École Doctorale des Sciences de l'Environnement d'Île-de-France

Année Universitaire 2012-2013

Modélisation Numérique de l'Écoulement Atmosphérique et Assimilation de Données

Olivier Talagrand

Cours 6

6 Mai 2013

Best Linear Unbiased Estimate

State vector x, belonging to *state space* $\mathcal{S}$ ($\dim \mathcal{S} = n$), to be estimated.

Available data in the form of

- A '*background*' estimate (*e. g.* forecast from the past), belonging to *state space*, with dimension n

  $x^b = x + \zeta^b$

- An additional set of data (*e. g.* observations), belonging to *observation space*, with dimension p

  $y = Hx + \varepsilon$

  H is known linear *observation operator.*

Assume probability distribution is known for the couple (ζ^b, ε).

Assume $E(\zeta^b) = 0$, $E(\varepsilon) = 0$, $E(\zeta^b \varepsilon^{\mathrm{T}}) = 0$ (not restrictive)

Set $E(\zeta^b \zeta^{b\mathrm{T}}) = P^b$ (also often denoted B), $E(\varepsilon\varepsilon^{\mathrm{T}}) = R$

Best Linear Unbiased Estimate (continuation 2)

Apply formulæ for Optimal Interpolation

$$\boldsymbol{x}^a = \boldsymbol{x}^b + P^b H^{\mathrm{T}} [H P^b H^{\mathrm{T}} + R]^{-1} (\boldsymbol{y} - H\boldsymbol{x}^b)$$
$$P^a = P^b - P^b H^{\mathrm{T}} [H P^b H^{\mathrm{T}} + R]^{-1} H P^b$$

x^a is the *Best Linear Unbiased Estimate* (*BLUE*) of x from x^b and y.

Equivalent set of formulæ

$$\boldsymbol{x}^a = \boldsymbol{x}^b + P^a H^{\mathrm{T}} R^{-1} (\boldsymbol{y} - H\boldsymbol{x}^b)$$
$$[P^a]^{-1} = [P^b]^{-1} + H^{\mathrm{T}} R^{-1} H$$

Matrix $K = P^b H^{\mathrm{T}} [H P^b H^{\mathrm{T}} + R]^{-1} = P^a H^{\mathrm{T}} R^{-1}$ is *gain matrix*.

If probability distributions are *globally* gaussian, *BLUE* achieves bayesian estimation, in the sense that $P(\boldsymbol{x} \mid \boldsymbol{x}^b, \boldsymbol{y}) = \mathcal{N}[\boldsymbol{x}^a, \boldsymbol{P}^a]$.

Best Linear Unbiased Estimate (continuation 3)

H can be any linear operator

Example : (scalar) satellite observation

$\boldsymbol{x} = (x_1, \ldots, x_n)^T$ temperature profile

Observation $y = \Sigma_i h_i x_i + \varepsilon = \boldsymbol{H}\boldsymbol{x} + \varepsilon$, $\boldsymbol{H} = (h_1, \ldots, h_n)$, $E(\varepsilon^2) = r$

Background $\boldsymbol{x}^b = (x_1^b, \ldots, x_n^b)^T$, error covariance matrix $P^b = (p_{ij}^b)$

$$\boldsymbol{x}^a = \boldsymbol{x}^b + P^b H^T [HP^b H^T + R]^{-1} (\boldsymbol{y} - H\boldsymbol{x}^b)$$

$[HP^b H^T + R]^{-1} (\boldsymbol{y} - H\boldsymbol{x}^b) = (y - \Sigma_\iota h_\iota x_\iota^b) / (\Sigma_{ij} h_i h_j p_{ij}^b + r) \equiv \mu$ scalar !

- $P^b = p^b \boldsymbol{I}_n$ $x_i^a = x_i^b + p^b h_i \mu$
- $P^b = \mathrm{diag}(p_{ii}^b)$ $x_i^a = x_i^b + p_{ii}^b h_i \mu$
  - General case $x_i^a = x_i^b + \Sigma_j p_{ij}^b h_j \mu$

Each level i is corrected, not only because of its own contribution to the observation, but because of the contribution of the other levels to which its background error is correlated.

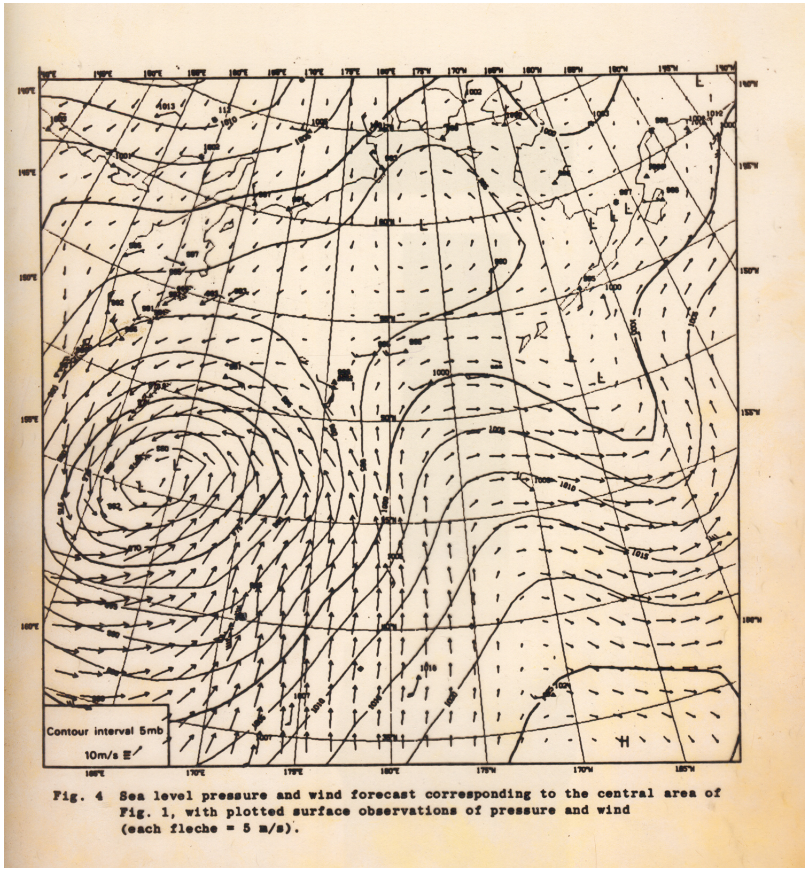

Fig. 4 Sea level pressure and wind forecast corresponding to the central area of Fig. 1, with plotted surface observations of pressure and wind (each fleche = 5 m/s).

