

Numerical Aspects of the Application of Recursive Filters to Variational Statistical Analysis. Part I: Spatially Homogeneous and Isotropic Gaussian Covariances

R. JAMES PURSER AND WAN-SHU WU

Science Applications International Corporation, Beltsville, Maryland

DAVID F. PARRISH

National Centers for Environmental Prediction/Environmental Modeling Center, Washington, D.C.

NIGEL M. ROBERTS

Joint Centre for Mesoscale Meteorology, University of Reading, Reading, and Met Office, Bracknell, United Kingdom

(Manuscript received 18 February 2002, in final form 4 October 2002)

ABSTRACT

The construction and application of efficient numerical recursive filters for the task of convolving a spatial distribution of “forcing” terms with a quasi-Gaussian self-adjoint smoothing kernel in two or three dimensions are described. In the context of variational analysis, this smoothing operation may be interpreted as the convolution of a covariance function of background error with the given forcing terms, which constitutes one of the most computationally intensive components of the iterative solution of a variational analysis problem.

Among the technical aspects of the recursive filters, the problems of achieving acceptable approximations to horizontal isotropy and the implementation of both periodic and nonperiodic boundary conditions that avoid the appearance of spurious numerical artifacts are treated herein. A multigrid approach that helps to minimize numerical noise at filtering scales greatly in excess of the grid step is also discussed. It is emphasized that the methods are not inherently limited to the construction of purely Gaussian shapes, although the detailed elaboration of methods by which a more general set of possible covariance profiles may be synthesized is deferred to the companion paper (Part II).

1. Introduction

There are many methods available for objectively analyzing the meteorological data required to initialize a numerical weather prediction model (e.g., see Daley 1991). Those methods based on formal statistical principles (e.g., Gandin 1963; Lorenc 1986, 1997; Parrish and Derber 1992; Courtier et al. 1998), which permit a proper accounting to be taken of multivariate aspects of the problem, have now largely superseded the overtly empirical methods of “successive corrections” (Bergthorsen and Döös 1955; Cressman 1959; Barnes 1964). Nevertheless, for specialized applications, the empirical methods continue to enjoy the advantages of greater computational efficiency and the ability to adapt more flexibly to the typically large inhomogeneities of density and quality of the available data. While the high efficiency of empirical methods becomes progressively less

of a critical factor as available computational power continues to increase, adaptivity remains a factor of considerable importance in circumstances where the day-to-day variability of data quality and quantity are hard to predict beforehand, such as occurs in the processing of satellite sounding data. In this context Hayden and Purser (1995), following up on the work of Purser and McQuigg (1982), developed a numerically efficient and spatially adaptive analysis scheme using spatial smoothers. Each spatial smoother was built up of more basic numerical operators consisting of rather simple recursive filters acting unidirectionally upon the gridded data residuals.

The numerical efficiency of these basic operators can also be turned to advantage within a statistical analysis scheme, specifically in the synthesis of the effective covariance-convolution operators needed by the descent algorithms of the large-scale linear solvers involved (Lorenc 1992). The Statistical Spectral Interpolation (SSI) of the National Centers for Environmental Prediction (NCEP) is an example of an analysis scheme in which the spectral representation of the background error covariance is employed directly (Parrish and Derber

Corresponding author address: R. James Purser, NOAA/NCEP, W/ NP2 Rm. 207, WWBG, 5200 Auth Rd., Camp Springs, MD 20746-4304.

E-mail: jim.purser@noaa.gov

1992). Methods of this type are inherently limited in their ability to deal conveniently with geographical inhomogeneities. Although one motivation of the present two-part study was to develop the tool of recursive filters to allow the operational three-dimensional variational analysis (3DVAR) scheme to accommodate spatial inhomogeneities in the background covariance, the inhomogeneous and anisotropic aspects of the filtering technique will be reserved for the companion paper (Purser et al. 2003, referred to henceforth as Part II); the focus of the present paper is to demonstrate the ability of appropriately constructed recursive filters to achieve acceptably isotropic covariance functions of Gaussian form. Part II will extend this study to more general non-Gaussian profiles and explore the case of spatially adaptive covariances of either locally isotropic or generally anisotropic forms.

A brief review of the ideas that underlie 3DVAR is given in section 2 in order to clarify the points at which the recursive filter plays a part. In section 3 we set forth the relevant theory pertaining to the construction of basic recursive filters capable of being forged into convolution operators reasonably representing the qualities desired by modeled covariance convolutions within an adaptive analysis scheme with a uniform Cartesian grid and with homogeneous covariances. Like the Gaussian covariances of Derber and Rosati (1989), which are obtained by multiple iterations of a diffusion operator, the basic recursive filters are crafted to produce approximately Gaussian smoothing kernels (but in fewer numerical operations than are typical in the explicit diffusion method). Some of the technicalities discussed in this section are treated in greater detail in the appendixes. In another paper, Wu et al. (2002), we provide examples of the applications of some of the techniques presented here to global variational analysis of meteorological data.

2. 3DVAR

In this section we attempt to follow the notation of Ide et al. (1997), writing the abstract vector representing the atmospheric state as \mathbf{x} , with “background” and “analysis” versions of this indicated by subscripts, that is, \mathbf{x}_b and \mathbf{x}_a . The component of error in the background is denoted $\boldsymbol{\epsilon}_b$:

$$\mathbf{x}_b = \mathbf{x} + \boldsymbol{\epsilon}_b. \quad (2.1)$$

The observational data are collected into another abstract vector \mathbf{y}_o whose components are related to the state vector \mathbf{x} through the application of a generalized, possibly nonlinear, interpolation operator \mathcal{H} together with an effective measurement error $\boldsymbol{\epsilon}_o$:

$$\mathbf{y}_o = \mathcal{H}(\mathbf{x}) + \boldsymbol{\epsilon}_o. \quad (2.2)$$

The statistical characteristics of the errors $\boldsymbol{\epsilon}_b$ and $\boldsymbol{\epsilon}_o$ are quite difficult to describe in complete detail owing to numerous complicating factors. Among the common

simplifying assumptions, we normally assume unbiasedness:

$$\langle \boldsymbol{\epsilon}_b \rangle = 0, \quad (2.3)$$

$$\langle \boldsymbol{\epsilon}_o \rangle = 0. \quad (2.4)$$

The covariance, $\mathbf{R} \equiv \langle \boldsymbol{\epsilon}_o \boldsymbol{\epsilon}_o^T \rangle$, of observational error is assumed to be diagonal, equivalent to assuming the observational errors are statistically independent. The corresponding covariance, $\mathbf{B} \equiv \langle \boldsymbol{\epsilon}_b \boldsymbol{\epsilon}_b^T \rangle$, of background error, however, is *never* assumed to be diagonal in its representation based on the state-space constructed from the gridded value components; the characteristically smooth form in space of background errors implies that neighboring points have errors, in fields of the same type, that are strongly positively correlated. Although the principles of variational analysis can accommodate strong nonlinearities if required, it is often numerically convenient to exploit the typically weak nonlinearity of \mathcal{H} by approximating the effects on $\mathcal{H}(\mathbf{x})$ of small increments of \mathbf{x} , using the linearization of \mathcal{H} :

$$\mathbf{H}d\mathbf{x} = d\mathcal{H}(\mathbf{x}). \quad (2.5)$$

The “primal” variational principle in 3DVAR seeks the minimum over \mathbf{x} of the penalty function $\mathcal{L}_1(\mathbf{x})$ defined by

