

Improving Weather Forecasting Accuracy by Using r-Adaptive Methods Coupled to Data Assimilation Algorithms

Chris Budd, Mike Cullen and Chiara Piccolo

Abstract Weather impacts all of our lives and we all take a close interest in it, with every news report finishing with a weather forecast watched by millions. Accurate weather forecasting is essential for the transport, agricultural and energy industries and the emergency and defence services. The Met Office plays a vital role by making 5-day forecasts, using advanced computer algorithms which combine numerical weather predictions (NWP) with carefully measured data (a process known as data assimilation). However, a major limitation on the accuracy of these forecasts is the sub-optimal use of this data. Adaptive methods, developed in a partnership between Bath and the Met Office have been employed to make better use of the data, thus improving the Met Office operational data assimilation system. This has led to a significant improvement in forecast accuracy as measured by the UK Index [9] with great societal and economic impact. Forecasts, of surface temperatures, in particular, are pivotal for the OpenRoad forecasting system used by local authorities to plan road clearing and gritting when snow or ice are predicted.

Data Assimilation

Data Assimilation is an essential part of the Met Office forecasting procedures, and involves combining the predictions of a numerical weather prediction calculation with data that is received from satellites, radiosondes and other observations [7]. Typically in a data assimilation calculation an optimal estimate for the initial conditions for a forecast (called the analysis) is obtained by comparing the forecasts obtained with that initial condition with data over a period of time (the assimilation

C. Budd (✉)
University of Bath, Bath BA2 7AY, UK
e-mail: mascjb@bath.ac.uk

M. Cullen · C. Piccolo
Met Office, FitzRoy Road, Exeter EX1 3PB, UK
e-mail: mike.cullen@metoffice.gov.uk

C. Piccolo
e-mail: chiara.piccolo@metoffice.gov.uk

window). The procedure to do this used by the Met Office is variational and since its implementation about 20 years, this has led to a significant increase in the Met Office forecasting accuracy [7].

However, a significant problem faced by the Met Office is that of assimilating data in the presence of atmospheric inversion layers or other fine structures. In an inversion layer for example, there is a very rapid temperature change from the cold air trapped below the warmer air above. Typically this change occurs on a length scale smaller than the computational grid. Misrepresenting these layers in the computations, leads to spurious spreading of the influence of the observations. This is because the errors in the data above and below the inversion layer are usually weakly correlated. This has a negative effect on the forecast, degrading the forecast performance in particular when trying to determine the weather close to the ground. The nature, and societal impact, of this problem can be severe both in the air and on the roads and is described in the following two quotes from publications authored by Chiara Piccolo and Mike Cullen [8, 9]:

A common problem in forecast case-studies is the misrepresentation of inversions and stratocumulus layers in the assimilation due to inappropriate background error covariances, e.g. smooth and broad vertical correlation functions which do not allow accurate fitting of high resolution radiosonde soundings. This inhibits the ability to diagnose realistic stratocumulus layers and boundary-layer structures which then results in poor forecasts. An example of the impact of this problem in the Met Office NWP system happened in December 2006 when poor visibility at Heathrow led to significant travel disruption during the Christmas period. In this instance radiosonde observations were not able to improve the analysis of the inversion and so the fog was not accurately forecast. Met Office Publication [8].

The accurate representation of the boundary layer in NWP models is important for instance in the forecasting of fog or icy roads. Met Office Publication [8].

