

Diabatic initialization using recursive filters

By PETER LYNCH^{1*} and XIANG-YU HUANG², ¹*Meteorological Service, Glasnevin Hill, Dublin 9, Ireland*; ²*Danish Meteorological Institute, Lyngbyvej 100, DK-2100 Copenhagen, Denmark*

(Manuscript received 22 June 1993; in final form 22 November 1993)

ABSTRACT

Several initialization schemes based on recursive filters are formulated and tested with a numerical weather prediction model, HIRLAM. These have an advantage over schemes which use non-recursive filters in that they derive the initialized values from a diabatic trajectory passing through the original analysis. The changes to the analysed fields are comparable in size to typical observational errors. A non-recursive implementation of the recursive filters makes the new initialization schemes as easy to use as the original non-recursive filter schemes. It also allows use of higher-order filters without added cost. An initialization method using a 6th order filter is compared to a scheme based on a non-recursive optimal filter, and is found to produce similar results for less than half the computational cost. If the sole aim is noise suppression, a filter whose output validates later than the initial time may be used. The advantage of this is that computation time is further reduced and phase error completely eliminated.

1. Introduction

Analysed data to be used as initial conditions for a numerical forecast must be adjusted to remove spuriously large gravity wave components. There are several reasons for this *initialization*: it enables us to generate realistic patterns of divergence, vertical velocity and other diagnostic fields; it eliminates high frequency noise which otherwise corrupts short-range forecasts; it facilitates the quality control and assimilation of observational data; it leads to more realistic precipitation patterns in the early hours of the forecast; finally, in the absence of initialization, excessive gravity wave activity can induce numerical instabilities, necessitating the use of a shorter time step.

Initialization of data for a limited area model has been successfully performed by application of a simple filter to time-series of dependent variables produced by the forecast model. In Lynch and Huang (1992), short-range forward and backward adiabatic integrations were carried out to produce sequences of values centered at the analysis time.

When these were used as input to a filter, the output was found to yield a noise-free integration. To allow for irreversible effects, Huang and Lynch (1993) followed a short backward adiabatic integration by a diabatic forecast of twice the duration to produce diabatic trajectories centered at the analysis time. The values yielded by this forecast were input to the filter and resulted in diabatically initialized fields.

The forward diabatic integration will not normally coincide at time $t=0$ with the original analysis. Ideally, the diabatic trajectory passing through the analysis point at the initial time should be used. We can generate this trajectory for positive time by integrating the model; but for negative time, irreversibility of the diabatic processes makes it difficult to calculate. In the present paper a method of filtering a forward trajectory to generate initialized fields is considered. This is achieved by means of a recursive filter.

Straightforward application of a second-order filter to a one-sided trajectory produces initialized fields which yield a noise-free integration, but causes a phase-shift so that the analysis and forecast changes are large. Several alternative formulations of the scheme are discussed and

* Corresponding author.

shown to be satisfactory when compared with the method of Huang and Lynch (1993). By implementing a recursive filter in a non-recursive way, no penalty in computer time or storage is incurred when a higher-order filter is used. It is shown that straightforward application of a 6th-order filter to a 1.5 h forecast produces initialized fields acceptably close to the analysis and eliminates high frequency oscillations while having minimal impact on a 24-h forecast. This scheme uses less than 50% of the computation required by that of Huang and Lynch (1993).

Filter output valid at the end of the span may be produced by reversing the sequence of coefficients. An integration may then be continued from this point, the original fields having been replaced by the filtered ones. The subsequent forecast is noise-free at virtually no extra cost in computation time.

2. Theory of recursive digital filters

General accounts of digital filters are available in publications on digital signal processing, for example Oppenheim and Schaffer (1989). Kulhánek (1976) reviews them from the standpoint of geophysics. For a recent discussion of the filters of interest in the present paper, see Lynch (1993a).

2.1. Definition of digital filters

Given a discrete function of time $\{x_n\}$ a *non-recursive* digital filter is defined by

$$y_n = \sum_{k=-N}^N a_k x_{n-k}. \quad (1)$$

The output y_n at time $n \Delta t$ depends on both past and future values of x_n , but not on other output values. A *recursive* digital filter of order N is defined by

$$y_n = \sum_{k=0}^N a_k x_{n-k} + \sum_{k=1}^N b_k y_{n-k}. \quad (2)$$

The output y_n at time $n \Delta t$ in this case depends on past and present values of the input and also on previous output values. Recursive filters are more powerful than non-recursive ones, but can also be more problematical, as the feedback of the output can give rise to instability. The response of a non-recursive filter to an impulse at $n=0$ vanishes for

$|n| > N$, giving rise to the alternative name *finite impulse response* or FIR filter. Recursive filters have longer memories: since the response of a recursive filter to an impulse input can persist indefinitely, it is known as an *infinite impulse response* or IIR filter.

The frequency response of a digital filter is easily found: let $x_n = \exp(in\theta)$, where θ is the *digital frequency*, and assume an output of the form $y_n = H(\theta) x_n$; substituting into (2), the transfer function $H(\theta)$ is

$$H(\theta) = \frac{\sum_{k=0}^N a_k e^{-ik\theta}}{1 - \sum_{k=1}^N b_k e^{-ik\theta}}. \quad (3)$$

For nonrecursive filters, the denominator reduces to unity. Defining $z = \exp(i\theta)$, it can be shown that a recursive filter is stable if the roots of the *characteristic polynomial*

$$z^N - \sum_{k=1}^N b_k z^{N-k} = 0$$

are inside the unit circle $|z| = 1$.

The output of the nonrecursive filter (1) for an input $x_n = z^n$ is of the form $y_n = H z^n$, where H is a polynomial in z (multiplied by z^{-N}). There are several design techniques for choosing the coefficients $\{a_n\}$ (see Huang and Lynch (1993), Appendix A, for two methods). The design of recursive filters is more complicated. The response function H is a rational function of z . The effectiveness of such filters is influenced not only by the forced response, but also by transient components whose rate of decay determines when a steady state response is reached. The recursive filter (2) can be regarded as a difference equation of order N . The general solution comprises two distinct parts, a component which is a solution of the homogeneous equation and a forced component due to the input. If the filter is stable the former component decays exponentially with increasing n and is called the transient response. The rate of decay depends on the distance of the poles z_p of H from the unit circle, that is on $|z_p|$. Therefore, *ceteris paribus*, we should arrange for the poles to have the smallest possible moduli; this is the idea which is used to devise the Quick-Start filter (see Appendix A).

