
IFS DOCUMENTATION – Cy41r1
Operational implementation 12 May 2015

**PART V: ENSEMBLE PREDICTION
SYSTEM**

Table of contents

Chapter 1 **Methodology**
Chapter 2 **Computational details: initial perturbations**
Chapter 3 **Computational details: non-linear integrations**
References

© Copyright 2015

European Centre for Medium-Range Weather Forecasts
Shinfield Park, Reading, RG2 9AX, England

Literary and scientific copyrights belong to ECMWF and are reserved in all countries. This publication is not to be reprinted or translated in whole or in part without the written permission of the Director. Appropriate non-commercial use will normally be granted under the condition that reference is made to ECMWF. The information within this publication is given in good faith and considered to be true, but ECMWF accepts no liability for error, omission and for loss or damage arising from its use.

Chapter 1

Methodology

Table of contents

- 1.1 Introduction**
- 1.2 Initial condition perturbations**
- 1.3 Singular vectors**
 - 1.3.1 Formulation of the singular vector computation
 - 1.3.2 Use of the total energy norm
 - 1.3.3 Use of the Hessian of the 3D-Var objective function
- Appendix A. Eigenvalue algorithms**
 - A.1 The Lanczos algorithm
 - A.2 The Jacobi–Davidson algorithm
 - A.3 Solution method

1.1 INTRODUCTION

Ensemble prediction is a method to predict the probability distribution of forecast states, given a probability distribution of random analysis error and model uncertainties. In ensemble systems, the impact of analysis and model uncertainties on forecasts states are simulated by starting the ensemble of forecasts from different initial conditions, and computing each forecast evolution using schemes designed to simulate model uncertainties. Since 1992, ensemble prediction systems have been part of the operational forecasting suite at ECMWF ([Palmer *et al.*, 1993](#)) and NCEP ([Toth and Kalnay, 1993, 1997](#)). A general description of the first version of the ECMWF EPS can be found in [Molteni *et al.* \(1996\)](#), and a description of a more recent version of the operational system is given by [Leutbecher and Palmer \(2008\)](#).

The implementations by ECMWF and NCEP in the early 90-s were followed in 1995 by MSC Canada ([Houtekamer *et al.*, 1996b,a](#)) and in 1998 by BoM Australia. The reader is referred to [Buizza *et al.* \(2005\)](#) for a comparison of the performance of the ECMWF, MSC and NCEP ensemble systems over the northern hemisphere, and to [Bourke *et al.* \(2005\)](#) for a comparison of the performance of the ECMWF and the BoM systems over the southern hemisphere. In the past decade, other meteorological centres have started producing routinely ensemble-based probabilistic forecasts, see for example [Park *et al.* \(2008\)](#) for comparison of the performance of eight global ensemble prediction systems archived within the TIGGE project. The reader is referred to [Hagedorn *et al.* \(2012\)](#) for a more recent comparison of the performance of the ECMWF EPS and multi-model ensemble systems defined by the TIGGE ensembles.

Since 12 September 2006, the ECMWF operational ensemble system has been running with a variable resolution (VAREPS, see [Buizza *et al.* \(2007\)](#) for a description of the system), with 51 members, one starting from unperturbed initial conditions (the control forecast) and 50 from perturbed initial conditions defined by adding small amplitude perturbations to the operational analysis.

Since 11 March 2008, the medium-range variable-resolution ensemble system (VAREPS) has been merged with the monthly ensemble system, which has been operational at ECMWF since October 2004 (see [Vitart \(2004\)](#) for a description of the system). The merged VAREPS/monthly system (hereafter simply called the Ensemble Prediction System, EPS; see [Vitart *et al.*, 2008](#)) runs twice-a-day, at 00 and 12 UTC. On Mondays and Thursdays, the 00 UTC forecasts are extended to day 32.

The atmospheric model IFS is coupled from initial time to the ocean model NEMO (see [Madec, 2008](#)). The coupling uses the single executable framework described by [Mogensen *et al.* \(2012b\)](#). The ocean initial

conditions come from the near-real-time component of the operational ocean analysis (NEMOVAR, see [Mogensen *et al.*, 2012a](#)) and from the Ocean ReAnalysis System 4 (ORAS4, see [Balmaseda *et al.*, 2013](#)) for re-forecasts. Further details on the coupling are given in [3.1](#).

The 51 EPS non-linear integrations are performed twice a day, with initial times 00 and 12 UTC, up to forecast day 15, with a (T639) resolution from day 0 to day 10, and with lower (T319) resolution from day 10 to day 15. In the vertical, all integrations have 91 levels, with the model top at 1 Pa.

On 11 March 2008, a new ensemble re-forecast suite has been implemented to estimate the model climatological distribution (see [Hagedorn, 2008](#) for a description of the system). The re-forecast suite is based on 11-member ensembles, integrated with the same model version used in the operational EPS, and run twice weekly (on Mondays and Thursdays) for the same calendar day over the past 20 years (in 2014, from 1994 to 2013). Thus, every week 440 ensemble members $((11 + 11) \times 20)$ can be accessed to estimate the model climatology and/or for further calibration purposes. The initial conditions for the re-forecasts are defined by ERA-Interim analyses ([Dee *et al.*, 2011](#)). Re-forecasts are used to generate some operational medium-range (e.g. the Extreme Forecast Index) and monthly products.

The first EPS configurations were designed on the notion that forecast uncertainty is dominated by error or uncertainty in the initial conditions (see [Harrison *et al.*, 1999](#)). This was consistent with studies that show that, when two operational forecasts differ, it is usually differences in the analyses rather than differences in model formulation that are critical to explaining this difference. Later on it was recognized that random model errors due to physical parameterizations and the effects of subgrid-scale processes could affect the forecast quality too. Two schemes are used to represent model uncertainties in the EPS since November 2010: (i) the Stochastically Perturbed Parameterization Tendencies scheme (SPPT) and (ii) the Stochastic Kinetic Energy Backscatter scheme (SKEB). They simulate the effect on forecast quality of random model errors due to physical processes and subgrid-scale effects.

The SPPT scheme was first implemented in October 1998 and referred to originally as *stochastic physics* ([Buizza *et al.*, 1999b](#)). The scheme perturbs the tendencies by terms that are given by a random pattern times the net parameterized physics tendencies. The random pattern varies horizontally and with time; each ensemble member uses a different realisation of the random pattern. Two major revisions of the SPPT scheme took place in September 2009 ([Palmer *et al.*, 2009](#)) and in November 2010. As in the original scheme, tendency perturbations are given by a multiplicative noise term based on the net parameterized physics tendencies. In the revised SPPT scheme, the random pattern that defines the perturbations varies smoothly in space and time in contrast to the original scheme. The random numbers at a grid point are given by a truncated Gaussian distribution. The pattern is defined as the sum of three independent random patterns with horizontal correlation scales of 500, 1000 and 2000 km and correlation time scales of 6 h, 3 d and 30 d, respectively. The standard deviation of the three patterns decreases with increasing spatial and temporal scale from 0.52 to 0.18 and finally to 0.06.

The SKEB scheme perturbs the vorticity tendencies with stochastic noise that is given by an evolving 3-dim pattern times a term that is proportional to the square root of an estimate of the kinetic energy dissipation rate in the numerical model [Berner *et al.* \(2009\)](#). The dissipation rate estimate consists of contributions from the numerical dissipation and a contribution from convection. The correlation time of the pattern is 7 h and the spatial correlations have been estimated from coarse-graining experiments [Palmer *et al.* \(2009\)](#).

There are two key advantages of using ensemble-based probabilistic forecasts instead of single forecasts. Firstly, they provide a more complete estimation of the future forecast states, since they not only provide the most likely scenario but also the probability that alternatives might occur. Secondly, they provide more consistent forecasts ([Zsoter *et al.*, 2009](#)). [Fig. 1.1](#) shows a schematic illustration of the phase-space evolution of the probability distribution function (PDF) of analysis error throughout the forecast range. A specific isopleth (e.g. the 1 standard deviation isopleth) is illustrated. It is assumed that the initial time PDF is a Gaussian distribution. At initial time ([Fig. 1.1\(a\)](#)), the isopleth is shown as isotropic, i.e. bounding an n -sphere, where the dimension of phase space n is of the order of $(O(10^8))$ for the ECMWF operational forecast model). In general, this error will not be isotropic—analysis error is likely to be larger along directions which are less well observed, and vice versa. However, at least in principle, it is straightforward to define an inner product on phase space, with respect to which the initial PDF is




Figure 1.1 A schematic illustration of the growth of an isopleth of the forecast error probability distribution function, from (a) initial phase, to (b) linear growth phase, to (c) nonlinear growth phase, to (d) loss of predictability. See text for further details.

isotropic. This inner product, defined from the analysis error covariance matrix, plays a fundamental role in the theory of singular vectors discussed below.

In the early part of the forecast, error growth is governed by linear dynamics. During this period an initially spherical isopleth of the PDF will evolve to bound an n -dimensional ellipsoidal volume (Fig. 1.1(b)). The major axis of the ellipsoid corresponds to a phase-space direction which defines the dominant finite-time instability of that part of phase space (relative to the analysis error covariance metric). The arrow shown in Fig. 1.1(b) points along the major axis of the ellipsoid. It can be thought of as evolving from the arrow shown in Fig. 1.1(a). Note that the arrows in Fig. 1.1(a) and Fig. 1.1(b) are not parallel to one another. This illustrates the non-modal nature of linear perturbation growth.

