A nonhydrostatic version of the model has been developed, but the results presented below are for the hydrostatic version (application to the nonhydrostatic version would simply involve filtering of the two additional prognostic variables). The grid-size here is approximately 18.3 km and there are 27 vertical levels. The time step is $\Delta t = 450$ s. It was found to be necessary to use a smaller time step, $\Delta t = 225$ s, for the backward adiabatic integration to ensure stability. The geographical area covered may be seen later in Fig. 7.

ALADIN does not have its own analysis scheme. The initialized analysis produced by the global spectral model ARPEGE is transformed to the resolution required for ALADIN. This transformation introduces spurious noise that must be removed by initialization on the limited domain. Four forecasts were performed starting from analyses valid at 0000 UTC 21 January 1996. One was from the transformed analysis without any subsequent initialization. The other three forecasts followed digital filtering initialization in three versions.

1) A Lanczos filter with span 6 h and a cutoff $\tau_c = 6$ h.

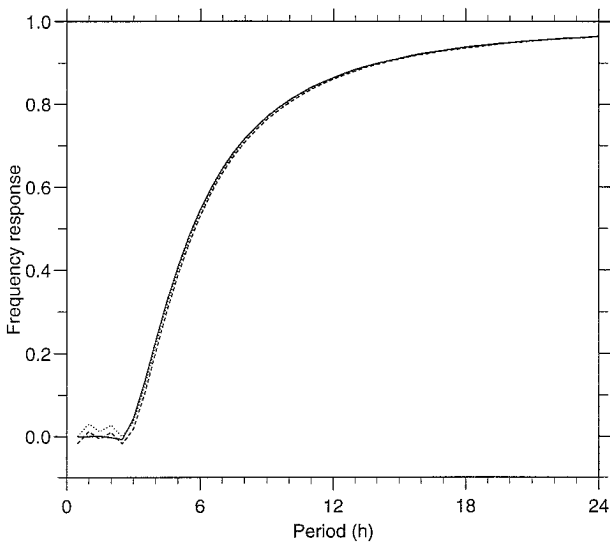


FIG. 5. The frequency response of the filters used in the three schemes described in the text. Solid line—scheme 1, Lanczos filter applied once; dashed line—scheme 2, Dolph–Chebyshev filter applied once; dotted line—scheme 3, Dolph–Chebyshev filter applied twice. Note that for scheme 3 the square of the response function is plotted.




FIG. 6. Evolution of the mean absolute surface pressure tendency for the first six forecast hours, starting from uninitialized data (solid) and from data filtered with schemes 1 (dotted), 2 (dashed), and 3 (dot-dashed).

A backward integration of 3 h was followed by a 6-h diabatic forecast. Filtering was applied once only, to the forward integration.

- 2) A Dolph–Chebyshev filter with span 4.5 h and stop-band edge $\tau_s = 3$ h. Filtering was again applied once only, to the forward integration.
- 3) A Dolph–Chebyshev filter with span 2.25 h and stop-band edge $\tau_s = 3$ h. Filtering was applied twice, to both the backward and forward integrations. This is the new scheme.

The frequency responses of the filters used in the three schemes are shown in Fig. 5. Note that for the new scheme it is the square of the response function that is plotted since the filter is applied twice. The Lanczos filter is precisely that used in LH92 and is known to be effective. Note that the abscissa is the period τ . It is clear that the responses of the three filtering schemes are very similar. The time spans used here are longer than those considered in section 3 and allow a better damping of high frequencies, with a ripple ratio lower than 0.05 instead of 0.10 for Dolph–Chebyshev filters. At the same time, the attenuation of periods longer than 12 h remains small. Note that the constraints on the filters are not as strong in this case as might be required for a global model (as discussed in Fillion et al. 1995) or a limited-area model with its own assimilation cycle.

The evolution of the absolute tendency of surface pressure averaged over the forecast domain is shown in Fig. 6. The excessive noise present in the uninitialized forecast (solid curve) is successfully removed by all the filtering schemes, and they appear to be equally effective in this regard. There is a small amount of residual noise but it is not of sufficient amplitude to cause concern.

