

The University of Reading

Weak Constraints in Four-dimensional Variational Assimilation

L.R. Watkinson¹, A.S. Lawless¹, N.K. Nichols¹ and
I. Roulstone²

NUMERICAL ANALYSIS REPORT 1/2007

Presented at the 7th International Workshop on Adjoint
Applications in Dynamic Meteorology, October 2006

¹*Department of Mathematics
The University of Reading
Whiteknights, PO Box 220
Reading, RG6 6AX*

²*Department of Mathematics
The University of Surrey
Guildford, GU2 7XH*

Department of Mathematics

1 Introduction

Data assimilation for numerical weather prediction (NWP) aims to use information from observations together with a numerical forecasting model to produce the best estimate of the state of the atmosphere. Many operational forecasting centres have moved towards variational assimilation techniques, in which the analysis is found by minimizing a function measuring the distance between observations and the model state. However since the observations themselves are insufficient to determine the atmospheric state, extra constraints must be incorporated into the variational problem. Such constraints fall into two categories, strong constraints, which must be satisfied exactly, and weak constraints, which need be only approximately satisfied (Sasaki 1970). Of the first type the most common approach is to incorporate the numerical model as a strong constraint, so that a sequence of states over a time interval of observations must satisfy the model equations. This leads to the method of four-dimensional variational assimilation (4D-Var), which is now operational in many forecasting centres (Rabier *et al.* 2000, Laroche *et al.* 2005, Rawlins 2005). Of the weak constraints the most common is the requirement that the analysed state be close to a previous short-range forecast or background field. This leads to the so-called ‘background term’ within the 4D-Var cost function and such a term is now standard in almost all implementations of 4D-Var.

Other constraints may also be added into the cost function in two ways. The most common way is the incorporation of linear constraints through the formulation of the background error covariance matrix. This method is particularly used to enforce balance conditions within the analysis. Such constraints are specified in the assimilation system by means of a variable transformation, which is designed to implicitly enforce relationships between the analysis increments (Parrish and Derber 1992, Cullen 2003). Alternatively weak constraints may be added as extra terms on the cost function, to enforce balance or other conditions only weakly. This method has been used to filter high frequency waves from the analysis or to smooth the final solution (Gauthier and Thépaut 2001, Lin *et al.*

assimilation system to exactly conserve such quantities, since if the background value is in error then the observations should be able to correct it. However we may expect that information can be gained from weakly constraining these quantities in the assimilation scheme, since they are unlikely to vary widely between one assimilation cycle and the next.

It may be argued that if the model itself is conserving and we constrain only the initial state by means of a background term, then implicitly we are also constraining the conserved quantities. In this paper we consider whether the explicit incorporation of conservation properties as a weak constraint can add information to a 4D-Var analysis, even if the model itself is conserving. To study this problem we use a simple Hamiltonian system, the two-body problem of celestial mechanics. This is a Hamiltonian system with known conservation properties and so we can test rigorously whether there is any extra benefit to be gained by making the constraint on the conserved quantities explicit. Initial results reported in Watkinson *et al.* (2005) and Watkinson *et al.* (2006) indicated that such a weak constraint behaves very differently from a standard background constraint. Here we provide further numerical evidence that this is the case and detailed analysis to explain this difference. The paper is organised as follows. In the next section we set out the mathematical formulation for the incorporation of weak constraints into 4D-Var. Section 3 presents the two-body problem that we use in this study in its continuous and discrete formulations. In section 4 we examine in detail the incorporation of a background constraint and two conservation weak constraints into 4D-Var applied to the two-body problem. The different formulations are compared numerically in section 5, before we present our conclusions in section 6.