The arrows at initial and forecast time define the dominant singular vector at initial and final time (with respect to the analysis error covariance metric). At forecast time, the dominant singular vector defines the dominant eigenvector of the forecast error covariance matrix. See Section 1.3.1 for more details.

The growth of the (isopleth of the) PDF between Fig. 1.1(b) and Fig. 1.1(c) describes a nonlinear evolution of the PDF. In Figs 1.1(c) the PDF has deformed from its ellipsoidal shape in Fig. 1.1(b). The nonlinear deformation will cause the PDF to evolve away from a Gaussian distribution. Put another way, in the nonlinear phase, the PDF in any given direction is partially determined by perturbations which, in the linear phase, were orthogonal to that direction. Finally, Fig. 1.1(d) shows (schematically) the situation where the evolved PDF has effectively become indistinguishable from the system's attractor, so that all predictability has been lost.

1.2 INITIAL CONDITION PERTURBATIONS

The operational EPS samples initial uncertainty using perturbations from an ensemble of data assimilations (EDA) and perturbations constructed from the leading singular vectors since June 2010. The perturbations from the EDA replace the perturbations that were based on evolved singular vectors prior to June 2010.

The observed values are perturbed with stochastic noise in the EDA. The variance of the observation perturbations is given by the observation error variance estimate that is used in the data assimilation. In addition, model uncertainties are represented in the EDA using the SPPT scheme. At the time of writing (April 2014), the EDA consists of 25 perturbed analyses with an outer loop resolution of T399. Further details about the data assimilation system and the EDA are provided in Part II of the IFS documentation and by Isaksen *et al.* (2010). The EDA also provides the operational high-resolution 4D-Var with an estimate of the flow-dependent variations of the background error variances.

If all relevant sources of uncertainty were represented accurately in the EDA, it should provide a reliable sample from the distribution of initial uncertainties. Ensemble prediction experiments using only EDA initial perturbations indicate that this goal has not been achieved yet. Therefore, [Buizza *et al.* \(2008\)](#) propose the hybrid approach of combining EDA perturbations with singular vector perturbations. The perturbations from the EDA sample errors that have grown in the past while the perturbations based on singular vectors sample structures that are likely to dominate forecast error in the future, say at Day 2 and beyond.

1.3 SINGULAR VECTORS

The number of degrees of freedom of the operational ECMWF model is many orders of magnitude larger than the largest practicable ensemble size. This raises the question of whether any particular strategy is desirable in sampling the initial PDF. If initial errors can occur independently in all the phase-space directions, then a strategy of random-under-sampling could lead to an EPS whose reliability was poor, especially for cases of small ensemble spread. In particular, if the spread from a randomly under-sampled ensemble was found to be small on a particular occasion, this could either be because the flow was especially predictable, or because the ensemble perturbations poorly sample the unstable subspace in which the analysis error lay. From a credibility perspective, it is important to try to minimize the latter type of occurrence.

An alternative strategy is to base the perturbations on the leading singular vectors. Clearly, by focusing on the unstable subspace, the cases of small spread being associated with large forecast error should be minimized, at least in the linear and weakly nonlinear range. In addition to this, there are several related reasons why the initial perturbations for the ECMWF EPS are based on the dominant singular vectors.

Firstly, as shown by [Rabier *et al.* \(1996\)](#) and [Gelaro *et al.* \(1998\)](#), the sensitivity of day-2 forecast error to perturbations in the initial state projects well into the space of dominant singular vectors. Rabier *et al.* have shown that cases of severe forecast failure can be dramatically improved if the analysis is modified using the sensitivity perturbations.

Secondly, provided the metric or inner product for the singular vectors is an accurate reflection of the analysis error covariance matrix then, as mentioned, the evolved singular vectors point along the largest eigenvectors of the forecast error covariance matrix. As such (see [Ehrendorfer and Tribbia, 1997](#)), perturbations constructed from the dominant singular vectors represent the most efficient means for predicting the forecast error covariance matrix, given a pre-specified number of allowable tangent model integrations.

Thirdly, singular vector perturbations may provide a relatively efficient means of sampling the forecast error PDF in the nonlinear range, particularly during transitions in weather regimes. For example, [Mureau *et al.* \(1993\)](#) have shown a case where the singular vector perturbations were successful in capturing a major transition to blocking, where random perturbations were inadequate. [Gelaro *et al.* \(1998\)](#) have documented further such cases. A more systematic study of the ability of ensembles to describe the probability of regime transitions in the weakly nonlinear forecast range has been made by [Trevisan *et al.* \(2001\)](#) using an intermediate-complexity model of the extratropical circulation. Relatively small ensembles initialised using firstly singular vectors, and secondly local Lyapunov vectors, were compared with a large Monte-Carlo ensemble based on random perturbations. It was found that small ensemble spread from the singular-vector ensemble was a reliable indicator of small ensemble spread from the Monte-Carlo ensemble. By contrast, small ensemble spread from the Lyapunov-vector ensemble was a much less reliable indicator of small spread from the Monte-Carlo ensemble.

Fourthly, in practice only imperfect information about the initial PDF is available. Hence it is difficult to even define a truly random initial sampling.

In the current system, each EPS perturbation is a linear combination of singular vectors with maximum growth computed using a total energy norm, over a 48 hour time period. The leading 50 singular vectors are used for each extra-tropical hemisphere, and the leading 5 singular vectors are used in each of up to 6 tropical areas. They are computed at T42 resolution and with 91 vertical levels. The assumption underlying the linear combination is that initial error is distributed in the space spanned by the singular

vectors according to an isotropic Gaussian distribution. A Gaussian sampling technique is used to sample realisations from this distribution (see [Leutbecher and Palmer \(2008\)](#) and [Section 2.3](#)).

The reader is referred to [Buizza and Palmer \(1995\)](#) for a description of the singular vector computation in the extra-tropics, and to [Barkmeijer *et al.* \(2001\)](#) for a discussion of the computation of the tropical singular vectors. The reader is also referred to [Palmer *et al.* \(1998\)](#) for a discussion of the sensitivity of the singular vector structure to the metric used to measure perturbation growth, to [Hoskins *et al.* \(2000\)](#) for an analysis of the nature of singular vector growth, and to [Barkmeijer *et al.* \(1998\)](#) and [Lawrence *et al.* \(2009\)](#) for comparisons of singular vectors computed using a total energy metric and singular vectors based on a Hessian metric. [Leutbecher and Lang \(2013\)](#) study the reliability of ensemble variances in subspaces spanned by the leading singular vectors. They also examine the sensitivity of this reliability to the dimension of the singular vector subspace and the horizontal resolution of the singular vectors.

1.3.1 Formulation of the singular vector computation

One way to define singular vectors is by means of a maximization problem. The scalar which has to be maximized can be written as

$$\frac{[\mathbf{E}\mathbf{x}, \mathbf{M}\mathbf{x}]}{[\mathbf{D}\mathbf{x}, \mathbf{x}]} \quad (1.1)$$

where $[\cdot, \cdot]$ denotes the Euclidean inner product, $[\mathbf{x}, \mathbf{y}] = \sum(x_i y_i)$, and \mathbf{D} and \mathbf{E} are positive definite operators. The operator \mathbf{M} is the propagator of the tangent model. It assigns to a particular vector \mathbf{x} the linearly evolved vector $\mathbf{M}\mathbf{x}$ for a given forecast time and with respect to a reference trajectory. Hence, the scalar defined by (1.1) is the ratio between the \mathbf{E} -norm of the evolved vector \mathbf{x} and the \mathbf{D} -norm of \mathbf{x} at initial time. Notice that the norm at initial and final time may differ. The leading singular vector has the property that it maximizes the scalar, the second singular vector maximizes the scalar in the space \mathbf{D} -orthogonal to the leading singular vector, and so forth. In this way, one obtains a set of singular vectors which are \mathbf{D} -orthogonal at initial time and \mathbf{E} -orthogonal at final time. The actual computation of the singular vectors in the IFS is done by solving an equivalent eigenvalue problem. Observe that the solutions of the maximization problem (1.1) also satisfy the generalized eigenvalue problem (1.2) given by

$$\mathbf{M}^* \mathbf{E} \mathbf{M} \mathbf{x} = \lambda \mathbf{D} \mathbf{x} \quad (1.2)$$

where \mathbf{M}^* is the adjoint of \mathbf{M} . The defining equation (1.2) can be generalized by activating operators in the singular vector computation, see [Section 1.3](#). It is, for instance, possible to set the state vector to zero outside a prescribed area at optimization time, by using a projection operator \mathbf{P} ([Buizza, 1994](#)). Consequently, the growth of singular vectors outside the target area is not taken into account in the actual computation. In using this projection operator, the eigenvalue problem (1.2) becomes $\mathbf{M}^* \mathbf{P}^* \mathbf{E} \mathbf{P} \mathbf{M} \mathbf{x} = \lambda \mathbf{D} \mathbf{x}$. To keep the notation as simple as possible, these additional operators will be omitted from the basic eigenvalue problem (1.2).

The operator \mathbf{D} determines the properties by which the singular vectors are constrained at initial time. As such, it can be interpreted as an approximation of the inverse of the analysis error covariance matrix \mathbf{P}^a . Currently, there are two methods to compute singular vectors, depending on the form of \mathbf{D} . Both methods will be discussed in [Subsections 1.3.2](#) and [1.3.3](#).