$$2\mathcal{L}_1(\mathbf{x}) = (\mathbf{x} - \mathbf{x}_b)^T \mathbf{B}^{-1} (\mathbf{x} - \mathbf{x}_b) + [\mathbf{y} - \mathcal{H}(\mathbf{x})]^T \mathbf{R}^{-1} [\mathbf{y} - \mathcal{H}(\mathbf{x})], \quad (2.6)$$

which may be justified by minimum-variance arguments. The solution, $\mathbf{x} = \mathbf{x}_a$, then obeys

$$\mathbf{B}^{-1}(\mathbf{x}_a - \mathbf{x}_b) = \mathbf{H}^T \mathbf{R}^{-1} [\mathbf{y} - \mathcal{H}(\mathbf{x}_a)] \quad (2.7)$$

and, hence, must define an increment of the form

$$\mathbf{x}_a - \mathbf{x}_b = \mathbf{B} \mathbf{H}^T \mathbf{f}, \quad (2.8)$$

where

$$\mathbf{f} = \mathbf{R}^{-1} [\mathbf{y} - \mathcal{H}(\mathbf{x}_a)]. \quad (2.9)$$

The inherent smoothness of the background field errors $\boldsymbol{\epsilon}_b$, and hence that of the covariance \mathbf{B} of these errors, is therefore imprinted on the analysis increments themselves.

If we neglect the effects of nonlinearity, we find that, instead of solving for \mathbf{x} directly, we can instead first solve for the smaller vector \mathbf{f} in the implied “dual” variational principle that minimizes $\mathcal{L}_2(\mathbf{f})$ defined by

$$2\mathcal{L}_2(\mathbf{f}) = \mathbf{f}^T (\mathbf{R} + \mathbf{H} \mathbf{B} \mathbf{H}^T) \mathbf{f} - 2\mathbf{f}^T \mathbf{d}, \quad (2.10)$$

where \mathbf{d} is the “innovation”:

$$\mathbf{d} = \mathbf{y}_o - \mathcal{H}(\mathbf{x}_b). \quad (2.11)$$

Treating this dual form of the analysis problem is essentially the approach adopted at the Data Assimilation Office of the National Aeronautics and Space Administration (NASA) Goddard Space Flight Center and described by da Silva et al. (1995) and by Cohn et al.

(1998). This “physical-space statistical analysis system” (PSAS) may be advantageous when using sophisticated preconditioning methods based on careful grouping of data, as discussed by Daley and Barker (2000). However, when only the simplest preconditioning strategies are employed, Courtier (1997) shows that the primal and dual forms imply essentially identical condition numbers for the linear inversion problems. Regardless of which form of 3DVAR is adopted, one must rely on iterative methods to converge toward a practical solution. The most costly part of each iterative step of a solution algorithm is the multiplication of some grid-space vector \mathbf{v} by the covariance matrix \mathbf{B} . This is required once per iteration whether treating the primal or the dual form of the problem. Even a single multiplication of this form becomes prohibitively expensive if carried out explicitly with the full matrix having the dimensionality of \mathbf{x} . One remedy, adopted by Gaspari and Cohn (1998, 1999) and extending earlier work of Oliver (1995), is to constrain the allowed covariance models to those possessing compact support. They show how families of approximately Gauss-shaped compact-support covariance models may be formulated and applied with much greater numerical efficiency than would be possible for more general functions. Our own approach to achieving reasonable efficiency without restricting the covariances to being of the compact-support type is to assume that an operation of the form $\mathbf{B}\mathbf{v}$ may be factored into smaller, less costly, factors.

In the first step, the multivariate structure of \mathbf{B} is broken apart by the judicious selection of a set of non-standard analysis variables for which the contributions from \mathbf{B} naturally separate out. For example, a single variable representing the quasigeostrophically balanced combination of mass and rotational wind fields can be attributed a univariate spatial covariance for its background errors quite independently of the corresponding spatial covariance for the residual unbalanced rotational wind component. Meanwhile, the divergent wind field can be treated independently of either. Further steps in the program of reducing the operator \mathbf{B} might be, next, to carry out a crude separation of a few additive components of the operator on the basis of their characteristic spatial scales. If this can be done to render the resulting operator components into Gaussian forms, then, in the absence of anisotropies obliquely oriented with respect to the grid, the Gaussians themselves may be factored into the three respective coordinate directions. [A more general discussion of radially symmetric covariance functions can be found in the work of Oliver (1995).] Finally, along each single dimension, a computational advantage may be gained by employing a spatially recursive filter, carefully constructed to mimic the required Gaussian convolution operator, but at a fraction of the still considerable cost of applying directly the explicit Gaussian convolution operator itself. It is the objective of the following sections to reveal precisely how such a recursive filter may be fabricated and

applied. We do *not* wish to imply that the Gaussian form is inherently desirable in data assimilation. On the contrary, careful investigation of the spatial profiles of forecast background error (Thiébaux 1976; Thiébaux et al. 1986; Hollingsworth and Lönnberg 1986) reveals covariance functions that cannot be reconciled with the Gaussian shape alone. However, by treating the two- or three-dimensional quasi-Gaussian filter combination as a relatively cheap “building block,” a far larger range of possible profile shapes becomes accessible, by the superposition of appropriately weighted combinations of quasi Gaussians of different sizes and by the application of the negative Laplacian operator to such components in order to induce the negatively correlated sidelobe characteristics of some components of background error. These aspects are dealt with in Part II. Thus, the motivating consideration for using recursive filters in this context is predominantly that of computational efficiency together with the recognition that much more general forms become available through the exploitation of superposition.

3. Homogeneous recursive filtering theory

The theory of digital filtering was initially developed in the context of time series analysis. However, many of the same techniques are equally applicable in two or more spatial dimensions when the numerical grid is of a sufficiently regular configuration, as it usually is in numerical weather analysis. While we attempt to keep the technical discussion of this section self-contained, other related aspects of the topic of digital filter design are well covered by standard texts such as Otnes and Enochson (1972) and Papoulis (1984).

a. Quasi-Gaussian recursive filters in one dimension

Let $K/\delta x^2$ denote the finite-difference operator:

$$K(\psi)_i/\delta x^2 = -(\psi_{i-1} - 2\psi_i + \psi_{i+1})/\delta x^2, \quad (3.1)$$

approximating the differential operator, $-d^2/dx^2$, on a line grid of uniform spacing δx . The spectral representation of the operator at wavenumber k (wavelength $2\pi/k$) is

$$\hat{K}(k) = \left[2 \sin\left(\frac{k\delta x}{2}\right) \right]^2.$$

Inverting this relationship, we obtain a formula for k^2 in terms of \hat{K} :

$$k^2 = \frac{4}{\delta x^2} \left[\arcsin\left(\frac{\hat{K}^{1/2}}{2}\right) \right]^2. \quad (3.2)$$

Clearly, the same formula relates operator $-d^2/dx^2$ to operator K ; in fact, the algebraic manipulations we set forth here can be regarded as an application of the “calculus of operators” (Dahlquist and Björck 1974, p. 311). Using the standard expansion,

TABLE 1. Coefficients b_{ij} for quasi-Gaussian filters up to degree six.