2.2. Design of recursive filters

In the simplest design method for recursive filters, a classical analog lowpass response is specified; a transformation of variables then converts this to discrete time in such a way that the required filter coefficients may be deduced. There are numerous classical filter functions which may be used as a basis for digital filter design. They are determined by the manner in which the ideal low-pass filter response is approximated in the pass- and stop-bands. It is important to choose the type of filter appropriate to the problem. All the classical types are optimal in one sense or another: for example, the Butterworth filter has an amplitude response which is as flat as possible at the low-frequency limit; the Bessel filter has a group delay that is as near constant as possible for a given filter order. In Appendix A we derive an expression for the N -pole filter whose transient solution decays most rapidly. It has all its poles on the negative real axis, at $s_p = -\sigma$:

$$H(s) = \frac{\sigma^N}{(s + \sigma)^N}.$$

The prototype version ($\omega_c = 1$) corresponds to the value $\sigma = \sqrt{1/(2^{1/N} - 1)}$. This is what we propose to call the N th-order Quick-Start filter.

There are several methods of deriving a digital transfer function from an analog response function. They all involve some mapping from the s -plane to the z -plane, chosen to preserve properties such as optimality of the filter. We will discuss only the bilinear transformation (for some alternatives see Lynch, 1993a). The theory behind this procedure can be found in the literature on digital signal processing. The bilinear transform is a mapping from the s -plane to the z -plane, given by

$$s = \frac{1}{\mu_c} \left[\frac{z-1}{z+1} \right], \quad (4)$$

where $\mu_c = \tan(\theta_c/2)$ with $\theta_c = \omega_c \Delta t = 2\pi \Delta t/\tau_c$ defining the cutoff frequency. Under this transformation the imaginary axis $s = i\omega$ of the s -plane maps into the unit circle $z = \exp(i\theta)$ in the z -plane

$$\omega = \frac{\tan(\theta/2)}{\tan(\theta_c/2)}$$

in such a way that $\omega = 1$, the cutoff frequency of

the prototype analog filter, corresponds to θ_c (this is the reason for the factor μ_c in (4)). Conversion from a prototype analog filter to a digital filter with cutoff frequency θ_c is implemented by the substitution of (4) into the transfer function $H(s)$ to obtain a function of z . The coefficients are deduced by comparing the result with the general expression (3).

The above may be made clearer by an example. The second-order prototype quick-start filter, normalized so that $H(0) = 1$, has a response function

$$H(s) = \frac{s_p^2}{(s - s_p)^2}, \quad (5)$$

where $s_p = (-\sqrt{\sqrt{2} + 1}, 0)$. Under the bilinear transformation this becomes

$$H(z) = G \frac{(z+1)^2}{(z-z_p)^2}, \quad (6)$$

where $G = (s_p \mu_c / (1 - s_p \mu_c))^2$ and z_p is the image of s_p . This response function corresponds to a filter of the form

$$y_{n+1} = (a_0 x_{n+1} + a_1 x_n + a_2 x_{n-1}) + (b_1 y_n + b_2 y_{n-1}),$$

and by comparing (6) with (3), the coefficients can be read off:

$$a_0 = \frac{1}{2} a_1 = a_2 = G, \quad b_1 = 2z_p, \quad b_2 = -z_p^2.$$

Graphs of the amplitude and phase response of several filters designed using the bilinear transform may be found in Lynch (1993a).

3. Application to initialization

The non-recursive filters employed by Lynch and Huang (1992) and Huang and Lynch (1993) have linear phase-error and produce output which is valid at the center of the span; thus, such filters applied to a forward integration yield data applicable at half the span. Recursive filters (IIR filters) have more complex phase-errors. Their use in initialization depends upon a simple idea: *if a filter with a group delay δ is applied for an integration over a time-span $T_s = \delta$, the output at $t = T_s$*

may be assumed to apply at the initial time $t=0$. This is what is required for initialization. Evidence in support of this idea is given in Subsection 3.3 below. An application of a recursive filter to initialization of a barotropic model is described in Lynch (1993b). Application to the HIRLAM model will be presented in the following sections.

3.1. HIRLAM model

The HIRLAM model (HIRLAM: High Resolution Limited Area Model) is the forecast component of a comprehensive numerical weather prediction system, which is under continuous development in a joint Nordic-Dutch-Irish research project and is operationally used in Denmark, Finland, Ireland, The Netherlands and Sweden. The model is documented in Kållberg (1989).

The HIRLAM model is a primitive equation model with horizontal wind (u, v), temperature (T), specific humidity (q) and surface pressure (p_s) as prognostic variables. The dependent variables are staggered on the Arakawa C-grid. The model has a rotated latitude/longitude grid, a hybrid sigma-pressure vertical coordinate, a second-order accuracy finite difference scheme, and a leapfrog semi-implicit scheme with Asselin time-filtering. The model physics includes a radiation scheme, a Kuo-type convection scheme, a stratiform condensation scheme, the ECMWF surface parameterization scheme, and a vertical diffusion scheme.

The model version we are using in this study has




Fig. 1. The model domain used in this study. The analysis of the mean sea level pressure at 00UTC 21 March 1993 is also shown.

a grid of 162×136 points and 31 levels in the vertical. The horizontal resolution is $0.21^\circ \times 0.21^\circ$. The time step for integration is 2.5 min. The forecast domain is shown in Fig. 1. The initial mean sea-level pressure is also given in the figure.

3.2. Nonrecursive implementation of recursive filters

A filter of the form (2) is to be applied to model output over a span $T_s = K \Delta t$. In signal processing the recursive nature of the filter is of great advantage in limiting the storage requirements when long series of output values are to be generated: only $2N+1$ values need to be stored (the latest values of x and y). But for initialization we require output at only a single time (the end of the span T_s) and storage can be further reduced by expressing the recursive filter in nonrecursive form. If we define input and output vectors $x = (x_0, x_1, \dots, x_K)^T$ and $y = (y_0, y_1, \dots, y_K)^T$, the filter (2) can be expressed as

$$y = Ax + By,$$

where A and B are lower triangular matrices. This may be solved for the output:

$$y = (I - B)^{-1} Ax = Fx. \quad (7)$$

The last row of this equation gives an expression for the output y at $t = T_s$ in terms of the inputs $x_0, x_1, \dots, x_K$:

$$y_K = \sum_{n=0}^K F_{K_n} x_n \quad (8)$$

This is formally similar to a nonrecursive filter (1) with the important distinction that the time series of x -values is no longer centered around the analysis time. Only the last row of the matrix F need be stored; this matrix is easily generated row-by-row so that matrix inversion is not required (see Appendix B).