1.3.2 Use of the total energy norm

When using the total energy norm, or any other simple operator, at initial time, the generalized eigenvalue problem (1.2) can be simplified to an ordinary eigenvalue problem. In this case the \mathbf{D} -norm of \mathbf{x} reads as

$$[\mathbf{x}, \mathbf{D}\mathbf{x}] = \frac{1}{2} \int_0^1 \int_{\Sigma} \left\{ \Delta^{-1} \zeta_{\mathbf{x}} \cdot \zeta_{\mathbf{x}} + \Delta^{-1} D_{\mathbf{x}} \cdot D_{\mathbf{x}} + \frac{c_p}{T_{\text{ref}}} T_{\mathbf{x}}^2 + w_q \frac{L_{\text{cond}}^2}{c_p T_{\text{ref}}} q_{\mathbf{x}}^2 \right\} d\Sigma \frac{\partial p}{\partial \eta} d\eta + \frac{1}{2} \int_{\Sigma} R_{\text{dry}} T_{\text{ref}} P_{\text{ref}} \ln \pi_{\mathbf{x}}^2 d\Sigma \quad (1.3)$$

where $\zeta_{\mathbf{x}}$, $D_{\mathbf{x}}$, $T_{\mathbf{x}}$ and $\ln \pi_{\mathbf{x}}$ stand for the vorticity, divergence, temperature, specific humidity and logarithm of the surface pressure component of the state vector \mathbf{x} , and c_p is the specific heat of dry air at constant pressure, L_{cond} is the latent heat of condensation at 0°C , R_{dry} is the gas constant for dry

air, $T_{\text{ref}} = 300 \text{ K}$ is a reference temperature and $P_{\text{ref}} = 800 \text{ hPa}$ is a reference pressure. The parameter w_q defines the relative weight given to the specific humidity term. The integration over the sphere and in the vertical are denoted by $\int_{\Sigma} d\Sigma$ and $\int_0^1 d\eta$, respectively.

Since the operator \mathbf{D} is a diagonal matrix, one can easily define a matrix \mathbf{C} so that $\mathbf{C}^2 = \mathbf{D}^{-1}$. Multiplying both sides of (1.2) to the left and right with \mathbf{C} , yields the equation

$$\mathbf{C}\mathbf{M}^*\mathbf{E}\mathbf{M}\mathbf{C}\mathbf{x} = \lambda \mathbf{x} \quad (1.4)$$

which can be solved using the Lanczos algorithm, see [Section A.1](#). The energy metric is believed to be a useful approximation to the analysis error covariance metric ([Palmer *et al.*, 1998](#)).

1.3.3 Use of the Hessian of the 3D-Var objective function

In the incremental formulation of 3D-Var, the Hessian of the objective function \mathfrak{S} can be used as an approximation of the inverse of the analysis error covariance matrix. The objective function has the form

$$\mathfrak{S}(\delta\mathbf{x}) = \frac{1}{2}\delta\mathbf{x}^T\mathbf{B}^{-1}\delta\mathbf{x} + \frac{1}{2}(\mathbf{H}\delta\mathbf{x} - \mathbf{d})^T\mathbf{R}^{-1}(\mathbf{H}\delta\mathbf{x} - \mathbf{d}) \quad (1.5)$$

and the increment $\delta\mathbf{x}^a$ where \mathfrak{S} attains its minimum, provides the analysis \mathbf{x}^a which is defined by adding $\delta\mathbf{x}^a$ to the background \mathbf{x}^b so that

$$\mathbf{x}^a = \mathbf{x}^b + \delta\mathbf{x}^a \quad (1.6)$$

The operators \mathbf{B} and \mathbf{R} are covariance matrices of the background and observation error respectively and \mathbf{d} is the innovation vector given by

$$\mathbf{d} = \mathbf{y} - \mathbf{H}\mathbf{x}^b \quad (1.7)$$

where \mathbf{y} is the observation vector and \mathbf{H} is a linear approximation of the observation operator in the vicinity of \mathbf{x}^b .

The Hessian $\nabla\nabla\mathfrak{S}$ of the objective function is given by

$$\nabla\nabla\mathfrak{S} = \mathbf{B}^{-1} + \mathbf{H}^T\mathbf{R}^{-1}\mathbf{H} \quad (1.8)$$

Provided that the background error ($\mathbf{x}^b - \mathbf{x}^t$) and the observation error ($\mathbf{y} - \mathbf{H}\mathbf{x}^t$) are uncorrelated, with \mathbf{x}^t the true state of the atmosphere, the Hessian $\nabla\nabla\mathfrak{S}$ is equal to the inverse of the analysis error covariance matrix \mathbf{P}^a . This follows by noting that the objective function is quadratic and attains its unique minimum at $\delta\mathbf{x}^a$ and consequently, by using (1.6),

$$\mathbf{B}^{-1}[\mathbf{x}^a - \mathbf{x}^t - (\mathbf{x}^b - \mathbf{x}^t)] + \mathbf{H}^T\mathbf{R}^{-1}[\mathbf{H}(\mathbf{x}^a - \mathbf{x}^t) + \mathbf{H}\mathbf{x}^t - \mathbf{y}] = 0 \quad (1.9)$$

Rewriting (1.9) gives

$$(\mathbf{B}^{-1} + \mathbf{H}^T\mathbf{R}^{-1}\mathbf{H})(\mathbf{x}^a - \mathbf{x}^t) = \mathbf{B}^{-1}(\mathbf{x}^b - \mathbf{x}^t) - \mathbf{H}^T\mathbf{R}^{-1}(\mathbf{H}\mathbf{x}^t - \mathbf{y}) \quad (1.10)$$

Using the assumption that the background and observation error are uncorrelated the above equation implies that

$$(\mathbf{B}^{-1} + \mathbf{H}^T\mathbf{R}^{-1}\mathbf{H})\mathbf{P}^a(\mathbf{B}^{-1} + \mathbf{H}^T\mathbf{R}^{-1}\mathbf{H})^T = (\mathbf{B}^{-1} + \mathbf{H}^T\mathbf{R}^{-1}\mathbf{H})^T \quad (1.11)$$

Now by multiplying each side of (1.11) to the right with its transpose, the desired result follows.

The defining eigenvalue problem for the singular vectors becomes

$$\mathbf{M}^*\mathbf{E}\mathbf{M}\mathbf{x} = \lambda \nabla\nabla\mathfrak{S} \mathbf{x} \quad (1.12)$$

Since the objective function is quadratic in the incremental formulation, the Hessian applied to a vector $\nabla\nabla\mathfrak{S} \mathbf{x}$ in (1.12) can be evaluated as finite difference of two gradients: $(\nabla\nabla\mathfrak{S})\mathbf{x} = \nabla\mathfrak{S}(\mathbf{x} + \mathbf{x}^b) - \nabla\mathfrak{S}(\mathbf{x}^b)$. The generalized eigenvalue problem (1.12) is solved by using the Jacobi–Davidson algorithm, see [Section A.2](#).

APPENDIX A. EIGENVALUE ALGORITHMS

A.1 The Lanczos algorithm

Algorithms based on Lanczos theory are very useful to solve an eigenvalue problem when only a few of the extreme eigenvectors are needed. It can be applied to large and sparse problems. The algorithm does not access directly the matrix elements of the operator that defines the problem, but it gives an estimate of the eigenvectors through successive application of the operator.

Consider the eigenvalue problem

$$\mathbf{A}\mathbf{x} = \sigma_i^2 \mathbf{x} \quad (\text{A.1})$$

where the matrix \mathbf{A} is $N \times N$ dimensional, and symmetric. Without loss of generality, we can also suppose that it is real.

If \mathbf{A} is a real, symmetric matrix, then there exists an orthogonal real matrix \mathbf{Q} such that

$$\mathbf{Q}^T \mathbf{A} \mathbf{Q} = \mathbf{D}(\lambda_1, \dots, \lambda_N) \quad (\text{A.2})$$

where $\mathbf{D}(\lambda_1, \dots, \lambda_N)$ is a diagonal matrix, and \mathbf{Q}^T denotes the transpose of \mathbf{Q} (Schur decomposition theorem).

The Lanczos algorithm does not directly compute the diagonal matrix \mathbf{D} , but it first computes a partial transformation of the matrix \mathbf{A} using a tridiagonal matrix \mathbf{T}

$$\mathbf{Q}^T \mathbf{A} \mathbf{Q} = \mathbf{T} \quad (\text{A.3})$$

with

$$\mathbf{T} = \begin{bmatrix} \alpha_1 & \beta_2 & 0 & \cdot & \cdot & \cdot & \cdot \\ \beta_2 & \alpha_2 & \beta_3 & 0 & \cdot & \cdot & \cdot \\ 0 & \beta_3 & \cdot & \cdot & \cdot & \cdot & \cdot \\ \cdot & 0 & \cdot & \cdot & \cdot & 0 & \cdot \\ \cdot & \cdot & \cdot & \cdot & \cdot & \beta_{J-1} & 0 \\ \cdot & \cdot & \cdot & 0 & \beta_{J-1} & \alpha_{J-1} & \beta_J \\ \cdot & \cdot & \cdot & \cdot & 0 & \beta_J & \alpha_J \end{bmatrix} \quad (\text{A.4})$$

and with

$$\mathbf{Q} = [\mathbf{q}_1, \dots, \mathbf{q}_J] \quad (\text{A.5})$$

where the vectors \mathbf{q}_j are column vectors, and where the number of iterations J is much smaller than the dimensionality of the problem, $J \ll N$. Then, the Lanczos algorithm finds the diagonal decomposition of \mathbf{T}