i	$j = 1$	$j = 2$	$j = 3$	$j = 4$	$j = 5$	$j = 6$
1	1	$\frac{1}{12}$	$\frac{1}{90}$	$\frac{1}{560}$	$\frac{1}{3150}$	$\frac{1}{16\ 632}$
2		1	$\frac{1}{6}$	$\frac{7}{240}$	$\frac{41}{7560}$	$\frac{479}{453\ 600}$
3			1	$\frac{1}{4}$	$\frac{13}{240}$	$\frac{695}{60\ 480}$
4				1	$\frac{1}{3}$	$\frac{31}{360}$
5					1	$\frac{5}{12}$
6						1

$$\arcsin(z) = \sum_{i=0}^{\infty} \gamma_i z^{2i+1}, \quad |z| < 1, \quad (3.3)$$

where

$$\begin{aligned} \gamma_i &= \frac{1}{(2i+1)} \cdot \frac{(2i-1)!!}{(2i)!!} \\ &\equiv 1, \frac{1}{3}, \frac{1}{2}, \frac{1}{5}, \frac{1.3}{2.4}, \frac{1}{7}, \frac{1.3.5}{2.4.6}, \dots, \end{aligned} \quad (3.4)$$

we may obtain a power expansion for $k^2 \delta x^2$, and thence, the expansions for the term $(k^2 \delta x^2)^i$:

$$(k^2 \delta x^2)^i = \sum_{j=i}^{\infty} b_{i,j} \hat{K}^j. \quad (3.5)$$

From (3.2) and (3.3) we find that

$$b_{1,j} \equiv \sum_{k=0}^{j-1} \frac{\gamma_k \gamma_{j-1-k}}{4^{(j-1)}} = \frac{2(j!)^2}{j^2(2j)!}, \quad (3.6)$$

and other coefficients then follow by the application of the inductive step:

$$b_{i,j} = \sum_{k=1}^{j+1-i} b_{1,k} b_{i-1,j-k}. \quad (3.7)$$

The coefficients, $b_{i,j}$, which are all positive and rational, are listed in Table 1 for $j \leq 6$.

Consider the differential operator, $\mathcal{D}_{(n)}$,

$$\mathcal{D}_{(n)} = 1 - \frac{a^2}{2} \frac{d^2}{dx^2} + \frac{a^4}{2!4} \frac{d^4}{dx^4} + \dots + \frac{1}{n!} \left(-\frac{a^2}{2} \frac{d^2}{dx^2} \right)^n, \quad (3.8)$$

whose spectral representation is

$$\hat{\mathcal{D}}_{(n)} \equiv 1 + \frac{a^2 k^2}{2} + \frac{a^4 k^4}{2!4} + \dots + \frac{1}{n!} \left(\frac{a^2 k^2}{2} \right)^n$$

or, with $\sigma = a/\delta x$,

$$\begin{aligned} \hat{\mathcal{D}}_{(n)} &= 1 + \sigma^2 \left(\frac{k^2 \delta x^2}{2} \right) + \frac{\sigma^4}{2!} \left(\frac{k^2 \delta x^2}{2} \right)^2 + \dots \\ &+ \frac{\sigma^{2n}}{n!} \left(\frac{k^2 \delta x^2}{2} \right)^n. \end{aligned} \quad (3.9)$$

Since

$$\lim_{n \rightarrow \infty} \hat{\mathcal{D}}_{(n)} = \exp\left(\frac{a^2 k^2}{2}\right), \quad (3.10)$$

the substitution of each power series (3.5) up to degree n for the powers of k^2 into (3.9) gives us a way of approximating this exponential function in terms of \hat{K} :

$$\begin{aligned} \hat{\mathcal{D}}_{(n)}^* &= 1 + \frac{\sigma}{2} \sum_{j=1}^n b_{1,j} \hat{K}^j + \frac{1}{2!} \left(\frac{\sigma^2}{2} \right)^2 \sum_{j=2}^n b_{2,j} \hat{K}^j + \dots \\ &+ \frac{1}{n!} \left(\frac{\sigma^2}{2} \right)^n b_{n,n} \hat{K}^n. \end{aligned} \quad (3.11)$$

Correspondingly, there is a finite-difference operator, $\mathcal{D}_{(n)}^*$, composed of the n th-degree expansion of K implied by this approximation, which, following a rearrangement of terms, we may write

$$\begin{aligned} \mathcal{D}_{(n)}^* &= 1 + b_{1,1} \frac{\sigma^2}{2} K + \left[b_{1,2} \left(\frac{\sigma^2}{2} \right) + \frac{b_{2,2}}{2!} \left(\frac{\sigma^2}{2} \right)^2 \right] K^2 \\ &+ \dots + \left[\sum_{j=1}^n \frac{b_{j,n}}{j!} \left(\frac{\sigma^2}{2} \right)^j \right] K^n. \end{aligned} \quad (3.12)$$

Note that, owing to the positivity of all the coefficients $b_{i,j}$, this operator is positive definite and therefore possesses a well-defined inverse. Note also that, for $\sigma \gg 1$, the only coefficients in (3.12) that remain significant are the ‘‘diagonal’’ ones, $b_{i,i} = 1$, yielding simply the truncated Taylor series for the exponential function of $\sigma^2 K/2$. Shortly, we shall examine the practical impact of the off-diagonal components, $b_{i,j}, j > i$, but first we describe the process of extracting from the above algebraic developments a practical class of smoothing filters.

The reciprocal of the function $\exp(a^2 k^2/2)$ in (3.10) is a Gaussian function in k and is the Fourier transform of a convolution operator (on the line x) whose kernel is also of Gaussian form. Provided we can find a practical way to invert the operator equation,

$$\mathcal{D}_{(n)}^* \mathbf{s} = \mathbf{p}, \quad (3.13)$$

for a given input distribution, \mathbf{p} , the resulting output, \mathbf{s} , will be an approximation to the convolution of \mathbf{p} by the Gaussian function whose spectral transform is the reciprocal of the right-hand side of (3.10). The approximation, $(\mathcal{D}_{(n)}^*)^{-1}$, to this convolution is what we refer to as a quasi-Gaussian filter. The common centered second moment of operator $\mathcal{D}_{(n)}$ and its approximation, $\mathcal{D}_{(n)}^*$, is




FIG. 1. Comparison of one-dimensional applications of recursive filters approximating a Gaussian (shown solid). Dashed curves show filter approximations: (a) order $n = 1$, (b) $n = 2$, and (c) $n = 4$, with (long dashes) and without (short dashes) the off-diagonal b coefficient refinements.

exactly $-a^2$, so a is a convenient measure of the intrinsic distance scale of the smoothing filter implied by the inversion of (3.13). A useful fact is that the square of the intrinsic scale of the composition of sequential smoothing filters is the sum of squares of the scales of the individual components. Also, as a consequence of the statisticians' "central limit theorem" (e.g., Wilks 1995) applied to convolutions in general, the effective convolution kernel of such a composition of several identical filter factors resembles a Gaussian more closely than does the representative factor. Thus, provided it becomes feasible to invert (3.13), we possess the means to convolve a gridded input distribution with a smooth quasi-Gaussian kernel, at least in one dimension.