The adiabatic digital filtering initialization scheme as described in Lynch and Huang (1992) is a standard option for initialization in the HIRLAM system. It is a straightforward matter to test recursive filters as expressed in (8): we simply redefine the filter coefficients. In practise, we perform the calculation of (8) for each model variable x at each gridpoint and level, accumulating the

sum as the forecast proceeds. Thus, the only extra storage needed is a single extra field for the final initialized values y_K .

3.3. Filter parameter values and delay

There are three parameters which determine the characteristics of an initialization scheme using the recursive filter: the filter order, the cutoff frequency and the filter span. In the experiments described in this paper, the Quick-Start filter is used with a cutoff period of 3 h. The filter span is 1.5 h, which needs 36 time-steps of model integrations, i.e., $K = 36$ in (8). The delay of the prototype second-order analog Quick-Start filter is $\delta_0 = 2/\sigma_0 = 1.287$ (see Appendix A). An approximate expression for the delay of the digital filter resulting from the bilinear transform is derived in Lynch (1993a); for the chosen parameters it gives $\Delta_0 = (\tau_c/2\pi) \delta_0 \approx 0.61$ h. Thus, the delay is less than the filter span, and there will be a phase error due to the application of the filter. The phase error can be reduced considerably by applying the filter only to the analysis increment, as described below.

If the variation of the group delay with frequency is large, different components will be phase-shifted to different degrees, resulting in distortion of the field. However, a graph of group delay (see Lynch, 1993a, Fig. 5.2) indicates that for the Quick-Start filter this variation is quite small for low frequencies. Moreover, if such a problem were to arise, the phase distortion could be reduced by using a Bessel filter, whose group delay is maximally flat. Another possible concern is that addition of an increment with a phase-shift to a first guess without one may result in significant errors. However, if both the increment and the phase distortion are small, this should be a second-order effect; in practice, no such difficulty has been observed.

Table 1. Position of pole σ , start-up time T_0 and delay δ_0 of quick-start filters of orders 2, 6, and 10; the dimensional delay Δ_0 is for a cutoff period of $\tau_c = 3$ h

N	σ	T_0	δ_0	Δ_0
2	1.554	0.644	1.287	0.614
6	2.858	0.350	2.100	1.003
10	3.733	0.268	2.679	1.279




Fig. 2. Input X and output Y for 2nd-order, 6th-order and 10th-order quick-start filters. Input is an oscillation with period 12 h. Outputs have the same period and (approximately) the same amplitude, but are phase-shifted to varying degrees. The cutoff period is $\tau_c = 3$ h.

An alternative method of minimizing the phase error is to use a higher order filter. The values of $\sigma = \sqrt{1/(2^{1/N} - 1)}$, $T_0 = 1/\sigma$ and $\delta_0 = N/\sigma$ for the prototype analog Quick-Start filters of orders 2, 6 and 10 are given in Table 1. The dimensional delay $\Delta_0 = (\tau_c/2\pi) \delta_0$ for $\tau_c = 3$ h is also tabulated. This shows how the delay increases with filter order. Results with a second-order filter are discussed in Section 5 and higher order filters are considered in Section 6.

A simple example of the effect of filter delay can be seen in Fig. 2. The input X is a pure sinusoid with a period of 12 h. Three outputs Y are shown; the filters used are second-order, sixth-order and tenth-order Quick-Start filters, each with a cutoff period of $\tau_c = 3$ h. In each case the oscillation is passed by the filter without undue attenuation. However, it is subject to a phase change which varies with the order. It is clear from the figure that the output is delayed relative to the input, and that this delay increases with filter order. Measurement of the actual delays show them to be in good agreement with the theoretical values. Note in particular that, for a time equal to the delay, the output is effectively frozen at its initial value: the solution more-or-less marks time for a period $t \approx \Delta_0$.

4. Design of the experiments

Several formulations of recursive filtering are described in this section. To illustrate the idea of

the methods, a schematic diagram is given in Fig. 3. The top diagram on the left shows adiabatic (A) and diabatic (D) trajectories for the time evolution of an "idealized" model variable. The difference between the adiabatic and diabatic processes is exaggerated to show the potential problem with diabatic digital filtering initialization and possible improvements of the new approaches. Other diagrams in Fig. 3 will be discussed in subsections below.

4.1. NOI: No initialization

In the 1st experiment, no initialization is performed. This experiment is referred to as NOI, and it provides one of the references to which the recursive filtering experiments are compared.

4.2. DOF: Diabatic optimal filter initialization

In the 2nd experiment the diabatic digital filtering initialization scheme of Huang and Lynch (1993) is used. The optimal filter with 3 h span,

15 h passband edge and 3 h stop band edge is used. The experiment is referred as DOF. It should be noted that DOF uses a *non-recursive* filter; the experiment is performed to provide another reference for comparison with *recursive* filters.

The way DOF is performed is shown by the centre left diagram of Fig. 3. An adiabatic integration backward in time is run first, following trajectory "A" for half of the filter span. A diabatic integration forward in time is then carried out following the trajectory "D". The optimal filter is applied to the values on "D" to yield the initialized field. A potential problem of DOF is that the diabatic trajectory "D" does not pass through the analysis state at the analysis time $t=0$. Thus, there may be a phase error due to the initialization scheme. Huang and Lynch (1993) did not find any serious consequences of this problem. Nevertheless, a formulation of the initialization scheme which removes it could be beneficial for theoretical studies and operational applications.

4.3. IIR: Infinite impulse response filter

In the 1st *recursive* filter application, referred as IIR, a forward diabatic model integration for 1.5 h is performed (lower left diagram, Fig. 3). The model integration follows the trajectory " x_i " and the filtered field follows the trajectory " y_i ". Due to the delay characteristics of IIRs, the output lags behind the input and, for a short span, one may assume the output to be valid at time $t=0$. However, as the span is somewhat longer than the delay for the chosen parameter values, there is a phase shift in the initialized field for this scheme.

4.4. INC: Incremental IIR filter

One means of minimizing the phase error is to apply IIR to the analysis increment (referred as INC), instead of to the analysis field itself. The basic assumption behind incremental initialization schemes is that the first-guess field (normally a 6-h forecast) is well balanced (e.g., Puri et al., 1982; Ballish et al., 1992). The analysis x^A is the sum of the first-guess x^F and an increment:

$$x^A = x^F + (x^A - x^F).$$

Two diabatic integrations are carried out, one from the analysis x^A and the other from the first-guess x^F , and two filtered fields y^A and y^F are




Fig. 3. Schematic diagrams of the diabatic digital filtering initialization techniques. See the text for details.

calculated using the IIR scheme. These are then used to construct the output:

$$y = x^F + (y^A - y^F).$$

If the analysis increment vanishes the filtering has absolutely no effect, irrespective of the span. Thus, any phase shift applies only to the analysis increment, not to the background field.