$$\mathbf{T} = \mathbf{S}^T \mathbf{D} \mathbf{S} \quad (\text{A.6})$$

The elements of the diagonal matrix \mathbf{D} are an estimate of the eigenvalues of \mathbf{A} , and an estimate of the eigenvectors are given by $\mathbf{Y} = [\mathbf{y}_1, \dots, \mathbf{y}_j]$, with

$$\mathbf{Y} = \mathbf{Q} \mathbf{S} \quad (\text{A.7})$$

The actual computation is performed by writing (A.3) as

$$\mathbf{A} \mathbf{Q} = \mathbf{Q} \mathbf{T} \quad (\text{A.8})$$

Equating columns of (A.8), it follows that

$$\mathbf{A} \mathbf{q}_j = \beta_{j-1} \mathbf{q}_{j-1} + \alpha_j \mathbf{q}_j + \beta_j \mathbf{q}_{j+1} \quad (\text{A.9})$$

for $j = 1, \dots, J$. The orthogonality of the vectors \mathbf{q}_j implies that

$$\alpha_j = \mathbf{q}_j^T \mathbf{A} \mathbf{q}_j \quad (\text{A.10})$$

Moreover, if

$$\mathbf{r}_j = (\mathbf{A} - \alpha_j \mathbf{I}) \mathbf{q}_j - \beta_{j-1} \mathbf{q}_{j-1} \quad (\text{A.11})$$

is non-zero, then

$$\mathbf{q}_{j+1} = \frac{\mathbf{r}_j}{\beta_j} \quad (\text{A.12})$$

where $\beta_j = \pm\sqrt{\langle \mathbf{r}_j; \mathbf{r}_j \rangle}$. An iterative application of these equations, with a randomly chosen starting vector \mathbf{q}_1 , defines the Lanczos iterative procedure. The total number of iterations J determines the accuracy of the computation. As this number increases, more eigenpairs can be separated from the others, independently from the choice of the starting vector \mathbf{q} . This separation starts from the boundaries of the eigenvalue spectrum. The accuracy of the eigenvectors is less than the accuracy of the singular values, say to order ε when the precision of the singular values is of order ε^2 .

The reader is referred to [Golub and van Loan \(1983\)](#) for a theoretical description of the Lanczos algorithm. The Lanczos code is available in NAG issue 17. In the IFS, the Lanczos code in subroutine **CONGRAD** is used to compute the singular vectors used in the operational EPS.

A.2 The Jacobi–Davidson algorithm

The generalized eigenproblem

$$\mathbf{A}\mathbf{x} = \lambda\mathbf{B}\mathbf{x} \quad (\text{A.13})$$

is usually handled by bringing it back to a standard eigenproblem

$$\mathbf{B}^{-1}\mathbf{A}\mathbf{x} = \lambda\mathbf{x} \quad (\text{A.14})$$

The matrix $\mathbf{B}^{-1}\mathbf{A}$ is in general nonsymmetric, even if both \mathbf{A} and \mathbf{B} are symmetric. However, if \mathbf{B} is symmetric and positive definite, the \mathbf{B} -inner product is well defined. The matrix $\mathbf{B}^{-1}\mathbf{A}$ is symmetric in this inner product if \mathbf{A} is symmetric:

$$[\mathbf{w}, \mathbf{B}^{-1}\mathbf{A}\mathbf{v}]_{\mathbf{B}} = [\mathbf{B}\mathbf{w}, \mathbf{B}^{-1}\mathbf{A}\mathbf{v}] = [\mathbf{w}, \mathbf{A}\mathbf{v}] = [\mathbf{A}\mathbf{w}, \mathbf{v}] = [\mathbf{B}^{-1}\mathbf{A}\mathbf{w}, \mathbf{v}]_{\mathbf{B}} \quad (\text{A.15})$$

The proposed method to solve (A.14) constructs a set of basis vectors \mathbf{V} of a search space v , c.f. the Lanczos method. The approximate eigenvectors are linear combinations of the vectors \mathbf{V} . The classical and most natural choice for the search space v , for instance utilized in the Lanczos method, is the so-called Krylov subspace, the space spanned by the vectors

$$\mathbf{v}, \mathbf{B}^{-1}\mathbf{A}\mathbf{v}, (\mathbf{B}^{-1}\mathbf{A})^2\mathbf{v}, \dots, (\mathbf{B}^{-1}\mathbf{A})^{i-1}\mathbf{v} \quad (\text{A.16})$$

This i -dimensional subspace is denoted by $\mathbf{K}^i(\mathbf{v}, \mathbf{B}^{-1}\mathbf{A})$. The vector \mathbf{v} is a starting vector that has to be chosen. The Krylov subspace is well suited for computing dominant eigenpairs since the vector $(\mathbf{B}^{-1}\mathbf{A})^{i-1}\mathbf{v}$ points more and more in the direction of the dominant eigenvector of $\mathbf{B}^{-1}\mathbf{A}$ for increasing i .

Given a search space v , the approximate eigenpair (θ, \mathbf{u}) of (A.15) is a linear combination of the basis vectors of v so that

$$\mathbf{u} = \mathbf{V}\mathbf{y} \quad (\text{A.17})$$

A suitable criterion for finding an optimal pair is the Galerkin condition that the residual

$$\mathbf{r} = \mathbf{B}^{-1}\mathbf{A}\mathbf{u} - \theta\mathbf{u} = \mathbf{B}^{-1}\mathbf{A}\mathbf{V}\mathbf{y} - \theta\mathbf{V}\mathbf{y} \quad (\text{A.18})$$

is \mathbf{B} -orthogonal to the search space v . Hence

$$\mathbf{V}^T\mathbf{B}\mathbf{r} = 0 \quad (\text{A.19})$$

and consequently, using (A.19),

$$\mathbf{V}^T\mathbf{A}\mathbf{V}\mathbf{y} - \theta\mathbf{V}^T\mathbf{B}\mathbf{V}\mathbf{y} = 0 \quad (\text{A.20})$$

Note that the resulting eigenproblem is of the dimension of the search space, which is generally much smaller than of the original problem. The basis vectors are usually orthonormalized so that $\mathbf{V}^T\mathbf{B}\mathbf{V} = \mathbf{I}$. Approximate eigenpairs that adhere to the Galerkin condition are called Ritz pairs.

It can be shown that the residuals $\mathbf{r}_1, \mathbf{r}_2, \dots, \mathbf{r}_i$ form a \mathbf{B} -orthogonal basis for $\mathbf{K}^i(\mathbf{r}_1, \mathbf{B}^{-1}\mathbf{A})$ when the approximate eigenpairs are computed according to (A.17).

A.3 Solution method

As was stated before, the natural search space for the generalized eigenvalue problem is the Krylov subspace $\mathbf{K}^i(\mathbf{v}, \mathbf{B}^{-1}, \mathbf{A})$. This basis can be generated by expanding the basis by new residual vectors. The problem in the construction of a basis for this space is that operations with \mathbf{B}^{-1} are needed. Since in our application the inverse of \mathbf{B} is not known explicitly, its action is approximated by the Conjugate Residual method (CR). To compute the vector

$$\mathbf{r} = \mathbf{B}^{-1}\tilde{\mathbf{r}} \quad (\text{A.21})$$

one iteratively solves the system

$$\mathbf{B}\mathbf{r} = \tilde{\mathbf{r}} \quad (\text{A.22})$$

Iterative solution methods require, apart from vector operations, only multiplications with \mathbf{B} . The vector \mathbf{r} can in principle be determined to high accuracy. This, however, may require many multiplications with \mathbf{B} and hence may be very expensive. Therefore, the action of \mathbf{B}^{-1} is approximated to low accuracy, by performing only a few steps with an iterative solution method. The number of iterations is controlled by NINNER. The subspace generated in this way is not a Krylov subspace and the basis vectors are not the residuals (A.19) but only approximations to it. As a consequence they are not perfectly \mathbf{B} -orthogonal. The orthogonalization has to be done explicitly.

The complete algorithm can be summarized as follows.

- (i) Choose a starting vector \mathbf{v} .
- (ii) Compute $\mathbf{B}\mathbf{v}$, \mathbf{B} -normalize \mathbf{v} .
- (iii) Repeat the steps (iv)–(x) NITERL times.
- (iv) Compute $\mathbf{A}\mathbf{V}$, $\mathbf{V}^T\mathbf{A}\mathbf{V}$
- (v) Solve small eigenproblem $\mathbf{V}^T\mathbf{A}\mathbf{V}\mathbf{y} = \theta\mathbf{y}$.
- (vi) Select Ritz value θ and \mathbf{y} .
- (vii) Compute Ritz vector $\mathbf{u} = \mathbf{V}\mathbf{y}$ and residual $\tilde{\mathbf{r}} = \mathbf{A}\mathbf{V}\mathbf{y} - \theta\mathbf{B}\mathbf{V}\mathbf{y}$.
- (viii) Compute approximately $\mathbf{v} = \mathbf{B}^{-1}\tilde{\mathbf{r}}$ with the CR-method.
- (ix) \mathbf{B} -orthonormalize new \mathbf{v} against \mathbf{V} .
- (x) Expand \mathbf{V} with the resulting vector.