One way to overcome the problem of insufficient resolution of inversion layers is simply to use a finer computational grid everywhere. However this is computationally inefficient and would entail a major redevelopment of the Met Office codes. An alternative possibility is to use a mesh which locally adapts to the fine structures, but is coarse elsewhere. Collaborative research between the group led by Prof. Chris Budd at the University of Bath and the Met Office, had been going on for some period of time on the possibility of using r-adaptive computational meshes to resolve local structures in numerical weather prediction algorithms. Whilst this research had originally been in the context of computing the dynamics of the solution, it was realised by Mike Cullen, then head of the variational data assimilation research group at the Met Office, that it might be possible to also use these r-adaptive methods to reduce the errors in data assimilation due to the inappropriate background error covariances in a manner which could then be coupled to existing data assimilation software. As a result, such adaptive methods were then applied to better resolve the troublesome inversion layers. The preliminary results were so successful in reducing computational errors, that they were developed by Chiara Piccolo at the Met Office into a fully operational adaptive data assimilation code which is used everyday to forecast the weather. We now describe the ideas behind this and the way that it led to impact on the Met Office system.

Adaptive Mesh Methods

The underpinning research originally conducted at Bath started as a systematic study of cheap, flexible, and robust, adaptive mesh redistribution methods with evolving mesh density which could be used to solve PDEs with structures evolving over small time and length scales. These computational meshes, coupled to appropriate methods for solving the PDE (such as finite element or finite volume methods), can then be used in numerical algorithms to compute the solutions of such evolutionary PDEs in several spatial dimensions. Such PDEs are typically discretised on a mesh and the discrete equations solved numerically. If features of the solution evolve on small time or length scales, conventional methods (based on nearly uniform meshes) may fail, whereas the adaptively redistributed meshes provide accurate robust solutions in a wide range of applications. The research led by Chris Budd has been centred on devising methods for moving the mesh so that a *fixed number* of mesh points are concentrated where they are most needed to resolve fine structures, such as the atmospheric inversion layers, without additional computational cost. The advantages of this approach (usually called r-adaptivity) over other adaptive approaches are that it is computationally simpler with straightforward and unchanging data structures, r-adaptive algorithms can be readily inserted into legacy software, the mesh regularity can be controlled a-priori and it can explicitly exploit the structures of the underlying PDE.

The Bath team, in collaboration with researchers Simon Fraser University (Canada), originally developed procedures for constructing r-adaptive moving mesh methods in one-dimension that could cope with specific singular PDEs which evolve over small time and length scales. In such r-adaptive methods a monitor function $m(x, t)$ is used to control the density of the mesh points, so that there are a large number of points concentrated when $m(x, t)$ is itself large. The monitor function can either be dictated by the mathematical analysis of the system, for example an estimate of the local truncation error of the numerical method, or can be determined directly from the physics described by the PDE. To be more precise, if the PDE with solution $u(x, t)$ has independent spatial variable $x \in [0, 1]$ with associated mesh points $X_i(t)$, with $i = 0 \dots N$, then the mesh points can be determined by solving the set of equidistribution equations

$$\int_{X_i(t)}^{X_{i+1}(t)} m(x, t) dx = \frac{\gamma(t)}{N} \tag{1}$$

Here $\gamma(t) = \int_0^1 m(x, t) dx$ is independent of x . The system (1) automatically concentrates the mesh points close to regions when $m(x, t)$ is large, however these equations are cumbersome to solve in practice. One approach to solve them is to consider the $X_i(t)$ to be point values of a continuous function $X(\xi, t)$, $\xi \in [0, 1]$ so that $X_i = X(i/N, t)$. Then a differentiation of (1) becomes the ‘equidistribution equation’

$$m X_\xi = \gamma(t) \quad \text{or on differentiation} \quad (m X_\xi)_\xi = 0. \tag{2}$$