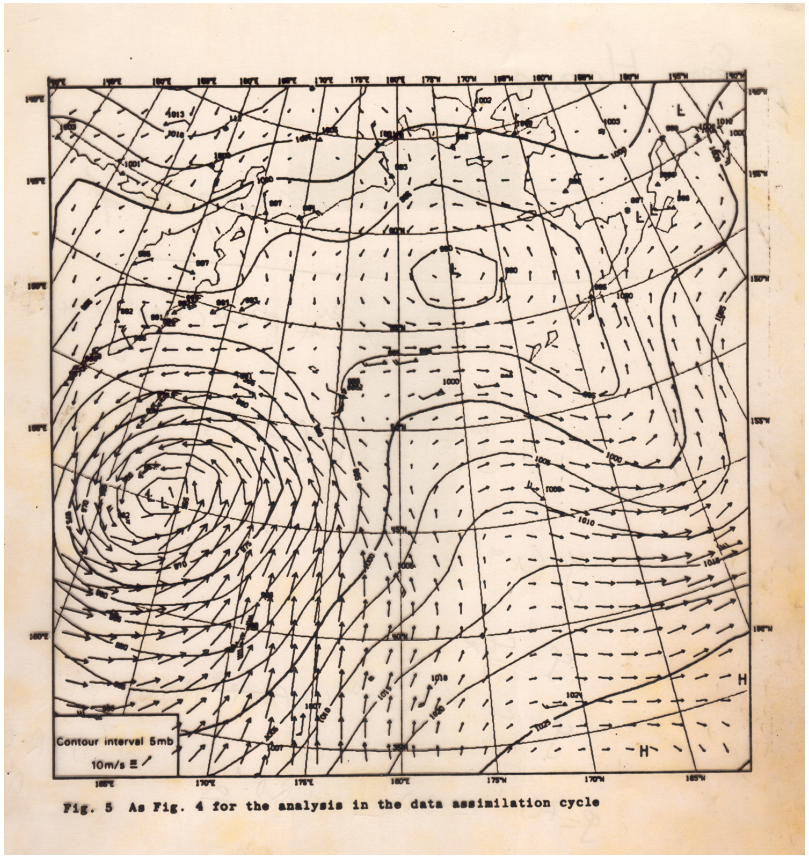

Fig. 5 As Fig. 4 for the analysis in the data assimilation cycle

After A. Lorenc

Best Linear Unbiased Estimate (continuation 4)

Variational form of the *BLUE*

BLUE x^a minimizes following scalar *objective function*, defined on state space

$\xi \in \mathcal{S} \rightarrow$

- $\mathcal{J}(\xi) = (1/2)\,(x^b - \xi)^{\mathrm{T}}\,[P^b]^{-1}\,(x^b - \xi) + (1/2)\,(y - H\xi)^{\mathrm{T}}\,R^{-1}\,(y - H\xi)$

$= \mathcal{J}_b + \mathcal{J}_o$

'*3D-Var*'

Can easily, and heuristically, be extended to the case of a nonlinear observation operator H.

Used operationally in USA, Australia, China, …

Question. How to introduce temporal dimension in estimation process ?

- Logic of Optimal Interpolation can be extended to time dimension.

- But we know much more than just temporal correlations. We know explicit dynamics.

Real (unknown) state vector at time k (in format of assimilating model) x_k. Belongs to state space $\mathcal{S}$ ($\dim \mathcal{S} = n$)

Evolution equation

$$x_{k+1} = M_k(x_k) + \eta_k$$

M_k is (known) model, η_k is (unknown) model error

Sequential Assimilation

- Assimilating model is integrated over period of time over which observations are available. Whenever model time reaches an instant at which observations are available, state predicted by the model is updated with new observations.

Variational Assimilation

- Assimilating model is globally adjusted to observations distributed over observation period. Achieved by minimization of an appropriate scalar *objective function* measuring misfit between data and sequence of model states to be estimated.

Sequential Assimilation

- Observation vector at time k

$y_k = H_k x_k + \varepsilon_k$ $\qquad k = 0, \ldots, K$

$E(\varepsilon_k) = 0 \quad ; \quad E(\varepsilon_k \varepsilon_j^{\mathrm{T}}) = R_k \, \delta_{kj}$

H_k linear

- Evolution equation

$x_{k+1} = M_k x_k + \eta_k$ $\qquad k = 0, \ldots, K\text{-}1$

$E(\eta_k) = 0 \quad ; \quad E(\eta_k \eta_j^{\mathrm{T}}) = Q_k \, \delta_{kj}$

M_k linear

- $E(\eta_k \varepsilon_j^{\mathrm{T}}) = 0$ (errors uncorrelated in time)

At time k, background x^b_k and associated error covariance matrix P^b_k known

- Analysis step

$$x^a_k = x^b_k + P^b_k H_k^T [H_k P^b_k H_k^T + R_k]^{-1} (y_k - H_k x^b_k)$$
$$P^a_k = P^b_k - P^b_k H_k^T [H_k P^b_k H_k^T + R_k]^{-1} H_k P^b_k$$

- Forecast step

$$x^b_{k+1} = M_k x^a_k$$
$$P^b_{k+1} = E[(x^b_{k+1} - x_{k+1})(x^b_{k+1} - x_{k+1})^T] = E[(M_k x^a_k - M_k x_k - \eta_k)(M_k x^a_k - M_k x_k - \eta_k)^T]$$
$$= M_k E[(x^a_k - x_k)(x^a_k - x_k)^T] M_k^T - E[\eta_k (x^a_k - x_k)^T] - E[(x^a_k - x_k)\eta_k^T] + E[\eta_k \eta_k^T]$$
$$= M_k P^a_k M_k^T + Q_k$$

At time k, background x^b_k and associated error covariance matrix P^b_k known

- Analysis step

$$x^a_k = x^b_k + P^b_k H_k^{\mathrm{T}} [H_k P^b_k H_k^{\mathrm{T}} + R_k]^{-1} (y_k - H_k x^b_k)$$
$$P^a_k = P^b_k - P^b_k H_k^{\mathrm{T}} [H_k P^b_k H_k^{\mathrm{T}} + R_k]^{-1} H_k P^b_k$$

- Forecast step

$$x^b_{k+1} = M_k x^a_k$$
$$P^b_{k+1} = M_k P^a_k M_k^{\mathrm{T}} + Q_k$$

Kalman filter (KF, Kalman, 1960)

Must be started from some initial estimate (x^b_0, P^b_0)

If all operators are linear, and if errors are uncorrelated in time, Kalman filter produces at time k the *BLUE* ${x^b}_k$ (resp. ${x^a}_k$) of the real state x_k from all data prior to (resp. up to) time k, plus the associated estimation error covariance matrix ${P^b}_k$ (resp. ${P^a}_k$).