The matrix \mathbf{V} contains the basis for the search space, the vector \mathbf{v} contains the new basis vector and (\mathbf{u}, θ) is an approximate eigenpair. In step (vi) the pair with the largest θ is selected because only the dominant part of the spectrum is of interest. However, if the eigenpair approximation reaches a certain accuracy ε (XKAPA in namelist NAMLCZ), i.e., $\tilde{\mathbf{r}} = \mathbf{A}\mathbf{u} - \theta\mathbf{B}\mathbf{u} \leq \varepsilon$, a smaller θ is selected. In step (viii) a few CR steps are performed to approximate the action of \mathbf{B}^{-1} . In step (ix), the Modified Gram–Schmidt procedure is used for reasons of numerical stability. For a more detailed description of the Jacobi–Davidson algorithm, the reader is referred to [Sleijpen and van der Vorst \(1996\)](#)

Chapter 2

Computational details: initial perturbations

Table of contents

- 2.1 Introduction**
- 2.2 Singular vectors**
 - 2.2.1 The singular vector code
 - 2.2.2 Computation of singular vectors
- 2.3 Initial perturbation component based on singular vectors**
- 2.4 EDA component of the initial perturbations**

2.1 INTRODUCTION

For a description of the evolution of the ECMWF Ensemble Prediction System (EPS), the reader is referred to [Molteni *et al.* \(1996\)](#), [Buizza *et al.* \(1998, 1999a, 2003, 2007\)](#), [Leutbecher and Palmer \(2008\)](#) and [Buizza *et al.* \(2008\)](#). As mentioned in [Section 1.1](#), for each initial date, an ECMWF ensemble comprises one ‘control’ forecast, which is a forecast started from the operational analysis, and N_{ens} perturbed forecasts. The initial conditions for the perturbed integrations are constructed by adding and subtracting to the operational analysis $N_{\text{ens}}/2$ perturbations. These perturbations are determined from perturbations originating from the Ensemble of Data Assimilations (EDA) and from the leading singular vectors. We begin with a description of the singular vector perturbations. The EDA perturbations are described in [section 2.4](#).

2.2 SINGULAR VECTORS

The part of the initial perturbations that is based on singular vectors is constructed with a Gaussian sampling technique ([Leutbecher and Palmer, 2008](#)). The sampling uses the leading singular vectors of each of the following sets of singular vectors:

- (i) Initial 50 leading northern hemisphere singular vectors, file ‘svifs’.
- (ii) Initial 50 leading southern hemisphere singular vectors, file ‘svifss’.
- (iii) Initial 5 leading singular vectors in up to 6 regions targeted on tropical cyclones, file(s) ‘svifs1’ to ‘svifs6’.

2.2.1 The singular vector code

The singular vector computation is called from **CUN3** in **CNT0** or from **CUN2** in **CUN1**, depending on whether the Hessian of the objective function is used or not. General routines related to the singular vector computation can be found in the directory ‘sinvect’. The technical routine needed for the Lanczos and Jacobi–Davidson algorithms is contained in the directory ‘var’ of IFS. Details of the set-up routines are given in [Table 2.1](#).

Subroutines **OPK** and **OPM** are the routines in which the main computational work of the singular vector computation takes place. The defining equation for the singular vectors [\(1.2\)](#) takes the form

$$\text{OPK}(\mathbf{x}) = \lambda \text{OPM}(\mathbf{x}) \tag{2.1}$$

The routine **OPM** is the identity in case a simple energy norm is used at initial time, or it is equal to the Hessian of the 4D-Var objective function. In the routine **OPK**, the propagators of the tangent and

Table 2.1 Constants in namelist *NAMLCZ* controlling the singular vector computation. Defaults refer to *LECMWF=TRUE*.

Name	Type	Purpose	Default
LANCZOS	LOGICAL	Activates singular vector computation with Lanczos algorithm (NCONF = 601)	TRUE
LJACDAV	LOGICAL	Activates singular vector computations with Jacobi–Davidson algorithm (NCONF = 131)	FALSE
L_USE_CONGRAD	LOGICAL	Use Lanczos algorithm in subroutine CONGRAD	TRUE
LOCNORM	LOGICAL	Switch to localize norm computation in grid space	TRUE
ALAT1	REAL	} NW corner of local area (defined by point 1)	90.0
ALON1	REAL		0.0
ALAT3	REAL		} SE corner of local area (defined by point 3)
ALON3	REAL	359.9	
NLEVMIN	INTEGER	Minimum level of local area	1
NLEVMAX	INTEGER	Maximum level of local area	NLEVG
LSPTRLC0	LOGICAL	Switch to truncate in spectral space at initial (0) time	FALSE
LSPTRLC1	LOGICAL	Switch to truncate in spectral space at final (1) time	FALSE
NWTRMIN0(1)	INTEGER	Spectral coefficients with total wavenumber outside window are set to zero	0
NWTRMAX0(1)	INTEGER		NXMAX
LNEWNORMT0	LOGICAL	Switch to re-define the norm at initial time	FALSE
NEWNORMT0	INTEGER	Re-defines initial norm 1: total energy 2: kinetic energy 3: vorticity squared 4: stream function squared 5: rotational kinetic energy	1
NITERL	INTEGER	Maximum number of Lanczos or Jacob–Davidson inner iterations	70
NWEIGL	INTEGER	Number of singular vectors that are requested	50
XKAPA	REAL	Relative accuracy of RITZ value to consider a singular vector as converged	0.01
XMIN_RITZ	REAL	Smallest accepted RITZ value in CONGRAD	1.0
NINNER	INTEGER	Number of Jacobi–Davidson inner iterations	2
NJDSTOP	INTEGER	Value of NSTOP when LJACDAV = TRUE	144
LEVOLC	LOGICAL	Switch to evolve singular vectors	TRUE
NEIGEVO	INTEGER	Number of singular vectors to evolve	35
NLANTYPE	INTEGER	Determines type of singular vectors 1: energy type norms are used 2 and 3: obsolete 4: same as 1 but now NCONF = 131 5: Hessian is used at initial time 6: for computing eigensystem of the Hessian	1
LSYMCHECK	LOGICAL	Switch to activate symmetry test for operator OPK	FALSE
L_SUBSPACE_SVS	LOGICAL	Compute singular vectors in subspace orthogonal to first N_DIM_SUBSPACE vectors in file 'sv_subspace'	FALSE
N_DIM_SUBSPACE	INTEGER	Number N_{sub} of vectors in 'sv_subspace' used to define subspace	
LSCALC	LOGICAL	Write inner product coefficients into grib-file 'scalw'	FALSE
LRNORMALIZE	LOGICAL	Re-normalise initial singular vectors using inner product defined by 'suscal'	FALSE

The routine **SULCZ** defines the constants as listed in Table 2.1 and determines the configuration of the singular vector computation.

adjoint model are evaluated. Depending on the setting of the constants in `NAMLCZ`, additional operators may be active, such as `SPTRLCZTL` when `LSPTRLC0 = TRUE`. It truncates the state vector in spectral space by setting all harmonic coefficients with total wavenumber smaller than `NWTRMIN0`, or larger than `NWTRMAX0`, to zero at initial time.

2.2.2 Computation of singular vectors

The computation of singular vectors for the EPS is organised in the following tasks in family (`/mc/main/HH/sv`). Task names are given in parentheses.

- (i) Retrieval of the 6-hour forecast which is used to initialise the singular vector trajectory (`getini`).
- (ii) Interpolation of the forecast to `T42L91`, the resolution of the singular vectors (`inidata_sv`).
- (iii) Computation of the extra-tropical initial singular vectors for the southern (`sh/sv`) and northern (`nh/sv`) hemispheres.
- (iv) Optimisation regions for the singular vector computations targeted on tropical cyclones are determined (`targets`).
- (v) Definition of the subspace of extra-tropical singular vectors used when computing the singular vectors targeted on a tropical cyclone (`subspace`).
- (vi) Computation of up to six sets of singular vectors targeted on tropical cyclones (`tc1/sv`), (`tc2/sv`), ... (`tc6/sv`).

At the time of revising this documentation (April 2014), the singular vector resolution is `T42L91`.

The nonlinear trajectory which defines the tangent-linear system is started from a 6-hour forecast from the delayed-cutoff stream. The use of a slightly less accurate trajectory does not significantly alter the singular vector structure and does not compromise the probabilistic skill of the EPS ([Leutbecher, 2005](#)). It has the advantage that more resources can be devoted to the singular vector computation and that the nonlinear forecasts can be started as soon as the final analysis is available.

For each hemisphere, the leading 50 singular vectors are computed. The tangent-linear model uses an adiabatic version of the forecast model (apart from vertical diffusion and surface friction).

Initial EPS perturbations in the tropics are included since cycle `Cy24r3` (introduced operationally in January 2002). The leading 5 singular vectors are computed in each of up to 6 target areas. Each tropical depression/storm (WMO class ≥ 1) reported on the GTS is potentially assigned its own optimisation region. When the tropical depression/storm could be tracked in the EPS run 12 hours before the initial time, the track information is used to determine the optimisation region; otherwise the optimisation region is centred on the reported position of the tropical storm. When the procedure generates more than 6 target regions, the regions that overlap the most or are closest to each other are merged. In addition, regions with a very large overlap are merged. If no tropical storm is present in the Caribbean region an additional optimisation region in the Caribbean is specified ($0^\circ - 25^\circ\text{N}$ and $100^\circ - 60^\circ\text{W}$). The computation of the singular vectors targeted on tropical cyclones uses a diabatic version of the tangent-linear model; i.e. large-scale condensation and cumulus convection are represented in the tangent-linear model ([Barkmeijer *et al.*, 2001](#); [Puri *et al.*, 2001](#)).