As a matrix, $\mathcal{D}_{(n)}^*$ is banded and, for an infinite domain, symmetric. Conventionally, the linear inversion of a system such as (3.13) might be effected by employing an LU factorization (Dahlquist and Björck 1974) of $\mathcal{D}_{(n)}^*$,

$$\mathcal{D}_{(n)}^* \equiv \mathbf{A}\mathbf{B}, \tag{3.14}$$

with lower-triangular band matrix, \mathbf{A} and upper-triangular band matrix, \mathbf{B} , allowing the solution to proceed as two steps of recursive substitution. On an infinite grid, the same principle pertains, but with the guaranteed simplification of (i) a translational symmetry ensuring that every row of \mathbf{A} is identical (allowing for the trivial translation) and every row of \mathbf{B} is identical and (ii) ordinary matrix symmetry by which we can ensure that \mathbf{B} is simply the transpose of \mathbf{A} . In this case, the LU decomposition of $\mathcal{D}_{(n)}^*$ is also of the symmetric, or Cholesky, type (Dahlquist and Björck 1974).

In the two stages of solution,

$$\mathbf{A}\mathbf{q} = \mathbf{p}, \tag{3.15}$$

$$\mathbf{B}\mathbf{s} = \mathbf{q}, \tag{3.16}$$

the explicit recursions of the back substitutions are the following basic recursive filters:

$$q_i = \beta p_i + \sum_{j=1}^n \alpha_j q_{i-j} \quad \text{and} \tag{3.17}$$

$$s_i = \beta q_i + \sum_{j=1}^n \alpha_j s_{i+j}, \tag{3.18}$$

which are conveniently referred to as the advancing and backing steps, respectively, since, in the first, index i must be treated in increasing order while, in the second, it must be treated in decreasing order, in order that the terms on the right are already available at each step. Note that the correspondences between notations of (3.15), (3.16) and of (3.17), (3.18) are

$$A_{i,i} = B_{i,i} = 1/\beta, \tag{3.19}$$

$$A_{i,i-j} = B_{i,i+j} = -\alpha_j/\beta, \quad j \in [1, n]. \tag{3.20}$$

Defining the total "substance" implied by the distribution \mathbf{p} to be $\sum_i \delta x p_i$, the operator $\mathcal{D}_{(n)}^*$, and hence its inverse, preserve this quantity. By symmetry, the factor operators, \mathbf{A} and \mathbf{B} must therefore also preserve substance, implying that

$$\beta = 1 - \sum_{j=1}^n \alpha_j. \tag{3.21}$$

The task of distilling the coefficients α_j from the parameters defining $\mathcal{D}_{(n)}^*$ is somewhat technical and is relegated to appendix A. Filters may be constructed at different orders, n , and with or without the refinements implied by the off-diagonal coefficients, $b_{i,j}$, for $j > i$. For $n = 1$ the filter response comprises back-to-back decreasing exponential functions, which Fig. 1a shows (dashed curve) in comparison to the Gaussian function (solid curve) of the same width of one grid unit. Better approximations to the Gaussian are obtained after application of the second-order filters, as shown in Fig. 1b, and fourth-order filters, shown in Fig. 1c, for the

case of the filter with only the diagonal coefficients b (short-dashed curves) and with all of the b coefficients (long-dashed curves). We see that the advantage of keeping all the coefficients is greater at higher order, where they make the resulting filter response a significantly better approximation to the intended Gaussian function. However, the alternative treatments of the b coefficients are virtually indistinguishable at smoothing scales of a few grid units, as the truncation errors of the component numerical derivative operators become insignificant in comparison to the error resulting from the finite truncation of the series for the Gaussian employed in the construction of the filter operator. The cost of applying the filters with or without the off-diagonal b coefficients is the same; therefore, we always adopt the more accurate formulation that includes the off-diagonal coefficients.

We have described the idealized case of operators acting on data extending indefinitely in both directions. In practice, we are confronted with geometrical constraints, either in the form of definite lateral boundaries to the domain, or as periodic conditions appropriate to a cyclic domain. Fortunately, it is possible to generalize the application of the advancing and backing recursive filters to both of these situations. Appendix B treats the case of lateral boundaries and shows how the effect of a continuation of the domain to infinity can be simulated by the imposition of appropriate “turning” conditions at the transition between the advancing and backing stages. Appendix C treats the case of periodic boundary conditions. In both of these special cases the main part of the filtering algorithm and the basic filter coefficients employed are the same as in the case of the infinite domain. By a generalization of the treatment used in the cyclic case, one may efficiently distribute the recursions across multiple processors of a massively parallel computer. This is discussed in detail in a related context by Fujita and Purser (2001).

b. Quasi-Gaussian filters in two dimensions

Let x and y be horizontal Cartesian coordinates, and k and l the associated wavenumber components. Then in two dimensions, we can exploit the factoring property of isotropic Gaussians:

$$\exp\left(-\frac{a^2\rho^2}{2}\right) = \exp\left(-\frac{a^2k^2}{2}\right) \exp\left(-\frac{a^2l^2}{2}\right), \quad (3.22)$$

where $\rho = (k^2 + l^2)^{1/2}$ is the total wavenumber. In terms of basic one-dimensional Gaussian smoothing filters, $\mathcal{D}_{(\infty)}^{(x)}$ and $\mathcal{D}_{(\infty)}^{(y)}$, operating in the x and y directions, a two-dimensional *isotropic* filter, G_a , also of Gaussian form, results from the successive application of the one-dimensional factors, $\mathcal{D}_{(\infty)}^{(x)}$ and $\mathcal{D}_{(\infty)}^{(y)}$. For example, an input field, χ , is smoothed to produce the output field, ψ , by the convolution:

$$\psi(\mathbf{x}_1) = \iint G_a(\mathbf{x}_1, \mathbf{x}_2)\chi(\mathbf{x}_2) dx_2 dy_2 \equiv G_a \star \chi, \quad (3.23)$$

where

$$G_a = \mathcal{D}_{(\infty)}^{(y)} \star \mathcal{D}_{(\infty)}^{(x)}.$$

The crucial significance of the Gaussian form for the one-dimensional filters is that this form is the *only* shape that, upon combination by convolution in the x and y directions, produces an *isotropic* product filter. In order to generalize our filters to alternative shapes, while preserving two-dimensional isotropy, we shall always attempt to base the construction of the more general filters on the building blocks supplied by the quasi-Gaussian products of the approximations, $\mathcal{D}_{(\infty)}^{(x)}$ and $\mathcal{D}_{(\infty)}^{(y)}$, to the true Gaussian smoothers. But we must first establish what is the minimum order of such a filter that will preserve the isotropy of the product combination, at least to a degree that ensures that any residual anisotropies are not obtrusively obvious.

Figure 2 depicts the results obtained by smoothing a delta function placed at the center of a square grid. Figure 2a shows the result of a single application of the first-order filter, $\mathcal{D}_{(1)}$, in the x and y directions. This result is clearly neither smooth nor even approximately isotropic. Figures 2b and 2c show the results obtained by using the filters of orders two and four. We see that the appearance of isotropy is not adequately attained until the order exceeds two, but the fourth-order filter shown in Fig. 2c seems to provide an excellent approximation to the isotropic Gaussian. For applications in data assimilation, it is usually worth the cost of applying a filter of at least fourth order if the filter is to be applied only once in each of the orthogonal grid directions. For a roughly equivalent cost, one may also apply the simple first-order filter four times in succession (but with a scale only a half as large in each instance). This has been the approach employed in earlier applications of recursive filtering to data analysis, such as Purser and McQuigg (1982), Lorenc (1992), and Hayden and Purser (1995). The result is shown in Fig. 2d, but is clearly inferior to the use of the single fourth-order filter.