If in addition errors are gaussian, the corresponding conditional probability distributions are the respective gaussian distributions $\mathcal{N}[{x^b}_k, {P^b}_k]$ and $\mathcal{N}[{x^a}_k, {P^a}_k]$.

A didactic example (Ghil et al.)

Barotropic model

$$\begin{cases} \dfrac{\partial \varphi}{\partial t} + \operatorname{div}(\varphi \underline{U}) = 0 \\ \dfrac{\partial \underline{U}}{\partial t} + \operatorname{grad}\left(\varphi + \frac{1}{2} U^2\right) + \underline{k} \times (f + \xi)\, \underline{U} = 0 \end{cases}$$

One dimension, periodic

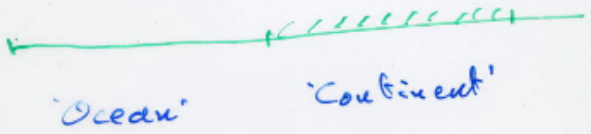

Linearized (conserves energy)

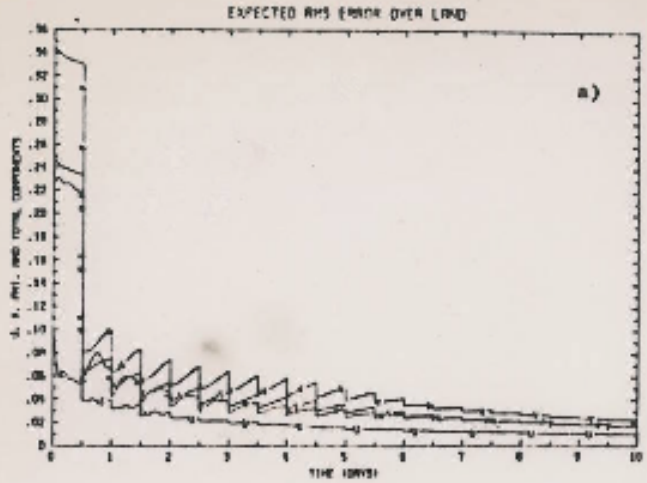

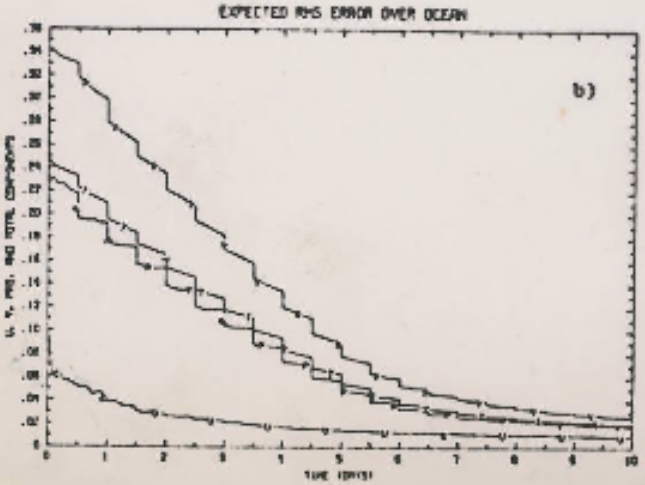

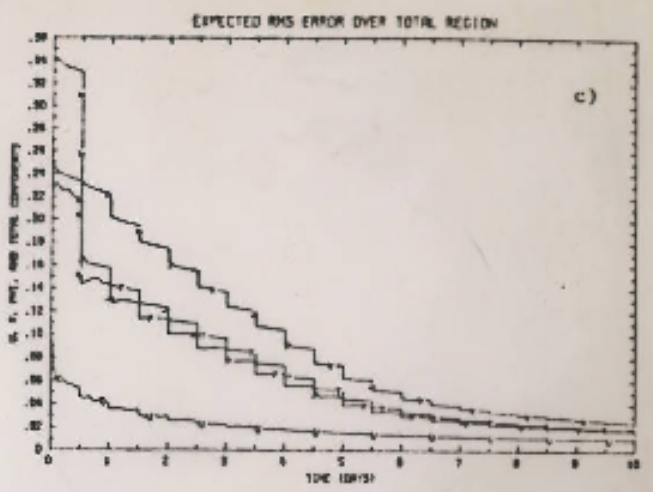

Fig. 2

The components of the total expected rms error (Erms), (trace: $P_k)^{1/2}$, in the estimation of solutions to the stochastic-dynamic system (Ψ,H), with Ψ given by (3.6) and $H = (I\ 0)$. System noise is absent, $Q = 0$. The filter used is the standard K-B filter (2.11) for the model.

a) Erms over land; b) Erms over the ocean; c) Erms over the entire L-domain

In each one of the figures, each curve represents one component of the total Erms error. The curves labelled U, V, and P represent the u component, v component and ϕ component, respectively. They are found by summing the diagonal elements of P_k which correspond to u, v, and ϕ, respectively, dividing by the number of terms in the sum, and then taking the square root. In a) the summation extends over land points only, in b) over ocean points only, and in c) over the entire L-domain. The vertical axis is scaled in such a way that 1.0 corresponds to an Erms error of v_{max} for the U and V curves, and of ϕ_0 for the P curve. The observational error level is 0.089 for the U and V curves, and 0.080 for the P curve. The curves labelled T represent the total Erms error over each region. Each T curve is a weighted average of the corresponding U, V, and P curves, with the weights chosen in such a way that the T curve measures the error in the total energy $u^2 + v^2 + \phi^2/\gamma$, conserved by the system (3.1). The observational noise level for the T curve is then 0.088. Notice the immediate error decrease over land and the gradual decrease over the ocean. The total estimation error tends to zero.