The computation is triggered for tropical cyclones located between $40^\circ\text{S}-40^\circ\text{N}$. Therefore, the optimisation regions can significantly overlap with the optimisation region of the extra-tropical singular vector computation. In order to avoid a duplication of perturbations present in the subspace of the extra-tropical singular vectors, the singular vectors targeted on a tropical cyclone are computed in the subspace orthogonal to the space spanned by the leading 50 extra-tropical singular vectors ([Leutbecher, 2007](#)).

2.3 INITIAL PERTURBATION COMPONENT BASED ON SINGULAR VECTORS

The perturbations of the initial conditions are obtained from the singular vectors via a multivariate Gaussian sampling technique described in more detail by [Leutbecher and Palmer \(2008\)](#). Preliminary experimentation with an earlier version of the Gaussian sampling software is reported by [Ehrendorfer and Beck \(2003\)](#).

The generation of the ensemble initial perturbations consists of 3 steps. Task names are given in parentheses.

- (i) retrieval of the estimate of analysis error (getea).
- (ii) singular vector retrieval and pre-processing (getsvs).
- (iii) generation of initial perturbations (rot). The initial condition perturbations are written to a grib-file 'pert.grib' and are later archived in MARS (task lag/./archive/icp/ml).

Hereafter a brief description of each step is given.

The analysis error estimate is retrieved and smoothed in task (mc/sv/getea). The 4D-Var estimate of analysis error standard deviation is obtained from the FDB (or ECMWF MARS archive). The standard deviation is retrieved on all model levels on a reduced Gaussian grid corresponding to the singular vector resolution. To avoid spuriously small local values of the analysis error standard deviation, the fields are smoothed by truncation to total wavenumber ≤ 21 .

The linear combination of the singular vectors into initial perturbations is done in task (mc/sv/rot) The relevant shell-script is `sample_svs` in scripts/gen. The Gaussian sampling is performed in the following three steps

- (i) Computation of the standard deviation BETA of the sampled Gaussian distribution for each set of singular vectors. This happens by comparing the singular vectors with an estimate of analysis error standard deviation provided by 4D-Var. See program `aev_norm` in prepdata/mc_tools.
- (ii) Generation of coefficients for the Gaussian sampling. See program `gen_sv_coeff` in prepdata/mc_tools.
- (iii) Linear combination of singular vectors. See program `sv_lin_combi` in prepdata/mc_tools.

The scaling of the width of the Gaussian distribution is performed independently for each set of singular vectors. A standard deviation β of the Gaussian distribution to be sampled is determined as follows. The singular vectors are compared with an estimate of the analysis error standard deviation using an analysis error variance norm. This is defined as the L^2 -norm of the perturbation scaled by the estimate of the analysis error standard deviation. For the re-forecasts, this estimate of the analysis error standard deviation is provided by the four-dimensional variational assimilation system used in ERA-Interim. In the real-time forecasts, scaled analysis error standard deviations obtained from the EDA (mars type=ses) are used as an estimate of the analysis error standard deviation. The summation in the norm extends over the entire model grid and involves the horizontal wind components, temperature and the logarithm of surface pressure, i.e. the variables defining the dry model state in grid point space. The standard deviation of the Gaussian for a set of singular vectors is set to

$$\beta = \gamma/\bar{\kappa}, \quad (2.2)$$

where $\bar{\kappa}$ denotes the average analysis error variance norm for the set of singular vectors. The parameter γ is determined empirically to yield adequate ensemble dispersion. Presently, the EPS uses $\gamma = 0.0125$ for the extra-tropical singular vectors and $\gamma = 0.01875$ for the singular vectors targeted on tropical cyclones. The reforecasts use $\gamma = 0.010$ and $\gamma = 0.015$, respectively, as the ERA-Interim analysis error estimate is different from that of the real-time analysis.

The scaling factor γ for the singular vector perturbation amplitude is determined by SHELL-variables: EPSGAMMA, EPSGAMMA_TC2INI. Table 2.2 describes how γ is set for each set of singular vectors depending on the two variables.

The coefficients for each set of singular vectors are generated by a call to program `gen_sv_coeff`. For each singular vector j and each odd member k random coefficients α_{jk} are drawn from the Gaussian distribution with mean zero and standard deviation β . The coefficients for the even members are set to minus the coefficients of the preceding member, i.e.

$$\alpha_{j\ k+1} = -\alpha_{j\ k}, \quad \text{for } k = 1, 3, 5, \dots \quad (2.3)$$

Table 2.2 *SHELL*-variables determining the scaling parameter γ which controls the standard deviation of the sampled Gaussian distribution according to Eq. (2.2).

γ	singular vector sets	singular vector files
EPSGAMMA	initial extra-tropical SVs	svifs, svifss
EPSGAMMA * EPSGAMMA_TC2INI	initial SVs targeted on TCs	svifs1, . . . , svifs6

A Gaussian distribution has no compact support, i.e. the coefficients can become arbitrarily large although this is unlikely. In order to avoid numerical instabilities, due to excessively large perturbation amplitudes the coefficients are limited to the range $[-3\beta, +3\beta]$.

The singular vectors are linearly combined using program `sv_lin_combi`. The linear combination is performed in spectral space using the truncation of the singular vector computation. For each set of singular vectors, the singular vectors and the previously generated coefficients are read. The perturbations are written to a single file `pert.grib`. The perturbations are then transferred to the FDB as type=`icp` (initial condition perturbation).

The final step in generating the perturbed initial conditions consists of adding the singular vector component of the initial perturbation to the analyses that have been perturbed with the EDA-based perturbations. The perturbations based on singular vectors are added in tasks `pert_ic/NN/pertinic`. Here NN refers to the EPS member number. The addition is performed with program `add_pert`.

2.4 EDA COMPONENT OF THE INITIAL PERTURBATIONS

The EDA perturbations are based on differences between a perturbed EDA member and the EDA mean. The perturbations are computed from 6-hour forecasts as a proxy for a sample of analysis errors because the EDA is scheduled in the operational system to run in delayed mode in order to guarantee sufficient computational resources for the EDA (Isaksen *et al.*, 2010).

Presently, the EDA contains $N_{\text{EDA}} = 25$ perturbed analyses. This provides $2N_{\text{EDA}} = 50$ different perturbations \mathbf{x}'_k

$$\mathbf{x}'_k = (-1)^{k-1} (\mathbf{x}_{\lfloor (k+1)/2 \rfloor} - \bar{\mathbf{x}})$$

with $k = 1, \dots, 50$ and $\mathbf{x}_j, \bar{\mathbf{x}}$ denoting the j -th perturbed member and the mean of all perturbed members of the EDA, respectively. The perturbation of an odd member $2j - 1$ is the negative of the perturbation of the subsequent even member $2j$:

$$\mathbf{x}'_{2j-1} = -\mathbf{x}'_{2j}$$

For instance, EDA member 1 is used to perturb EPS members 1 and 2 while EDA member 25 is used to perturb EPS members 49 and 50.

Now, the tasks are described that generate the EDA perturbations. The tasks are grouped in three families. In the first family, the high-resolution analysis that defines the centre of the EPS initial condition distribution is retrieved at full resolution (T1279L137) from MARS and interpolated to the EPS resolution (Tasks `inigroup/an/getini` and `inidata`). In the second family (`inigroup/get_eda`), the EDA data is retrieved from MARS and interpolated to the resolution of the EPS (T639L91, Tasks `get_eda/NNN/getini` and `inidata`). In addition, the mean of the EDA forecasts and the EDA perturbations are computed in task `get_eda/eda_mean` using program `comp_mean_pert`. In the third family (`inigroup/eda_pert`), the EDA perturbations are added to the centre analysis. Tasks `eda_pert/NNN/trans_an` generate the perturbed analyses $\mathbf{x}_c + \mathbf{x}'_{\text{NNN}}$. Here \mathbf{x}_c refers to the centre-analysis. Program `add_pert` sums unperturbed analysis and EDA perturbations. The perturbed surface fields are then constrained via the program `adjust_pertsfc` to physical values and to be consistent with the HTESSEL IFS surface parametrisation scheme (see Part IV, Chapter 7).

The generation of the perturbed analyses can proceed even if some EDA members are not available. Tasks `get_eda/get/NNN/edamemberok` and `get_eda/get/check_edamemberok` generate the file `file_eda_list_ok`, which contains a list with the numbers of the available EDA members. If one or more perturbed EDA members are unavailable, N_{EDA} is reduced accordingly.

Chapter 3

Computational details: non-linear integrations

Table of contents

- 3.1 Introduction**
- 3.2 Family legA**
 - 3.2.1 Control forecast family legA/fc/cf
 - 3.2.2 Perturbed forecast family legA/fc/pf
 - 3.2.3 Calibration/validation forecast family legA/fc/cv
- 3.3 Family legB**
 - 3.3.1 Control forecast family legB/fc/cf
 - 3.3.2 Perturbed forecast family legB/fc/pf
- 3.4 Family legC**
 - 3.4.1 Control forecast family legC/fc/cf
 - 3.4.2 Perturbed forecast family legC/fc/pf

3.1 INTRODUCTION

At the time of writing this documentation for model cycle Cy40r1 (April 2014), the model resolution is T639L91 between forecast day 0 and 10, and T319L91 between forecast day 10 and 15 (or 32 on Mondays and Thursdays at 00 UTC). The T639 and T319 resolutions use time steps of 20 min and 45 min, respectively. The merged VAREPS/monthly system (hereafter simply called the Ensemble Prediction System, EPS) runs twice-a-day, at 00 and 12 UTC. On Mondays and Thursdays the 00 UTC forecasts are extended to day 32.