4.5. RAD: Recursive adiabatic/diabatic scheme

Another approach to minimize the phase difference associated with IIR is to use an adiabatic integration backward to counteract the phase error produced by the diabatic forward integration. This experiment is referred to as RAD, as it uses the recursive filter and requires both adiabatic (backward) and diabatic (forward) integrations.

The idea of RAD is illustrated in the three diagrams on the right hand side of Fig. 3. The first step is to run the adiabatic model backward (x_i) and compute the filtered trajectory (y_i). The "final" value y_K is then used as the initial state for the diabatic integration. There is a small phase-shift as the delay characteristics of the filter are less than perfect. During the forward, diabatic, integration the filtering is applied again. The final value y_K is assumed to be valid at the analysis time ($t=0$) and yields the initialized field. There is another phase error introduced here, which is in the opposite direction to the adiabatic one, due to the direction of the integrations. The total phase error is therefore expected to be much smaller than that of IIR.

The RAD scheme has some resemblance to the DOF scheme. However, there are fundamental differences between them. First of all, the valid time on the filtered trajectory has conceptually never left the analysis time as the filtered field may be assumed to apply at $t=0$ as long as the span is less than the filter delay; therefore, the phase shift is very small. Secondly, the adiabatic backward integration is used not to centre the trajectory values, but only to reduce the phase error due to the imperfect filter.

4.6. IAD: Incremental adiabatic/diabatic scheme

Finally, the idea of RAD is also applied to the analysis increment. This scheme, denoted IAD, combines the ideas of INC and RAD, with a view to further reducing the phase error. It involves

4 short integrations, a reverse adiabatic and a forward diabatic one for both the analysis and first-guess fields.

5. Results with 2nd-order filters

5.1. Reduction of noise

The mean absolute surface pressure tendency N_1 is chosen to measure the global noise level. It is defined as

$$N_1 = \frac{1}{IJ} \sum_{i=1}^I \sum_{j=1}^J \left| \frac{\partial p_s}{\partial t} \right|_{ij},$$

where the summation is calculated over the whole model domain. In Fig. 4, the variation of N_1 is shown for all experiments. All the curves converge as the time integration is carried on. The distinct differences between the uninitialized forecast and the initialized ones in the first 6 h convincingly show that the initial noise due to the imbalance between mass and wind fields is effectively removed by all the initialization schemes.

A more localized look at the noise is provided in Fig. 5, where the surface pressure p_s at a model point (indicated by "D" in Fig. 1) is plotted for the first 6 h for all experiments. It is seen that, in agreement with the results shown in Fig. 4, the large oscillations in p_s in NOI do not appear in other experiments. It may be noted that the IIR curve seems to follow a different path from the other initialized runs, which cluster together; this is




Fig. 4. Time evolution of the mean absolute surface pressure tendency N_1 averaged over the interior model domain (hPa/3 h) for various schemes.




Fig. 5. Time evolution of the mean sea level pressure p_s (hPa) at a model point, which is indicated in Fig. 1 by “D”.

presumably due to the presence of a phase shift resulting from the span being larger than the delay.

While the variation in p_s is a good indicator for external noise, the variation in the vertical velocity ω at the middle level reflects the presence of internal noise. In Fig. 6, ω at level 16 (about 500 hPa) at the same model point is shown for all experiments. Very high frequency oscillations in ω with period much shorter than 2 h, appear in the uninitialized run (NOI). These do not appear in other experiments; however, in INC and IAD there are oscillations of about a 2-h period in ω which are not filtered out. These oscillations are present not only in NOI, but also in the forecast from the first-guess fields (not shown). Note that the IIR curve appears to lead the other curves in phase by about 1 h.




Fig. 6. Time evolution of the vertical velocity ω (hPa/s) at level 16 (about 500 hPa) at a model point, which is indicated in Fig. 1 by “D”.

Some comments can be made here on the incremental initialization schemes. As the basic assumption of all incremental initialization schemes is that the first-guess field is well balanced, high frequencies in the forecast from the first-guess (if they exist) should be allowed and not filtered out. If these oscillations are physically significant, as assumed in incremental initialization schemes, both INC and IAD give very good results: they preserve those oscillation implicit in the first-guess while removing those due to the observations. On the other hand, in an operational context, noise may persist longer than the analysis cycle time; under these circumstances, the incremental initialization schemes may give rise to problems, while filtering of the complete model fields (as in DOF, IIR and RAD) can provide noise free initialized fields.

5.2. Changes in initial fields and forecasts

Apart from controlling the high frequency noise, an initialization scheme should cause only minor changes to the initial fields and should not degrade the forecast which follows the initialization. One motivation for the present investigation is to reduce the changes made by DOF in initial fields and forecasts.

From Figs. 5 and 6, it can be seen that the noise in NOI is damped during the model integration and that the model integration may be considered free from noise after 6 h. The best initialization scheme should thus produce a solution which coincides closely with NOI after a 6-h integration. Using this as a criterion, a subjective comparison of the experiments has been made, with the following result in order of decreasing merit: (1) IAD, (2)

Table 2. Root-mean-square (rms) and maximum (max) differences in temperature, wind components and surface pressure between the initialized and uninitialized analyses

EXP	T (K)		u (ms ⁻¹)		v (ms ⁻¹)		p _s (hPa)	
	max	rms	max	rms	max	rms	max	rms
DOF	6.68	0.28	8.56	0.63	7.27	0.59	1.53	0.42
IIR	4.45	0.45	16.41	1.16	10.29	1.24	1.83	0.62
INC	2.75	0.29	6.63	0.76	7.54	0.81	1.76	0.48
RAD	6.33	0.21	7.78	0.43	4.98	0.39	1.21	0.31
IAD	2.46	0.13	4.15	0.30	2.66	0.27	0.81	0.24

RAD, (3) DOF, (4) INC and (5) IIR. Careful comparisons have also been made for three more model points, indicated by A, B and C in Fig. 1. On the two "ocean" points (A and C), the differences between different experiments are smaller than those on the two "land" points (D is a land point). Nevertheless, the results are qualitatively similar for all the points.