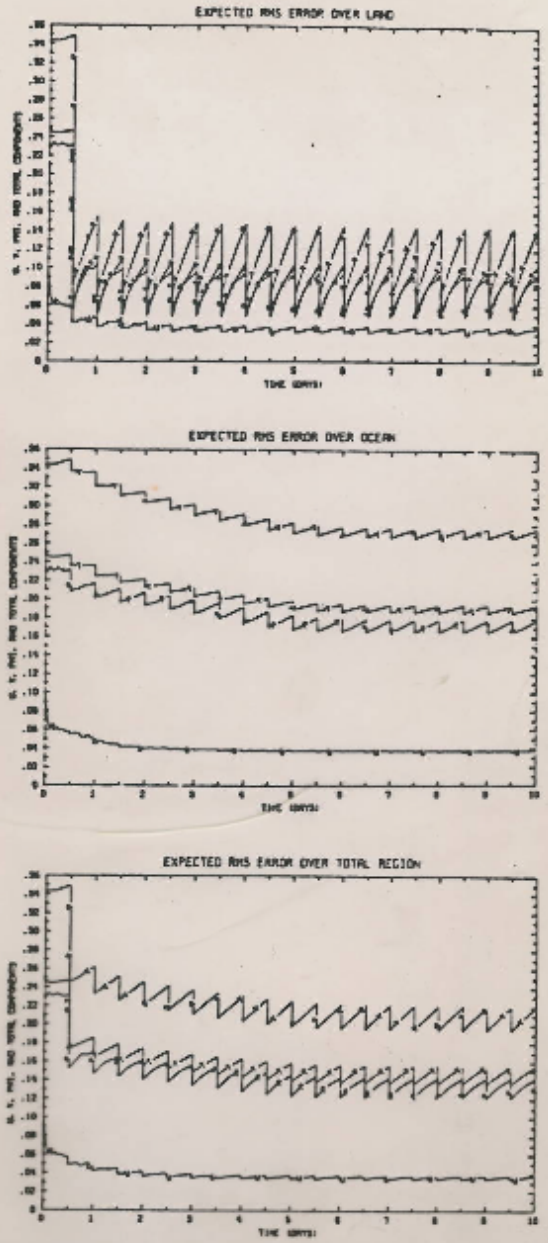

Fig. 6 This figure and the following ones show the properties of the estimated algorithm (2.11) in the presence of system noise, $Q \neq 0$. This figure gives the Erms estimation error, and is homologous to Fig. 2. Notice the sharper increase of error over land between synoptic times, and the convergence of each curve to a periodic, nonzero function.

Nonlinearities ?

Model is usually nonlinear, and observation operators (satellite observations) tend more and more to be nonlinear.

- Analysis step

$$x^a_k = x^b_k + P^b_k H_k{}'^{\mathrm{T}} [H_k' P^b_k H_k{}'^{\mathrm{T}} + R_k]^{-1} [y_k - H_k(x^b_k)]$$
$$P^a_k = P^b_k - P^b_k H_k{}'^{\mathrm{T}} [H_k' P^b_k H_k{}'^{\mathrm{T}} + R_k]^{-1} H_k' P^b_k$$

- Forecast step

$$x^b_{k+1} = M_k(x^a_k)$$
$$P^b_{k+1} = M_k' P^a_k M_k{}'^{\mathrm{T}} + Q_k$$

Extended Kalman Filter (EKF, heuristic !)

Costliest part of computation

$$P^{b}_{k+1} = M_k P^{a}_{k} M_k^{\mathrm{T}} + Q_k$$

Multiplication by M_k = one integration of the model between times k and $k+1$.
Computation of $M_k P^{a}_{k} M_k^{\mathrm{T}} \approx 2n$ integrations of the model

Need for determining the temporal evolution of the uncertainty on the state of the system is the major difficulty in assimilation of meteorological and oceanographical observations

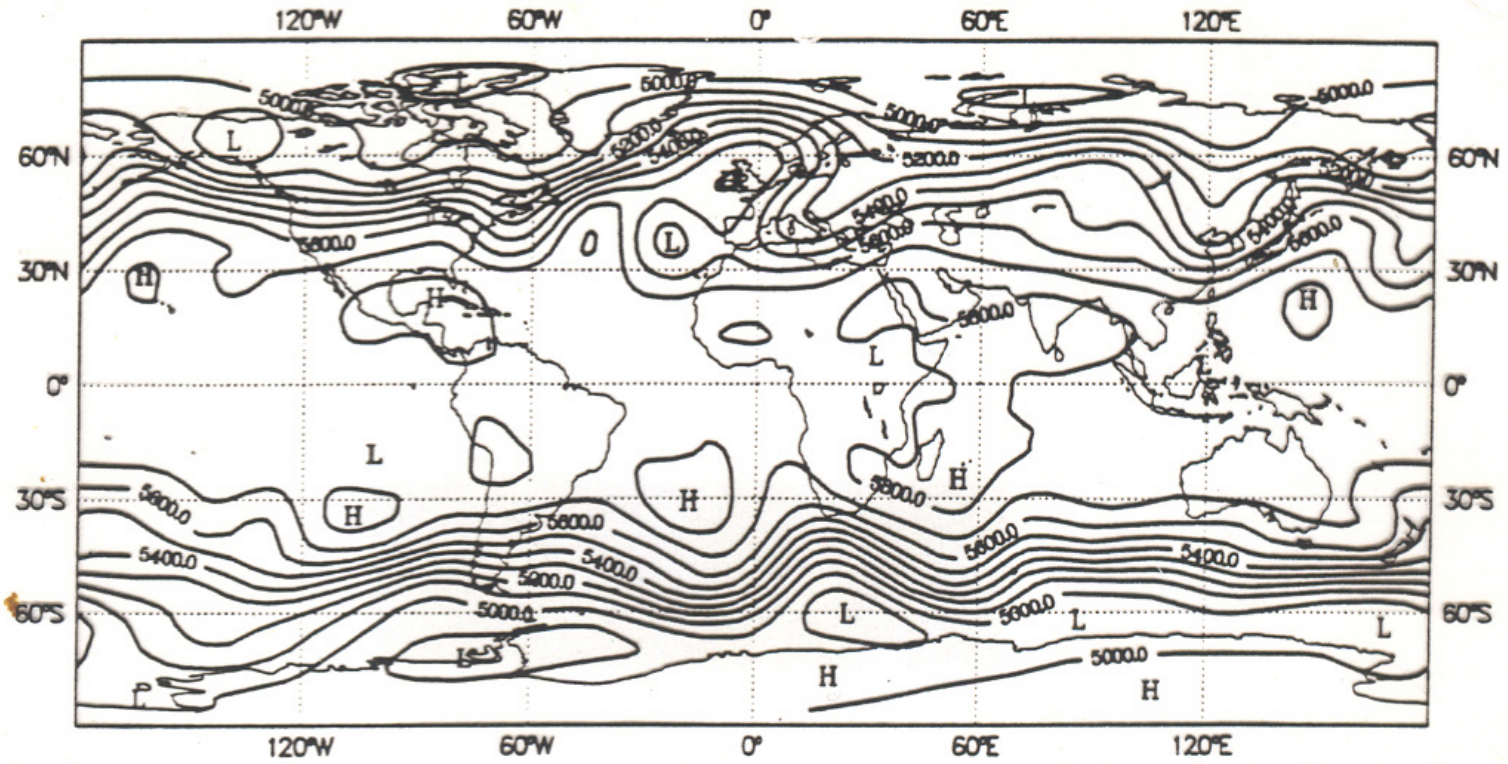

Analysis of 500-hPa geopotential for 1 December 1989, 00:00 UTC (ECMWF, spectral truncation T21, unit *m*. After F. Bouttier)