Very often, the physical variables of interest in an analysis are derivatives of the variables it is convenient to base the covariance model on. For example, covariances of the streamfunction or velocity potential (scalars) are often more convenient to handle than the derived covariances among velocity components at two locations. Since we may wish to employ the results of our filters as building blocks of such differentiated covariances, it is as well to examine the derivatives of fields analogous to those of Fig. 2. In order to permit any departures from isotropy to stand out more clearly, we take the Laplacian of the result of smoothing the delta function. Figure 3 shows three such results (with a slightly smaller scale than was used in Fig. 2), involving single applications (in x and y) of $\mathcal{D}_{(n)}$ with n being 2,




FIG. 2. Sequential application of quasi-Gaussian recursive filters of order n in two dimensions: (a) $n = 1$, (b) $n = 2$, (c) $n = 4$, and (d) four applications of filters with $n = 1$ with scale parameter adjusted to make the result comparable with the other single-pass filters.

4, and 6 in panels a, b, and c, respectively. Even more so than in Fig. 2, we see that it is not until we adopt at least fourth-order filtering that we obtain an acceptable degree of isotropy. For reference, the “right answer” obtained using the Laplacian of the true Gaussian, G_a , is shown in Fig. 3d.

c. Numerical robustness and multigrid refinements

A recognized problem with high-order recursive filters (e.g., Otnes and Enochson 1972) is their susceptibility to numerical noise derived from the amplification of roundoff effects, especially as the filtering scale becomes significantly larger than the grid scale. A suggestion of this problem is barely visible in the raggedness of the outer contours of Fig. 3c, where 32-bit precision was used for both the calculation of the filter coefficients and the application of the filters themselves. The problem is deferred to larger filtering scales when all calculations are performed at 64-bit precision. Also, even when the filtering itself is performed at 32-bit precision, it is still beneficial to employ 64-bit precision in

calculating the coefficients used. However, a more satisfactory and permanent solution is to avoid the problem altogether, which can only be done by keeping the scale of the computational grid for the recursive filter operations commensurate with the desired filtering scale in each direction. A natural remedy, in cases where the grid dimensions permit it, is to employ a “multigrid” strategy. General discussions of such methods can be found in Brandt (1977) and an excellent introductory exposition of the method is Briggs (1987). Essentially, the field to be smoothed at a certain filtering scale is first transferred (by adjoint interpolation from the initial fine grid) to a grid whose coarseness is comparable with, but still sufficiently resolves, this smoothness scale. The smoothing is performed by the high-order recursive filter on the generally somewhat coarser grid, now without risk of numerical noise, and at a numerical cost that is usually significantly less than the cost of the equivalent operation applied to the original fine grid. The resulting smooth field is finally interpolated back to the fine grid. The implied operator representing this combination of steps remains self-adjoint and, provided the order of

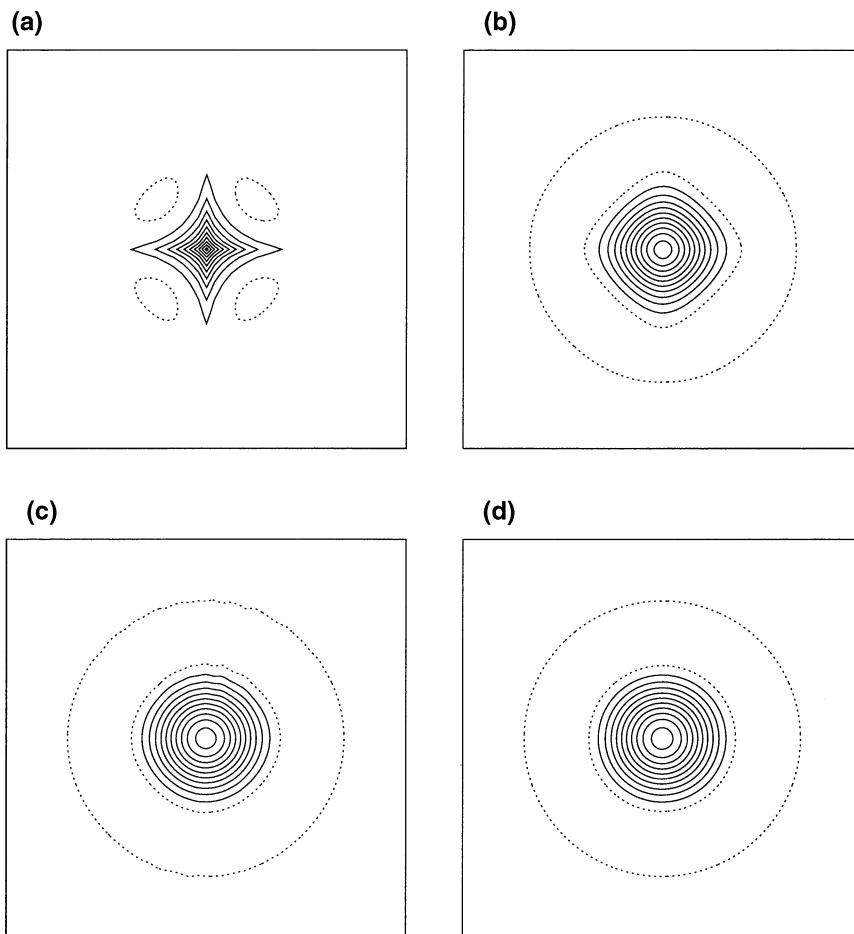


FIG. 3. Negative Laplacian applied to quasi-Gaussian recursive filters with (a) $n = 2$, (b) $n = 4$, (c) $n = 6$, and (d) corresponding contours for the exact Gaussian. Contours are at multiples of odd numbers with negative contours shown as dotted curves.

accuracy of the interpolations is large enough, no discernible hint of roughness appears in the resulting smooth output. For the higher-order filters, which are susceptible to the amplification of roundoff noise, the multigrid remedy should be applied whenever the smoothing scale exceeds about six of the basic grid spaces.

The simplest multigrid structure is one in which the spacing in the successive grids doubles. Then, except for the possible overlaps (which are desirable in the case of bounded domains in order to preserve the same centered interpolation operators everywhere), each coarse grid is a subset of its finer predecessor. For cyclic domains, this simplification obviously works only when the periodic grid dimensions are divisible by powers of two. For bounded domains, the judicious use of overlaps of the coarser grids far enough beyond the finer grid boundaries to accommodate the interpolation stencils enables one to adopt the grid-scale-doubling arrangement without numerical restrictions on the grid dimensions. Interpolation is assumed to occur only between adjacent grids of the scale hierarchy. Interpolation from

a coarse grid to the next finer grid in two dimensions is accomplished by passing through an intermediate stage in which one dimension is “fine” and the other is “coarse.” In this way, only one-dimensional interpolation operators need be considered. Assuming each coarse grid overlaps the next finer grid by a sufficient margin, all interpolations can be performed using the same standard centered formula. Table 2 lists the weights w_p associated with source points distant p grid units from the target point for midpoint interpolation from a uniform grid at (even) orders of accuracy between 2 and 12. [Analytic expressions for these interpolation coefficients may be found in the appendix A of Purser and Leslie (1988).] Experience suggests that sixth-order interpolations are adequate for most purposes.