2 Four-dimensional variational assimilation with weak \mathfrak{d}

we define a function $F(t, \mathbf{x}, \dot{\mathbf{x}})$ which measures the distance between the model trajectory and the observations. The 4D-Var problem, with only the observation term, can then be formulated in the following way: Minimize the cost function

$$J_o(\mathbf{x}) = \int_{t_0}^{t_N} F(t, \mathbf{x}, \dot{\mathbf{x}}) dt, \quad (2)$$

subject to the strong constraint (1). Using standard methods of the calculus of variations (Courant and Hilbert, 1953) we incorporate the strong constraint by defining a Lagrangian

$$\mathcal{L} = \int_{t_0}^{t_N} \{F(t, \mathbf{x}, \dot{\mathbf{x}}) + \boldsymbol{\lambda}(t)^T (\dot{\mathbf{x}} - \mathbf{f}(\mathbf{x}))\} dt, \quad (3)$$

where $\boldsymbol{\lambda}(t)$ is a vector of Lagrange multipliers. Then the solution of the variational problem is given by the value of \mathbf{x} which ensures that the first variation $\delta \mathcal{L} = 0$. If we define $G(\mathbf{x}, \boldsymbol{\lambda}, \dot{\mathbf{x}}, \dot{\boldsymbol{\lambda}}) \equiv F(t, \mathbf{x}, \dot{\mathbf{x}}) + \boldsymbol{\lambda}(t)^T (\dot{\mathbf{x}} - \mathbf{f}(\mathbf{x}))$, then necessary conditions to ensure $\delta \mathcal{L} = 0$ are given by Euler-Lagrange equations

$$G_{\mathbf{x}} - \frac{d}{dt} G_{\dot{\mathbf{x}}} = 0, \quad (4)$$

$$G_{\boldsymbol{\lambda}} - \frac{d}{dt} G_{\dot{\boldsymbol{\lambda}}} = 0, \quad (5)$$

where subscripts indicate derivatives. The second of these simply recovers the equations of the dynamical system (1). The first Euler-Lagrange equation implies the *adjoint equations*

$$\dot{\boldsymbol{\lambda}} = - \left[\frac{\partial \mathbf{f}}{\partial \mathbf{x}} \right] \boldsymbol{\lambda} + \frac{\partial F}{\partial \mathbf{x}} - \frac{d}{dt} \frac{\partial F}{\partial \dot{\mathbf{x}}}, \quad (6)$$

where $\partial \mathbf{f} / \partial \mathbf{x}$ is the Jacobian of \mathbf{f} with respect to \mathbf{x} . We also must satisfy the boundary conditions

$$\boldsymbol{\lambda}(t_N) = 0, \quad \boldsymbol{\lambda}(t_0) = 0. \quad (7)$$

We impose the first of these at the final time t_N . The second boundary condition specifies that the gradient of the Lagrangian is zero at the initial time.

We now consider the effect on the problem of adding a weak constraint. We consider a general weak constraint $\Phi(\mathbf{x}(t_0))$, which is a function of \mathbf{x} at the initial time. Such a constraint can be incorporated into the variational problem by defining a new Lagrangian

$$\mathcal{L} = \int_{t_0}^{t_N} G dt + \Phi(\mathbf{x}(t_0)), \quad (8)$$

where the Lagrange multiplier for the constraint is constant and can be taken to be equal to one without loss of generality. Since this constraint is not integrated over time, the Euler-Lagrange equations remain unchanged. However the constraint changes the boundary

planetary motion is not exactly described by a reduced system, in the case where a small body such as the Earth orbits a large body such as the sun, the reduced problem and the full system are very similar and Kepler's laws apply to the reduced problem. These laws provide extra qualitative information about the system being modelled and can be summarized as follows (cf. Acheson 1997, p.71):

- The shape of an orbit is a conic section with the centre of mass at one focus.
- A line from the orbitting body to the fixed body will sweep out equal areas in equal times.
- The square of the period of the orbit is proportional to the cube of its semi-major axis.