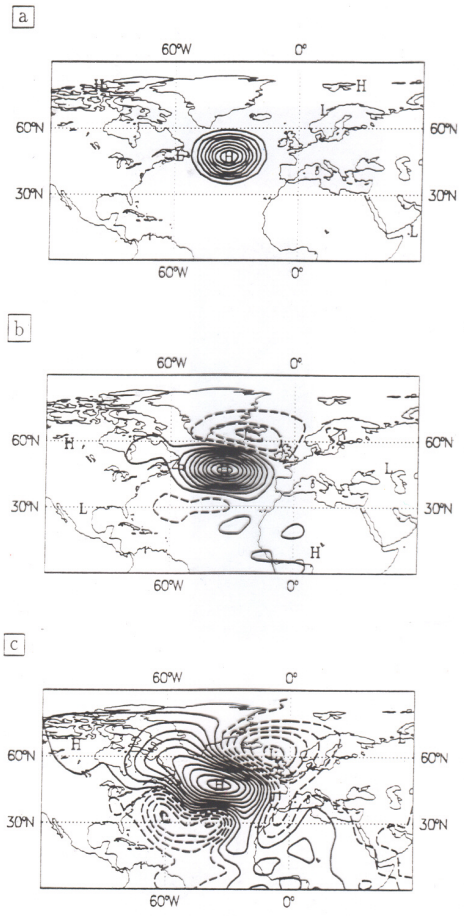

Temporal evolution of the 500-hPa geopotential autocorrelation with respect to point located at 45N, 35W. From top to bottom: initial time, 6- and 24-hour range. Contour interval 0.1. After F. Bouttier.

Two solutions :

- *Low-rank filters*

  Use low-rank covariance matrix, restricted to modes in state space on which it is known, or at least assumed, that a large part of the uncertainty is concentrated (this requires the definition of a norm on state space).

  Reduced Rank Square Root Filters (*RRSQRT*, Heemink)

  Singular Evolutive Extended Kalman Filter (*SEEK*, Pham)

  ....

Second solution :

- *Ensemble filters*

  Uncertainty is represented, not by a covariance matrix, but by an ensemble of point estimates in state space that are meant to sample the conditional probability distribution for the state of the system (dimension $L \approx O(10\text{-}100)$).

  Ensemble is evolved in time through the full model, which eliminates any need for linear hypothesis as to the temporal evolution.

  Ensemble Kalman Filter (*EnKF,* Evensen, Anderson, …)

How to update predicted ensemble with new observations ?

Predicted ensemble at time k : $\{x^b_l\}$, $\quad l = 1, \ldots, L$
Observation vector at same time : $y = Hx + \varepsilon$

- Gaussian approach

  Produce sample of probability distribution for real observed quantity Hx
  $y_l = y - \varepsilon_l$
  where ε_l is distributed according to probability distribution for observation error ε.

  Then use Kalman formula to produce sample of 'analysed' states

  $$x^a_l = x^b_l + P^b H^{\mathrm{T}} [HP^b H^{\mathrm{T}} + R]^{-1} (y_l - Hx^b_l) , \qquad l = 1, \ldots, L \qquad (2)$$

  where P^b is covariance matrix of predicted ensemble $\{x^b_l\}$.

  Remark. If P^b was exact covariance matrix of background error, (2) would achieve Bayesian estimation, in the sense that $\{x^a_l\}$ would be a sample of conditional probability distribution for x, given all data up to time k.

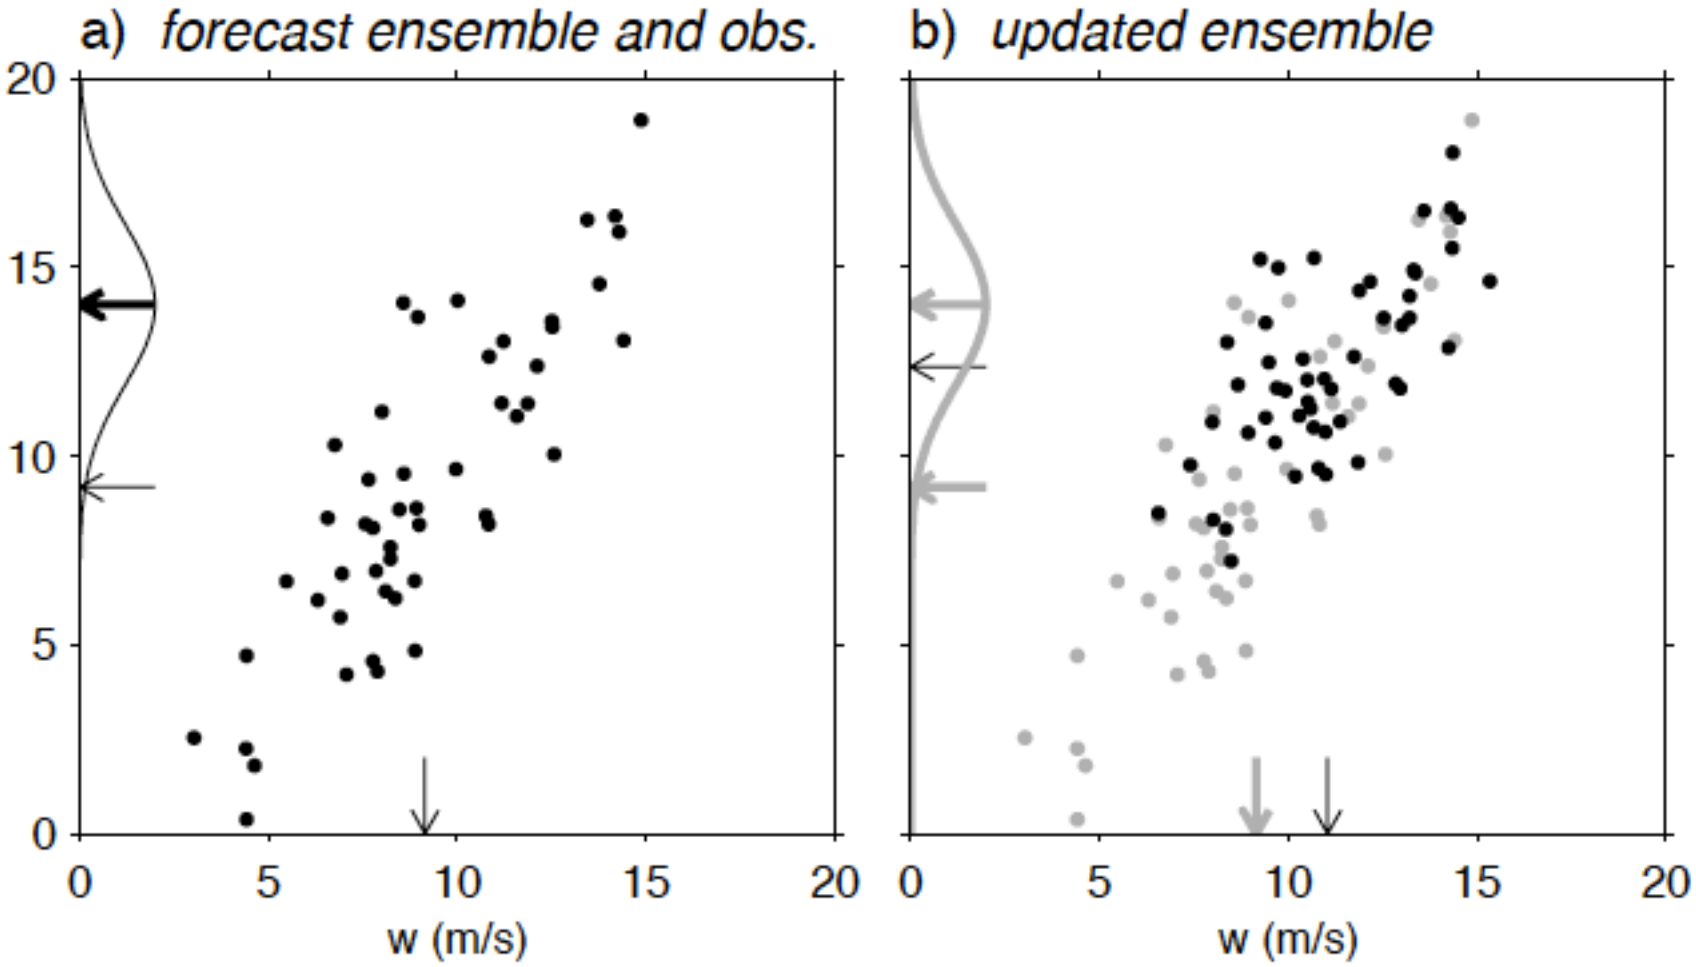
a) forecast ensemble and obs.
b) updated ensemble
20
15
10
5
0
0
5
10
15
20
w (m/s)
0
5
10
15
20
w (m/s)