4. Discussion

The problem of efficiently accommodating accurate approximations to isotropic covariance functions in a variational analysis has been solved using recursive nu-

TABLE 2. Coefficients w_j for uniform grid midpoint interpolation at orders of accuracy, n , up to 12, where index j is the half-integer distance, in grid units, between the target point and the source point to which weight w_j applies.

n	γ	$\gamma^{w_{1/2}}$	$\gamma^{w_{3/2}}$	$\gamma^{w_{5/2}}$	$\gamma^{w_{7/2}}$	$\gamma^{w_{9/2}}$	$\gamma^{w_{11/2}}$
2	2	1					
4	16	9	-1				
6	256	150	-25	3			
8	2048	1225	-245	49	-5		
10	65 536	39 690	-8820	2268	-405	35	
12	524 228	320 166	-76 230	22 869	-5445	847	-63

merical filters. The covariances are never explicitly computed; instead, their effects as convolution operators are represented. This is achieved through a sequence of applications of carefully designed recursive filters operating along the various lines of the appropriately chosen computational grids. In a regional analysis, there is no reason not to use the grid of the intended numerical prediction model. While the building blocks of a covariance operator generated by recursive filtering are always of quasi-Gaussian form as required by the factorization property (3.22), the final covariances themselves may be given more general radial profiles through the superposition of varyingly scaled component Gaussians, or by combining these components with the application of Laplacian operators, as we discuss in Part II.

A global model requires additional refinements to the technique we have described here for purely geometrical reasons. For one thing, the presence of polar coordinate singularities makes the generation of even approximately isotropic covariances impossible by the method of sequentially applying the quasi-Gaussian recursive filters sequentially along the two principal grid directions. But even if we adopt a patchwork of overlapping rectangular grids [which is essentially the approach we have adopted for the global analysis described in Wu et al. (2002)], the use of the present filters, which are homogeneous in scale only relative to the grid to which they are applied, translates to the application of smoothing operations, which, on grids of global extent, cannot be everywhere even approximately of uniform scale in the degree to which they smooth the given data. In order to address this need to generalize the filters to obtain control over their spatial smoothing scale from one geographical location to another, Part II will also describe technical extensions that not only allow geographical control of the characteristic scale and amplitude of the covariance, but that also allow geographically adaptive three-dimensionally anisotropic stretching of the quasi-Gaussian building blocks from which the covariances are constructed.

Acknowledgments. The authors would like to thank Drs. John Derber, Dezso Devenyi, and Andrew Lorenc for many helpful discussions; Dr. Wanqiu Wang for valuable comments made during internal review; and three anonymous reviewers for their valuable comments

and suggestions. We also thank Prof. Eugenia Kalnay and Drs. Stephen Lord and Roger Daley for their encouragement and support. This work was partially supported by the NSF/NOAA Joint Grants Program of the U.S. Weather Research Program. This research is also in response to requirements and funding by the Federal Aviation Administration (FAA). The views expressed are those of the authors and do not necessarily represent the official policy or position of the FAA.

APPENDIX A

Obtaining Filter Coefficients for a Given Scale

In section 3, $\mathcal{D}_{(n)}^*$ was defined as a real-coefficient monic polynomial of K , so we may use this polynomial's roots κ_p (whose complex members come in conjugate pairs) to perform the operator factorization:

$$\mathcal{D}_{(n)}^* = \prod_{p=1}^n \left(1 - \frac{K}{\kappa_p} \right). \quad (\text{A.1})$$

In terms of the shift operator, Z defined

$$(Z\psi)_i = \psi_{i+1}, \quad (\text{A.2})$$

we have

$$K = -Z + 2 - Z^{-1}, \quad (\text{A.3})$$

and each factor in (A.1) is expressible:

$$1 - \frac{K}{\kappa_p} = \frac{(Z - 2\omega_p + Z^{-1})}{2 - 2\omega_p},$$

where

$$\omega_p = 1 - \kappa_p/2.$$

Therefore, the smaller of the two possible roots,

$$\zeta_p = [\omega_p + (\omega_p^2 - 1)^{1/2}]^{\pm 1} \quad (\text{A.4})$$

of the quadratic,

$$z^2 - 2\omega_p z + 1 = 0,$$

allows the factorization of the term

$$1 - \frac{K}{\kappa_p} = \left(\frac{1 - \zeta_p Z^{-1}}{1 - \zeta_p} \right) \left(\frac{1 - \zeta_p Z}{1 - \zeta_p} \right). \quad (\text{A.5})$$

The operation,

$$\psi = \left(\frac{1 - \zeta_p}{1 - \zeta_p Z^{-1}} \right) \chi,$$

describes the application of the right-moving, complex coefficient, first-order recursive filter,

$$\psi_i = (1 - \zeta_p)\chi_i + \zeta_p\psi_{i-1},$$

whose stability is guaranteed by the property $|\zeta_p| < 1$ [the other root of (A.4) is its reciprocal]. Likewise, the inverse of the other operator factor on the right of (A.5) describes the operation of a left-moving stable recursive filter. By a similar decomposition for all the $p \in \{1, \dots, n\}$, we deduce that the operators of (3.15) and (3.16), comprising the inverses of the right-moving and left-moving high-order filters, are formally

$$\mathbf{A} = \prod_{p=1}^n \left(\frac{1 - \zeta_p Z^{-1}}{1 - \zeta_p} \right) \quad \text{and} \quad (\text{A.6})$$

$$\mathbf{B} = \prod_{p=1}^n \left(\frac{1 - \zeta_p Z}{1 - \zeta_p} \right), \quad (\text{A.7})$$

whose operator kernels are real by virtue of the occurrences of the complex κ_p , and hence of the ζ_p , in conjugate pairs, and whose coefficients are constructible by the explicit convolution products of Z^{-1} and Z prescribed by (A.6) and (A.7).

We summarize the practical steps required in order to obtain these coefficients. First, one locates the complex roots, κ_p , of the real-coefficient polynomial in K that defines $\mathcal{D}_{(n)}^*$. Then, the corresponding quantities ω and, in each case, the smaller of the two roots ζ_p are obtained. Then the convolution polynomial (A.6) is constructed using these complex (in general) ζ_p . Finally, we invoke (3.17) and (3.18) to get the algorithmically convenient coefficients α_i and γ of this filtering scheme.

The same kind of root-finding method is also used in the construction of other homogeneous recursive filters, as described in Hamming (1977) or Otnes and Enochson (1972). However, the complex arithmetic can be avoided by an alternative iterative method of obtaining the convolution kernels of \mathbf{A} and its adjoint \mathbf{B} . According to (3.14), these are the rows of the Cholesky factors of an infinite symmetric band matrix representation of $\mathcal{D}_{(n)}^*$, but in practice, given a sufficiently large *finite* symmetric matrix having the same generic rows as $\mathcal{D}_{(n)}^*$, one finds that the last rows of its Cholesky lower-triangular factor will be numerically indistinguishable, both from each other, and from the generic row of the Cholesky factor of the infinite matrix representation of $\mathcal{D}_{(n)}^*$. By an adaptation of the regular Cholesky algorithm,

$$\mathbf{A}_{i,j} = \left(\mathcal{D}_{i,j}^* - \sum_{k<j} \mathbf{A}_{i,k} \mathbf{A}_{j,k} \right) / A_{j,j}, \quad j < i, \quad (\text{A.8})$$

$$\mathbf{A}_{i,i} = \left(\mathcal{D}_{i,i}^* - \sum_{k<i} \mathbf{A}_{i,k} \mathbf{A}_{i,k} \right)^{1/2}, \quad (\text{A.9})$$

one can compute successive rows until, within roundoff, the subsequent differences from row to row are negligible. This iterative approach is equally valid and never fails in practice to reproduce the same numerical coefficients as is given by the formal root-finding procedure.