I

3.2 The discrete model

When discretizing a Hamiltonian system it is important not only to reduce the local truncation error, but to preserve the global qualitative features of the system being modelled, such as conservation laws or symmetries. Symplectic integration schemes for Hamiltonian problems are designed to ensure that the conservation properties of the continuous system are preserved in the discrete model. The scheme we use for modelling the two-body problem is the Stormer-Verlet scheme, which is a second-order accurate symplectic method. This scheme has previously been used to model similar Hamiltonian systems (Budd and Piggott 2000, Leimkuhler and Reich 2001). We define the position vector of the discrete system $\mathbf{Q} = (Q_1, Q_2)^T$ and the discrete momentum vector $\mathbf{P} = (P_1, P_2)^T$. Then applying the Stormer-Verlet scheme to the equations of the reduced two-body problem (14) and (15) we obtain the discrete system

$$\mathbf{P}^{n+\frac{1}{2}} = \mathbf{P}^n - \frac{h}{2} \frac{\mathbf{Q}^n}{(Q_1^{n^2} + Q_2^{n^2})^{\frac{3}{2}}}; \quad (17)$$

$$\mathbf{Q}^{n+1} = \mathbf{Q}^n + h\mathbf{P}^{n+\frac{1}{2}}; \quad (18)$$

$$\mathbf{P}^{n+1} = \mathbf{P}^{n+\frac{1}{2}} - \frac{h}{2} \frac{\mathbf{Q}^{n+1}}{(Q_1^{n+1^2} + Q_2^{n+1^2})^{\frac{3}{2}}}, \quad (19)$$

where h is the non-dimensional time step.

In order to demonstrate the conservation properties of this scheme we run the model from the initial conditions $\mathbf{Q} = (1, 0), \mathbf{P} = (0, 1)$, taken from Budd and Piggott (2003). The analytical solution with these starting values is a circular orbit with a period of 2π and a radius of one. For a time step of $h = 0.001$ we verify that the error in the energy compared to the constant analytical value remains small. In Figure 1 we show the error in the energy over two orbits for a circular orbit (eccentricity = 0) and an orbit with eccentricity equal to 0.5. We note that the scales on the two plots are different. For the circular orbit the energy is almost exactly conserved, with the error of the order of machine precision. As the eccentricity is increased there is a greater error in the energy conservation, of the order 10^{-7} , except at one point in the orbit. This corresponds to the point of closest approach, which for this experiment is the point of the initial conditions. At this point the energy returns to its initial value, due to the symmetry of the problem. For elliptical orbits started away from the point of closest approach the error increases slightly at the point of closest approach on the orbit. Schemes with varying time steps must be then used if the error is to be reduced at all points on the orbit. These results are discussed in more detail in Watkinson *et al.* (2005) and Watkinson (2006). For the

purpose of studying the behaviour of weak constraints in 4D-Var we do not need the complication of large ellipticities and so we may restrict ourselves to circular orbits. We

We now consider the effect of including the Hamiltonian, or total energy, as a weak constraint. We define a weak-constraint term

$$J_E = \frac{1}{2} \alpha_2 (E(\mathbf{q}, \mathbf{p}) - E(\mathbf{q}_b, \mathbf{p}_b))^2 \quad (24)$$

where $E(\mathbf{q}, \mathbf{p})$ is defined by (12) and α_2 is a scalar. Again this is of the form (10), where $\mathbf{g}(\mathbf{x})$ is now the scalar function $E(\mathbf{q}, \mathbf{p})$ and $\mathbf{C}^{-1} = \alpha_2 \mathbf{I}$. The gradient equation for this constraint then contains the term

$$\nabla J_E = \alpha_2 (E(\mathbf{q}, \mathbf{p}) - E(\mathbf{q}_b, \mathbf{p}_b)) \begin{pmatrix} \partial E / \partial \mathbf{q} \\ \partial E / \partial \mathbf{p} \end{pmatrix}. \quad (25)$$

Using the canonical equations (13) we see that this is equal to

$$\nabla J_E = \alpha_2 (E(\mathbf{q}, \mathbf{p}) - E(\mathbf{q}_b, \mathbf{p}_b)) \begin{pmatrix} -\dot{\mathbf{p}} \\ \dot{\mathbf{q}} \end{pmatrix}, \quad (26)$$

where the dot indicates the total time derivative. We see therefore that when the Hamiltonian is used as a constraint the gradient depends not only on the energy difference from the background, but also on the time derivative of the current model solution. The gradient with respect to the position variables depends on the rate of change of momentum and similarly the gradient with respect to the momentum variables depends on the rate of change of position. Whereas the gradient term arising from a background constraint depends on variables at one particular time only, that arising from the Hamiltonian constraints depends on the model time tendency at that time. Hence we may expect these two constraints to act differently within the data assimilation system. From the structure of (25) and the canonical equations we see that for 4D-Var applied to any Hamiltonian system, adding a weak constraint on the Hamiltonian will give gradient components in the direction of the model tendencies.