C. Snyder

Month-long Performance of EnKF vs. 3Dvar with WRF

— EnKF — 3DVar (prior, solid; posterior, dotted)

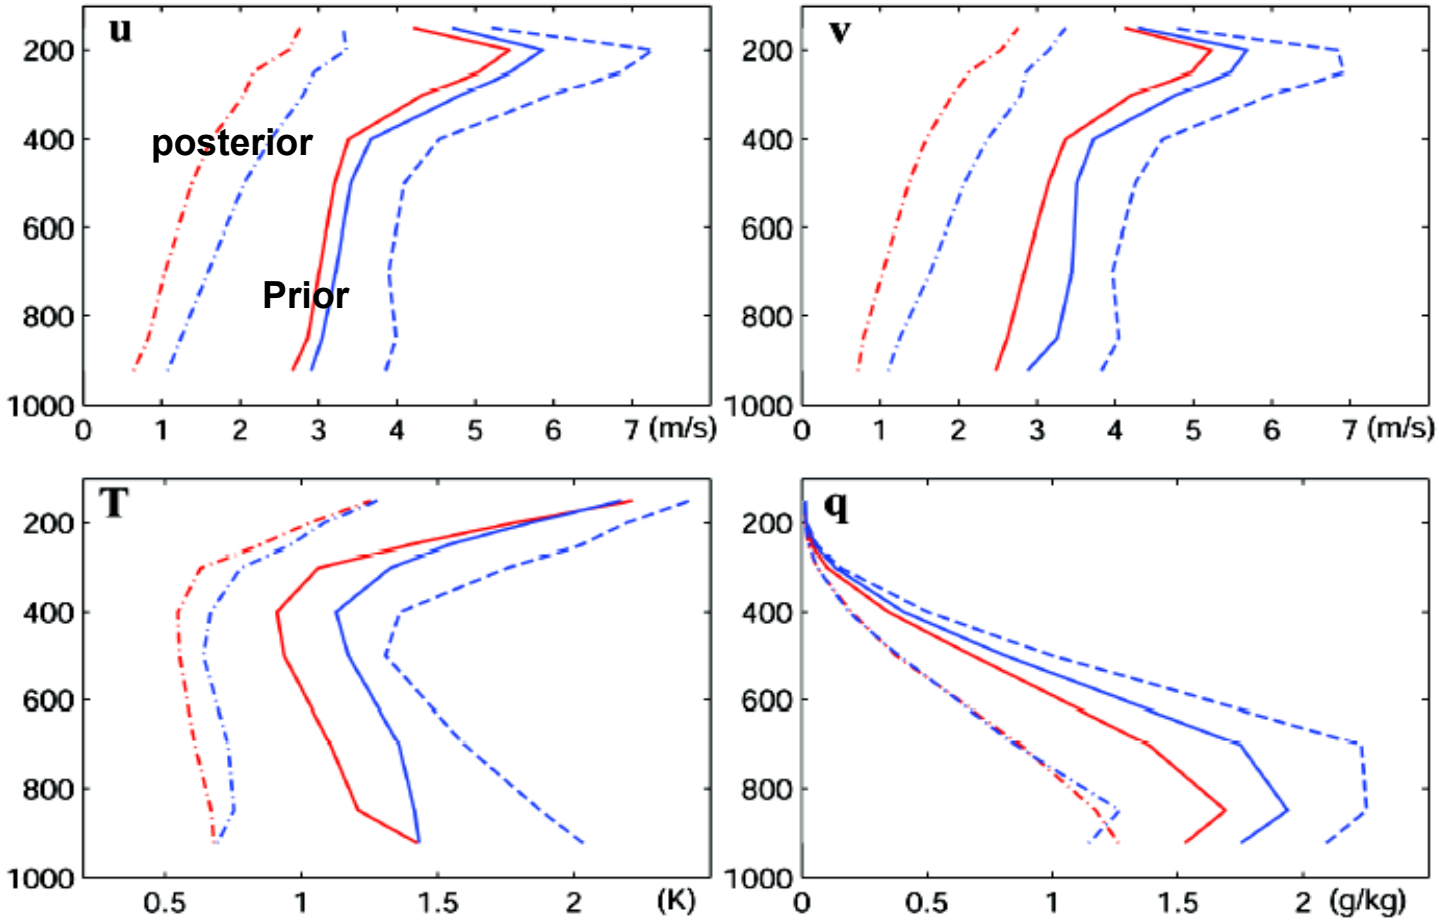

Better performance of EnKF than 3DVar also seen in both 12-h forecast and posterior analysis in terms of root-mean square difference averaged over the entire month

(Meng and Zhang 2007c, MWR, in review)

Nonlinear Observation Operators

The cases in which $\mathbf{y} = h(\mathbf{x}) + \epsilon$ or $\mathbf{x}$ is non-Gaussian, or both, can be handled easily.