The average absolute surface pressure tendency is an indicator of the level of noise in the external gravity wave components. It is also necessary to study the impact of filtering on spurious internal gravity wave energy. For this purpose, the 500-hPa vertical velocity at the initial time was examined for the four forecasts (results not shown). It was found that large amplitude, small-scale noise in the uninitialized run, particularly in the region of the Alps, was effectively removed by all the filtering schemes.

The changes induced by the three filtering schemes were examined and compared. It was found that the new scheme consistently led to smaller changes than either of the schemes involving single filtering. The impact on the 500-hPa temperature analysis is shown in Fig. 7. Figures 7a,b,c show, respectively, the increments due to filtering schemes 1, 2, and 3. It is clear that the changes brought about by the new scheme are significantly smaller than those of the alternatives. Notwithstanding this, the noise reduction is equally effective. It is reasonable to ascribe the reduction in initialization increments for the new scheme to the reduced diabatic discrepancy associated with the shorter span since, as we have seen, the filter frequency response is virtually identical for the three schemes.

The saving in computation time for the new scheme may easily be estimated. Let Δt_B and c_B be the time step and cost per time step for the backward adiabatic run, Δt and c the corresponding values for the forward diabatic run, and $2T$ the total time span. We may suppose $\Delta t_B = \Delta t/2$ and $c_B = c/2$. Then the cost of the old schemes (single filtering) is