The final constraint we consider is a weak constraint on the angular momentum, of the form

$$J_L = \frac{1}{2} \alpha_3 (L(\mathbf{q}, \mathbf{p}) - L(\mathbf{q}_b, \mathbf{p}_b))^2, \quad (27)$$

where L is the angular momentum defined by (16) and α_3 is a scalar. The gradient term introduced by this constraint is equal to

$$\nabla J_L = \alpha_3 (L(\mathbf{q}, \mathbf{p}) - L(\mathbf{q}_b, \mathbf{p}_b)) \begin{pmatrix} p_2 \\ -p_1 \\ -q_2 \\ q_1 \end{pmatrix}. \quad (28)$$

For this constraint the gradient depends on the angular momentum difference from the background and on the actual values of the trajectory. The gradient with respect to position has a component in the direction defined by the momentum coordinates and the gradient with respect to momentum has a component in the direction defined by the position coordinates. We now study the effect of these differences in some numerical experiments.

5 Numerical experiments

To examine how the different weak constraints affect a 4D-Var analysis we set up a 4D-Var system for the two-body problem presented in Section 3. We consider a discrete formulation of the cost function introduced in Section 2, which we write in the standard form

$$J_o(\mathbf{x}_0) = \frac{1}{2} \sum_{i=0}^n (\mathbf{y}_i - \mathbf{h}_i(\mathbf{x}_i))^T \mathbf{R}_i^{-1} (\mathbf{y}_i - \mathbf{h}_i(\mathbf{x}_i)), \quad (29)$$

where \mathbf{y} are the observations at time t_i , \mathbf{h}_i is the observation operator and \mathbf{R} is the observation error covariance matrix. The cost function is minimized with respect to the initial state \mathbf{x}_0 , subject to the discrete model equations described in Section 3.2. Discrete equivalents of the weak constraints described in Section 4 can be added in turn to this cost function. The adjoint model for the system is generated from the discrete numerical model using the approach of automatic differentiation (Giering and Kaminski 1998). The code is verified by using the standard adjoint test and gradient test (Watkinson 2006, Li *et al.* 1994). The discrete cost function is minimized using the quasi-Newton method of the CONMIN algorithm (Shanno and Phua 1980)

For each of the experiments described the true solution is found by running the model with the same initial conditions as in Section 3.2, with a model time step of 0.001. The solution trajectory is a circular orbit of radius one. Identical twin assimilation experiments are carried out using observations taken from the true trajectory with a random Gaussian error of variance 10^{-4} . The assimilation time window is taken to be π time units, which corresponds to half of an orbit, and a forecast is run for a further four complete orbits from the end of the assimilation window. Observations are taken of the whole model state at the end of the assimilation window and halfway through the window, so that the observation operator is equal to the identity matrix at each of these times. As in Watkinson *et al.* (2005) the 4D-Var minimization is stopped when all of the following conditions hold (Gill

et al. 1986, p.306):

$$\frac{|J_{k-1} - J_k|}{1 + |J_k|} < \epsilon, \quad \frac{\|\mathbf{x}_{k-1} - \mathbf{x}_k\|_2}{1 + \|\mathbf{x}_k\|_2} < \epsilon^{\frac{1}{2}}, \quad \frac{\|\nabla J_k\|_2}{1 + |J_k|} < \epsilon^{\frac{1}{3}}, \quad (30)$$

where k is the iteration number, J_k is the sum of J_o and the constraint term evaluated at the iterate \mathbf{x}_k and ϵ is a tolerance which we choose to be 10^{-6} .