Let $\hat{h} = (N_e)^{-1} \sum h(\mathbf{x}^i)$. Define $\mathbf{X}$ as before and

$$\mathbf{Y} \quad = \quad \text{matrix with columns } (N_e - 1)^{-1/2} \left(h(\mathbf{x}^i) + \epsilon^i - \hat{h} - \hat{\epsilon} \right),$$

EnKF update is as before, but with nonlinear $h(\mathbf{x}^i)$:

$$\xi^i = \mathbf{x}^i + \hat{\mathbf{K}} \left(\mathbf{y}^o - (h(\mathbf{x}^i) + \epsilon^i) \right)$$

$$\hat{\mathbf{K}} = \mathbf{X}\mathbf{Y}^T \left(\mathbf{Y}\mathbf{Y}^T \right)^{-1}$$

Nonlinear Observation Operators (cont.)

Note relation to the BLUE:

- $\mathbf{X}\mathbf{Y}^T$ is a sample estimate of $\mathbf{P}_{xy} = \mathrm{cov}(\mathbf{x}, \mathbf{y})$
- $\mathbf{Y}\mathbf{Y}^T$ is a sample estimate of $\mathbf{P}_{yy} = \mathrm{cov}(\mathbf{y})$
- $\hat{\mathbf{K}} \rightarrow \mathbf{P}_{xy}\mathbf{P}_{yy}^{-1}$ for large N_e

Thus, sample mean converges to the BLUE.

$$\begin{aligned} \hat{\mathbf{x}}^a &= \hat{\mathbf{x}}^f + \hat{\mathbf{K}}\left(\mathbf{y}^o - (\hat{h} + \hat{\epsilon})\right) \\ &\rightarrow E(\mathbf{x}) + \mathbf{P}_{xy}\mathbf{P}_{yy}^{-1}(\mathbf{y}^o - E(\mathbf{y})) \end{aligned}$$

EnKF for nonlinear observation operators approximates the BLUE.

But problems

- Collapse of ensemble for small ensemble size (less than a few hundred). Empirical 'covariance inflation'

- Spurious correlations appear at large geographical distances. Empirical 'localization' (see Gaspari and Cohn, 1999, *Q. J. R. Meteorol. Soc*)

- In formula

$$x^a{}_l = x^b{}_l + P^b H^{\mathrm{T}} [HP^b H^{\mathrm{T}} + R]^{-1} (y_l - Hx^b{}_l) , \qquad l = 1, \ldots, L$$

P^b, which is covariance matrix of an L-size ensemble, has rank L-1 at most. This means that corrections made on ensemble elements are contained in a subspace with dimension L-1. Obviously very restrictive if $L \ll p$, $L \ll n$.

Houtekamer and Mitchell (1998) use two ensembles, the elements of each of which are updated with covariance matrix of other ensemble.

There exist many variants of Ensemble Kalman Filter

Ensemble Transform Kalman Filter (***ETKF***, Bishop *et al.*, *Mon. Wea. Rev.*, 2001)

Requires a prior 'control' analysis x_c^a, emanating from a background x_c^b. An ensemble is evolved about that control without explicit use of the observations (and without feedback to control)

More precisely, define $L \times L$ matrix T such that, given $P^b = ZZ^T$, then $P^a = ZTT^TZ^T$ (not trivial, but possible). Then the background deviations $x^b{}_l - x_c^b$ are transformed through $Z \rightarrow ZT$ into an ensemble of analysis deviations $x^a{}_l - x_c^a$.

(does not avoid collapse of ensembles)

Local Ensemble Transform Kalman Filter (***LETKF***, Hunt *et al.*, *Physica D*, 2007)

Each gridpoint is corrected only through the use of neighbouring observations.

Other variants of Ensemble Kalman Filter

'*Unscented*' *Kalman Filter* (Wan and van der Merve, 2001, Wiley Publishing)

Weighted Kalman Filter (Papadakis *et al*., 2010, *Tellus A*)

Inflation-free Ensemble Kalman Filters (Bocquet and Sakov, 2012, *Nonlin. Processes Geophys.*)

Situation still not entirely clear.

In any case, optimality always requires errors to be independent in time. In order to relax that constraint, it is necessarily to augment the state vector in the temporal dimension.

Bayesian properties of Ensemble Kalman Filter ?

Very little is known.

Le Gland *et al*. (2011). In the linear and gaussian case, the discrete pdf defined by the filter, in the limit of infinite sample size L, tends to the bayesian gaussian pdf.

No result for finite size (note that ensemble elements are not mutually independent)

In the nonlinear case, the discrete pdf tends to a limit which is in general not the bayesian pdf.

Variational Assimilation

Variational approach can easily be extended to time dimension.

Suppose for instance available data consist of

- Background estimate at time 0

  $x_0^b = x_0 + \zeta_0^b \qquad E(\zeta_0^b \zeta_0^{bT}) = P_0^b$

- Observations at times $k = 0, \ldots, K$

  $y_k = H_k x_k + \varepsilon_k \qquad E(\varepsilon_k \varepsilon_j^T) = R_k \delta_{kj}$

- Model (supposed for the time being to be exact)

  $x_{k+1} = M_k x_k \qquad k = 0, \ldots, K\text{-}1$

Errors assumed to be unbiased and uncorrelated in time, H_k and M_k linear

Then objective function

$\xi_0 \in \mathcal{S} \rightarrow$

$$\mathcal{J}(\xi_0) = (1/2)\,(x_0^b - \xi_0)^T [P_0^b]^{-1} (x_0^b - \xi_0) + (1/2)\,\Sigma_k [y_k - H_k \xi_k]^T R_k^{-1} [y_k - H_k \xi_k]$$

subject to $\xi_{k+1} = M_k \xi_k\,, \qquad k = 0, \ldots, K\text{-}1$

$$\mathcal{J}(\xi_0) = (1/2)\,(x_0^b - \xi_0)^{\mathrm{T}}\,[P_0^b]^{-1}\,(x_0^b - \xi_0) + (1/2)\,\Sigma_k [y_k - H_k \xi_k]^{\mathrm{T}}\,R_k^{-1}\,[y_k - H_k \xi_k]$$

Background is not necessary, if observations are in sufficient number to overdetermine the problem. Nor is strict linearity.

How to minimize objective function with respect to initial state $u = \xi_0$ (u is called the *control variable* of the problem) ?