APPENDIX B

Nonperiodic End Conditions

Consider a mutually adjoint pair of n th-order basic recursive filters \mathbf{A}^{-1} (“advancing”) and \mathbf{B}^{-1} (“backing”) on a uniform grid and with constant coefficients such that, if $\mathbf{q} = \mathbf{A}^{-1}\mathbf{p}$, and $\mathbf{r} = \mathbf{B}^{-1}\mathbf{p}$, then, for a generic grid point i ,

$$q_i = \beta p_i + \sum_{j=1}^n \alpha_j q_{i-j} \quad \text{and} \quad (\text{B.1})$$

$$r_i = \beta p_i + \sum_{j=1}^n \alpha_j q_{i+j}. \quad (\text{B.2})$$

On an infinite grid, assuming all data remain bounded at infinity, one can easily verify for the symmetrized output, \mathbf{s} :

$$\mathbf{s} = \mathbf{B}^{-1}\mathbf{A}^{-1}\mathbf{p} \equiv \mathbf{A}^{-1}\mathbf{B}^{-1}\mathbf{p}. \quad (\text{B.3})$$

We ask, given a *finite* domain, $i \in [1, N]$, and assuming this interval contains the support of input \mathbf{p} , how does one “prime” those values of \mathbf{s} just inside the boundary at $i = N$ in order to enable the backing filter, \mathbf{B}^{-1} when applied to $\mathbf{q} = \mathbf{A}^{-1}\mathbf{p}$, to simulate implicitly the effect of a continuation of the gridded values of \mathbf{q} beyond this boundary? The solution of this puzzle is found by exploiting the commutativity (B.3). Let $\hat{\mathbf{s}}_j$ denote the n vector of $(s_{j+1-n}, \dots, s_j)^T$. Then $\hat{\mathbf{s}}_N$ are the last n values of s belonging to the actual domain while $\hat{\mathbf{s}}_{N+n}$ is the vector of n values one would have obtained just beyond, if the grid were continued. Define a lower-triangular $n \times n$ matrix, \mathbf{L} with elements

$$L_{i,i} = 1 \quad \text{and} \quad (\text{B.4})$$

$$L_{i+i,i} = -\alpha_j, \quad (\text{B.5})$$

and an upper-triangular $n \times n$ matrix, \mathbf{U} , with elements

$$U_{i,i+j} = \alpha_{N-j}. \quad (\text{B.6})$$

Then, since $\hat{\mathbf{r}}_{N+n} = \mathbf{0}$, it must follow from $\mathbf{s} = \mathbf{A}^{-1}\mathbf{r}$ that

$$\mathbf{L}\hat{\mathbf{s}}_{N+n} = \mathbf{U}\hat{\mathbf{s}}_N \quad (\text{B.7})$$

and, from $\mathbf{s} = \mathbf{B}^{-1}\mathbf{q}$, that

$$\mathbf{L}^T\hat{\mathbf{s}}_N = \mathbf{U}^T\hat{\mathbf{s}}_{N+n} + \hat{\mathbf{q}}_N/\beta. \quad (\text{B.8})$$

Thus, eliminating $\hat{\mathbf{s}}_{N+n}$, we obtain the turning conditions that prime the backing filter:

$$(\mathbf{L}^T - \mathbf{U}^T\mathbf{L}^{-1}\mathbf{U})\hat{\mathbf{s}}_N = \hat{\mathbf{q}}_N/\beta. \quad (\text{B.9})$$




FIG. B1. Comparison of filter responses to initial data composed of impulses placed two units and eight units away from the right-hand boundary with (solid) and without (short dashed) the special end conditions described in appendix B. (a) Results for the first-order filter with a characteristic scale of five units, and (b) for the fourth-order filter with the same scale.

For the simplest filter pair, $n = 1$, $\alpha_1 = \alpha$, this formula reduces to

$$(1 - \alpha^2)s_N = \beta q_N \tag{B.10}$$

or, since for the conservative filter, $\beta = 1 - \alpha$,

$$s_N = \frac{1 - \alpha}{1 + \alpha} q_N. \tag{B.11}$$

To summarize, the bounded-domain algorithm will consist of the following three steps: (i) apply advancing filter \mathbf{A}^{-1} , (ii) prime the end conditions according to the procedure we have defined in this appendix, and (iii) apply backing filter \mathbf{B}^{-1} .

By way of illustrating the effects of these end conditions, Fig. B1 compares first-order (panel a) and fourth-order (panel b) filter responses with and without these end conditions for applications to initial data consisting of impulses two and eight grid units from the right-hand end. In each case, the smoothing scale is 4.0 grid units. The short-dashed curves show the output without applying the special end conditions (i.e., applying the backing filter as if all data beyond the boundary are reset to zero) while the solid curves show the “correct” responses using the end conditions we have described. In both cases the neglect of the proper end conditions has a serious effect on the response associated with the impulse placed two units from the end, being much larger in the case of the higher-order filter shown in the second panel, while, for the first-order filter, the differences for the case of the impulse placed eight units in become negligible. However, for the higher-order filter of Fig. B1b, we see that significant differences persist even when the initial impulse is as much as eight units away from the boundary and that, there-

fore, it is not wise to neglect the proper end conditions for the higher-order filters.

APPENDIX C

Periodic End Conditions

Preserving the notation of appendix B, we consider the problem of defining the correct priming vector $\hat{\mathbf{q}}_0$ for the advancing filter on a cyclic domain with period N , such that the values obtained are consistent with the wraparound condition, $\hat{\mathbf{q}}_N = \hat{\mathbf{q}}_0$. In the recursion (3.17), the sensitivity of $\hat{\mathbf{q}}_N$ to $\hat{\mathbf{q}}_{N-1}$, given that input element p_N is unchanged, is expressed as

$$d\hat{\mathbf{q}}_N = \mathbf{T}d\hat{\mathbf{q}}_{N-1},$$

where

$$\mathbf{T} = \begin{pmatrix} 0 & 1 & 0 & \cdots & 0 \\ 0 & 0 & 1 & \cdots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ \alpha_n & \alpha_{n-1} & \alpha_{n-2} & \cdots & \alpha_1 \end{pmatrix}.$$

Applying the chain rule, we deduce that the sensitivity of $\hat{\mathbf{q}}_N$ to $\hat{\mathbf{q}}_0$ when the intervening input elements, p_1, \dots, p_N , remain unchanged, is

$$d\hat{\mathbf{q}}_N = \mathbf{T}^N d\hat{\mathbf{q}}_0 \tag{C.1}$$

or, for some linear n -vector function $\hat{\mathbf{h}}$ of all the p_1, \dots, p_N ,

$$\hat{\mathbf{q}}_N = \mathbf{T}^N \hat{\mathbf{q}}_0 + \hat{\mathbf{h}}(\mathbf{p}). \tag{C.2}$$

The identification of $\hat{\mathbf{q}}_0$ with $\hat{\mathbf{q}}_N$ therefore requires that

$$\hat{\mathbf{q}}_0 = (\mathbf{I} - \mathbf{T}^N)^{-1} \hat{\mathbf{h}} \tag{C.3}$$

We do not need an explicit formula for the vector function $\hat{\mathbf{h}}(\mathbf{p})$ because we can construct it simply by running a preliminary trial of the recursive filter with the starting values, $\hat{\mathbf{q}}_0^*$, of this trial solution \mathbf{q}^* set to zero. Then, from (C.2), the final n values $\hat{\mathbf{q}}_N^*$ of the trial solution are simply the vector $\hat{\mathbf{h}}$ itself. A second sweep of the recursive filter with the proper initial setting computed from (C.3) completes the advancing filter in a self-consistent way.