The latter equation can then be solved by relaxation, through solving the moving mesh PDE (MMPDE)

$$X_\tau = (m X_\xi)_\xi . \quad (3)$$

Typically this MMPDE is solved either simultaneously with the underlying PDE, or alternately with it. The research programme at Bath extended these MMPDE methods to two and three dimensions, by using ideas from geometry and fluid mechanics. This work facilitated the development of the Parabolic Monge-Ampere (PMA) algorithm [3, 4]. This method for mesh redistribution combined the equidistribution of an appropriate monitor of the solution with optimal transport methods and the solution of an associated Monge-Ampere equation. The PMA algorithm was first implemented by a PhD student, John Williams (2000–2004), and proved to be effective on model problems. In a more developed form, it was the basis of an invited paper [5], which described in detail how the PMA method could either be used to solve PDEs or to derive meshes to better represent the fine structure in their solutions. This paper was of significant interest to the Met Office as many meteorological phenomena occur on small length scales relative to the overall scale of the Earth. In 2006, an EPSRC/Met Office CASE student at Bath (Emily Walsh), started a programme of research developing the PMA algorithm specifically for meteorological problems. The PMA algorithm was applied to improve the numerical prediction of severe storms associated with rapid variations in wind speed and temperature. Intensive research in this context led to the identification of appropriate monitors $m(x, t)$ of the atmospheric state, based upon estimates of the potential vorticity, which in turn were invoked to obtain effective computational meshes [6]. Chris Budd and a Met Office sponsored PDRA Phil Browne, in collaboration with Mike Cullen and Chiara Piccolo, have continued this process to develop a fast, general purpose adaptive 3D adaptive mesh redistribution algorithm based on PMA [1] which is useable both for the UK Area weather forecast and many other applications. This algorithm is now being extended in a NERC funded collaboration with Reading and Imperial College, to determine adaptive meshes on the sphere.

As described above, resolving the fine structures associated with rapid temperature changes across inversion layers is important for accurate data assimilation calculations and this led to the next application of adaptive mesh methods to these problems.

Coupling r-Adaptive Methods to Data Assimilation Algorithms

Data assimilation matches the sequence of weather forecasts u_i (with $i = 0 \dots N$) from an initial state u_0 to data points y_i which are measurements, with error, of functions $H(u)$ of the true state at different times t_i . One of the main objectives of data assimilation is to find the best, unbiased, estimate of the initial state (called the

analysis), which is consistent with this data. In a typical meteorological forecast the state u_0 has around 10^9 degrees of freedom, and this is compared to about 10^6 data values. This is thus a significantly ill conditioned and underdetermined problem. To regularise the system some a-priori information must be included. This is usually an estimate u_B of the background state at the initial time, usually given by a numerical forecast, again with an associated error. For computational purposes it is usually assumed that the data errors are Gaussian with (known) covariance R and that the background error is also Gaussian with covariance B . A powerful algorithm for determining u_0 which is in constant use at the Met Office in its operational codes, is the variational procedure for which we have

$$u_0 = \operatorname{argmin}(J), \quad \text{where} \quad J = \alpha \|u_0 - u_B\|_B^2 + \beta \sum_{i=0}^N \|y_i - H(u_i)\|_R^2. \quad (4)$$

If $N = 0$ this is the 3d-Var method used for the UKV forecasts, and if $N > 0$ it is the 4d-Var method used for global forecasts. In the 4d-Var formulation, the variational procedure is coupled to the strong constraint given by the NWP procedure for which

$$u_{i+1} = M_i(u_i),$$

with M_i a complex nonlinear function representing the evolution of the weather between sample points.

In either case it is important to have good estimates of the error covariances. Usually the procedure works well, but in the case of inversion layers and other fine structures, these errors can be misrepresented. To address this problem we used the adaptive mesh transformations to rescale the spatial coordinates used in the data assimilation calculation. A key breakthrough in this work was the incorporation of an appropriate monitor function $m(x)$ into the rescaling algorithm. It was found that an effective monitor function was a smoothing of the static stability estimate

$$m(x, z, t) = \sqrt{a^2 + b^2 \theta_z^2} \quad (5)$$

where θ is the *potential temperature* which changes rapidly across the inversion layer, z is the vertical coordinate and x the horizontal coordinate. In the implementation by Piccolo and Cullen described in [9] a two-stage process is used to calculate the mesh, with the first mesh generated using the background state, and the second using a preliminary analysis. The mesh calculated using this preliminary analysis state were the used in the assimilation procedure to estimate the analysis. An example of the results of this is illustrated in Fig. 1, taken from Figs. 2 and 3 of [9]. The rescaling of the vertical coordinates using this monitor function meant that the vertical correlations of the background error covariance matrix in the inversion or ground boundary layers were much better resolved. In particular this has improved the ability of the assimilation system to accurately use high-resolution information like radiosonde soundings. This algorithm has proved especially appropriate, flexible and robust for this proce-