Use iterative minimization algorithm, each step of which requires the explicit knowledge of the local gradient $\nabla_u \mathcal{J} \equiv (\partial \mathcal{J}/\partial u_i)$ of $\mathcal{J}$ with respect to u.

How to numerically compute the gradient $\nabla_u \mathcal{J}$?

Direct perturbation, in order to obtain partial derivatives $\partial \mathcal{J} / \partial u_i$ by finite differences ? That would require as many explicit computations of the objective function $\mathcal{J}$ as there are components in u. Practically impossible.

Gradient computed by *adjoint method*.

Adjoint Method

Input vector $\boldsymbol{u} = (u_i)$, $\dim \boldsymbol{u} = n$

Numerical process, implemented on computer (*e. g.* integration of numerical model)

$$\boldsymbol{u} \rightarrow \boldsymbol{v} = \boldsymbol{G}(\boldsymbol{u})$$

- $\boldsymbol{v} = (v_j)$ is *output vector* , $\dim \boldsymbol{v} = m$
- Perturbation $\delta\boldsymbol{u} = (\delta u_i)$ of input. Resulting first-order perturbation on $\boldsymbol{v}$
  - $\delta v_j = \Sigma_i \, (\partial v_j / \partial u_i) \, \delta u_i$
- or, in matrix form
  - $\delta \boldsymbol{v} \; = \; \boldsymbol{G}' \, \delta \boldsymbol{u}$
- where $\boldsymbol{G}' \equiv (\partial v_j / \partial u_i)$ is local matrix of partial derivatives, or jacobian matrix, of $\boldsymbol{G}$.

Adjoint Method (continued 1)

$$\delta \boldsymbol{v} \;=\; \boldsymbol{G}'\,\delta \boldsymbol{u} \qquad \text{(D)}$$

- Scalar function of output

$$\mathcal{J}(\boldsymbol{v}) \;=\; \mathcal{J}[\boldsymbol{G}(\boldsymbol{u})]$$

Gradient $\nabla_{\boldsymbol{u}} \mathcal{J}$ of $\mathcal{J}$ with respect to input $\boldsymbol{u}$?

'Chain rule'

$$\partial \mathcal{J}/\partial u_i = \Sigma_j \, \partial \mathcal{J}/\partial v_j \, (\partial v_j/\partial u_i)$$

or

- $$\nabla_{\boldsymbol{u}} \mathcal{J} \;=\; \boldsymbol{G}'^{\mathrm{T}} \nabla_{\boldsymbol{v}} \mathcal{J} \qquad \text{(A)}$$

Adjoint Method (continued 2)

- $\boldsymbol{G}$ is the composition of a number of successive steps

$$\boldsymbol{G} = \boldsymbol{G}_N \circ \ldots \circ \boldsymbol{G}_2 \circ \boldsymbol{G}_1$$

'Chain rule'

$$\boldsymbol{G}' = \boldsymbol{G}_N' \ldots \boldsymbol{G}_2' \boldsymbol{G}_1'$$

Transpose

- $$\boldsymbol{G}'^{\mathrm{T}} = \boldsymbol{G}_1'^{\mathrm{T}} \boldsymbol{G}_2'^{\mathrm{T}} \ldots \boldsymbol{G}_N'^{\mathrm{T}}$$

Transpose, or *adjoint*, computations are performed in reversed order of direct computations.

If $\boldsymbol{G}$ is nonlinear, local jacobian $\boldsymbol{G}'$ depends on local value of input $\boldsymbol{u}$. Any quantity which is an argument of a nonlinear operation in the direct computation will be used again in the adjoint computation. It must be kept in memory from the direct computation (or else be recomputed again in the course of the adjoint computation).

If everything is kept in memory, total operation count of adjoint computation is at most 4 times operation count of direct computation (in practice about 2).

Adjoint Approach

$$\mathcal{J}(\xi_0) = (1/2)\,(x_0^b - \xi_0)^{\mathrm{T}}\,[P_0^b]^{-1}\,(x_0^b - \xi_0) + (1/2)\,\Sigma_k [y_k - H_k\xi_k]^{\mathrm{T}}\,R_k^{-1}\,[y_k - H_k\xi_k]$$

subject to $\xi_{k+1} = M_k\xi_k\,, \qquad k = 0, \ldots, K\text{-}1$

Control variable $\qquad \xi_0 = \boldsymbol{u}$

Adjoint equation

$$\lambda_K = \qquad H_K^{\mathrm{T}}\,R_K^{-1}\,[H_K\,\xi_K - y_K]$$

$$\lambda_k = M_k^{\mathrm{T}}\lambda_{k+1} + H_k^{\mathrm{T}}\,R_k^{-1}\,[H_k\,\xi_k - y_k] \qquad k = K\text{-}1, \ldots, 1$$

$$\lambda_0 = M_0^{\mathrm{T}}\lambda_1 \quad + H_0^{\mathrm{T}}\,R_0^{-1}\,[H_0\,\xi_0 - y_0] \quad + \; [P_0^b]^{-1}\,(\xi_0 - x_0^b)$$

$$\nabla_u \mathcal{J} = \lambda_0$$

Result of direct integration (ξ_k), which appears in quadratic terms in expression of objective function, must be kept in memory from direct integration.

Adjoint Approach (continued 2)

Nonlinearities ?

$$\mathcal{J}(\xi_0) = (1/2)\,(x_0^b - \xi_0)^{\mathrm{T}}\,[P_0^b]^{-1}\,(x_0^b - \xi_0) + (1/2)\,\Sigma_k [y_k - H_k(\xi_k)]^{\mathrm{T}}\,R_k^{-1}\,[y_k - H_k(\xi_k)]$$

subject to $\xi_{k+1} = M_k(\xi_k)\,, \qquad k = 0, \ldots, K\text{-}1$

Control variable $\xi_0 = \boldsymbol{u}$

Adjoint equation

$$\lambda_K = H_K{}'^{\mathrm{T}}\,R_K^{-1}\,[H_K(\xi_K) - y_K]$$

$$\lambda_k = M_k{}'^{\mathrm{T}}\lambda_{k+1} + H_k{}'^{\mathrm{T}}\,R_k^{-1}\,[H_k(\xi_k) - y_k] \qquad k = K\text{-}1, \ldots, 1$$

$$\lambda_0 = M_0{}'^{\mathrm{T}}\lambda_1 + H_0{}'^{\mathrm{T}}\,R_0^{-1}\,[H_0(\xi_0) - y_0] + [P_0^b]^{-1}\,(\xi_0 - x_0^b)$$

$$\nabla_u \mathcal{J} = \lambda_0$$

Not heuristic (it gives the exact gradient $\nabla_u \mathcal{J}$), and really used as described here.