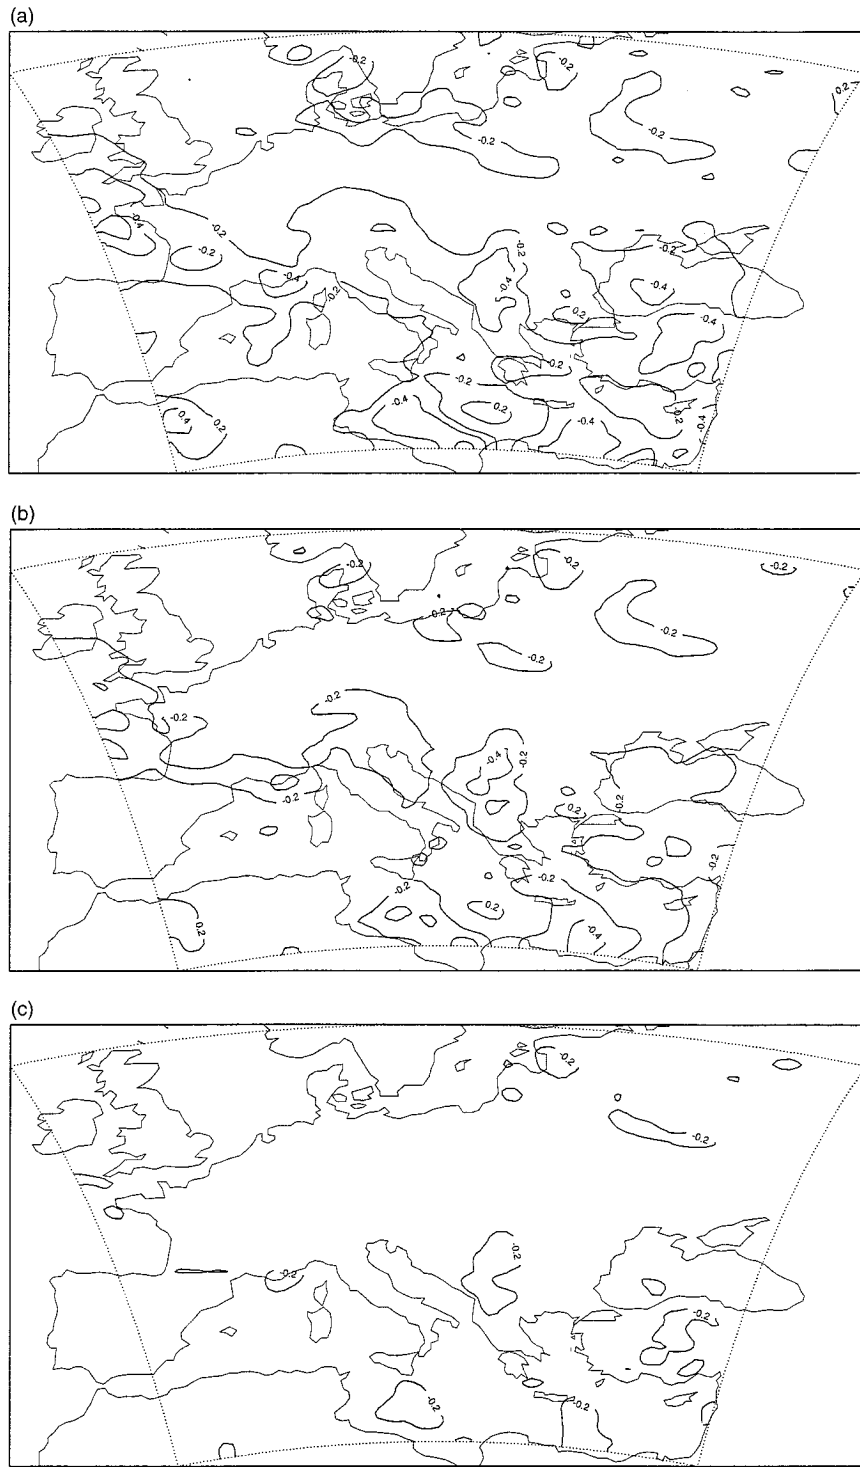


FIG. 7. Impact on the 500-hPa temperature analysis: (a), (b), and (c) show the increments due to filtering schemes 1, 2, and 3. The contour interval is 0.2 K.

$$C = \frac{T}{\Delta t_B} c_B + \frac{2T}{\Delta t} c = \frac{3Tc}{\Delta t}.$$

If the filter span of the new scheme is half as long, the cost is

$$C = \frac{T}{\Delta t_B} c_B + \frac{T}{\Delta t} c = \frac{2Tc}{\Delta t},$$

which results in a saving of 33%. This was the case in the above example: the computation time (seconds of

CPU usage) for scheme 1 was 1292 s; for scheme 2 it was 955 s; and for scheme 3 it was 662 s. The reduction is of practical benefit in an operational context.

The above results demonstrate that a significant gain in efficiency can be made by means of the new scheme. This scheme results in smaller increments to the initial fields but is equally effective in reducing high frequency noise in the forecast. The new scheme with Dolph–Chebyshev filters has been implemented in the operational ALADIN-LACE suite (focusing on central and eastern Europe) and in the preoperational ALADIN-FRANCE suite at Météo-France. It has also been implemented in the global system ARPEGE, in the framework of incremental initialization, so as to further reduce the diabatic discrepancy and preserve tidal modes (as suggested in LH94). In order to reasonably damp the large-scale gravity components introduced by analysis,

a more selective Dolph–Chebyshev filter has been used (with $\tau_s = 5$ h as the stop-band edge).

REFERENCES

- Bubnová, R., G. Hello, P. Bénard, and J.-F. Geleyn, 1995: Integration of the fully elastic equations cast in the hydrostatic pressure terrain-following coordinate in the framework of the ARPEGE/Aladin NWP system. *Mon. Wea. Rev.*, **123**, 515–535.
- Fillion, L., H. L. Mitchell, H. Richie, and A. Staniforth, 1995: The impact of a digital filter finalization technique in a global data assimilation system. *Tellus*, **47A**, 304–323.
- Huang, X.-Y., and P. Lynch, 1993: Diabatic digital-filtering initialization: Application to the HIRLAM model. *Mon. Wea. Rev.*, **121**, 589–603.
- Lynch, P., 1997: The Dolph–Chebyshev window: A simple optimal filter. *Mon. Wea. Rev.*, **125**, 655–660.
- , and X.-Y. Huang, 1992: Initialization of the HIRLAM model using a digital filter. *Mon. Wea. Rev.*, **120**, 1019–1034.
- , and —, 1994: Diabatic initialization using recursive filters. *Tellus*, **46A**, 583–597.