The above evaluation is based upon a subjective judgement. Objective comparisons between different experiments have also been made by computing the root-mean-square (rms) and maximum (max) differences in the model fields between the initialized and uninitialized experiments. In Table 2, the rms and max differences in temperature (T), winds (u, v) and surface pressure (p_s) between the initialized and uninitialized analyses are shown. The numbers in the table are reasonably small and generally less than the observational errors. Compared with DOF, the rms and max differences are generally larger in IIR and INC and smaller in RAD and IAD (the max differences in temperature T are somewhat erratic). Similar calculations are made for the 24-h forecasts. The results in Table 3 show very similar characteristics to Table 2. In general, the rms differences are smaller for the forecasts than for the analyses. Using the figures in Table 2 and Table 3 to place the 5 schemes in order of merit leads to the same result as that from the subjective comparisons: (1) IAD, (2) RAD, (3) DOF, (4) INC and (5) IIR (ordered from the smallest to the largest changes).

The size of the changes induced by initialization should be determined by the degree of imbalance

in the analysis. In the experiments discussed above, while rms changes are acceptably small, maximum changes are larger than might be expected. These are due in part to phase shifting of the fields by the filtering process. In the following sections we discuss some alternative formulations of the filtering procedure which reduce the phase error so as to diminish the impact of the initialization on the analysis and forecast fields, while maintaining a satisfactory level of noise suppression.

6. Higher order filters

The order of the Quick-Start filter was set to 2 in the experiments discussed above. The simplest scheme using a recursive filter is that designated IIR: it involves only a single integration forward in time for a span of 1.5 h. However, the scheme resulted in a significant phase error. In this section we show how this problem may be overcome by using a filter of a higher order. Due to the non-recursive implementation of the recursive filter (Subsection 3.2), there is no extra demand on computer memory as the filter order is increased. A higher order filter should give better initialization results, as the ratio between the start-up and delay times is inversely proportional to the order of the filter (see Appendix A). The delays for Quick-Start filters of several orders, given in Table 1, show how Δ_0 increases with the order N . So, for a fixed span, the phase error should diminish as N is increased.

The values in Table 1 also show that the start-up time decreases with increasing order. This suggests that a shorter span may suffice for higher filter orders. However, the quantity T_0 does not truly reflect the time required for the transient response to die away: this response is not purely exponential, but involves products of powers of t and exponentials. An elementary example illustrates the problem: the function $f(t) = t \exp(1 - t)$ has an e -folding time $t_e = 1$. But the function grows from zero to unity as time t varies between these values, and it does not fall to $1/e$ times its maximum value until $t = 3.15$. The e -folding time governs the long-term decay but not the initial behaviour. It is found that higher order filters require a span longer than might be expected from the magnitude of T_0 .

The initialization experiments were repeated with the IIR scheme, using filter orders ranging

Table 3. *Root-mean-square (rms) and maximum (max) differences in temperature, wind components and surface pressure between the 24-h forecasts starting from the initialized data sets and the uninitialized forecast*

EXP	T (K)		u (ms^{-1})		v (ms^{-1})		p_s (hPa)	
	max	rms	max	rms	max	rms	max	rms
DOF	3.84	0.19	9.27	0.46	6.73	0.43	0.83	0.16
IIR	6.04	0.31	10.07	0.80	12.72	0.84	1.15	0.19
INC	6.11	0.27	7.59	0.70	9.11	0.69	1.30	0.20
RAD	6.11	0.17	6.65	0.44	6.70	0.41	0.99	0.12
IAD	6.15	0.16	8.25	0.38	6.68	0.36	0.71	0.08

Table 4. Root-mean-square (rms) and maximum (max) differences in temperature, wind components and surface pressure between the initialized and uninitialized analyses; results are for 2nd, 6th and 10th order filters

N	T (K)		u (ms ⁻¹)		v (ms ⁻¹)		p _s (hPa)	
	max	rms	max	rms	max	rms	max	rms
2	4.45	0.45	16.41	1.16	10.29	1.24	1.83	0.62
6	4.03	0.30	10.55	0.75	6.45	0.80	1.66	0.53
10	3.64	0.22	7.22	0.51	4.47	0.54	1.76	0.50

from 1 to 10. The span and cutoff were unchanged throughout, at 1.5 h and 3 h respectively. It was found that the changes induced in the initial fields become smaller as the order is increased; furthermore, the impact on the 24-h forecast also decreases. However, the residual noise becomes larger with increasing order. Thus, the choice of filter order is made so as to achieve adequate suppression of noise while simultaneously keeping the impact on the analysed and forecast fields at a reasonable level.

Table 4 shows the rms and maximum changes resulting from initialization with filters of orders 2, 6 and 10. Table 5 gives the corresponding figures for 24-h forecasts. Clearly, the higher order filters result in smaller changes. Fig. 7 is a plot of the mean absolute surface pressure tendency N_1 for the 3 forecasts. It is seen that the 10th-order filter does not remove high-frequency noise to a sufficient degree. The 6th-order filter would appear to be an excellent compromise: noise suppression is satisfactory, while changes to the analysis are at a

Table 5. Root-mean-square (rms) and maximum (max) differences in temperature, wind components and surface pressure between the 24-h forecasts starting from the initialized data sets and the uninitialized forecasts; results are for 2nd, 6th and 10th order filters

N	T (K)		u (ms ⁻¹)		v (ms ⁻¹)		p _s (hPa)	
	max	rms	max	rms	max	rms	max	rms
2	6.04	0.31	10.07	0.80	12.72	0.84	1.15	0.19
6	3.23	0.19	7.52	0.51	9.26	0.53	0.69	0.12
10	3.24	0.14	6.27	0.37	6.20	0.37	0.45	0.08




Fig. 7. Time evolution of the mean absolute surface pressure tendency N_1 (hPa/3 h). The parameters are the same as in the experiment IIR, except for the filter order: N02: 2nd order; N06: 6th order; N10: 10th order.

level comparable to observational error and impact on the 24-h forecast is acceptable. Note, in particular, the rms difference in surface pressure, which is only 0.12 hPa in this case. Plots of pressure and vertical velocity (not shown) confirm the efficacy of the 6th-order filter in eliminating gravity wave noise.