This merged EPS/monthly configuration has been running in operations since 11 March 2008 (see [Buizza *et al.* \(2007\)](#) and [Vitart *et al.* \(2008\)](#) for a description of the operational configuration, and [Buizza *et al.* \(2003\)](#) for a description of the benefits of increasing the spatial resolution for the EPS). The EPS horizontal resolution was last changed on 26 January 2010.

Since the introduction of the merged medium-range/monthly EPS on 11 March 2008, each of the 51 15-day non-linear integrations have been performed in tasks grouped in three families, legA, legB and legC. This facilitates the production of the ensemble forecast products in different batches: the first one which includes only products up to forecast day 10, the second one which includes also the 15-day products, and the third one (on Mondays and Thursdays) which includes also the monthly products. A detailed description of the VAREPS technical implementation, including a detailed discussion on how to retrieve forecast data across the resolution change, can be found in [Buizza *et al.* \(2006\)](#).

An important addition to Ensemble in cycle Cy40r1 has been the introduction of ocean-atmosphere coupling from day 0, instead of from day 10 as in the previous cycles. The ocean model is NEMO (see [Madedec, 2008](#)) version 3.4.1 in the ORCA1 configuration (1 degree horizontal resolution) and 42 levels in the vertical. It is coupled to IFS using a single executable framework ([Mogensen *et al.*, 2012b](#)). The model timestep of NEMO is 1 hour with a coupling timestep to IFS every 3 hours. The ocean initial conditions come from the near real-time component of the operational ocean analysis (NEMOVAR, see [Mogensen *et al.*, 2012a](#)) or from the Ocean ReAnalysis System 4 (ORAS4, see [Balmaseda *et al.*, 2013](#)) for re-forecasts. The atmosphere is forced by the sea-surface temperatures (SSTs) from the ocean model plus an additional term which is the difference between the high resolution SST from OSTIA and the

low resolution SST from the ocean analysis at step 0 during the first 5 days of the forecast. After day 5, this additional term is relaxed to reach 0 by day 9. This additional term ensures that the SSTs during the first days of the forecasts still contains the high resolution information of the OSTIA SST analysis.

3.2 FAMILY LEGA

Family 'legA' includes the generation of the initial perturbations (see [Section 2.3](#)) and the first 10-day forecast integration. These forecasts are all coupled to the NEMO ocean model. Under sub-family 'legA/fc', three types of forecasts are currently run: the ensemble control forecast, the ensemble perturbed forecasts and the two extra calibration/validation forecasts.

3.2.1 Control forecast family legA/fc/cf

The control forecast family includes four tasks: task legA/fc/cf/getinileg copies the initial conditions in the forecast working directory, while task legA/fc/cf/modeleps performs the day 0 to day 10 integration at resolution T639 and with 91 vertical levels. The family oceanini prepares the ocean initial conditions. Tasks modeleps_tidy and spt2fdb deal with the restart data for the ocean model and the SPPT scheme, respectively. Control forecast data are archived in MARS (the ECMWF Meteorological Archival and Retrieval System) as type 'cf' of stream 'enfo'.

3.2.2 Perturbed forecast family legA/fc/pf

This family includes 50 sub-families, named legA/fc/pf/001, legA/fc/pf/002, Each sub-family includes the same two tasks as the control forecast: task legA/fc/cf/getinileg copies the perturbed initial conditions in the forecast working directory, while task legA/fc/cf/modeleps performs the day 0 to day 10 integration at resolution T639 and with 91 vertical levels. Perturbed forecast data are archived in MARS as type 'pf' of stream 'enfo'.

3.2.3 Calibration/validation forecast family legA/fc/cv

This family includes two sub-families, named legA/fc/cv/control_1 and legA/fc/cv/control_2: these two sub-families include tasks to run two extra reference forecasts up to 15 days. These two extra forecasts are performed with constant resolution for calibration and validation purposes (note that these two forecasts are not used to generate any ensemble product). Forecast data from these two calibration/validation forecasts are archived in MARS as type 'cv' of stream 'enfo'. Each sub-family includes the same two tasks as the control forecast: task legA/fc/cf/getinileg copies the perturbed initial conditions in the forecast working directory, while task legA/fc/cf/modeleps performs the 15-day integration, at resolution T639 and with 91 vertical levels for control_1, and at resolution T319 and with 91 vertical levels for control_2.

3.3 FAMILY LEGB

Under sub-family 'legB/fc', the ensemble control and the perturbed forecasts are extended from forecast day 10 to 15 with a T319 resolution and with 91 vertical levels. These forecasts are fully coupled to the ocean model.

Note that technically each legB forecast starts from the legA 9-day forecast: when the legB forecast has passed forecast day-1, which coincides with the legA forecast day-10, its data are archived in MARS as part of the 'enfo' stream. In other words, EPS forecast data in MARS stream 'enfo' include legA T639L91 forecasts from day 0 to day 10 and legB/legC T319L91 forecasts for lead times exceeding 10 days.

The accumulated fields in legB are re-initialized at day 10 from interpolated fields obtained from the legA integration. A conserving interpolation scheme available in IFS is used. The accumulated fields that have been interpolated to lower resolution at day 10 are archived in MARS in stream 'efov' and as the corresponding 'variable resolution' parameter. For instance, total precipitation (param=128) is archived as 'Total precipitation (variable resolution)'. It is identified as param=128.230. These variable resolution

fields at the truncation time are required to correctly compute accumulations that start at or before 10 d and end after 10 d.

3.3.1 Control forecast family legB/fc/cf

The control forecast family includes five tasks: task legB/fc/cf/getvarepsdata retrieves all 9-day forecast fields from legA required to start the legB forecast. Task legB/fc/cf/intHtoL interpolates the initial conditions from the legA to the legB resolution. Task legB/fc/cf/modeleps performs the day 9 to day 15 integration at resolution T319 and with 91 vertical levels. Tasks modeleps_tidy and sppt2fdb deal with the restart data for the ocean model and the SPPT scheme, respectively.

3.3.2 Perturbed forecast family legB/fc/pf

This family includes 50 sub-families named legB/fc/pf/001, legB/fc/pf/002, . . . Each sub-family handles the integration of a perturbed forecast of the ensemble. The subfamilies include the same tasks as the control forecast.

3.4 FAMILY LEGC

Under sub-family 'legC/fc', twice-a-week (on Mondays and Thursdays) the ensemble control and the perturbed forecasts are extended from forecast day 15 to day 32 with a T319 resolution and with 91 vertical levels, coupled to the ocean model as in LegB.

3.4.1 Control forecast family legC/fc/cf

The control forecast family includes three tasks: task legC/fc/cf/getvarepsdata retrieves all 15-day forecast fields from legB required to start the legC forecast. Task legC/fc/cf/intHtoL interpolates the initial conditions from the legB to the legC resolution. Task legC/fc/cf/modeleps performs the day 15 to day 32 integration at resolution T319 and with 91 vertical levels

3.4.2 Perturbed forecast family legC/fc/pf

This family includes 50 sub-families named legC/fc/pf/001, legC/fc/pf/002, . . . Their structures mirror that of the control forecast family legC/fc/cf.

References

- Balmaseda, M., Mogensen, K. and Weaver, A. T. (2013). Evaluation of the ECMWF ocean reanalysis ORAS4. *Q. J. R. Meteorol. Soc.*, **139**, 1132–1161.
- Barkmeijer, J., Buizza, R., Palmer, T. N., Puri, K. and Mahfouf, J.-F. (2001). Tropical singular vectors computed with linearized diabatic physics. *Q. J. R. Meteorol. Soc.*, **127**, 685–708.
- Barkmeijer, J., van Gijzen, M. and Bouttier, F. (1998). Singular vectors and estimates of the analysis error covariance metric. *Q. J. R. Meteorol. Soc.*, **124**, 1695–1713.
- Berner, J., Shutts, G. J., Leutbecher, M. and Palmer, T. N. (2009). A spectral stochastic kinetic energy backscatter scheme and its impact on flow-dependent predictability in the ECMWF ensemble prediction system. *J. Atmos. Sci.*, **66**, 603–626.
- Bourke, W., Buizza, R. and Naughton, M. (2005). Performance of the ECMWF and the BoM ensemble systems in the southern hemisphere. *Mon. Wea. Rev.*, **132**, 2338–2357.
- Buizza, R. (1994). Localization of optimal perturbations using a projection operator. *Q. J. R. Meteorol. Soc.*, **120**, 1647–1682.
- Buizza, R., Barkmeijer, J., Palmer, T. N. and Richardson, D. S. (1999a). Current status and future developments of the ECMWF ensemble prediction system. *Meteorol. Appl.*, **6**, 1–14.
- Buizza, R., Bidlot, J. R., Wedi, N., Fuentes, M., Hamrud, M., Holt, G., Palmer, T. N. and Vitart, F. (2006). The new ECMWF variable resolution ensemble prediction system (VAREPS): methodology and validation. *ECMWF Tech. Memo. No. 499*.
- Buizza, R., Bidlot, J. R., Wedi, N., Fuentes, M., Hamrud, M., Holt, G. and Vitart, F. (2007). The new ECMWF VAREPS (variable resolution ensemble prediction system). *Q. J. R. Meteorol. Soc.*, **133**, 681–695.
- Buizza, R., Houtekamer, P. L., Toth, Z., Pellerin, G., Wei, M. and Zhu, Y. (2005). A comparison of the ECMWF, MSC and NCEP global ensemble prediction systems. *Mon. Wea. Rev.*, **133**, 1076–1097.
- Buizza, R., Leutbecher, M. and Isaksen, L. (2008). Potential use of an ensemble of analyses in the ECMWF ensemble prediction system. *Q. J. R. Meteorol. Soc.*, **134**, 2051–2066.
- Buizza, R., Miller, M. and Palmer, T. N. (1999b). Stochastic representation of model uncertainties in the ECMWF ensemble prediction system. *Q. J. R. Meteorol. Soc.*, **125**, 2887–2908.
- Buizza, R. and Palmer, T. N. (1995). The singular vector structure of the atmosphere global circulation. *J. Atmos. Sci.*, **52**, 1434–1456.
- Buizza, R., Petroliaigis, T., Palmer, T. N., Barkmeijer, J., Hamrud, M., Hollingsworth, A., Simmons, A. and Wedi, N. (1998). Impact of model resolution and ensemble size on the performance of an ensemble prediction system. *Q. J. R. Meteorol. Soc.*, **124**, 1935–1960.
- Buizza, R., Richardson, D. S. and Palmer, T. N. (2003). Benefits of increased resolution in the ECMWF ensemble system and comparison with poor-man’s ensembles. *Q. J. R. Meteorol. Soc.*, **129**, 1269–1288.