A similar procedure is used for the backing filter on the cyclic domain, so the overall cost is double that of the recursive filter on a nonperiodic domain of N points. For this reason, when performance is critical, it may be preferable in practice to employ a generous overlap and the nonperiodic version of the filter instead. However, the extra overhead of computing the proper cyclic condition is a factor for consideration only for serial processing; when the domain is divided into segments, whether it is periodic or not, the recursions *always* need to be repeated, at least for a significant fraction of the domain, in order to achieve intersegment consistency. Fujita and Purser (2001) show how a generalization of this idea can form the basis for one of the ways in which recursive processes may be implemented efficiently in a distributed computing environment.

REFERENCES

- Barnes, S. L., 1964: A technique for maximizing details in numerical weather map analysis. *J. Appl. Meteor.*, **3**, 396–409.
- Bergthorssen, P., and B. Döös, 1955: Numerical weather map analysis. *Tellus*, **7**, 329–340.
- Brandt, A., 1977: Multilevel adaptive solutions of boundary value problems. *Math. Comput.*, **31**, 333–390.
- Briggs, W. L., 1987: *A Multigrid Tutorial*. SIAM Publications, 90 pp.
- Cohn, S. E., A. da Silva, J. Guo, M. Sienkiewicz, and D. Lamich, 1998: Assessing the effects of data selection with the DAO physical-space statistical analysis system. *Mon. Wea. Rev.*, **126**, 2913–2926.
- Courtier, P., 1997: Dual formulations of four-dimensional variational assimilation. *Quart. J. Roy. Meteor. Soc.*, **123**, 2449–2461.
- , and Coauthors, 1998: The ECMWF implementation of three-dimensional variational assimilation (3D-Var). I: Formulation. *Quart. J. Roy. Meteor. Soc.*, **124**, 1783–1807.
- Cressman, G. P., 1959: An operational analysis scheme. *Mon. Wea. Rev.*, **87**, 367–374.
- Dahlquist, G., and A. Björck, 1974: *Numerical Methods*. Prentice Hall, 573 pp.
- Daley, R. A., 1991: *Atmospheric Data Assimilation*. Cambridge University Press, 457 pp.
- , and E. Barker, 2000: NAVDAS 2000 source book. NRL Publ. NRL/PU/7530-00-418, Naval Research Laboratory, Monterey, CA, 153 pp.
- da Silva, A., J. Pfaendtner, J. Guo, M. Sienkiewicz, and S. Cohn, 1995: Assessing the effects of data selection with DAO's physical-space statistical analysis system. *Second International Symposium on Assimilation of Observations in Meteorology and Oceanography*, WMO/TD 651, 273–278.
- Derber, J., and A. Rosati, 1989: A global oceanic data assimilation system. *J. Phys. Oceanogr.*, **19**, 1333–1347.
- Fujita, T., and R. J. Purser, 2001: Parallel implementation of compact numerical schemes. NOAA/NCEP Office Note 434, 34 pp. [Available from NCEP, 5200 Auth Rd., Camp Springs, MD 20746.]
- Gandin, L. S., 1963: *Objective Analysis of Meteorological Fields* (in Russian). Gidrometeoizdat, 212 pp.
- Gaspari, G., and S. E. Cohn, 1998: Construction of correlation functions in two and three dimensions. Office Note Series on Global Modeling and Data Assimilation, DAO Office Note 96-03R1, DAO, GSFC, 53 pp.
- , and —, 1999: Construction of correlation functions in two and three dimensions. *Quart. J. Roy. Meteor. Soc.*, **125**, 723–757.
- Hamming, R. W., 1977: *Digital Filters*. Prentice Hall, 226 pp.
- Hayden, C. M., and R. J. Purser, 1995: Recursive filter objective analysis of meteorological fields: Applications to NESDIS operational processing. *J. Appl. Meteor.*, **34**, 3–15.
- Hollingsworth, A., and P. Lönnberg, 1986: The statistical structure of short-range forecast errors as determined from radiosonde data. Part I: The wind field. *Tellus*, **38A**, 111–136.
- Ide, K., P. Courtier, M. Ghil, and A. C. Lorenc, 1997: Unified notation for data assimilation: Operational, sequential and variational. *J. Meteor. Soc. Japan*, **75**, 181–189.
- Lorenc, A. C., 1986: Analysis methods for numerical weather prediction. *Quart. J. Roy. Meteor. Soc.*, **112**, 1177–1194.
- , 1992: Iterative analysis using covariance functions and filters. *Quart. J. Roy. Meteor. Soc.*, **118**, 569–591.
- , 1997: Development of an operational variational assimilation scheme. *J. Meteor. Soc. Japan*, **75**, 339–346.
- Oliver, D., 1995: Moving averages for Gaussian simulation in two and three dimensions. *Math. Geology*, **27**, 939–960.
- Otnes, R. K., and L. Enochson, 1972: *Digital Time Series Analysis*. Wiley, 467 pp.
- Papoulis, A., 1984: *Probability, Random Variables, and Stochastic Processes*. McGraw-Hill, 576 pp.
- Parrish, D. F., and J. C. Derber, 1992: The National Meteorological Center's Spectral Statistical-Interpolation Analysis System. *Mon. Wea. Rev.*, **120**, 1747–1763.
- Purser, R. J., and R. McQuigg, 1982: A successive correction analysis scheme using recursive numerical filters. Met Office Tech. Note 154, British Meteorological Office, 17 pp.
- , and L. M. Leslie, 1988: A semi-implicit semi-Lagrangian finite-difference scheme using high-order spatial differencing on a non-staggered grid. *Mon. Wea. Rev.*, **116**, 2067–2080.
- , W.-S. Wu, D. F. Parrish, and N. M. Roberts, 2003: Numerical aspects of the application of recursive filters to variational statistical analysis. Part II: Spatially inhomogeneous and anisotropic general covariances. *Mon. Wea. Rev.*, **131**, 1536–1548.
- Thiébaux, H. J., 1976: Anisotropic correlation functions for objective analysis. *Mon. Wea. Rev.*, **104**, 994–1002.
- , H. L. Mitchell, and D. W. Shantz, 1986: Horizontal structure of hemispheric forecast error correlations for geopotential and temperature. *Mon. Wea. Rev.*, **114**, 1048–1066.
- Wilks, D. S., 1995: *Statistical Methods in the Atmospheric Sciences: An Introduction*. Academic Press, 467 pp.
- Wu, W.-S., R. J. Purser, and D. F. Parrish, 2002: Three-dimensional variational analysis with spatially inhomogeneous covariances. *Mon. Wea. Rev.*, **130**, 2905–2916.