5.1 Perfect back round

For the first set of experiments we take the background field \mathbf{x}_b to be equal to the true solution at the initial time. Although this is unrealistic, since in practice it is the truth we are trying to estimate, these experiments allow us to study cleanly the different behaviours of the weak constraints. The initial guess for the 4D-Var minimization is a small perturbation to the truth, perturbed in such a way as to have an incorrect energy and angular momentum. We examine first the effect of a weak constraint on the background field itself, as described by (20). In Figure 2 we plot the root mean square error of the state vector during the assimilation and forecast period for experiments with only the observation term and with various weightings of the background weak constraint. For the case in which no weak constraint is applied we see that the error has a tendency which increases with time, with an oscillating error pattern superimposed. As the background constraint is applied with increasing weight both of these errors are reduced. However it appears that the background constraint is more efficient at suppressing the oscillations; the divergence with time is still present with the highest value of the weighting α_1 .

In Figure 3 we see the same experiment performed with the background constraint replaced by a weak constraint on the energy, as defined by (24). Different weightings are chosen from those used for the background constraint to ensure that the relative weight of the observation and constraint term is similar for both constraints. For the energy constraint the behaviour is different from the background constraint as the weighting is increased. In this case the oscillations are not smoothed out as they were for the background constraint, even for high values of α_2 . However we see that for the highest value of the weighting the divergence of the error with time is removed. If we use a weak constraint on the angular momentum (27) then the error plot is almost identical to that of the energy constraint shown in Figure 3.

To explain the difference in behaviour between the constraints we must consider the sources of the two different types of error we see in the unconstrained solution, the growth over time and the oscillations. As discussed in Watkinson *et al.* (2005), the increase in

the error over time is a result of the analysed and true orbits being out of phase with each other. Thus we can conclude that the period of the analysed orbit is in error from the true solution. From Kepler's third law this implies an error in the length of the semi-major axis. On the other hand the oscillations in the error curves can be explained by considering the comparison of two solutions with the same semi-major axes but different eccentricities. Two such solutions would have the same period and therefore will always be at the same point on the orbit, but the difference between them will be smaller as they pass the points of closest and furthest approach. This can be seen more clearly from the diagram in Figure 4.

Having understood these different sources of error we can now see more clearly why the different weak constraints behave in the way they do. Tw

p

for the energy constraint, the period of the orbit is constrained.

From these results with a perfect background constraint we can understand how the different formulations of the weak constraint affect the analysis in different ways. As we presented in Section 4 a weak constraint on the background field introduces a term in the gradient which depends on the difference in components of the state vector and this acts to control the shape of the final orbit. The energy constraint was shown in Section 4 to give a gradient term which depends on the difference in the total energy and this acts to control the period of the orbit. In a similar way a constraint on the angular momentum also controls the period, by means of a gradient term dependent upon the error in angular momentum.

In some ways these results may seem surprising. The model itself conserves energy and angular momentum and so we may expect that by constraining the initial state, we are also constraining the energy and angular momentum over the time window. Nevertheless it appears that the use of an explicit constraint on the conserved quantities gives extra information which is not provided by the use of the model alone. The performance of the 4D-Var assimilation system can be improved by making explicit use of the information which is implicitly contained within the system dynamics. However the experiments presented so far assume that the background state is equal to the true state, which would not be the case in practice. We now look at the effect of the constraints where we have a more realistic background field.

5.2 Noisy back round

We now consider adding the different weak constraints to the cost function when the background field is not equal to the truth. We use the same experimental design as in the

where $\mathbf{B} = \sigma_b^2 \mathbf{I}$.

For the energy and angular momentum constraints we must calculate the variance from the variance of the background. Thus the energy constraint becomes

$$J_E = \frac{1}{\sigma_E^2} (E(\mathbf{q}, \mathbf{p}) - E(\mathbf{q}_b, \mathbf{p}_b))^2, \quad (32)$$

where the variance σ_E^2 is defined by

$$\sigma_E^2 = \nabla_{\mathbf{x}_b} E(\mathbf{x}_b(t_0))^T \quad (33)$$

and $= (\sigma_b^2, \sigma_b^2, \sigma_b^2, \sigma_b^2)^T$. The angular momentum constraint has a similar form to the energy constraint. We add a term to the cost function of the form

$$J_L = \frac{1}{\sigma_L^2} (L(\mathbf{x}_b(t_0)) - L(\mathbf{x}(t_0)))^2, \quad (34)$$

where the variance σ_L^2 is defined by

$$\sigma_L^2 = \nabla_{\mathbf{x}_b} L(\mathbf{x}_b(t_0))^T. \quad (35)$$

We consider the effect of these constraints individually and their combined effects.