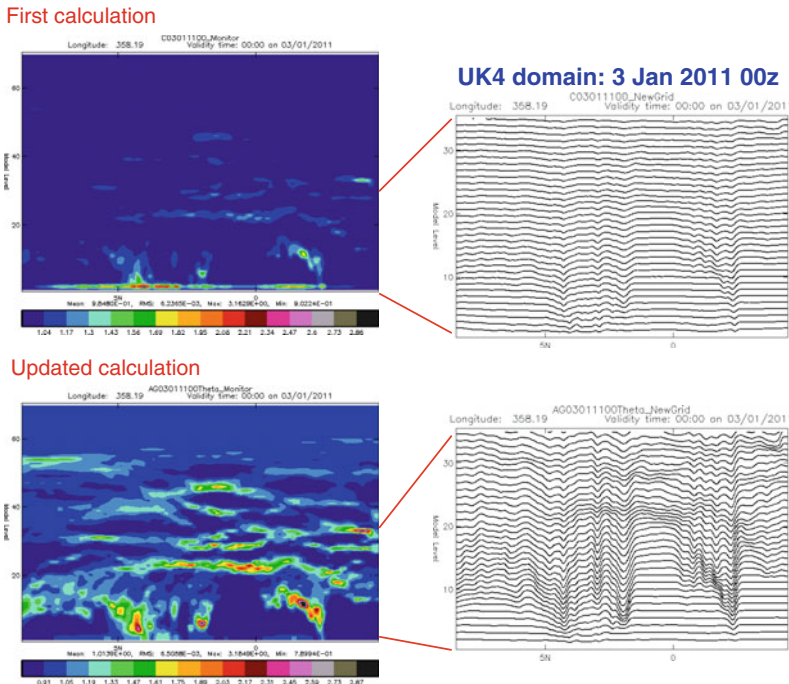


Fig. 1 This figure (replica of Figs. 2 and 3 of [9]) shows the two stages of the adaptive process. In the figure *top left* we see the monitor function for the background state with the adapted mesh *top right*. On the *bottom left* we see the monitor function for a preliminary analysis, on the *bottom right* we see the associated mesh

ture, and has been particularly suitable when dealing with real meteorological data, which can be very noisy.

Results

Adaptive data assimilation software, based directly on the research described above, was first incorporated by Chiara Piccolo into the 3d-Var algorithm used in the operational data assimilation code for the Met Office 4 km grid UK models in November 2010 [8, 9]. Operational codes make forecasts every six hours, and the operational codes, incorporating the adaptive algorithm, have been used to forecast the UK weather for the last five years. A direct consequence of this work is an improvement of the Met Office forecasting skill in terms of the so-called UK Index which is a measure of the forecasting skill of limited-area NWP models over the UK and is based on forecasts of selected parameters and for a selected set of positions verified by comparison with available station observations across the UK at

3–6-hourly intervals. One of the most significant results was an improvement in the error of forecasting the temperatures at 2 m above ground level.