7. Filtered start-up or launch

There are several reasons for performing initialization. Foremost among these is the elimination of spurious high-frequency gravity wave noise. This is essential for assimilation of observational data in an analysis/forecast cycle. Initialization also produces realistic diagnostic fields, such as divergence and vertical velocity, at the analysis time. Digital filtering has also been shown to yield realistic fields of prognostic variables, such as cloud water content, for which no observations are available (Huang and Sundqvist, 1993). Thus, initialization enhances and extends the analysis. However, if the sole requirement is noise suppression, initialization may be replaced by a filtering process which produces balanced fields valid at a time later than $t = 0$. This has the following advantages: overall computation time may be reduced (Subsection 7.1) or phase error may be completely eliminated (Subsection 7.2). To distinguish it from pure initialization, we will refer to this procedure as filtered launching.

7.1. Reversal of filter coefficients

The recursive filters discussed above have been so constructed that the output applies to the initial time. By a simple reversal of the coefficients (replacing F_{K_n} by F_{KK-n} in (8)) the output is valid at the end of the filter span $T_s = K \Delta t$. We may now replace the model variables at this point by the filter output and proceed with the integration. The advantage of this launching process is that there is no overlap between the initialization and the forecast, so that CPU time is practically the same as for an uninitialized forecast.

We present here the results of applying the same sixth-order filter as in Section 6: the cutoff is $\tau_c = 3$ h and the span $T_s = 1.5$ h, but the filter coefficients are reversed. The fields at $t = 1.5$ h are replaced by the filtered fields, a forward step is made and the forecast proceeds normally thereafter. The noise parameter N_1 is shown in Fig. 8. The curves NOI and N06 are for the uninitialized forecast and that filtered as in Section 6. The curve DFL results from the launching method: it coincides with NOI until time $t = T_s$, at which point the noise drops from around 7 hPa/3 h to about 3. The difference between the 24-h forecasts using the two alternative methods and the uninitialized forecast are given in Table 6. The impact on the forecast scores is about the same for the two schemes. The advantage of the filtered launch with reversed coefficients is its economy: there is virtually no additional computational cost in applying this scheme.




Fig. 8. Time evolution of the mean absolute surface pressure tendency N_1 (hPa/3 h). NOI: no initialization; N06: recursive initialization with 6th order filter; DFL: filtered launch with 6th order filter.

Table 6. Root-mean-square (rms) and maximum (max) differences in temperature, wind components and surface pressure between the 24-h forecasts starting from the initialized data sets and the uninitialized forecasts; results are for initialization with 6th-order filter (N06) and filtered start-up, with coefficients reversed (DFL)

EXP	T (K)		u (ms ⁻¹)		v (ms ⁻¹)		p _s (hPa)	
	max	rms	max	rms	max	rms	max	rms
N06	3.32	0.19	7.52	0.51	9.26	0.53	0.69	0.12
DFL	3.19	0.20	7.06	0.52	8.86	0.55	0.77	0.13

7.2. Phase – corrected launching

The 2nd-order filter applied in Section 5 was found to suppress noise effectively if applied over a span $T_s = 1.5$ h. But the delay of this filter, for a cutoff period $\tau_c = 3$ h, was $\delta_0 \approx 0.6$ h. When the output was assumed to be valid at $t = 0$ and a 24-h forecast made, the magnitude of the rms and maximum differences from the uninitialized forecast suggested a phase error; the evolution of p_s and ω at a particular point reinforced this suspicion (see curves marked IIR in Figs. 5, 6). If our interpretation of the delay parameter δ_0 is correct, the filter output should validate not at $t = 0$ but at $t = T_s - \delta_0$.

To confirm this argument, the second-order filter was applied over a span of $T_s = 1.5$ h, as in




Fig. 9. Time evolution of the mean absolute surface pressure tendency N_1 (hPa/3 h). NOI: no initialization; IIR: recursive initialization with 2nd order filter; PCL: phase-corrected launch with 2nd order filter.

Section 5, and the time-step reset to $T_s - \delta_0$; the model variables were then replaced by the filter output, a forward step made and the integration continued normally thereafter. The noise parameter N_1 is shown in Fig. 9. The curves NOI and IIR are for the uninitialized forecast and that initialized as in Section 5. The curve PCL results from the launching method: it coincides with NOI until time $t = T_s$, at which point the noise level is about 7 hPa/3 h. The filter output, with a value of N_1 around 2 hPa/3 h, is then used to continue the forecast from time $t = T_s - \delta_0 \approx 0.9$ h. It is similar to the curve for standard initialization (dotted line) shifted by 0.9 h. The difference between the 24-h forecasts using the two alternative methods and the uninitialized forecast are given in Table 7. The impact on the forecast scores is much smaller for the launching scheme than for the initialization. All the errors in the row marked PCL are significantly less than those marked IIR. Note, in particular, the rms error in surface pressure for PCL: it is only 0.04 hPa; for practical purposes, this forecast is identical to the reference run.

One advantage of the filtered launch is its economy: it requires only the computation time for an integration of duration δ_0 , since this is the amount by which we step back the forecast; in the present case it is equivalent to a forecast of about a half-hour. The other advantage is the minute impact on the forecast: the differences in Table 7 confirm the absence of phase error. The disadvantage of launching is the loss of other benefits of initialization, such as the production of realistic diagnostic fields valid at the analysis time. The choice depends on operational requirements.

Table 7. Root-mean-square (rms) and maximum (max) differences in temperature, wind components and surface pressure between the 24-h forecasts starting from the initialized data sets and the uninitialized forecasts; results are for initialization with 2nd-order filter (IIR) and phase-corrected launch with the same filter (PCL)

EXP	T (K)		u (ms ⁻¹)		v (ms ⁻¹)		p _s (hPa)	
	max	rms	max	rms	max	rms	max	rms
IIR	6.04	0.31	10.07	0.80	12.72	0.84	1.15	0.19
PCL	3.08	0.10	6.05	0.23	3.93	0.20	0.42	0.04

8. Conclusions

In this paper, a number of recursive filters have been investigated. Initialization schemes based on a particular feature of these filters (the phase-lag or delay of output with respect to input) are formulated and tested with a sophisticated numerical weather prediction model, HIRLAM. When a low order recursive filter is applied to the full fields produced by a one-sided integration, a phase error is incurred; this is due to the need to use a span considerably longer than the delay to achieve adequate filtering. It was shown that this problem could be overcome by applying the filter to the analysis increment (the INC scheme), by preceding the forward diabatic filtering by a reverse adiabatic filtering (the RAD scheme) or by a combination of these methods (the IAD scheme).

Filters of higher order were also investigated. It was shown that no extra storage was necessary if the non-recursive implementation (8) was used. The impact on analysis and forecast is less with higher order filters (for the same cutoff and span) but the degree of filtering is also less. It was found that a 6th-order filter represents a good compromise: the changes made to the analysed fields are of the order of observational errors; the impact on the 24-h forecast is very small; and elimination of high frequency noise is satisfactory.