- Dee, D. P., Uppala, S. M., Simmons, A. J., Berrisford, P., Poli, Kobayashi, S., Andrae, U., Balmaseda, M. A., Balsamo, G., Bauer, P., Bechtold, P., Beljaars, A. C. M., van de Berg, L., Bidlot, J., Bormann, N., Delsol, C., Dragani, R., Fuentes, M., Geer, A. J., Haimberger, L., Healy, S. B., Hersbach, H., Hólm, E. V., Isaksen, L., Kållberg, P., Köhler, M., Matricardi, M., McNally, A. P., Monge-Sanz, B. M., Morcrette, J.-J., Park, B.-K., Peubey, C., de Rosnay, P., Tavolato, C., Thépaut, J.-N. and Vitart, F. (2011). The ERA-Interim reanalysis: configuration and performance of the data assimilation system. *Q. J. R. Meteorol. Soc.*, **137**, 553–597.
- Ehrendorfer, M. and Beck, A. (2003). Singular vector-based multivariate normal sampling in ensemble prediction. *ECMWF Tech. Memo. No. 416*.
- Ehrendorfer, M. and Tribbia, J. J. (1997). Optimal prediction of forecast error covariances through singular vectors. *J. Atmos. Sci.*, **53**, 286–313.
- Gelaro, R., Buizza, R., Palmer, T. N. and Klinker, E. (1998). Sensitivity analysis of forecast errors and the construction of optimal perturbations using singular vectors. *J. Atmos. Sci.*, **55**, 1012–1037.
- Golub, G. H. and van Loan, C. F. (1983). *Matrix Computation*. North Oxford Academic Publ. Co. Ltd.
- Hagedorn, R. (2008). Using the ECMWF reforecast dataset to calibrate EPS forecasts. *ECMWF Newsletter No. 117*, pp. 8–13.
- Hagedorn, R., Buizza, R., Hamill, M. T., Leutbecher, M. and Palmer, T. N. (2012). Comparing TIGGE multi-model forecasts with re-forecast calibrated ECMWF ensemble forecasts. *Q. J. R. Meteorol. Soc.*, **138**, 1814–1827.
- Harrison, M. S., Palmer, T. N., Richardson, D. and Buizza, R. (1999). Analysis and model dependencies in medium-range ensembles: Two transplant case studies. *Q. J. R. Meteorol. Soc.*, **126**, 2487–2515.
- Hoskins, B. J., Buizza, R. and Badger, J. (2000). The nature of singular vector growth and structure. *Q. J. R. Meteorol. Soc.*, **126**, 1565–1580.
- Houtekamer, P. L., Lefaiivre, L. and Derome, J. (1996a). The RPN ensemble prediction system. In *Proc. ECMWF Seminar on Predictability*, Vol. II, pp. 121–146, 4–8 September 1995, Reading, UK.
- Houtekamer, P. L., Lefaiivre, L., Derome, J., Ritchie, H. and Mitchell, H. L. (1996b). A system simulation approach to ensemble prediction. *Mon. Wea. Rev.*, **124**, 1225–1242.
- Isaksen, L. M. B., Buizza, M. F., Haseler, J., Leutbecher, M. and Raynaud, L. (2010). Ensemble of data assimilations at ECMWF. *Technical Report 636*, ECMWF, Reading, UK.
- Lawrence, A. R., Leutbecher, M. and Palmer, T. N. (2009). The characteristics of hessian singular vectors using an advanced data assimilation scheme. *Q. J. R. Meteorol. Soc.*, **135**, 1117–1132.
- Leutbecher, M. (2005). On ensemble prediction using singular vectors started from forecasts. *Mon. Wea. Rev.*, **133**, 3038–3046.
- Leutbecher, M. (2007). On the representation of initial uncertainties with multiple sets of singular vectors optimised for different criteria. *Q. J. R. Meteorol. Soc.*, **133**, 2045–2056.
- Leutbecher, M. and Lang, S. T. K. (2013). On the reliability of ensemble variance in subspaces defined by singular vectors. *Q. J. R. Meteorol. Soc.*, doi:10.1002/qj.2229.
- Leutbecher, M. and Palmer, T. N. (2008). Ensemble forecasting. *J. Comp. Phys.*, **227**, 3515–3539.
- Madec, G. (2008). *NEMO reference manual, ocean dynamics component: NEMO_OPA. Preliminary version*. Institut Pierre-Simon Laplace (IPSL), France, note du pole de modelisation 27.
- Mogensen, K., Balmaseda, M. and Weaver, A. (2012a). The NEMOVAR ocean data assimilation system as implemented in the ECMWF ocean analysis for System 4. *Tech. Memo. 668*, ECMWF, available at www.ecmwf.int/publications/library/do/references/show?id=90389.

- Mogensen, K., Keeley, S. and Towers, P. (2012b). Coupling of the NEMO and IFS models in a single executable. *Tech. Memo. 673*, ECMWF, available at www.ecmwf.int/publications/library/do/references/show?id=90460.
- Molteni, F., Buizza, R., Palmer, T. N. and Petroliagis, T. (1996). The ECMWF Ensemble Prediction System: methodology and validation. *Q. J. R. Meteorol. Soc.*, **122**, 73–119.
- Mureau, F., Molteni, F. and Palmer, T. N. (1993). Ensemble prediction using dynamically conditioned perturbations. *Q. J. R. Meteorol. Soc.*, **199**, 299–298.
- Palmer, T. N., Buizza, R., Doblas-Reyes, F., Jung, T., Leutbecher, M., Shutts, G. J., Steinheimer, M. and Weisheimer, A. (2009). Stochastic parameterization and model uncertainty. *ECMWF Tech. Memo. No. 598*, pp. 1–42.
- Palmer, T. N., Gelaro, R., Barkmeijer, J. and Buizza, R. (1998). Singular vectors, metrics and adaptive observations. *J. Atmos. Sci.*, **55**, 633–653.
- Palmer, T. N., Molteni, F., Mureau, R., Buizza, R., Chapelet, P. and Tribbia, J. (1993). Ensemble Prediction. In *Proc. of the ECMWF Seminar on Validation of Models over Europe: Vol. 1*, pp. 21–66, 7–11 September 1992, Reading, UK.
- Park, Y.-Y., Buizza, R. and Leutbecher, M. (2008). TIGGE: preliminary results on comparing and combining ensembles. *ECMWF Tech. Memo. No. 548*.
- Puri, K., Barkmeijer, J. and Palmer, T. N. (2001). Ensemble prediction of tropical cyclones using targeted diabatic singular vectors. *Q. J. R. Meteorol. Soc.*, **127**, 709–734.
- Rabier, F., Klinker, E., Courtier, P. and Hollingsworth, A. (1996). Sensitivity of forecast errors to initial conditions. *Q. J. R. Meteorol. Soc.*, **122**, 121–150.
- Sleijpen, G. L. G. and van der Vorst, H. A. (1996). A Jacobi–Davidson iteration method for linear eigenvalue problems. *SIAM J. Matrix Anal. Appl.*, **17**, 401–425.
- Toth, Z. and Kalnay, E. (1993). Ensemble forecasting at NMC: the generation of initial perturbations. *Bull. Am. Meteorol. Soc.*, **74**, 2317–2330.
- Toth, Z. and Kalnay, E. (1997). Ensemble forecasting at NCEP and the breeding method. *Mon. Wea. Rev.*, **125**, 3297–3319.
- Trevisan, A., Pancotti, F. and Molteni, F. (2001). Ensemble prediction in a model with flow regimes. *Q. J. R. Meteorol. Soc.*, **127**, 343–358.
- Vitart, F. (2004). Monthly forecasting at ecmwf. *Mon. Wea. Rev.*, **132**, 2671–2779.
- Vitart, F., Buizza, R., Balmaseda, M. A., Balsamo, G., Bidlot, J., Bonet, A., Fuentes, M., Hofstadler, A., Molteni, F. and Palmer, T. N. (2008). The new VAREPS-monthly forecasting system: a first step towards seamless prediction. *Q. J. R. Meteorol. Soc.*, **134**, 1789–1799.
- Zsoter, E., Buizza, R. and Richardson, D. (2009). Jumpiness of the ecmwf and uk met office eps control and ensemble-mean forecasts. *Mon. Wea. Rev.*, **137**, 3823–3836.