The adaptive mesh transformation led to positive impact in the forecast skill of UK models both in winter and summer. Analysis RMS errors are reduced with respect to radiosonde, aircraft, SEVIRI and ground GPS observations for both periods. Background RMS errors are reduced with respect to aircraft, surface and ground GPS observations for both periods and also with respect to radiosonde observations for relative humidity in the lower part of the troposphere and for potential temperature around the inversions. These results are consistent with the change in the monitor function structures coming from the updated normalization procedure and recalculation of the adaptive mesh within the nonlinear minimization procedure. These led also to improvement of the background state in the full cycled analysis/forecast system and therefore to better representation of the vertical structure of the boundary layer. For these reasons this new version of the adaptive mesh transformation was implemented operationally in the Met Office data assimilation system in July 2011 for UKV and UK4 models, [9].

Societal Impact

Obviously, the improvement of the Met Office forecasting skill has had significant economic and societal impacts. According to Mike Cullen

The new method of adapting computational grids to the expected solution is now being exploited in the high resolution analyses used to drive the short-range forecasts for the UK. Particular benefit has been found in predicting low-level temperatures, which is very important for maintaining the road network in a safe condition and for predicting fog.

Temperature predictions are used, for example, in the Met Office OpenRoad software [10] that is employed to provide a 24h forecast for road state companies and to advise local councils on ice hazards and the need (or not) for road gritting, to help maintain essential road services, mainly in winter. In the winter of 2011/12, the Met Office provided OpenRoad based forecasts for over 350 routes in the UK. The use of OpenRoad reduces the impact of cold weather on road networks, in particular on road safety, and, via more accurate forecasting of road temperatures, leads to a more cost-effective use of grit supplies (gritting can cost a council in the order of £10k to £15k per day). Moreover, since salt is a corrosive substance, avoidance of gritting when it is not necessary, leads to savings for road users in general and to a reduction of damage to the transport infrastructure in particular.

References

1. Browne, P.A., Budd, C.J., Piccolo, C., Cullen, M.: Fast three dimensional r-adaptive mesh redistribution. *J. Comput. Phys.* **123**, (2014)
2. Budd, C.J., Huang, W., Russell, R.D.: Moving mesh methods for problems with blow-up. *SIAM J. Sci. Comput.* **17**, 305–327 (1996)

3. Budd, C.J., Williams, J.F.: Parabolic Monge-Ampere methods for blow-up problems in several spatial dimensions. *J. Phys. A* **39**, 5425–5463 (2006)
4. Budd, C.J., Williams, J.F.: Moving mesh generation using the parabolic Monge-Ampere equation. *SIAM J. Sci. Comput.* **31**, 3438–3465 (2009)
5. Budd, C.J., Huang, W.-Z., Russell, R.: Adaptivity with moving grids. *Acta Numerica* **18**, 1–131 (2009)
6. Budd, C.J., Cullen, M., Walsh, E.: Monge-Ampere based moving mesh methods for numerical weather prediction, with applications to the Eady problem. *J. Comp. Phys* **236**, 247–270 (2013)
7. Lorenc, A.C., Ballard, S.P., Bell, R.S., Ingleby, N.B., Andrews, P.L.F., Barker, D.M., Bray, J.R., Clayton, A.M., Dalby, T., Li, D., Payne, T.J., Saunders, F.W.: The met: office global three-dimensional variational data assimilation scheme. *Q. J. R. Meteorol. Soc.* **126**, 2991–3012 (2000)
8. Piccolo, C., Cullen, M.: Adaptive mesh method in the met office variational data assimilation system. *Q. J. R. Meteorol. Soc.* **137**, 631–640 (2011)
9. Piccolo, C., Cullen, M.: A new implementation of the adaptive mesh transform in the met office 3D-var system. *Q. J. R. Meteorol. Soc.* **138**, 1560–1570 (2012)
10. The OpenRoad Software <http://www.metoffice.gov.uk/roads/openroad>



<http://www.springer.com/978-3-319-25452-4>

UK Success Stories in Industrial Mathematics
Aston, P.J.; Mulholland, A.J.; Tant, K.M.M. (Eds.)
2016, XIV, 303 p. 81 illus. in color., Hardcover
ISBN: 978-3-319-25452-4