The non-recursive implementation of the recursive filters makes the new initialization schemes as easy to use as the original scheme of Huang and Lynch (1993). In addition, some CPU time is saved by using the new schemes, since recursive filters are in general more effective than non-recursive ones. The sixth-order IIR scheme yields results comparable to the method (DOF) of Huang and Lynch (1993) at less than half the cost: the latter requires a 1.5-h backward adiabatic integration and a 3-h diabatic forecast, whereas the former requires only a 1.5-h forecast.

The INC scheme requires two 1.5-h integrations, but the forecast from the first-guess field can be performed as soon as this field is available, so it involves no extra delay in an operational setting. The added cost of the mixed schemes, RAD and IAD, may be justified if their superior noise elimination characteristics are required. These schemes were tested with a 2nd-order filter, but should yield even better results with the 6th-order filter.

An alternative way of applying the filter was briefly described: if the sequence of coefficients is reversed, the output may be assumed to apply at the end of the span, so that the forecast can be continued from the end of the span and no extra integration time is required. This filtered launching provides a very economical means of eliminating high frequency noise if the other advantages of initialization are not required. A version of the launching described in Subsection 7.2 enables virtually complete removal of phase error, while still being very economical.

Finally, it may be noted that recursive filters of the type considered may be applied repeatedly during a forecast if noise control is desirable, for example in climate modelling. With the order of the filter coefficients reversed, no extra integration time is required. These filters may also be useful in the formulation of constraints for four-dimensional data assimilation; if they are used in a forward-backward cycle, the nett phase-error vanishes.

9. Acknowledgments

It is a pleasure to thank Nils Gustafsson, Erland Källén and Peter Lönnberg for their helpful comments and reviews of an earlier version of this paper.

10. Appendix A

Analog filters; the quick-start filter

A.1. General properties of analog filters

Analog filters are closely associated with linear differential equations. Let us consider the linear second order o.d.e. governing a forced, damped harmonic oscillator:

$$\ddot{y} + 2\sigma_0 \dot{y} + (\sigma_0^2 + \omega_0^2) y = x(t). \quad (\text{A1})$$

The coefficient $\sigma_0 > 0$ determines the damping and ω_0 is the natural frequency of the system. The Laplace transform of (A1) is

$$(s^2 + 2\sigma_0 s + (\sigma_0^2 + \omega_0^2)) \hat{y}(s) = \hat{x}(s) + (s + 2\sigma_0) y(0) + \dot{y}(0) \quad (\text{A2})$$

(hats denote transforms). The solution for zero initial conditions, $y(0) = \dot{y}(0) = 0$, is

$$\hat{y}(s) = H(s) \hat{x}(s), \quad (\text{A3})$$

where the *transfer* function or *system* function $H(s)$ is given by

$$H(s) = \frac{1}{s^2 + 2\sigma_0 s + (\sigma_0^2 + \omega_0^2)}. \quad (\text{A4})$$

For $\omega_0^2 > 0$ this function has complex conjugate poles in the s -plane, at $s_p = -\sigma_0 \pm i\omega_0$. The response function (A4) has the character of a low-pass filter: for large $|s|$, $H(s)$ behaves like s^{-2} , so its amplitude tends to zero as $\omega \rightarrow \pm\infty$, and high frequency inputs are effectively damped. For a sinusoidal input, $x(t) = \exp(i\omega t)$, the steady-state response is

$$y(t) = \left[\frac{1}{(s_p - i\omega)(\bar{s}_p - i\omega)} \right] e^{i\omega t} = H(i\omega) \cdot x(t).$$

Thus, pure sinusoidal input gives rise to an output of the same frequency, whose amplitude and phase are determined by $H(i\omega)$, the transfer function evaluated on the imaginary axis of the s -plane. We can write $H(i\omega)$ in terms of its amplitude and phase:

$$H(i\omega) = M(\omega) \exp(i\delta(\omega)),$$

and expressions for $M(\omega)$ and $\varphi(\omega)$ follow from (A4):

$$M(\omega) = \sqrt{\frac{1}{[\sigma_0^2 + (\omega_0 - \omega)^2][\sigma_0^2 + (\omega_0 + \omega)^2]}}, \quad (\text{A5})$$

$$\varphi(\omega) = -\tan^{-1} \left(\frac{\omega - \omega_0}{\sigma_0} \right) - \tan^{-1} \left(\frac{\omega + \omega_0}{\sigma_0} \right). \quad (\text{A6})$$

The special case of linear phase $\varphi = -\delta_0 \omega$ corresponds to delay in time for all components of the input irrespective of frequency. The time delay is given in this case by $\delta_0 = -d\varphi/d\omega$. Generalising this idea, we define the group delay, δ , as the derivative of the phase response (A6) with respect to frequency: $\delta = -d\varphi/d\omega$. We will refer to the

value of δ at $\omega = 0$ simply as the *delay*; from (A6) we get:

$$\delta_0 = \frac{2\sigma_0}{\sigma_0^2 + \omega_0^2}. \quad (\text{A7})$$

The point $s = i\omega_c$ where the magnitude-squared equals half the value at $\omega = 0$ determines the *cutoff frequency*:

$$|H(i\omega_c)|^2 = \frac{1}{2} |H(0)|^2. \quad (\text{A8})$$

It is convenient to consider *prototype* filters, for which $\omega_c = 1$; the general case is recovered by the transformation $s \rightarrow s/\omega_c$. Using (A4) and (A5), (A8) implies:

$$(\sigma_0^2 + \omega_0^2)^2 = 2(\sigma_0^2 - \omega_0^2) + 1. \quad (\text{A9})$$

This is the equation of a curve in the s -plane (Fig. 10). It is a particular case of a family of curves called Cassinian ovals. Since all prototype filters have parameters (σ_0, ω_0) lying on the curve (A9), there is in effect one free parameter: we may impose one further condition on the filter. Three special choices were considered in Lynch (1993a):




Fig. 10. This peanut-shaped curve is the locus in the s -plane of the poles $s_p = -\sigma_0 \pm i\omega_0$ of prototype 2nd-order filters. The 3 special cases considered in the text are indicated (BW, FD, QS). The curve is described by the equation $(\sigma_0^2 + \omega_0^2)^2 = 2(\sigma_0^2 - \omega_0^2) + 1$. The unit circle is shown for reference.

the Butterworth filter, whose amplitude response is maximally flat in the pass-band; the Bessel or flat-delay filter, with a phase response as nearly linear as possible in the pass-band; and the Quick-Start filter, with a transient response which decays as rapidly as possible. The points on (A9) corresponding to these three cases are indicated in Fig. 10 by BW, FD and QS respectively; the last of these will be described below.