In Figure 5 we show the evolution of the error for the experiments with only a background constraint, only an energy constraint and the combination of both these constraints, for a noisy background field. The curve for the experiment with no constraints is exactly the same as in the previous section. We consider first the background constraint.

We see that as with the point constraint, the evolution of the error is significantly reduced with the inclusion of the energy constraint. The error decreases rapidly at first and then levels off, reaching a minimum value of approximately 0.01 after about 400 iterations.

is not providing any extra information which is not present in the energy constraint. To test this hypothesis we compare the effect of using all three constraints together with an experiment using the background constraint plus twice the energy constraint. The resulting error plot is shown in Figure 6. We see that the errors for these two experiments are equal. The addition of the angular momentum constraint is equivalent to doubling the weight on the energy constraint. Thus, consistent with our previous experiments, we conclude that the angular momentum and energy constraints act in the same way, that is both act to constrain the period of the analysis.


6 Conclusions


When assimilating data into a dynamical system it is important to incorporate as much knowledge of the system as possible into the assimilation process. Most modern assimilation methods make good use of the system dynamics and of dynamical balance relationships. However knowledge of the conservation laws of the system is usually incorporated only implicitly through the system equations. In this study we have investigated the value of imposing information about these conservation laws explicitly as a weak constraint. A theoretical analysis of such weak constraints in the context of the two-body problem showed that they lead to terms in the gradient equations quite different from those pro-

- (2005). Evaluation of the operational 4D-Var at the Meteorological Service of Canada. Preprints, *Proceedings of the Fourth WMO International Symposium on Assimilation of Observations in Meteorology and Oceanography*, Prague, WMO, 139.
- Leimkuhler, B. and Reich, S., (2001). A reversible averaging integrator for multiple time-scale dynamics. *J. Comp. Phys.*, **171**, 95–114.
- Li, Y., Navon, I.M., Yang, W., Zou, X., Bates, J.R., Moorthi, S. and Higgins, R.W., (1994). Four-dimensional variational data assimilation experiments with a multilevel semi-Lagrangian semi-implicit general circulation model. *Mon. Wea. Rev.*, **122**, 966–983.
- Lin, C.-L., Chai, T. and Sun, J., (2002). On the smoothness constraints for four-dimensional data assimilation. *J. Comp. Phys.*, **181**, 430–453.
- Nichols, N.K., (2003). Treating model error in 3-D and 4-D data assimilation. In *Proceedings of the NATO Advanced Study Institute on Data Assimilation for the Earth System, Maratea, Italy, June 2002*, 127–135.
- Parrish, D.F. and Derber, J.C., (1992). The National Meterological Center's spectral statistical-interpolation analysis scheme. *Mon. Wea. Rev.*, **120**, 1747–1763.
- Rabier, F., Jarvinen, H., Klinker, E., Mahfouf, J.-F. and Simmons, A., (2000). The ECMWF operational implementation of four-dimensional variational assimilation. I: Experimental results with simplified physics. *Quarterly Journal of the Royal Meteorological Society*, **126**, 1143–1170.
- Rawlins, R., (2005). Operational implementation of 4D-Var in global model of the Met Office, U.K. Preprints, *Proceedings of the fourth WMO international symposium on assimilation of observations in meteorology and oceanography*, Prague, WMO, 138.
- Sasaki, Y., (1970). Some basic formalisms in numerical variational analysis. *Mon. Wea. Rev.*, **98**, 875–883.
- Shanno, D.F. and Phua, K.H., (1980). Remark on algorithm 500 - a variable method subroutine for unconstrained nonlinear minimization. *ACM Trans. on Mathematical Software*, **6**, 618–622.

Tr'

(a)








Figure 4: Illustration of the effect of comparing orbits with two different eccentricities.