A.2. The Quick-Start filter

Without forcing, the solution of (A1) is evanescent for $\sigma_0 > 0$; the transient response decays with an e -folding time $1/\sigma_0$. Thus, the most rapid decay obtains for the largest possible value of σ_0 . For the prototype second-order filter (A4) this is the point where the curve (A9) intersects the horizontal axis, with coordinates

$$\begin{aligned} \sigma_0 &= \sqrt{\sqrt{2} + 1} \\ &= \sqrt{\frac{1}{\sqrt{2} - 1}}, \quad \omega_0 = 0. \end{aligned} \quad (\text{A10})$$

The output of the filter based on this choice should reach a steady state faster than for any other prototype filter. It corresponds to the case of critical damping for a harmonic oscillator. We shall call the filter based on this choice the Quick-Start (QS) filter (no alternative name having been found in the literature). The time-constant $1/\sigma_0$ is a measure of the decay-rate of the transient response. We shall call this the *start-up* time $T_0 = 1/\sigma_0$. We define the non-dimensional ratio of the start-up and delay times:

$$\Gamma = \frac{\text{start-up}}{\text{delay}} = \frac{T_0}{\delta_0}. \quad (\text{A11})$$

For application to initialization it is desirable that Γ be as small as possible, so as to achieve maximum attenuation of the transient within the delay period. The minimum value $\Gamma = \frac{1}{2}$ is attained by the Quick-Start filter. Generalization to arbitrary order is straightforward: the N -pole filter whose transient solution decays most rapidly has all its poles on the negative real axis, at $s_p = -\sigma$:

$$H(s) = \frac{\sigma^N}{(s + \sigma)^N}. \quad (\text{A12})$$

The prototype version ($\omega_c = 1$) corresponds to the value $\sigma = \sqrt{1/(2^{1/N} - 1)}$. We will refer to this as the N th-order quick-start filter. It is easy to show that the delay is $\delta_0 = N/\sigma$, so that

$$\Gamma = \frac{1/\sigma}{N/\sigma} = \frac{1}{N}, \quad (\text{A13})$$

which decreases indefinitely with increasing order.

11. Appendix B

Inductive construction of the filter matrix F

The filter matrix F in (7) can be constructed row-by-row without matrix inversion. Let us define input and output vectors (of order $n+1$) by

$$\mathbf{x}_n = (x_0, \dots, x_n), \quad \mathbf{y}_n = (y_0, \dots, y_n)$$

and suppose they are related by the following expression:

$$\mathbf{y}_n = \mathbf{F}_n \mathbf{x}_n. \quad (\text{B1})$$

We wish to derive a similar relationship at step $n+1$. The recursive filter relationship (2) is

$$\begin{aligned} y_{n+1} &= \sum_{k=0}^N a_k x_{n+1-k} \\ &+ \sum_{k=1}^N b_k y_{n+1-k}. \end{aligned} \quad (\text{B2})$$

We define a coefficient vector $\mathbf{a}_n$ (of order $n+1$) by

$$\mathbf{a}_n = (\underbrace{0, \dots, 0}_{n+1-N \text{ terms}}, \underbrace{a_N, \dots, a_1}_N),$$

and $\mathbf{b}_n$ in a similar way. Then (B2) may be written

$$y_{n+1} = \mathbf{a}_n \mathbf{x}_n + \mathbf{b}_n \mathbf{y}_n + a_0 x_{n+1}. \quad (\text{B3})$$

This provides the next row of the matrix: using (B1) we can write

$$\mathbf{y}_{n+1} = \mathbf{F}_{n+1} \mathbf{x}_{n+1}, \quad (\text{B4})$$

where the filter matrix $\mathbf{F}_{n+1}$ is given in block form:

$$\mathbf{F}_{n+1} = \begin{pmatrix} \mathbf{F}_n & \mathbf{0} \\ \mathbf{a}_n + \mathbf{b}_n \mathbf{F}_n & a_0 \end{pmatrix}.$$

The initial rows are normally constructed by means of lower order filters: thus, $y_0 = x_0$,

$$\begin{aligned} y_1 &= (a_0 x_1 + a_1 x_0) + b_1 y_0 \\ &= a_0 x_1 + (a_1 + b_1) x_0 \end{aligned}$$

and so on until $n = N$, after which (B2) is used. In this paper the coefficients for each order n are those of a quick-start filter having the specified cutoff period τ_c . These are easily generated from (A12).

REFERENCES

- Ballish, B., Cao, X., Kalnay, E. and Kanamitsu, M. 1992. Incremental nonlinear normal-mode initialization. *Mon. Wea. Rev.* **120**, 1723–1734.
- Huang, X.-Y. and Lynch, P. 1993. Diabatic digital filtering initialization: Application to the HIRLAM model. *Mon. Wea. Rev.* **121**, 589–603.
- Huang, X.-Y. and Sundqvist, H. 1993. Initialization of cloud water content and cloud cover for numerical prediction models. *Mon. Wea. Rev.* **121**, 2719–2726.
- Kulhánek, O. 1976. *Introduction to digital filtering in geophysics*. Elsevier Sci. Publ. Co., Amsterdam, 168 pp.
- Kållberg, P. (ed.) 1989. *The HIRLAM level 1 system. Documentation manual*. Available from SMHI, S-601 76 Norrköping, Sweden.
- Lynch, P. 1993a. Digital filters for numerical weather prediction. *HIRLAM Technical Report no. 10*. (Available from FMI, PB 503, SF-00101 Helsinki, Finland.)
- Lynch, P. 1993b. Initialization using a recursive digital filter. In: *Research activities in atmospheric and ocean modeling*, CAS/JSC Working Group on Numerical Experimentation. Report No. 18, WMO Secretariat, Geneva, pp. 1.21–1.22.
- Lynch, P. and Huang, X.-Y. 1992. Initialization of the HIRLAM model using a digital filter. *Mon. Wea. Rev.* **120**, 1019–1034.
- Oppenheim, A. V. and Schaffer, R. W. 1989. *Discrete-time signal processing*. Prentice-Hall International, Inc., 879 pp.
- Puri, K., Bourke, W. and Seaman, R. 1982. Incremental linear normal mode initialization in four-dimensional data assimilation. *Mon. Wea. Rev.* **110**, 1773–1785.